TAMPEREEN TEKNILLINEN YLIOPISTO
TAMPERE UNIVERSITY OF TECHNOLOGY

Eero Lihavainen

# Image Analysis Methods for the Characterization of Mitochondrial Morphology and Dynamics

Julkaisu 1410 • Publication 1410

Tampere 2016

Tampereen teknillinen yliopisto. Julkaisu 1410
Tampere University of Technology. Publication 1410

Eero Lihavainen

# Image Analysis Methods for the Characterization of Mitochondrial Morphology and Dynamics

Thesis for the degree of Doctor of Science in Technology to be presented with due permission for public examination and criticism in Tietotalo Building, Auditorium TB109, at Tampere University of Technology, on the 30th of September 2016, at 12 noon.

Tampereen teknillinen yliopisto - Tampere University of Technology
Tampere 2016

**Supervisor:** Associate Professor Andre S. Ribeiro
Tampere University of Technology
Finland

**Pre-examiners:** Professor Ivo F. Sbalzarini
Technische Universität Dresden
Germany

Senior Lecturer Boguslaw Obara
Durham University
United Kingdom

**Opponent:** Professor Carolina Wählby
Uppsala University
Sweden

ISBN 978-952-15-3799-8 (printed)
ISBN 978-952-15-3810-0 (PDF)
ISSN 1459-2045

# Abstract

Microscopy is the primary tool for analyzing dynamical processes in live cells. With modern fluorescence microscopy techniques, it is possible to image organelles and other subcellular structures at high frame rates in single cells, and even inside living animals. One organelle of particular interest is the mitochondrion, in which abnormalities in dynamics and morphology have been associated with many diseases. Because of this, imaging methods have been, and will continue to be, indispensable for developing an understanding of the morphology and dynamics of mitochondria, and how deficiencies in them can contribute to diseases.

In microscopy studies of mitochondria, it is still common for researchers to rely on visual inspection, as opposed to thorough computational image analysis, for drawing conclusions from image data. This not only limits the amount of data that can be analyzed in a reasonable time, but also introduces human error to the study due to subjective judgement. In the case of mitochondria, the present lack of automatic analysis methods is partly explained by the complex dynamics and morphology of this organelle, which can cause mitochondria to appear vastly different in different environments. Due to this variation in appearance, methods for segmenting and tracking small intracellular particles are often not directly applicable to analyzing mitochondrial images, and instead, special methods may need to be crafted for each application. This thesis is an attempt to facilitate studies on mitochondria by presenting new tools and methods specifically developed for the quantitative analysis of mitochondrial morphology and dynamics from fluorescence microscope images.

This thesis has three main outcomes. First, we developed a software tool, Mytoe, for quantifying morphological features of mitochondria and estimating their motion from time-lapse fluorescence microscopy. Second, we developed a method for detecting the tips of mitochondria, and demonstrated how these tips can be tracked using a general particle tracking framework. Finally, we propose a novel method for quantifying mitochondrial fragmentation from two-photon microscope images of brain tissue where mitochondria have been fluorescently labeled. We expect these contributions to help provide insights about mitochondrial dynamics and structure in both single-cell imaging and animal disease models.

# Preface

This study was carried out at the Laboratory of Biosystem Dynamics in the Department of Signal Processing, Faculty of Computing and Electrical Engineering, Tampere University of Technology under the supervision of Associate Professor Andre S. Ribeiro.

First and foremost, I would like to express my sincere gratitude to my supervisor Andre Ribeiro for his support, guidance and perseverance during my studies.

Next, I would like to thank all of my co-authors for making this work possible. In particular, I want to thank Jarno Mäkelä for spending evenings with the microscope, to produce the data that much of this thesis has relied on. I am also grateful to Leonard Khiroug and his research group, especially Mikhail Kislin, for the fruitful collaboration that we have had.

I also wish to thank all the present and past members of the Laboratory of Biosystem Dynamics that I have had the pleasure to work with. Besides the science, it has been the discussions with you on different cultures, philosophy of science and video games that have made me grow the most as a person during these years.

I would also like to thank Emil Aaltonen foundation and the Vilho, Yrjö and Kalle Väisälä Foundation for financial support during my doctoral studies.

I am grateful to the pre-examiners Boguslaw Obara and Ivo Sbalzarini, for their insightful suggestions which have helped improve the thesis.

Finally, I am deeply grateful to my family, for all their support during the years. And of course Elina, who supported me as much as she could while writing her own thesis.

Tampere, July 2016

Eero Lihavainen

# Contents

# Acronyms

| | |
|---|---|
| AC | Active contour |
| ATP | Adenosine triphospate |
| BM3D | Block-matching and 3-D filtering |
| DWT | Discrete wavelet transform |
| DT-CWT | Dual-tree complex wavelet transform |
| EPE | Expected prediction error |
| GLCM | Graylevel co-occurrence matrix |
| GLM | Generalized linear model |
| LAP | Linear assignment problem |
| MHT | Multiple hypothesis tracking |
| mtDNA | Mitochondrial DNA |
| NA | Numerical aperture |
| OF | Optical flow |
| PSF | Point-spread function |
| ROS | Reactive oxygen species |
| SNR | Signal-to-noise ratio |
| SOAC | Stretching open active contour |
| SVM | Support vector machine |

# List of Publications

This thesis is a compound of the following publications. These are referred to in the text as Publication I, Publication II, and Publication III.

I Eero Lihavainen, Jarno Mäkelä, Johannes N. Spelbrink, Andre S. Ribeiro, "Mytoe: Automatic analysis of mitochondrial dynamics," *Bioinformatics*, vol 28, no. 7, pp. 1050 – 1051 Feb. 2012.

II Eero Lihavainen, Jarno Mäkelä, Johannes N. Spelbrink, Andre S. Ribeiro, "Detecting and tracking the tips of fluorescently labeled mitochondria in U2OS cells," *Lecture Notes in Computer Science*, 9280:363 – 372, Aug. 2015

III Eero Lihavainen, Mikhail Kislin, Dmytro Toptunov, Leonard Khiroug, Andre S. Ribeiro, "Automatic quantification of mitochondrial fragmentation from two-photon microscope images of mouse brain tissue," *Journal of Microscopy*, vol 260, no. 3, pp. 338 – 351 Dec. 2015.

In all publications, the author of this thesis designed and implemented all image analysis methods, performed their evaluation, and was the main responsible in writing the manuscripts. Publications I and II were conceived by Andre S. Ribeiro and the author of this thesis jointly. Publication III was conceived by the author of this thesis, following collaboration efforts between our group and Leonard Khiroug's group.

Publication III has been used by Mikhail Kislin in his PhD dissertation.

# 1 Introduction

The research field of *bioimage informatics* has emerged in recent decades from the need of automated, quantitative analysis of image data in the biomedical sciences. This need has itself emerged from important advances in the field of microscopy, where not only the quality of imaging has improved, but also multiple novel methods have emerged. In addition, these instruments have become more affordable, which has made multi-modal cell microscopy a standard method for research groups working in the field of Cell Biology. Furthermore, this was also supported by recent findings on fluorescent proteins and means to use them to gain information on a multitude of cellular processes, not possible otherwise.

To study such cellular processes and, in particular, to obtain quantitative analyzes, many cells need to be observed and thus many images are required. This is made possible by modern high-throughput imaging techniques. These techniques further enhance the need for bioimage informatics, since they make manual handling and analysis of data impractical, requiring instead efficient algorithms for automatic image analysis.

The development of novel imaging and cell biology techniques results in new image analysis problems, as the methods developed to handle older images in general are not directly applicable to new ones. Furthermore, in a cell imaging study, the appropriate choice of image analysis methods depends on factors such as the research question, the experimental setup and microscopy technique used, and the type and physiological state of the cells. Novel studies often require the adaptation of previously used image analysis methods to the problem at hand, or in some cases the development of new methods.

One particular object of study for which novel image analysis methods are currently required is the mitochondrion. Mitochondria are highly dynamic organelles that are involved in vital cellular processes such as energy production, and whose dysfunctions are related to various diseases. Because of that, they have been the topic of many recent studies in cell biology. From the image analysis perspective, they are challenging, in that they exhibit widely varying morphologies, from small particles to networked structures; this morphological variation is achieved through the processes of fusion and fission. Due to the unique and dynamic nature of the structures formed by mitochondria, existing methods for the analysis of

subcellular structures (Ruusuvuori et al., 2010) are often not directly applicable to mitochondrial image analysis. Indeed, several studies have recently proposed methods for problems such as detecting and tracking mitochondria from single-cell fluorescence microscope images.

In addition to single-cell imaging, thanks to recent methodological advances, mitochondria can be imaged inside tissues of living animals; for example, one can image mitochondria in the neocortex of live mice, which has the potential to result in important findings about dysfunctions of mitochondria in neurons and their consequences. This method, and similar ones, produce images qualitatively different from single-cell imaging, where cells are neatly separated on a microscope slide. Consequently, it requires the image analysis methods to be specifically tailored for the problem. For the particular problem of quantifying the fragmentation of mitochondria, i.e. the breakdown of the mitochondrial network due to an imbalance in rates of fusion and fission, some methods have been proposed prior to this thesis (Chen et al., 2011; Peng et al., 2011), but to the best of our knowledge, no method has provided an absolute quantification that requires no reference images.

## 1.1 Objectives of the thesis

We set out to develop tools and methods for automatically quantifying mitochondrial morphology and dynamics from fluorescence microscopy images, with potential applications in cell biology, drug screening and neuroscience. In particular, we focused on confocal microscope images of single cells and on two-photon microscope images of mitochondria in the mouse neocortex. For single-cell imaging, we focused on time-lapse images of a human osteosarcoma cell line (U2OS), which exhibits a sufficiently flat morphology so that the analysis can be performed from two-dimensional images. We also wanted to assess the applicability of machine learning methods to mitochondrial image analysis, in particular for the detection of the tips of individual mitochondria, to be used as targets for tracking, and for quantifying mitochondrial fragmentation in tissue.

In summary, this thesis has two main objectives:

1. To develop methods and software tools for analyzing the morphology and dynamics of mitochondria from time-lapse fluorescence microscopy images of single cells. This involves developing methods for detecting mitochondria and quantifying their shape, as well as estimating their motion. This objective was completed and then published in Publications I and II.

2. To develop a method for quantifying the degree of mitochondrial fragmentation in two-photon microscope images of cortical tissue. Due to the qualitative differences of these images to single-cell images, this requires

methods that are separate from those achieved in Objective 1. This objective was completed and then published in Publication III.

## 1.2 Thesis outline

This thesis is organized as follows: In Chapter 2 we review the biology of mitochondria and discuss their functions and dysfunctions that motivate the present thesis. We also discuss current techniques that are used to image mitochondria in live cells and live animals. Chapter 3 provides a theoretical background on statistical learning, which has an important role in the bioimage analysis methods proposed in the thesis. In Chapter 4 we provide an overview of the image analysis methods that are relevant to our studies of mitochondria. Next, a summary of the results obtained in the publications is presented in Chapter 5. Finally, we present our conclusions and discuss the outcomes of the thesis in Chapter 6.

# 2 Biological background

## 2.1 Mitochondria

### 2.1.1 Function and dynamics

Mitochondria are organelles that are found in all eukaryotic organisms. They are best known for their role in animal cells as the providers of energy in the form of adenosine triphosphate (ATP). In addition to this, mitochondria perform other important functions in metabolism: for example, in conditions of starvation, they are responsible for oxidizing amino acids for ATP production. Meanwhile, in conditions of excess, they participate in biosynthesis processes by accelerating the production of fatty acids and sterols, which are needed for building cell membranes (Alberts et al., 2008). Mitochondria also participate actively in apoptosis, or programmed cell death, by releasing cytochrome C, which activates the self-destruction machinery of the cell (Alberts et al., 2008).

The morphology of mitochondria varies from small particles to filamentous, networked structures (see Figure 2.1). This variation in morphology is possible because mitochondria are capable of fusing and dividing. The processes of fusion and fission allow for restructuring the mitochondrial network for optimal energy dissipation (see Westermann (2010)). Interestingly, during apoptosis, mitochondria become fragmented by rapid fission. There is no evidence that this change in mitochondrial morphology would be a cause of apoptosis, but it appears that some components of the mitochondrial fusion and fission machinery are involved in apoptosis (Suen et al., 2008).

Recent research has unveiled a function for the constant mitochondrial fusion-fission dynamics. Namely, it appears to be part of a quality control mechanism where mitochondria use fission to exclude parts of the network that contain damaged mitochondrial DNA (mtDNA). According to recently published works (see Fischer et al. (2012) for a review), this quality control works by inhibiting fusion for mitochondria with low membrane potential, which is a sign of damage. The damaged mitochondria are then degraded by the process of mitophagy. The role of fusion in mitochondrial quality control, on the other hand, is that of enabling the mixing of contents inside the mitochondrial network, thus diluting damage-inducing molecules.

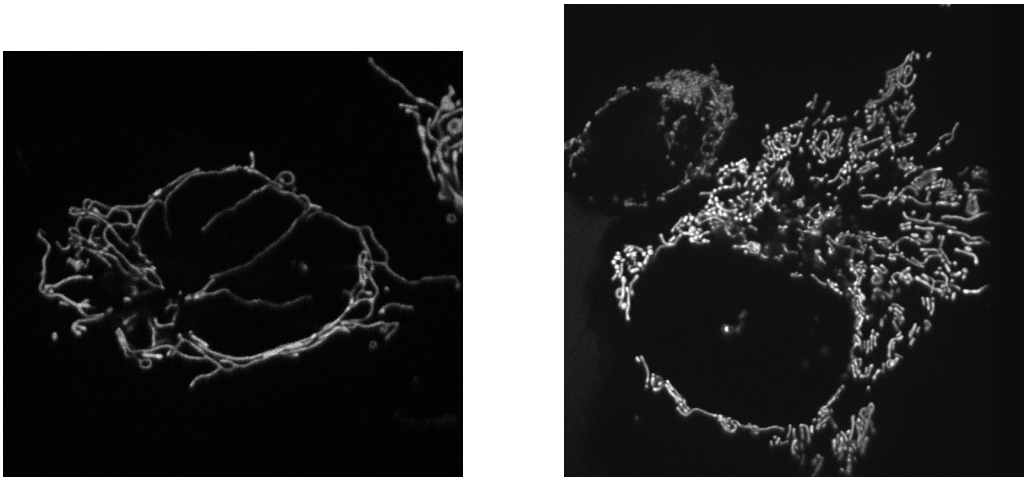

**Figure 2.1:** Mitochondria can exhibit varying morphologies. The image on the left shows a cell with elongated and networked mitochondria, whereas on the right, the mitochondria are much shorter and disconnected.

In many cell types, mitochondria are spread out uniformly inside the cytoplasm (Rafelski, 2013) and are relatively static (Saunter et al., 2009). However, in some situations, they need to be actively transported to a location in the cell requiring high energy consumption; this occurs in neurons, in which mitochondria are actively transported along the axons and dendrites to the peripheral regions of the cell. For this, there exists specialized transport machinery that has been extensively studied in recent years (Sheng and Cai, 2012). In brief, mitochondria attach to motor proteins that provide either anterograde (towards the periphery) or retrograde (back towards the cell body) transport, along either microtubules or actin filaments.

### 2.1.2 Mitochondria in disease and ageing

Given their importance in cell metabolism, it is not surprising that the malfunctioning of mitochondria is associated to various diseases. *Mitochondrial diseases* constitute a group of diseases that are caused by failures in mitochondrial functioning, more specifically in the complexes responsible for oxidative phosphorylation. These failures can be due to mutations in mtDNA or nuclear DNA (Ylikallio and Suomalainen, 2012). Symptoms of mitochondrial diseases vary based on the complexes that are affected (Ylikallio and Suomalainen, 2012); they include myopathy, lactic acidosis and neurologic symptoms such as mental regression, among others (Debray et al., 2008).

Dysfunctions of mitochondria are also associated with neurodegenerative diseases, such as Parkinson's, Alzheimer's, Huntington's and Charcot-Marie Tooth diseases. In neurons, it is vital that mitochondria can travel from the cell body to the dendritic and axonal termini. Disruption of such transport leads to neuronal

dysfunctions, and can have several causes (see Chen and Chan (2009)): (i) impaired mitochondrial fission, causing crowding of mitochondria; (ii) fragmentation of mitochondria due to impaired fusion, which is associated with a deficiency in transport ability; or (iii) defects in the mitochondrial transport mechanisms. In addition to problems in transport, impaired mitophagy of damaged mitochondria will lead to the accumulation of dysfunctional mitochondria, which may impair neuronal functionality.

It has long been hypothesized that accumulated damage in mtDNA could be a major contributing factor in the ageing of organisms (see e.g. Alberts et al. (2008)). The so-called free radical theory of aging hypothesized that reactive oxygen species (ROS), which are mutation-causing side products of mitochondrial functioning, accelerate the aging of tissue by damaging the mtDNA, causing in turn more ROS to be produced. However, even though excessive damage in mtDNA has been shown to cause premature ageing in mice (Trifunovic et al., 2004), it is not yet clear whether the natural mtDNA mutation rate in humans is large enough to cause ageing (Payne and Chinnery, 2015).

Although the free radical theory of ageing is unsupported by evidence, some studies have suggested that damaged mtDNA in stem cells leads to defects in somatic stem cell homeostasis (Ahlqvist et al., 2015). This, in turn, could lead to premature ageing. A recent study suggested that stem cells may have a mechanism for rejuvenating themselves by segregating older – and thus potentially more damaged – mitochondria to the differentiating daughter cell in division (Katajisto et al., 2015). Furthermore, the authors showed that this segregation is required for the stem-like cells to keep their stemness. Given these results, an important direction for future research is to establish the role of mitochondrial damage in stem cells in ageing.

## 2.2 Imaging mitochondria in live cells

### 2.2.1 Fluorescence microscopy

Imaging individual organelles in live cells can be done via fluorescence microscopy. A fluorescence microscope works by illuminating the specimen with light (the excitation) at a specific wavelength, which will cause fluorescent molecules, or *fluorophores*, to emit light at a different wavelength. The simplest type of fluorescence microscope is the *epi-fluorescence* microscope. It relies on *wide-field* illumination, where the whole specimen is illuminated at once using filtered light from a lamp, and the emitted light is filtered and then collected using a camera (Stephens, 2003). This is illustrated in Figure 2.2.

In fluorescence microscopy, the specimen is visualized with a fluorescent *label*, a fluorophore that localizes in the region to be labeled, such as to an organelle, the DNA or the cytoplasm. Labels are also referred to as *stains* or *dyes*. In

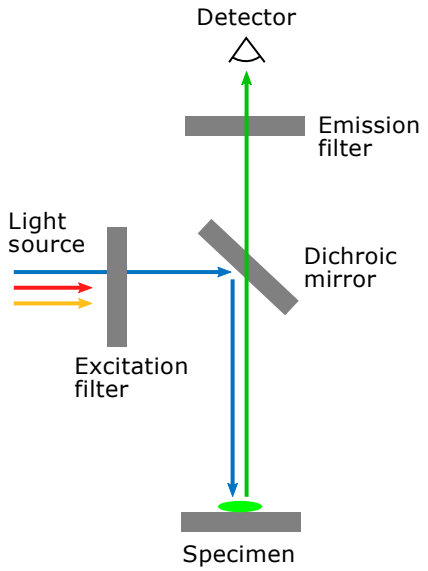


**Figure 2.2:** A schematic view of a widefield microscope. The illumination light is filtered using the excitation filters, so as to match the excitation spectrum of the target fluorophores. The light is then directed to the whole specimen using a dichroic mirror. The emitted light is then filtered to remove unwanted wavelengths, and finally collected by a camera.

live-cell imaging, particularly over large time scales, it is often required that the fluorophore not hinder cellular functions. For example, DAPI is a stain that binds to the DNA, and is commonly used for visualizing the nucleus. The drawback of such a stain is that it interferes with vital cellular processes, such as transcription, thus potentially affecting experimental results. Fluorescent proteins on the other hand, are ideal in this regard; they were originally extracted from animals such as jellyfish, which express them naturally, and they do not interfere with the growth of the host cells (Andersen et al., 1998).

In the case of mitochondria, the label is usually a fusion of a protein that localizes in the mitochondrial matrix or in the intermembrane space, and of a fluorescent protein. One such label, which we have used in Publications I and II, is mitoDsRED2, a fusion of a red fluorescent protein DsRed2 and a mitochondrial-targeting component of the human cytochrome c oxidase. To study mitochondria in live animals, the animals can be genetically modified to express fluorescent label proteins. For example, a transgenic mouse line that expresses mitochondrially targeted fluorescent proteins selectively in neurons was introduced in (Misgeld et al., 2007).

Exposure to light causes changes in cell homeostasis; this damage is referred to as *photodamage*. It is believed to mainly result from light interacting with fluorescent molecules, causing them to become chemically reactive, which leads to the generation of ROS in cells (Frigault et al., 2009). Well-known effects of

photodamage are delayed mitosis, or a complete arrest of the cell cycle. These can be used to control photodamage effects in the experimental setup. However, these effects can vary between cell types, as well as between individual cells. Due to this, photodamage can, in a sense, amplify the cell-to-cell variability, thus skewing the experimental results. Furthermore, cells can become more sensitive to photodamage when affected by other stressors; this needs to be taken into account particularly when studying cells under stress conditions (Magidson and Khodjakov, 2013).

When imaging organelles inside cells, epi-fluorescence microscopy has suboptimal resolution. Since all the light emitted from the specimen is collected, much of the light from outside the focal plane is present in an image. This is problematic especially if the specimen extends in the axial dimension. Such out-of-focus light can be either from the targeted fluorophores, or from fluorophores native to the cell that cause *autofluorescence*. Because of this, epi-fluorescence imaging is often supplemented with deconvolution, which attempts to increase the resolution by mathematically inverting the blurring described by the point spread function (PSF) (Shaw, 2006).

*Confocal* microscopy is a more sophisticated technique that achieves better resolution when compared to epi-fluorescence. It illuminates the specimen by focusing a laser beam to one small volume at a time. Superior resolution is achieved by passing the emission signal through a pinhole aperture before reaching the detector. This has the effect of discarding much of the light from outside of the illumination volume. Due to this, confocal microscopy is effective at *optical sectioning*, i.e. at imaging three-dimensional objects as a series of thin sections at different focal planes (Stephens, 2003). A schematic illustrating the principle of confocal microscopy is shown in Figure 2.3.

While confocal microscopy is effective for single-cell imaging, it has some drawbacks when imaging tissues. First, the depth that can be reached by imaging is limited to approximately 40 $\mu m$, because of the scattering of light from the tissue (Norman, 2005). Second, since the excitation light passes through the tissue, the amount of photodamage is pronounced. Because of this, a technique more suitable for imaging mitochondria in living animals is *multiphoton* microscopy. The operating principle is that multiple light pulses are used to excite the fluorophore simultaneously, which will cause emission at a wavelength higher than that of the individual pulses. This allows for imaging deeper inside the tissue, because there is less scattering of light. Additionally, the excitation pulses can be of much lower energy when compared to confocal microscopy, and thus it will cause less photodamage in the out-of-focus regions. Because of this, in Publication III, we used two-photon microscopy to image the neocortex of mice.

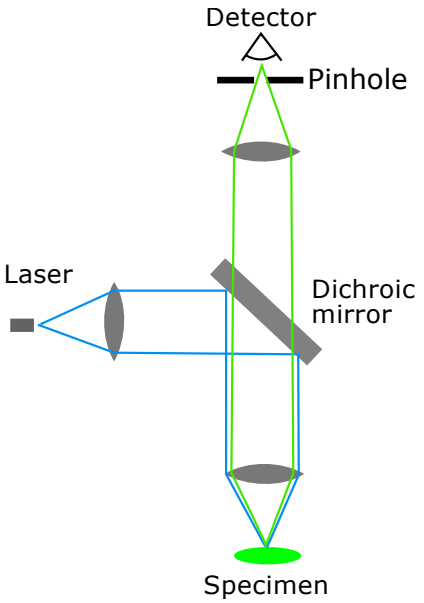


**Figure 2.3:** A schematic of a laser scanning confocal microscope. The illuminating laser beam is directed to the specimen and focused to one small volume at a time. The excitation light is directed to the detector, typically a photomultiplier tube, through a pinhole.

### 2.2.2 The tradeoff between image quality and cell health

Poor image quality will increase the errors made by any image analysis method. Thus, and given the relatively invasive nature of fluorescence microscopy, an important consideration when designing an experiment is to satisfy the image quality requirements of the analysis methods, while keeping minimal the damage caused to the cells.

There are several factors that affect the visual quality of a fluorescence microscope image. The most important are resolution, which is limited by the physical properties of the imaging system, and noise, i.e. the random fluctuations in recorded intensity. In addition to these, the uniformity of the distribution of the fluorescent label and the possible autofluorescence level of the cell can also affect the results.

The emission of photons is an inherently noisy process. This photon noise can be counteracted by increasing the amount of fluorophores or by using more light to excite the cells, with the obvious drawbacks mentioned above. In addition to this, some of the noise is generated in the detector. Here, the main noise sources are the conversion of photons to electrons, which also results in photon noise, the readout, i.e. the conversion of the resulting charge to a voltage, and thermal noise, i.e. the generation of electrons due to thermal vibration. The latter can be minimized by cooling down the detector (Vliet et al., 1998). In practice, in

fluorescence microscopy, the two sources of photon noise affect the most the image quality. The experimentalist will be able to control such amount of noise to an extent, by varying the amount of light used to illuminate the specimen, which can be done either by adjusting the light intensity or by tuning the exposure time.

The maximum resolution of a fluorescence microscope depends on the numerical aperture (NA) of the lens, which in turn depends on the refractive index $n$ of the medium between the specimen and the objective front lens, and on the half-angular aperture $\alpha$ of the objective, via the relation $\text{NA} = nsin(\alpha)$. The highest numerical apertures in commonly used microscopy objectives are around 1.35; these are oil-immersion objectives, which benefit from a higher refractive index compared to dry objectives where the medium is air.

Denoting the emission wavelength as $\lambda$, the resolving power of such a lens is the diffraction limit $d = 0.5\lambda/\text{NA}$. That is, two point sources closer than $d$ to each other will appear as a single PSF -shaped blob in the image. As an example, for a lens with $\text{NA} = 1.35$ and a red fluorescent protein with an emission wavelength of 700 $nm$, the imaging system can theoretically resolve two objects at 0.26 $\mu m$ apart. This is approximately half the width of a mitochondrion (see e.g. Westermann (2010)). It should be noted that this is not a limit on how precisely objects can be localized; for example, the location of an isolated fluorophore can be estimated with sub-pixel precision as the centroid of a Gaussian function fitted to the image of the molecule (see e.g. Mortensen et al. (2010)).

The diffraction limit can be overcome by super-resolution (SR) microscopy techniques (see e.g. Fornasiero and Opazo (2015)). These can resolve objects a few nanometers apart, making them effective at visualizing subcellular structures such as the actin cytoskeleton (Fornasiero and Opazo, 2015). So far, SR methods have had limited application in live cell imaging. A major obstacle has been that they require high illumination intensities, making it difficult to avoid the effects from photodamage from being to intrusive to the measurement results (Fornasiero and Opazo, 2015). However, there is some evidence that by using light from the red spectrum, along with other optimizations to the experimental design, one can attenuate photodamage considerably (Wäldchen et al., 2015).

Given that the main parameters used to control the image quality are those that determine the amount of photodamage taken by the cell, it is important to devise image analysis methods that are robust to noise and, simultaneously, are adapted to the objects under study.

# 3 Supervised machine learning

Statistical learning is at the heart of most data analysis. Methods from statistical estimation theory and pattern recognition have been important parts of image processing and analysis for decades (Geman and Geman, 1984; Otsu, 1979), and recent developments in artificial neural networks have resulted in learning methods outperforming previous methods in problems such as object detection (Krizhevsky et al., 2012) and cell segmentation in microscopy (Ronneberger et al., 2015) by a large margin. In this chapter, we provide a brief overview of statistical learning theory, and discuss supervised learning methods that are relevant in bioimage analysis. We do not attempt to provide a comprehensive review of the vast array of methods that have been successful, but limit the discussion to the methods used in this thesis.

## 3.1 Theory of machine learning

The problem of *supervised learning* can be formulated as a function approximation problem: given some vector of explanatory variables, or features, $\mathbf{x}$, and a response variable $y$, with an assumed functional relationship between them, we try to approximate a function $f$ such that $f(\mathbf{x}) = y$. This is done by generalizing from a set of examples referred to as the *training data* $\mathcal{D}_{train} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$. Having found an approximation $\hat{f}$ for $f$, we can predict the response $\hat{y}$ for a previously unseen $\mathbf{x}$ as $\hat{y} = \hat{f}(\mathbf{x})$. If $y$ is a qualitative variable, i.e. it represents discrete classes, we call the problem *classification*, and the resulting model is a *classifier*. If $y$ is quantitative, the problem is called *regression*. We note that supervised learning is conceptually different from *unsupervised learning*, which attempts to find structure in a data set without any response variable. The most common class of unsupervised learning methods is clustering, in which a data set is partitioned into $K$ subsets, where the members in each cluster are similar to each other in some statistical sense (Hastie et al., 2009).

In supervised learning, the goodness of the approximation is measured with a *loss function* $L$, which is most commonly the squared error

$$L(y, \hat{f}(\mathbf{x})) = (y - \hat{f}(\mathbf{x}))^2 \tag{3.1}$$

for regression problems, and the zero-one loss

$$L(y, \hat{f}(\mathbf{x})) = \begin{cases} 0, & \text{if } y = \hat{f}(\mathbf{x}) \\ 1, & \text{otherwise} \end{cases} \tag{3.2}$$

in the case of classification. The problem of learning can be stated as finding $\hat{f}$ such that the expected loss, or *expected prediction error*

$$\text{EPE}(\mathbf{x}) = \mathbb{E}(L(y, \hat{f}(\mathbf{x}))) \tag{3.3}$$

is minimized. In practice, the EPE has to be approximated by averaging over the training samples. Given this framework, learning methods can be seen to differ in the form of the model assumed for $f$, the choice of loss function, and possible regularization methods incorporated in the minimization of the expected error, that encode prior assumptions about the structure of the solutions (Hastie et al., 2009).

In the case of the squared-error loss, the error of Eq. 3.3 can be written as

$$\text{EPE(x)} = (\mathbb{E}(\hat{f}(\mathbf{x})) - y)^2 + \mathbb{E}\left[(\hat{f}(\mathbf{x})) - \mathbb{E}(\hat{f}(\mathbf{x}))^2)\right]. \tag{3.4}$$

In this expression, the first summand is the squared bias of the prediction, and the second expression is its variance. The bias is the difference between the true response and the expected estimate, and the variance measures the dispersion of the prediction about its mean. In practice, there is a tradeoff between these components of the error: increasing a model's complexity decreases the bias but increases the variance, and vice versa (Hastie et al., 2009).

An important aspect of statistical learning theory is the generalization ability of a learning algorithm. For example, it is important to know whether a learning method minimizes the EPE given enough samples, and how fast the generalization improves with increasing sample size. These questions are answered by the statistical learning theory of Vapnik and Chervonenkis, which has been the basis for developing successful learning methods such as the support vector machine (SVM); the details of the theory are outside the scope of this thesis, but we refer the reader to Vapnik's introduction to the theory in (Vapnik, 1999).

## 3.2 Supervised learning methods

### 3.2.1 Linear models

*Linear models* constitute a simple yet powerful class of models for $f$. They model the response as a linear combination of the features, i.e.

$$f(x_1, x_2, \ldots, x_p) = \beta_0 + \sum_{i=1}^{p} \beta_i x_i \tag{3.5}$$

where $\beta_i$ are the model parameters. Using the more convenient notations $\beta = [\beta_0, \beta_1, \ldots, \beta_p]^T$ and $\mathbf{x} = [1, x_1, \ldots, x_p]^T$, this can be written

$$f(\mathbf{x}) = \mathbf{x}^T\beta. \tag{3.6}$$

If we choose the squared error loss, a solution approximating the minimization of EPE is given by

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \frac{1}{N}\sum_{i=1}^{N}(y_i - \mathbf{x}_i^T\beta)^2, \tag{3.7}$$

which is also known as linear regression (Hastie et al., 2009).

It is useful to view the linear model from a statistical modeling perspective: If one assumes that $y$ has a normal distribution with variance $\sigma^2$, the model can be stated as

$$y \sim \mathcal{N}(\mathbf{x}^T\beta, \sigma^2), \tag{3.8}$$

or equivalently as

$$y = \mathbf{x}^T\beta + \epsilon, \tag{3.9}$$

where the *error* $\epsilon$ is a zero-mean normal random variable with variance $\sigma^2$. This latter formulation highlights the interpretation that the observed values for $y$ are noisy measurements of an underlying true value, and that we are modeling the error. In this model, the least-squares estimate of Eq. 3.7 corresponds to the maximum likelihood estimate (Hastie et al., 2009).

Many of the most commonly used classifiers are linear; while not linear models in the sense of Eq. 3.5, they model the response as depending on the features through a linear combination. More specifically, a linear classifier is of the form

$$f(\mathbf{x}) = h(\mathbf{x}^T\beta + \beta_0), \tag{3.10}$$

where $h$ is a function that maps its argument to one of the possible classes. Note that for classifiers we use here the notation $\mathbf{x} = [x_1, x_2, \ldots, x_p]$ and $\beta = [\beta_1, \beta_2, \ldots, \beta_p]$. In the case of a binary classifier, i.e. where there are two possible classes that can be encoded as 0 and 1, $\beta$ defines a hyperplane - the *decision boundary* - that divides the feature space into two subspaces corresponding to the different classes; this is illustrated in Figure 3.1. A simple linear classifier is obtained from the linear model of 3.5 by setting

$$h(x) = \begin{cases} 0, & \text{if } x > 0 \\ 1, & \text{if } x \leq 0 \end{cases}. \tag{3.11}$$

This is the classification model in methods such as the perceptron (Rosenblatt, 1958), linear discriminant analysis (Fisher, 1936) and the support vector machine (SVM) (Cortes and Vapnik, 1995); the differences between these classifiers lie in how the location of the decision boundary is set. For example, the SVM is a maximum-margin classifier, that is, if the classes are linearly separable, it selects the hyperplane that is maximally and equally far from the closest training sample of each class.

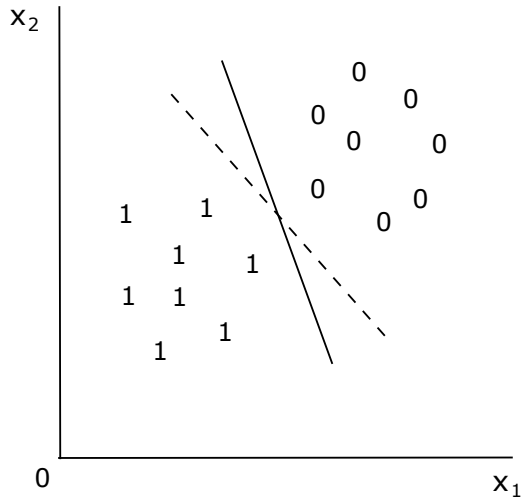


**Figure 3.1:** An example training data set with linearly separable classes. Zeros and ones correspond to different classes, and the two lines are possible decision boundaries that classify the training data correctly.

### 3.2.2 Generalized linear models

In some applications, $y$ takes values in a subset $S$ of $\mathbb{R}$, and using a linear model would yield predictions outside $S$; for example, the linear model can predict negative responses, which is nonsensical if the measurements of $y$ correspond to, e.g., concentrations of molecules. *Generalized linear models* (GLMs) (J. A. Nelder, 1972) maintain the linear relationship between the coefficients and the features, while allowing the modeling of the errors with a more general family of distributions. This is achieved through using a *link function* $g : S \mapsto \mathbb{R}$, which is defined such that $\mathbb{E}(y) = g^{-1}(\mathbf{x}^T\beta)$ holds. The choice of the link function depends thus on the error distribution. Although the principle can be applied to other distributions, GLMs are specifically defined for the exponential family of distributions; in particular, the linear model of Eq. 3.8 is a GLM with the indentity function as $g$. Another example of a GLM is logistic regression: the response is assumed to be Bernoulli distributed, and the model prediction can be interpreted as the probability of the response having value 1. Thus, it is in fact a linear classifier. For logistic regression, the *logit* link $g(p) = ln(p/(1-p))$ is the most common choice, although other choices exist, such as the probit link, which is the inverse cumulative density function of the standard normal distribution (Hastie et al., 2009).

Beta regression (Ferrari and Cribari-Neto, 2004) is a GLM-like model that is useful when the response takes values in the open interval $(0, 1)$. The response is modeled by a Beta distribution, whose probability density function (PDF) is defined as

$$\pi(y; p, q) = \frac{\Gamma(p+q)}{\Gamma(p)\Gamma(q)} y^{p-1}(1-y)^{q-1} \tag{3.12}$$

where $0 < y < 1$, $p, q > 0$ and $\Gamma$ is the gamma function. Neither of the parameters $p$ and $q$ correspond to the expected value of the distribution, thus there is no obvious GLM-like formulation. However, it turns out that the distribution can be reparametrized to have a mean parameter $\mu$ and a dispersion parameter $\phi$, by setting $\mu = p/(p+q)$ and $\phi = p+q$. Now, for a Beta distributed $y$, we have $\mathbb{E}(y) = \mu$, and $\text{var}(y) = \mu(1-\mu)/(1+\phi)$; $\phi$ measures dispersion in the sense that increasing $\phi$ decreases the variance. Figure 3.2 shows the shape of the Beta distribution for different values of $\mu$ and $\phi$. Through this new parametrization we can express $\mu$ in terms of $\mathbf{x}^T\beta$ via a link function, which can be e.g. the logit or the probit. A notable feature of this model compared to e.g. the linear model of Eq. 3.8 is that the variance depends on the mean. In practice, the parameters $\beta$ and $\phi$ can be estimated by numerical maximization of the likelihood functions.

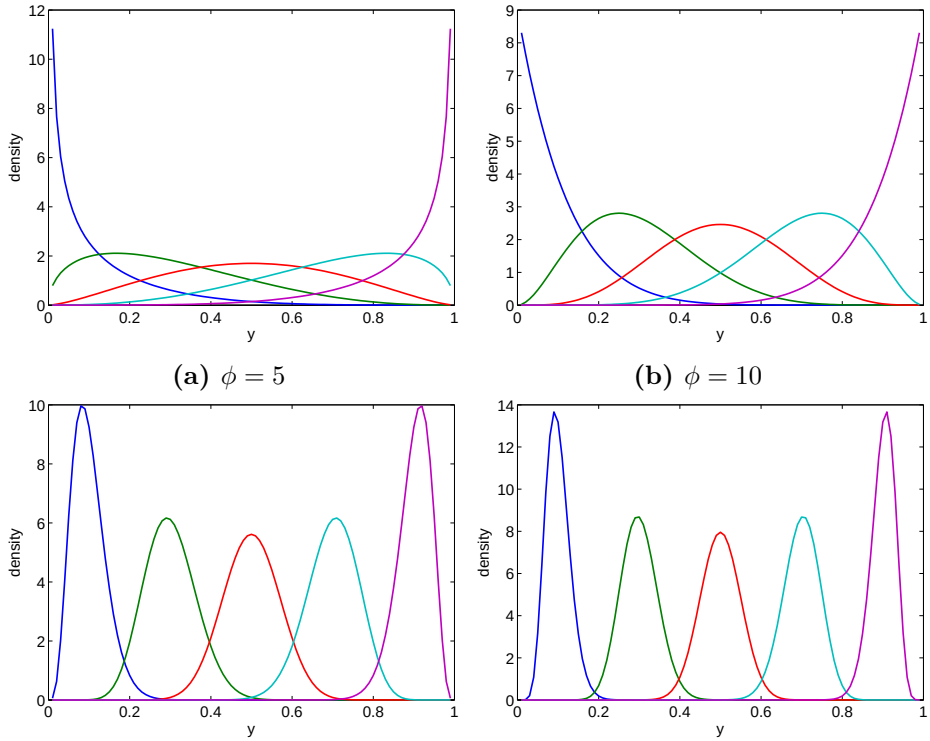


**(a)** $\phi = 5$ **(b)** $\phi = 10$

**(c)** $\phi = 50$ **(d)** $\phi = 100$

**Figure 3.2:** Probability density function of the Beta distribution for different values for the parameters $\mu$ and $\phi$. The colors correspond to different values of $\mu$ as follows: blue: 0.1, green: 0.3, red: 0.5, cyan: 0.7, purple: 0.9. It can be observed that the variance of the distribution decreases when $\mu$ moves towards the edges of the range $(0, 1)$.

### 3.2.3 Regularization methods

The more complex a supervised learning model is, the more prone it is to *overfitting*. That is, it is able to achieve a low training error by fitting exactly to the data

set, but suffers from poor generalization. This is exemplified in Figure 3.3, which shows an example of how increasing the order, and therefore the complexity, of a polynomial model results in overfitting a set of data points. In machine learning problems, overfitting is more likely to occur the less training samples there are; in particular, if the dimensionality $p$ of the feature vectors is much larger than the number $N$ of training samples, methods like least-squares are inadequate Hastie et al. (2009).

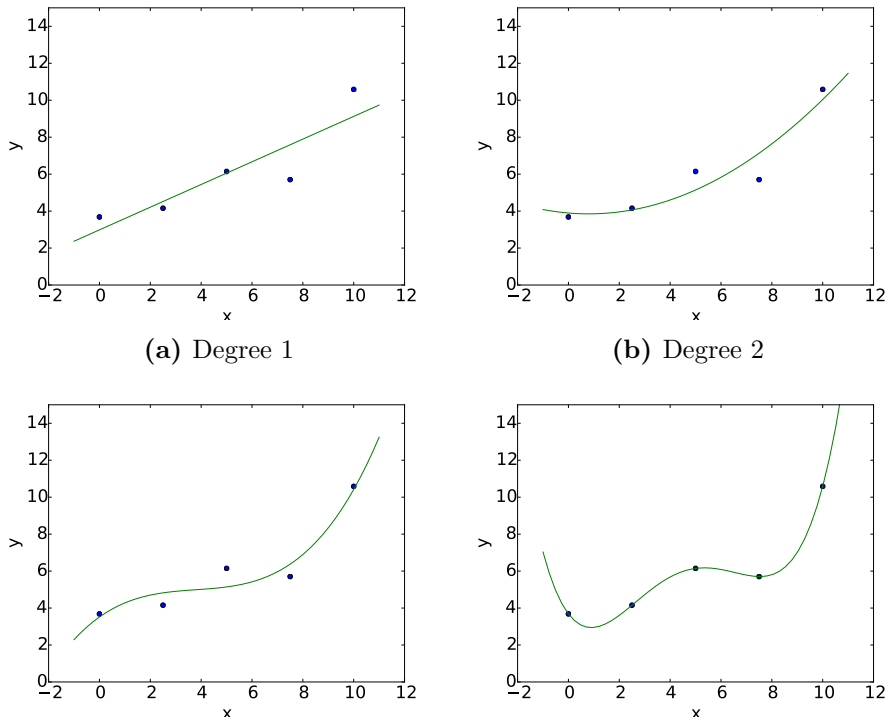


**(a)** Degree 1 **(b)** Degree 2

**(c)** Degree 3 **(d)** Degree 4

**Figure 3.3:** Increasing the degree of a polynomial fitted to linear data quickly results in overfitting. The points are data generated from a line $y = x + 1$ with Gaussian errors with $\sigma^2 = 1.22$, and the lines are polynomial curves of different degrees fitted to the points. Note that all of the models with a degree greater than unity would make poor predictions outside the range of the training points, and the model of degree 4 is complex enough to fit the training data perfectly.

Regularization methods attempt to reduce overfitting by restricting the space of allowed solutions to the estimation problem. In the case of GLMs, this can be done by adding a penalty term to the minimization problem solved in least-squares or maximum-likelihood estimation. In particular, the most generally useful penalty

terms restrict an $L_p$ norm of the parameter vector:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^{N} (y_i - \mathbf{x}_i^T \beta)^2 + \lambda \|\beta\|_p, \tag{3.13}$$

where $\lambda$ is a parameter controlling the strength of the regularization; the larger the $\lambda$, the stronger the regularization. Using the $L_2$ norm yields *ridge regression* (Hoerl and Kennard, 1970). In ridge regression, increasing the value of $\lambda$ moves the elements of $\beta$ towards zero and towards each other; thus the method avoids solutions where the coefficients vary wildly and have large magnitudes. If the penalty is instead on the $L_1$ norm, the problem is called the *lasso* (Tibshirani, 1996). The lasso also penalizes the magnitudes of the individual coefficients, but it has the attractive property of yielding sparse solutions, that is, solutions where some coefficients are zero. Increasing $\lambda$ will increase the sparsity of the resulting model, and for this reason the lasso is also an effective method of *feature selection*.

Another issue when trying to predict with $p \gg N$, is that there are bound to be correlations between the features. Fitting a GLM model without regularization will result in the inclusion of all of these features, thus the model will be redundant. This problem is also present when using ridge regression. The lasso, on the other hand, selects at most $N$ variables even when $p > N$, and it tends to select only one of a set of correlated features. Thus, while the lasso results in often desirable sparsity, in some applications, it may hamper the interpretability of the results, because interesting features may be excluded from the model. Designed to alleviate this problem, the *elastic net* (Zou and Hastie, 2005) is a regularization method that combines the penalty terms of ridge regression and the lasso.

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^{N} (y_i - \mathbf{x}_i^T \beta)^2 + \lambda_1 \|\beta\|_1 + \lambda_2 \|\beta\|_2. \tag{3.14}$$

Here, by controlling the ratio of the strenghts of the penalty terms, it is possible to control the tradeoff between sparsity and grouping – i.e. the inclusion of groups of correlated features. We used the elastic net in Publication III to perform feature selection from a large set of image features.

### 3.2.4 Support vector machines

For a linear classifier to be useful, its input features need to be at least approximately linearly separable. Nevertheless, the application of linear classifiers extends to nonlinearly separable datasets, because any data set can be made linearly separable by mapping it nonlinearly to a high enough dimension Cover (1965). The problem with such a data transformation is that the required dimension may be prohibitively high from a storage and computational standpoint. However, in methods where the training depends only on inner products of the feature vectors, it turns out that no explicit transformation is needed; the high-dimensional inner

products can be computed efficiently through the use of a *kernel function*. This so-called *kernel trick* is behind the success of the SVM.

The SVM is a linear classifier that finds the optimal separating hyperplane to separate the classes in the training data. In the linearly separable case, this is simply the hyperplane that is maximally far from the closest training sample of each class. This method has its origins in the *generalized portraits* proposed in (Vapnik and Lerner, 1963). The concept of margin and the optimal separating hyperplane are illustrated in Figure 3.4. Since real data is rarely linearly separable, the separating hyperplane is selected such that it is as close as possible to the misclassified training samples, as proposed in Cortes and Vapnik (1995). The loss function that is used in training the SVM is the *hinge loss*:

$$L(y) = \max(0, 1 - y(\mathbf{x}^T\beta - \beta_0)), \tag{3.15}$$

where the responses $y \in \{-1, 1\}$. The loss is zero for points correctly classified by the hyperplane, and proportional to the distance to the hyperplane for misclassified points. The optimization problem solved by the training algorithm can be expressed as

$$\min_{\beta,\beta_0} \left[ \frac{1}{N} \sum_{i=1}^{N} \max(0, 1 - y_i(\mathbf{x}_i^T\beta - \beta_0)) \right] + \lambda\|\beta\|^2. \tag{3.16}$$

Here, $\lambda$ is a regularization term that restricts the the amount of misclassifications allowed in the training data, at the expense of margin size. We note that although we focus here only on classification, the SVM can also be used to solve regression problems (Cortes and Vapnik, 1995; Schölkopf et al., 2000).

The optimization algorithm that is used to train the SVM depends only on the inner products of the feature vectors, and thus the kernel trick can be used. For two vectors $\mathbf{a}, \mathbf{b}$, a kernel $K(\mathbf{a}, \mathbf{b})$ computes the inner product $\psi(\mathbf{a})^T\psi(\mathbf{b})$, where $\psi$ is a function that maps its input to a higher dimensional space. The feature transformation can thus be implemented by replacing the inner products $\mathbf{x}_i^T\mathbf{x}_j$ with $K(\mathbf{x}_i, \mathbf{x}_j)$. Examples of the most common kernel functions are the *polynomial kernel*

$$K(\mathbf{a}, \mathbf{b}) = (1 + \mathbf{a}^T\mathbf{b})^d, \tag{3.17}$$

where $d$ is the order of the polynomial mapping; and the *radial-basis function kernel*

$$K(\mathbf{a}, \mathbf{b}) = \exp(-\gamma\|\mathbf{a} - \mathbf{b}\|^2), \tag{3.18}$$

where $\gamma > 0$. The latter is particularly interesting, because it corresponds to a mapping to an infinite-dimensional space.

In Publication III we applied an SVM to automatically detect regions in a tissue image that were of too low quality to be used for further analysis.

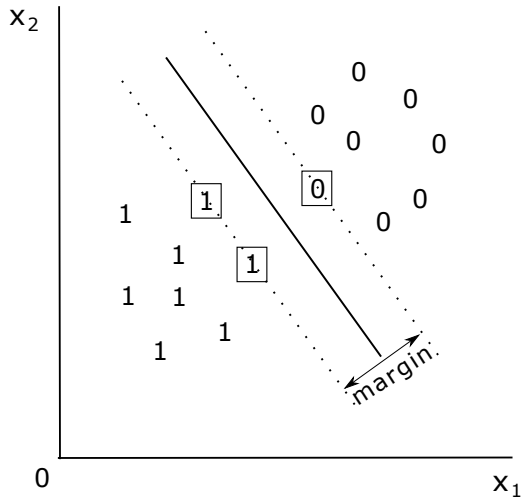


**Figure 3.4:** A linearly separable dataset for which the optimal separating hyperplane has been determined. The hyperplane is defined by the three boxed points that lie on the margin, which are referred to as the *support vectors*.

### 3.2.5 Random forests

*Decision tree* classifiers are tree models in which each node performs a binary test on a feature value, and branches into two edges that correspond to the different outcomes of the test; the leaf nodes correspond to different classes. A sample is classified by starting at the root node, performing the test and moving down the edge indicated by the test, and continuing traversing the tree until a leaf node is reached, which gives the predicted class. This is illustrated in Figure 3.5. The training of a decision tree starts with the whole training data set and finds the test that best separates the classes in the data, according to a chosen *impurity* measure; it then splits the data to the new nodes, and continues this process at each node that receives samples from different classes. When the training is finished, each leaf node contains samples from only one class, which is the class assigned to that leaf node (Hastie et al., 2009).

Decision trees in this simplest form overfit the data, since they always result in a perfect classification for the training set. Hence, training methods used in practice apply some rule to prune the tree, i.e. to avoid splits that do not improve the classification significantly, based on e.g. a statistical test. Another popular approach is to use *ensemble methods* to reduce variance: The *random forest* (Breiman, 2001) is an example of a broader class of *bagging* methods for classification, where multiple models are trained and their outputs are combined, usually by a majority vote. A random forest is an ensemble of decision trees, in which each tree is trained with a bootstrap sample, i.e. a random subset, of the training data. An important distinction to traditional decision trees is that each split is based on a random subset of the features. This makes random forests

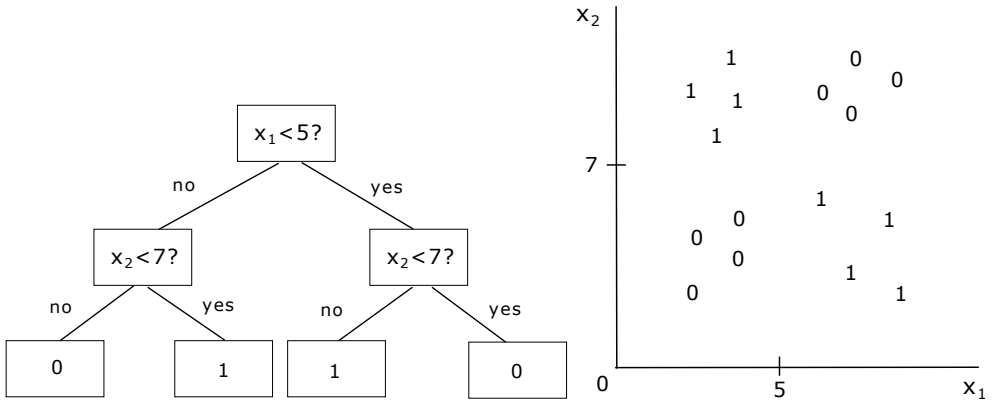


**Figure 3.5:** The diagram on the left depicts a decision tree that correctly classifies the data shown on the right. Zeros and ones correspond to samples from different classes. This simplistic classification problem is a pathological case for linear classification; indeed, the split on the root node of the decision tree can be thought of as a linear classifier, and it does not separate the classes in the data. The splits on the next level, however, separate the data perfectly.

effective at reducing variance; the correlation between the outputs of the different trees is decreased, which makes each tree more informative in the majority voting. In Publication II we trained a random forest to detect tips of mitochondria.

# 4 Analysis of bioimages

## 4.1 Segmentation of subcellular structures

Most image analysis tasks require a knowledge of the exact location of an object in an image. This can be achieved via *segmentation* - the partitioning of pixels into sets corresponding to different objects – e.g. cells, organelles and background. The partitions can be represented as images where all pixels have value zero, except those that correspond to the object – such an image is referred to as a *mask*. A related concept is *detection*, which generally refers to a process whose output is a single point that marks the approximate location of the object.

Most methods for segmenting subcellular structures in fluorescence microscopy images follow a general framework (adopted from Smal et al. (2010)) that consists of three steps: (i) noise reduction, to prevent errors in further processing steps; (ii) signal enhancement, to increase the contrast between the background and the objects to be segmented; and (iii) intensity thresholding, to produce a binary mask. Figure 4.1 illustrates these steps. We will next describe each step with examples of how recent methods implement them and, in particular, how these steps are used for obtaining segmented mitochondria from images. After that, we will consider alternative segmentation and detection approaches that do not fit this framework.

### 4.1.1 Noise reduction

Due to the limitations discussed in Chapter 2, fluorescence microscope images always contain noise. It is often desirable to reduce the amount of noise before further processing. Image noise is commonly assumed to be additive and Gaussian, which is a good approximation with reasonably high signal-to-noise ratios (SNRs) (Vliet et al., 1998). The measured image $I$ can then be modelled as

$$I(x, y) = f(x, y) + \eta(x, y), \tag{4.1}$$

where $f$ is the emission signal being measured. We assume that $\eta(x, y) \sim \mathcal{N}(0, \sigma^2)$ and that $\eta(x, y)$ are independent and identically distributed. The goal of noise reduction is to estimate $f$ given the measurement $I$.

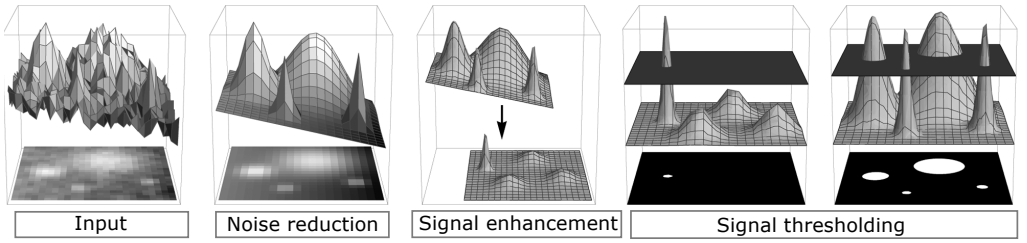


**Figure 4.1:** A general framework for detecting subcellular objects. Starting with a noisy input image, the first step is to apply a noise reduction method, such as Gaussian or median filtering. The next step, signal enhancement, subtracts possible uneven background signal from the image, and increases the contrast between objects and background. Finally, the signal thresholding step converts the enhanced image into a binary mask, and possibly removes objects that are too small or too large. Adapted with permission from Smal et al. (2010).

Traditionally, the most commonly used noise reduction methods are local filters, i.e. they estimate the value $f(x, y)$ from pixels of $I$ in a neighborhood $N$ of $(x, y)$, which is most often a square region centered at $(x, y)$. The simplest noise reducing filter is the mean filter, which computes the mean of the values in $N$. The mean filter is effective at removing Gaussian noise in neighborhoods where $f$ is approximately constant; in regions with sharp details such as edges, however, it will result in a blurry image. A more common version, the Gaussian filter, computes a weighted mean, where the weights are based on the distance to the center pixel, according to a 2-D Gaussian function; this offers slightly better edge preservation. On the other hand, the median filter, which estimates $f(x, y)$ as the median of the neighborhood, is effective at preserving edges. In addition, the median filter is fairly robust against photon shot noise, which is present in images with a low SNR. For these reasons, in Publications I and III, we used median filtering prior to segmenting mitochondria from the images.

More recently, research on noise reduction methods has focused on *non-local* methods. These are based on the assumption that small neighborhoods, or patches, of an image lie on a low-dimensional manifold, which implies that for any neighborhoods, multiple similar neighborhoods are present in the same image. The pioneering work of this approach was the *non-local means* (NLM) filter, which estimates $f(x, y)$ as the weighted mean of the center points of patches in $I$, where the weight increases with similarity to the patch centered $(x, y)$. The similarity is defined in terms of the euclidean distance between the patches interpreted as vectors. The current state of the art in image denoising is the non-local method *block-matching and 3-D filtering* (BM3D) (Dabov et al., 2007), which has been used as a preprocessing method in cell segmentation (Chowdhury et al., 2013). In Figure 4.2 we show examples of different noise reduction filters applied to a microscope image.

The assumption of Gaussian noise is not always appropriate: the SNR in fluores-

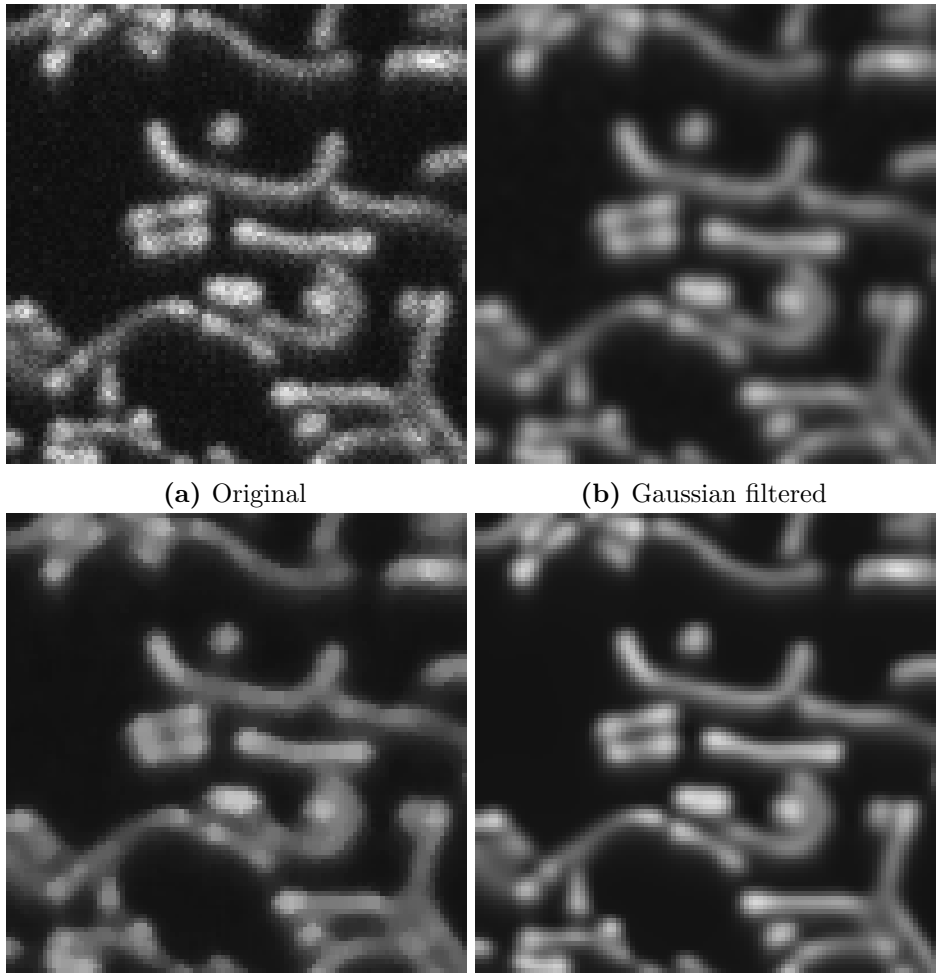

**(a)** Original **(b)** Gaussian filtered

**(c)** Median filtered **(d)** BM3D filtered

**Figure 4.2:** Outputs of different noise reduction techniques on an example image of mitochondria. The parameters were selected to remove the noise while preserving as much detail as possible: for the Gaussian filter, the standard deviation $\sigma = 1.2$ was selected, and for the median filter, the window size was $5 \times 5$. The BM3D required no parameters. Note that the median filter tends to exaggerate the edges in the image, while the BM3D produces a smooth image without blurring details as much as the Gaussian filter.

cence microscopy images may be low due to, e.g., low copy numbers of fluorescent proteins, or using weak illumination in order to avoid photodamage. In such cases, the noise is better modeled as Poissonian, or as a combination of Poissonian and Gaussian noise; methods based on the assumption of Gaussian noise will perform suboptimally. A natural way to approach this problem is to derive the solution to the denoising problem assuming the specific noise model, as was done, for example, in (Paul et al., 2013) in the context of segmentation. An alternative solution is to use a *variance-stabilizing* transformation on the image, which results in an image

whose noise has an approximately Gaussian distribution, on which standard denoising methods can be applied. For Poisson noise, one simple variance-stabilizing transformation is the square root, and a more general transformation has been proposed by Anscombe (1948); for mixed Poisson-Gaussian noise, a generalized Anscombe transformation has been derived Murtagh et al. (1995). After denoising, an inverse transformation has to be applied to obtain the final denoised image; finding a suitable inverse transformation has been an important research topic in recent years (Mäkitalo and Foi, 2013).

### 4.1.2 Signal enhancement

The goal of the signal enhancement step is to separate the objects from the background in their intensity, so as to make it possible to reliably segment them by thresholding. To this end, signal enhancement can consist of various processing substeps to subtract background fluorescence arising from, e.g., out-of-focus objects, and to decrease intensity variations between the objects of interest.

Similar to noise reduction, background subtraction can be stated as an estimation problem, where the background intensity for a pixel is estimated from a neighborhood of the pixel, and then subtracted from the pixel's intensity. For this, linear and nonlinear filters have been used; For example, Rizk et al. (2014) used a nonlinear rolling-ball filter which estimates the background intensity as the mode of the intensities in $N$. On the other hand, the linear Laplacian-of-Gaussian filter is commonly used especially for detecting spot-like objects (see (Smal et al., 2010)). For detecting spots in 3-D images, Allalou et al. (2010) proposed a method that enhances spots using a cosine filter.

For detecting curvilinear objects, several signal enhancement methods have been proposed. The method of Steger (1998) is based on two observations about the geometry of a typical line in an image: first, the derivative of the line profile, that is, the intensity profile along a line perpendicular to the line, should vanish in the center of the line, and second, the magnitude of its second derivative should be large. In order to assess these features at a specific location in an image, the direction normal to the hypothetical line must be estimated; this can be done by finding the direction in which the second directional derivative is maximized. Frangi et al. (1998) proposed a *vesselness* criterion for detecting blood vessels in medical imaging, which measures the similarity of the relations between eigenvalues of the Hessian matrix to those of a curvilinear structure, and involves a scale selection operation to account for lines of different sizes. Obara et al. (2012a) proposed to decrease the dependence of such measures on the local image contrast, by using a phase congruency tensor instead of the Hessian matrix, which is based on the notion that the local phase varies little with the scale at salient image structures.

For segmenting mitochondria in Publications I and III, we used morphological top-hat filtering. This was chosen as a more noise-resistant alternative to a previously

proposed method, which used a linear top-hat filter followed by median filtering to remove artifacts (Koopman et al., 2006). Peng et al. (2010) used adaptive local normalization, which refers to local centering and scaling in a adaptively sized window; in their method the window size is chosen, based on the local variance, to be large enough to contain both mitochondria and background. Although not specifically designed for segmenting mitochondria, the method of Rizk et al. (2014) is also shown to be able to quantify changes in mitochondrial length.We note that these approaches were designed for two-dimensional images, but they are each extendable to three dimensions. Recent works have addressed the problem of segmenting mitochondria from 3-D images: In (Nikolaisen et al., 2014), the authors proposed to use a deconvolution method for signal enhancement, and found that it compared favorably to their previous approach (Koopman et al., 2006). Recently, Viana et al. (2015) proposed detecting mitochondria using a combination of background subtraction via derivative filters, along with a multiscale vesselness filter.

### 4.1.3 Thresholding

Thresholding refers to choosing an intensity value $T$ such that all pixels brighter than $T$ will be considered to belong to an object, while the rest is considered to belong to the background. The threshold value can be selected manually; if the signal enhancement method is designed such that the same value for $T$ produces consistent results across images, $T$ can be used as a parameter for the method. For example, if the intensity values can be interpreted as likelihoods that an object is present, then the choice of $T$ reflects the required level of confidence to accept a detection. Nevertheless, not all methods produce such output, thus it is often desirable to select the threshold value automatically.

One of the most widely used methods for automatic threshold selection is *Otsu's method*. This method is a computationally efficient algorithm for partitioning the image pixels into two classes so as to maximize the inter-class variance as defined in linear discriminant analysis (Otsu, 1979). Otsu's method has been used for segmenting small subcellular particles (see Ruusuvuori et al. (2010)), and mitochondria (McClatchey et al., 2015; Vowinckel et al., 2015). Our method for segmenting mitochondria, used in Publications I and III uses Otsu's method for thresholding. Several other automatic threshold selection procedures have been proposed (see Sezgin and Sankur (2004) for a review); histogram-based methods select the threshold either based on the histogram shape, for example by locating a local minimum, or, like Otsu's method, by applying a clustering method to the pixel values. Threshold selection can be also done adaptively, by selecting different threshold values for different regions of the image based on the local intensity statistics. However, such methods can often be decomposed into a signal enhancement step and a global thresholding step.

In the common case that the image contains multiple objects to be segmented,

different objects will ideally not be connected, in which case they can be trivially separated by a labeling algorithm. However, in practice, thresholding results in errors of either undersegmentation, where one mask corresponds to multiple objects, or oversegmentation, where multiple mask correspond to a single object. These errors can be corrected to some extent by splitting and merging procedures; for example, undersegmented fluorescent spots can be further segmented by fitting a mixture of Gaussian functions to the pixels, and considering each component of the mixture a separate object (Ruusuvuori et al., 2010). Furthermore, due to imperfections in the image or errors in the processing steps, the result may contain segmented objects that do not correspond to objects. Often such spurious detections are smaller than the true objects, and can be filtered out by setting a threshold on their size.

### 4.1.4 Other segmentation approaches

Not all methods for detecting or segmenting subellular objects fall strictly under the framework presented above. For instance, Rizk et al. (2014) segment the objects in an image in several stages in a coarse-to-fine fashion, Obara et al. (2012b) apply the watershed algorithm instead of thresholding to segment fungal networks, and Allalou et al. (2010) apply thresholds to the outputs of several filters to obtain the final segmentation result. Next, we will discuss two important classes of segmentation methods, namely, *learning-based* methods and *deformable models*.

#### 4.1.4.1 Learning-based segmentation

Segmentation and detection can be seen as classification problems, where the labeling of the pixels is learned. In contrast to the thresholding framework presented above, learning-based segmentation methods often lack distinct steps of noise reduction or thresholding. A separate thresholding step is unnecessary if the output of a classifier is used as the segmentation. Regarding noise reduction, machine learning techniques do not necessarily benefit from such preprocessing, and instead require large amounts of training data to ensure resistance to noise (Ronneberger et al., 2015).

The class of a pixel should intuitively be decidable based on its local context; that is, features should be computed from a neighborhood of the pixel. For example, in Publication II, we used the pixel intensities of the neighborhood of a pixel as features in a Random Forest classifier to detect the presence of a tip of a mitochondrion. Other possible features include histograms and summary statistics of the pixel intensities, and descriptors of image texture.

In addition to classifiers, unsupervised learning is also used in segmentation. For example, Otsu's method (Otsu, 1979) solves automatic threshold selection as a clustering problem. Methods based on Bayesian estimation (Geman and Geman,

1984) or function approximation Mumford and Shah (1989) interpret the image as a degraded version of an underlying piecewise smooth (or constant) image, in which each unique value corresponds to a class, and the problem is the estimation of the underlying image. This is done by minimizing a distance measure between the image and the estimate, regularized by terms that encode assumptions about the structure of the solution. For example, the model of Mumford and Shah (1989) includes regularization terms to encourage smoothness of the image, and to penalize the length of the boundary separating the segments. From the statistical estimation perspective of e.g. (Geman and Geman, 1984), the choice of distance measure reflects the noise model assumed for the data; Lecellier et al. (2009) presented a general method for segmentation for the exponential family of noise models. Recently, this was extended by a general formulation that models noise via GLMs, and allows the modeling of linear degradations such as convolution by a PSF, thus combining deconvolution and other image restoration problems with segmentation (Paul et al., 2013).

One advantage of learning-based segmentation is that one can avoid the need for designing segmentation methods *ad hoc*. For example, Sommer et al. (2011) proposed a tool, ilastik, for interactive segmentation of biological images, the goal of which is to be easy to use without expert knowledge on machine learning or image analysis. In practice, the performance of such classification-based segmentation depends on the selection of informative features. However, methods based on convolutional neural networks, such as (Ronneberger et al., 2015), embed feature extraction and selection in the training process, thus showing promise of a true general purpose learning-based segmentation method.

#### 4.1.4.2 Deformable models

One approach for segmenting large objects is to use deformable models such as active contours (ACs). The first AC model, the *snake*, was proposed in (Kass et al., 1988). Snakes are closed parametric curves that are iteratively fitted to structures in the image, while penalizing unlikely shapes with a smoothness term. Another way to model ACs is to define them as a level set of a higher dimensional function; this segmentation approach was proposed by (Chan and Vese, 2001). In contrast to the snake model, the level-set approach is not sensitive to the initialization of the AC, and does not require knowledge of the number of objects in the image. It should be noted that AC models such as the Chan-Vese model (Chan and Vese, 2001) can be interpreted in the statistical estimation framework discussed above.

For segmenting relatively round objects, such as cells, the closed curve model of (Kass et al., 1988) is appropriate. However, filamentous structures like cytoskeletal filaments or mitochondria are better modelled as open contours: Smith et al. (2010) introduced the stretching open active contour (SOAC) model, which includes a stretching force in the ends of each contour. SOACs have been applied in the

segmentation of individual actin fibers (Smith et al., 2010) and actin networks (Xu et al., 2011). Figure 4.3 shows an example of a SOAC applied to segment an actin filament. A level-set based method for filament segmentation was presented in (Xiao et al., 2016); the method models the filament shapes through two level-set functions fit simultaneously. As opposed to the SOAC model, the method is not sensitive to initialization, and does not require the knowledge of the number of objects in the image; a major drawback of the model is that it cannot represent filaments that are overlapping or networked.

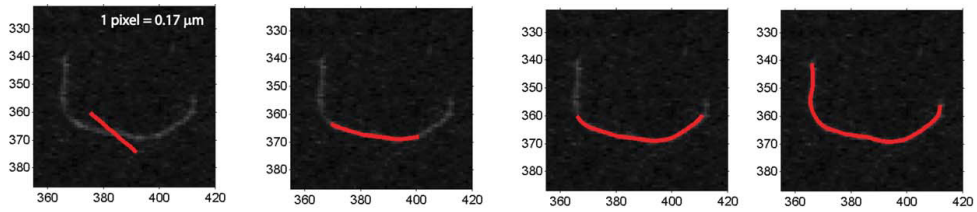


**Figure 4.3:** On the left, the SOAC (red) is initialized on a fluorescence microscope image of an actin filament. The rest of the images are snapshots of the curve being iteratively fitted to the shape of the filament. Adapted with permission from Smith et al. (2010).

## 4.2 Motion analysis

### 4.2.1 Tracking

Given a proper segmentation of the objects in each frame in a movie, their dynamics can be studied by tracking them. This entails finding an association between the segments in different frames that correspond to the same object. The simplest strategy for this *data association* problem is to link objects between each consecutive pair of frames, by minimizing a cost function that penalizes, e.g., the distance between the associated objects or differences in their spatial intensity distributions. This corresponds to the linear assignment problem (LAP) for which fast algorithms are available (Jonker and Volgenant, 1987). However, this approach is insufficient if occlusions are present in the movie, because an occlusion will lead to a trajectory being split; this is illustrated in Figure 4.4. Thus, some works have extended this approach to link the detections over multiple frames (Sbalzarini and Koumoutsakos, 2005).

The optimal solution to the data association problem is given by multiple hypothesis tracking (MHT) (Reid, 1979), which simply considers each possible combination of associations between objects, and chooses the one with minimal cost. This is computationally infeasible, and approximations need to be used. A popular solution is to start with the greedy pairwise matching outlined above, or a multi-frame extension, which may result in incomplete track segments, and then build complete tracks by combining these segments (Wu et al., 2011). One example

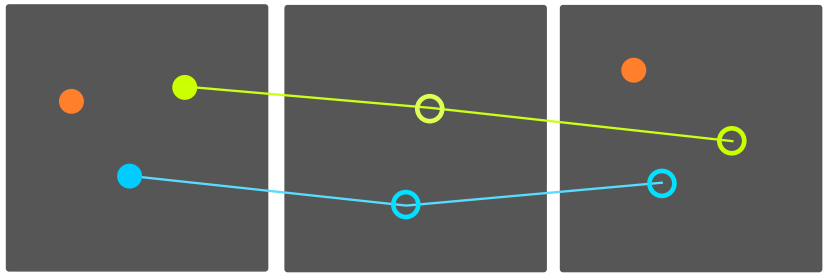

**Figure 4.4:** An illustration of the problem of occlusions when using only subsequent frame pairs for tracking. Each gray box represents a frame of a movie, and circles of different colors represent different objects; filled circles represent objects that are considered newly appearing by the tracking algorithm, and the lines mark the associations between objects in different frames. In this example, the orange object is occluded in the second frame, and because of that it is considered to be a different object in the first and the last frames.

of this is the method of (Jaqaman et al., 2008), which we used in Publication II to track detected tips of mitochondria. The method works by solving another LAP where the assignments are formed between starting points and endpoints of track segments, which are thus connected into full tracks. By allowing the linking of two starting points or two endpoints, the method also naturally accounts for splitting and merging events.

In such tracking methods, a vital component is the definition of the cost function. Simplest costs for linking two objects might penalize the euclidean distance or the amount of overlap between the objects. If the objects are not well separated throughout the movie, such cost functions are prone to erroneous assignments. A more robust approach is to formulate the cost functions based on a model of the objects' dynamics; in the case of intracellular structures, this is often a combination of Brownian motion and directed transport. Depending on the specimens and the experimental setup, other possible components for the cost function are intensities and shapes of the objects.

Not all methods require an explicit segmentation before tracking. In particular, active contours can be used for simultaneous segmentation and tracking. This is achieved by initializing the contours in the current frame by those found in the previous frame. The SOAC model has been used to track actin filaments Smith et al. (2010). Although this approach should be applicable to the analysis of mitochondria, such applications have not been reported; to our knowledge, no method has been proposed for the automated tracking the whole mitochondrial network.

### 4.2.2 Optical flow

It is possible to estimate object motion without explicit tracking. The concept of *optical flow* (OF) refers to the apparent motion of the intensity mass between frames in a movie. In the context of microscopy, the optical flow corresponds to the motion of the specimen, given that the surrounding environment is static relative to the microscope. For estimating OF in a movie, two assumptions are generally made. The first, so-called *optical flow assumption*, posits that any point in a frame have a corresponding point in the next frame. The second assumption is that of *brightness constancy*: the intensity of the apparently moving point remains unchanged. From these, one can formulate the *optical flow equation*:

$$I(x, y, t) = I(x + \delta x, y + \delta y, t + \delta t). \tag{4.2}$$

The problem is then to estimate the displacements $(\delta x, \delta y)$ for each point in the image.

In the *differential* class of methods for OF estimation, it is assumed that the flow is locally linear. This allows replacing the right side of eq. 4.2 with its first order Taylor approximation:

$$I(x, y, t) = I(x, y, t) + \delta x I_x(x, y, t) + \delta y I_y(x, y, t) + \delta t I_t(x, y, t). \tag{4.3}$$

Setting $\delta t = 1$ without loss of generality, this further simplifies to

$$\delta x I_x(x, y, t) + \delta y I_y(x, y, t) = -I_t(x, y, t), \tag{4.4}$$

which shows that the OF can be estimated based on only the spatial and temporal derivatives of the image. This problem is ill-posed, because there is only one equation but two parameters to estimate; thus, further assumptions are needed to regularize it.

Most modern OF estimation methods are based on the traditional *local* approach of Lucas and Kanade (Lucas and Kanade, 1981) and the *global* approach of Horn and Schunk (Horn and Schunck, 1981). The former regularizes the problem by assuming that the flow is constant inside a local window, and solves the flow for each pixel as a least-squares estimation problem inside the neighborhood defined by the window. The latter imposes a smoothness constraint on the flow field, while solving the least-squares estimation for each pixel simultaneously.

In Publication I, we applied the Lucas-Kanade method to quantify instantaneous velocities of mitochondria. The method is suitable because the assumption of locally constant flow holds reasonably well at small scales. A limitation is that the displacements must be small, which requires a fast rate of imaging. OF estimation has also been used for quantifying mitochondrial velocities in axons (Gerencser and Nicholls, 2008), and for other subcellular objects (Ulman, 2010).

## 4.3 Image-based measurements

Segmentation of objects enables morphometry, i.e. measurements of the shapes of the object, as well as the quantification of fluorescence intensity inside the objects. Fluorescence intensities can be used to estimate fluorescent protein counts; this principle has been applied in, e.g., counting fluorescently tagged mRNA molecules (Golding et al., 2005) and in estimating mitochondrial membrane potential (Koopman et al., 2006). Finally, given objects labeled with different fluorescent labels, their colocalization can be studied.

In addition to segmentation-based measurements, image *texture* can be described computationally. In the context of grayscale images, texture can be loosely defined as the local spatial distribution of intensity in an image, and can be thought as an arrangement of low-level structures such as edges and lines at different scales and orientations. Texture in a fluorescence microscope image can inform about the state of the specimen; for example, fragmented mitochondria exhibit a different texture from elongated mitochondria (see Figure 4.5), which we took advantage of in Publication III. This difference in texture is apparent even in images with relatively small resolution and low SNR, where individual mitochondria are not necessarily detectable by a segmentation algorithm.

An important application for image-based measurements is to use them as features in a learning method that, for example, classifies cells based on their phenotype. This is a key approach in image-based screening studies (Shariff et al., 2010). We will next outline methods that are used to quantify the morphology and texture of objects in images, and review applications of these in machine learning problems.

### 4.3.1 Morphometry

The theory of *mathematical morphology* gives tools for quantifying the morphology of segmented objects. For example, the length of an object can be approximated using morphological thinning (Guo and Hall, 1989), which "peels" the mask to form a single-pixel-wide mask; a rough estimate of the length is the number of pixels scaled by the distance between two objects one pixel apart. However, this underestimates distances between diagonally-connected pixels, and a better estimate is obtained by considering different arrangements of neighboring pixels and assigning different lengths to them Gonzalez and Woods (2006). Similarly, the area of the object is proportional to the number of pixels in the original mask, and the length of the perimeter can be computed by finding the edge pixels. More accurate estimates can be obtained by modeling the object shape with a spline, such as an active contour (Kass et al., 1988).

Besides elementary measurements of lengths and areas, more complex shape analysis is often necessary for describing object morphologies. For this, several shape descriptors have been proposed in the literature (Mingqiang et al., 2008). These include features that quantify the deviation of the shape from a simple

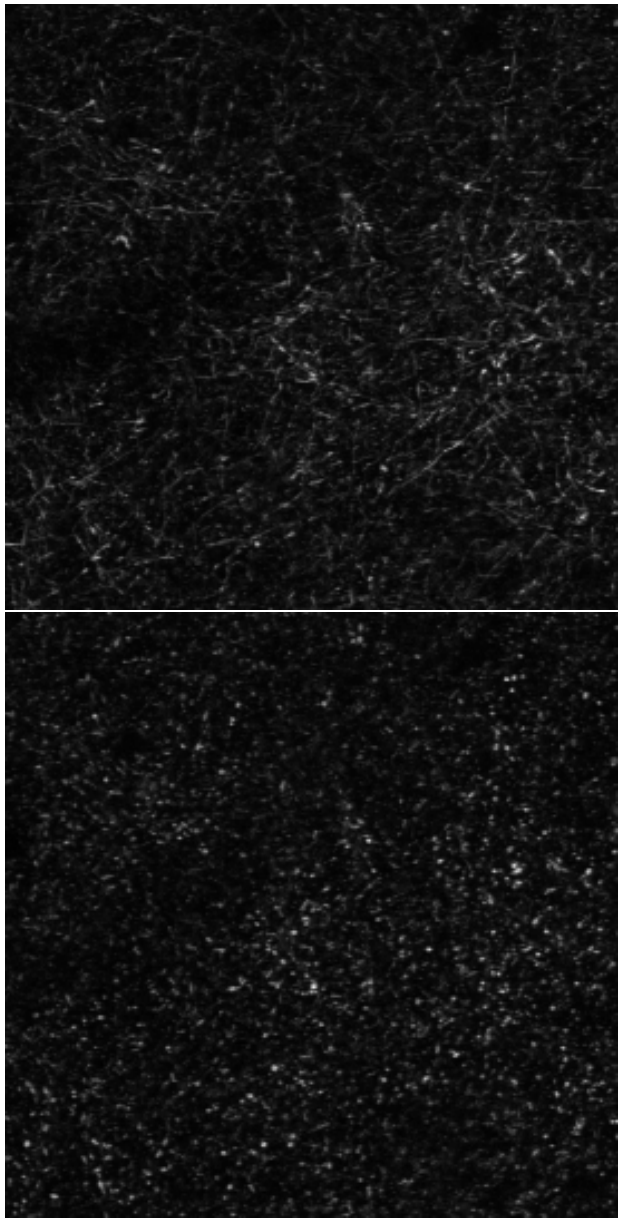

**Figure 4.5:** In two-photon microscopy microscope images of fluorescently labeled mitochondria in the mouse neocortex, the image texture changes when the mitochondrial morphology changes from elongated (top) to fragmented (down).

shape such as a circle, and various statistical descriptors based on mapping a shape to one dimension. Such descriptors can be useful in detecting abnormal shapes in cells; for example, Bakal et al. (2007) proposed several morphological features for describing cell shape. Nikolaisen et al. (2014) proposed a list of quantities to capture variations in mitochondrial morphology, including measurements based

on a decomposition of the mitochondrial network into individual branches similar to Publication I.

### 4.3.2 Texture analysis

The goal of texture analysis is to extract a set of features from an image that characterizes different aspects of texture. Depending on the scale one is interested in, the image can be a complete image of, e.g., a cell, or a subimage depicting a region of a cell or tissue. It should be noted that the concept of texture, and thus the statistical features developed to describe it, assume some spatial homogeneity; therefore, texture analysis is mostly meaningful in regions that exhibit such homogeneity.

Various statistical descriptors for texture have been proposed in the literature. An early example of texture features is that of the Haralick features, which are based on statistics of the *graylevel co-occurrence matrix* (GLCM) (Haralick et al., 1973). The GLCM is a two-dimensional histogram, in which each element $GLCM_{ij}$ contains the number of occurrences where a pixel and another pixel at a specified offset have values $i$ and $j$. Thus, the use of the GLCM for texture analysis is based on the assumption that image texture can be described by the statistical relationship between pairs of nearby pixels. From the GLCM, various summary can be computed that describe different aspects of texture; in the original work by Haralick et al., 14 features were defined, including correlations, variances, and entropy-based quantities. Haralick features are commonly used as a part of larger feature sets in bioimage analysis (Newberg et al., 2009).

Many successful texture feature extraction methods are based on filter banks. One such case are Gabor filters, which have been widely used since they were first applied for texture-based image retrieval by Manjunath and Ma (1996). This is mainly because they respond to edges and linear structures and can be tuned in orientation and scale. A Gabor function is defined as the product of a 2D complex sinusoidal wave multiplied by a Gaussian function, which yields the following expression:

$$g(x, y) = \left( \frac{1}{2\pi\sigma_x\sigma_y} \right) \exp\left[ -\frac{1}{2} \left( \frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) + 2\pi j\omega x \right], \tag{4.5}$$

where $\sigma_x^2$ and $\sigma_y^2$ are the horizontal and vertical spreads of the Gaussian function, and $\omega$ is the frequency of the complex sinusoid. Figure 4.6 shows examples of Gabor functions with different parameters.

The dual-tree complex wavelet transform (DT-CWT) can be used to efficiently compute texture features which share some similarity to Gabor features Selesnick et al. (2005), and which have achieved similar discrimination performance on a face recognition task (Eleyan et al., 2008). The DT-CWT can be seen as a complex-valued extension of the discrete wavelet transform (DWT), which is widely

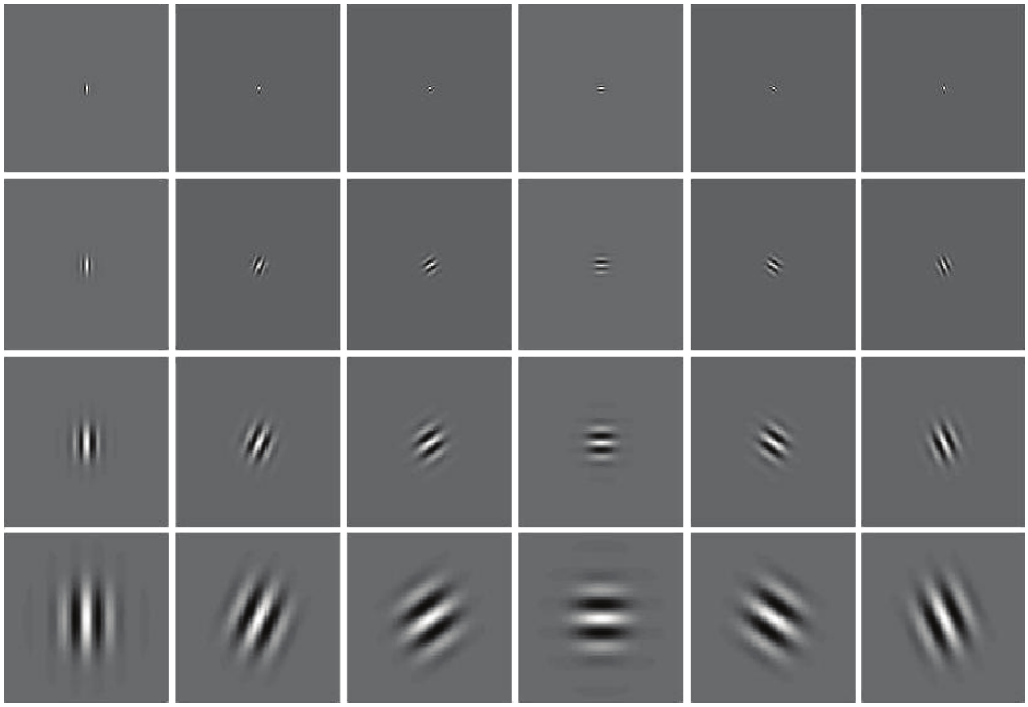

**Figure 4.6:** Real parts of Gabor filters at four scales and six orientations. Adapted with permission from Eleyan et al. (2008).

used in image processing and has been applied in the detection of subcellular structures (Olivo-Marin, 2002). The DWT is implemented as a tree-structured filter bank, which outputs three sub-bands at each scale, whereas the DT-CWT is implemented as two trees whose outputs correspond to the real and imaginary parts of the transform coefficients. One of the main drawbacks of the DWT from the point of view of texture analysis is that it does not represent arbitrarily oriented edges efficiently. In the DT-CWT, by contrast, the basis functions are directionally selective in six orientations, as shown in Figure 4.7. This results from the basis functions of the transform being complex-valued (Selesnick et al., 2005).

A filter bank for texture feature extraction outputs a $K$-dimensional feature vector for each pixel, where $K$ is the number of filters in the image. It is often desirable to summarize the features over the image pixels, so as to make them independent of image size, and to reduce the dimensionality of the feature vector. For Gabor features, Manjunath and Ma (1996) used the mean and standard deviation of the magnitude of each filter response. Recently, Zujovic et al. (2013) proposed a set of sub-band statistics for an image similarity metric for image retrieval; this includes means, standard deviations and first-order horizontal and vertical autocorrelations of the filter responses, as well as cross-correlations between responses of filters. Specifically, the cross-correlations are computed between filters with different

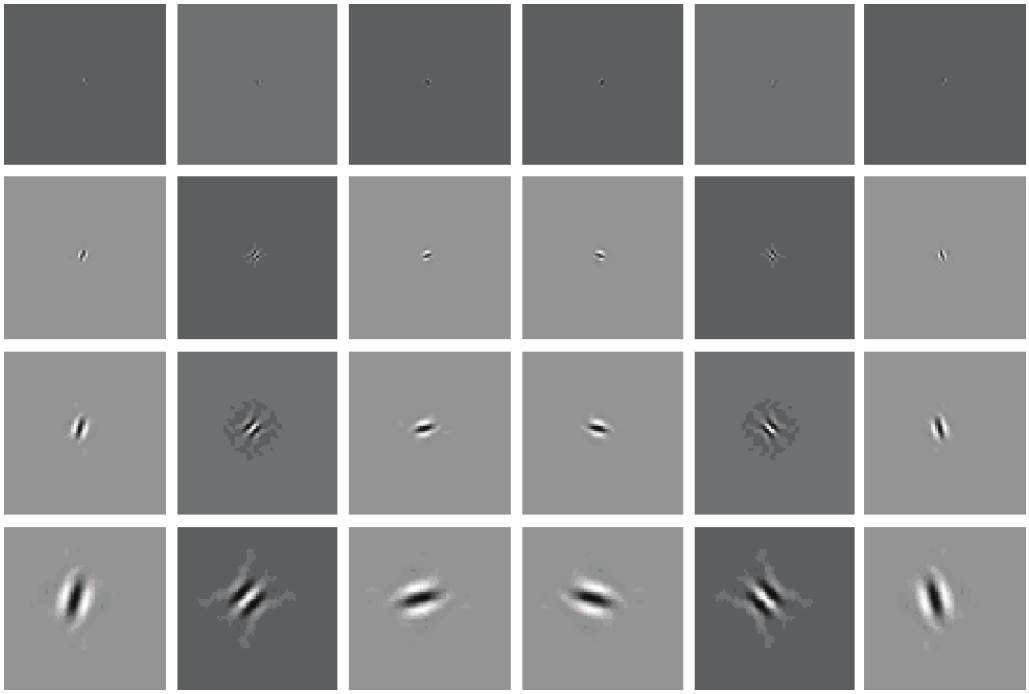

**Figure 4.7:** Real parts of the filters associated with the DT-CWT at four scales and six orientations. Adapted with permission from Eleyan et al. (2008).

scales and the same orientation, and between filters with the same scale and different orientations.

Another approach to summarizing texture features is to assume that for each pixel, the vector of filter responses is a noisy version of one of a relatively small set of prototype vectors. Such prototype vectors are called *textons*, and they can be estimated by applying K-means clustering to the filter responses of pixels in a training data set (Leung and Malik, 1999). In a test image, each pixel can then be assigned to the nearest texton via vector quantization, and the texture of the image can be described with the histogram of textons composing the image. This approach has been successfully applied in, for example, screening mammography images for malignant masses (Li et al., 2015).

### 4.3.3 Applications of machine learning in high content imaging

Machine learning methods are key components of high content imaging experiments (Shariff et al., 2010). From images of cells or cell populations, supervised learning can be used to identify, for example, cell cycle states or dead cells (Walter et al., 2010), protein localization patterns (Newberg et al., 2009), or abnormal morphologies (D'Ambrosio and Vale, 2010). On the other hand, unsupervised learning can be used to identify distinct phenotypes without manual labeling. For example, Bakal et al. (2007) identified cellular signaling networks by clustering

cellular phenotypes, as quantified by a large set of morphometric features. Peng et al. (2011) used clustering to identify a set of subgroups of mitochondria based on morphometric features and texture features. Using the same features, the authors then trained an SVM classifier to classify segmented mitochondria into these subgroups. This allowed using the histogram of subtype counts in a cell as a compact descriptor of the overall mitochondrial morphology in the cell. Another study (Zhong et al., 2012) proposed a method to identify cell cycle states from time series of cell images by a temporal clustering of feature vectors extracted from the images.

In Publication III, we used DT-CWT texture features in conjunction with morphometric features to predict the degree of fragmentation of mitochondria from two-photon microscopy images of the mouse neocortex. The texture features were defined as summary statistics of the different filter responses (Zujovic et al., 2013). Both feature sets explained some of the variance in the degree of fragmentation, but we found that the texture features were more informative with respect to the degree of fragmentation. Due to the low SNR in the images, we also trained a SVM classifier to exclude locations where mitochondrial fragmentation was difficult to assess. Similar methods have been previously proposed for characterizing mitochondrial shape in single-cell images: Reis et al. (2012) classified mitochondria in cells as fragmented, networked or swollen, based on a set of morphometric features and Zernike texture features. Lin et al. (2010) quantified mitochondrial fragmentation using a nearest neighbor regression method trained on morphometric features and Haralick features. Although our method was developed with tissue images in mind, it was shown to perform better than that of Lin et al. (2010) on single-cell images in III.

## 4.4 Validation of methods using simulated microscopy images

A common problem in developing methods for microscopy image analysis is that no objective ground truth is available; traditionally, methods have been validated against manual analysis by human experts, which is suspectible to human error. An attractive alternative that is often used is to algorithmically generate images that simulate microscope images. However, this requires careful design of the simulators to produce realistic images, in order that the results can be generalized to real images. In particular, it is important to avoid *inverse crimes*, that is, the use of the underlying assumptions of the analysis method to guide the simulation. Some recent work has focused on creating general tools for image simulation: (Lindén et al., 2016) developed a tool that simulates fluorescence microscope imaging of single molecules. The tool combines a complex model of molecule diffusion inside a cell, and realistic models of noise and image formation. Other simulators have focused on creating images of cell populations, both static (Lehmussola et al., 2007) and dynamic with divisions (Ulman et al., 2015). The

CellOrganizer tool (Murphy, 2016) generates three-dimensional models of cells and organelles by learning generative models for object shape and subcellular organization from images.

In Publications II and III, we generated simulated movies of fluorescently labeled mitochondria by modeling the mitochondria filaments as splines. The control points of the splines were subjected to random movement, constrained to maintain a level of smoothness in the filament shape, and to directed displacements that were equal for each control point in a filament. Figure 4.8 shows example frames from a simulated movie. It should be noted that this model is simplified in that it does not simulate the mitochondrial network structure nor the cell environment. However, the modeled filaments are allowed to overlap, resulting in images with apparent network-like structures. Thus, methods that are based on segmentation from static images can be validated using these images. For validating methods to analyze mitochondrial fusion and fission, the modeling of mitochondrial dynamics should be extended with fusion-fission events.

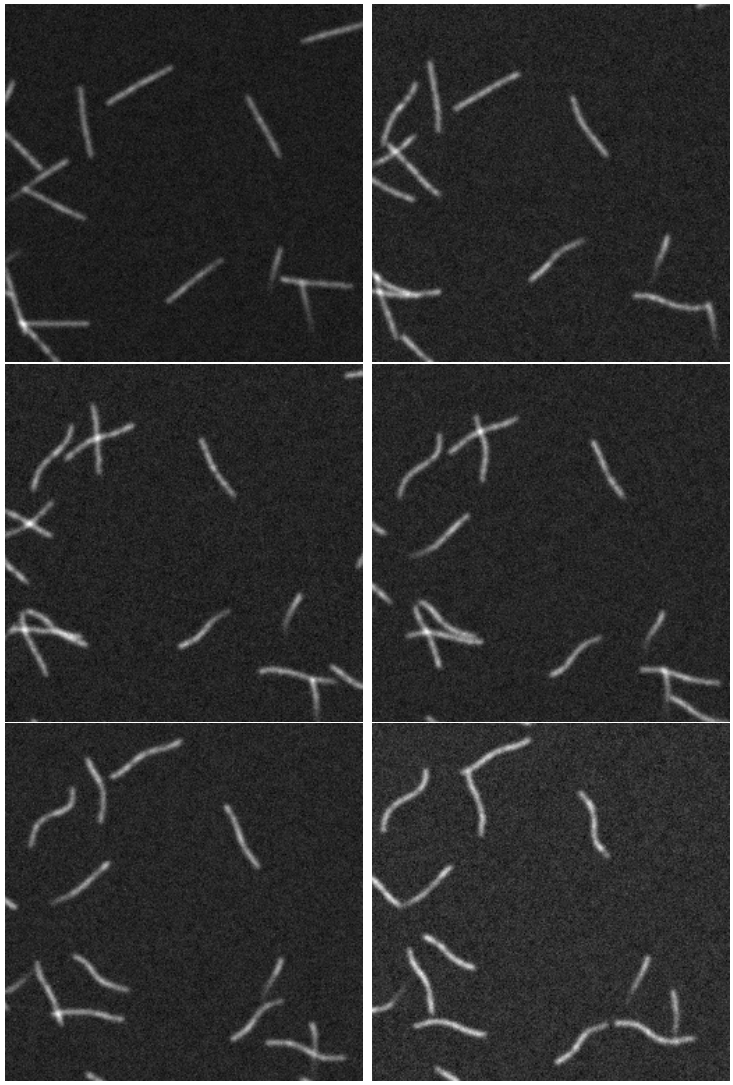

**Figure 4.8:** From left to right and top to bottom: six frames from a simulated movie which was used for validation in Publication II.

# 5 Summary of the results

This study has produced methods and tools for the analysis of images of mitochondria, obtained both from single-cell microscopy and from two-photon imaging of tissues. Our novel methods are expected to have numerous applications in basic and applied research on the morphology and dynamics of mitochondria, due to the widespread use of these imaging techniques to study the level of health of the cells and the effects of perturbations on the mitochondria structure and dynamics.

First, in Publication I, we proposed a software tool, Mytoe, for analyzing mitochondrial morphology and motion from time-lapse fluorescence microscopy images of single cells with fluorescently labeled mitochondria. The segmentation method is based on a previous approach (Koopman et al., 2006); for the quantification of morphological features we introduced a graph-based approach where the segmented objects are decomposed into a graph structure from which properties of individual edges, or branches, can be computed separately. Finally, the motion analysis is based on optical flow estimation following the Lucas-Kanade method. The analysis was tested by introducing simulated motion on real images, and it showed that our methods are robust to image degradation by noise.

In Publication II, we presented a novel method for detecting mitochondrial tips. The method is based on supervised learning techniques, and does not depend on a pre-segmentation of the mitochondria. This study was motivated by the need for a method for tracking mitochondria, to obtain a more detailed analysis of their motion that can be achieved using optical flow estimation, which only provides measurements of instantaneous velocities. The main advantage of a segmentation-free approach is that mitochondria that move outside the focus plane are not mistakenly detected as tips.

The tip detection method is based on a random forest classifier that classifies a small image patch as positive (tip) or negative (non-tip), and it was trained on manually selected positive and negative samples, as well as on randomly selected negative samples. We demonstrated that this novel method can detect the tips of fluorescently labeled mitochondria with better accuracy than a segmentation-based approach. We then applied a particle tracking method (Jaqaman et al., 2008) to track the detected tips; testing our methodology on movies of U2OS cells treated with Nocodazole, which disrupts the microtubules thus affecting mitochondrial

motility, we were able to measure the expected decrease in the mitochondrial speed, in agreement with the literature on the subject.

Finally, the new method proposed in Publication III allows the quantification of the degree of fragmentation in mitochondria from in vivo two-photon fluorescence microscope images of cell tissues. This method assigns a score between zero and one for a small patch extracted from an image of a tissue. The quantification is based on a Beta regression model that was fitted to scores, which were given by multiple experts for a set of training images. As features for the model, we used morphometric features based on the lengths and widths of segmented mitochondria, as well as texture features based on the DT-CWT. Furthermore, an SVM classifier was trained on the texture features to identify locations where the degree of fragmentation could not be reliably assessed – as measured by the percentage of experts that left the region unscored. Such locations can correspond to, e.g., blood vessels or too sparsely distributed mitochondria.

We found that the best predictions were given by the model which used both morphometric features and texture features. Of these, texture features were more informative. We first validated the method on synthetic images, and showed that it performs well on a wide range of image resolutions and signal-to-noise ratios. We then applied the method to images of mice undergoing cardiac arrest, and demonstrated that the method is sufficiently sensitive in order to detect the changes that cardiac arrest causes in mitochondrial morphology.

# 6 Conclusions and discussion

This thesis focused on developing tools and methods for the automatic analysis of the morphology and dynamics of mitochondria from fluorescence microscopy images. Recent years have witnessed increasing interest for analyzing properties of mitochondria from microscope images, and the contributions of this thesis are expected to be of use in such studies. Publication III showed this, by verifying that our methods were able to distinguish between the structure of mitochondria in cells prior to and following a cardiac arrest.

In the future, the method proposed in Publication II could be used to guide the segmentation and tracking of mitochondria with the SOAC method (Smith et al., 2010), which would allow a more complete characterization of mitochondrial dynamics. This could be useful, for example, in studying the interactions of mitochondria with the cytoskeleton.

Meanwhile, the method proposed in Publication III could be applied for the detection and quantification of the degree of mitochondrial fragmentation that is known to occur following various stresses and pathologies. Here, we used it for assessing the level of damage to neurons due to cardiac arrest. Furthermore, we expect the approach proposed here to be applicable to quantify the fragmentation of curvilinear structures in general.

All the methods proposed in this thesis were developed to be applied for two-dimensional images, although both confocal microscopy and two-photon microscopy are capable of producing three-dimensional images of cells and subcellular structures. This limits the applications of the methods to images where the mitochondrial structure is approximately two-dimensional, or to situations where the two-dimensional optical sections are sufficiently informative about the problem at hand. Nevertheless, the methods in Publication I are easy to extend to three dimensions, as the individual filtering and processing steps are also defined for three-dimensional images. Indeed, similar methods have been published recently that work well on three-dimensional images (Nikolaisen et al., 2014; Viana et al., 2015).

Meanwhile, extending the machine learning based methods of Publications II and III to three dimensions would not be as straightforward, as the dimensions of the feature spaces would grow. This would likely require considerably more training

data, and possibly the application of dimensionality reduction techniques. In the method of Publication III, the texture feature extraction could be based on a three-dimensional version of the DT-CWT. One interesting avenue for research in this regard would be the application of three-dimensional convolutional neural networks, which lack the need for feature extraction, but require large amounts of training data to avoid overfitting (Ji et al., 2013).

Another limitation that is present in all current methods for segmenting mitochondria is that when using standard fluorescent labels and diffraction-limited fluorescence microscopy, it is not possible to identify fused mitochondria from ones that are separate but near each other, at least without analyzing their motion. Furthermore, it is known that mitochondria can fuse partially by fusing only their outer membranes (Liu and Murphy, 2009). Since the resolution of a confocal microscope is at best approximately 100 $nm$ with deconvolution (Sedarat et al., 2004), whereas the distance between the inner and outer membranes is 40 $nm$ (Walther et al., 2009), present segmentation results are necessarily ambiguous about the connectivity of the mitochondrial network. One current solution for this problem that allows addressing specific research questions in this regard, is to make use of experimental setups that employ multiple fluorescent labels specifically setup for visualizing the mixing of mitochondrial contents, which is interpreted as fusion.

Although significant progress has been made in the recent years, mitochondrial image analysis is not a solved problem. In particular, to our knowledge, no method has demonstrated accurate tracking of the whole mitochondrial network, particularly the detection of mitochondrial fusion and fission events. Furthermore, mitochondrial appearance is dependent on the type of the cell, and the results of one method are not necessarily generalizable to other cell types than the few ones they were tested on.

We expect the solutions to these problems to arise in the near future. Namely, they will emerge from the combined efforts of improvement of super-resolution imaging techniques, from which more detailed images of mitochondria will result, improvement of fluorescent labeling techniques, and development of novel, computationally more efficient, image analysis methods.

# Bibliography


Ahlqvist, K. J., Suomalainen, A., and Hämäläinen, R. H., "Stem cells, mitochondria and aging," *Biochimica et Biophysica Acta (BBA) - Bioenergetics*, vol. 1847, no. 11, pp. 1380–1386, 2015.

Alberts, B., Johnson, A., Lewis, J., Raff, M., Roberts, K., and Walter, P., *Molecular biology of the cell*, 5th ed. New York: Garland Science, 2008.

Allalou, A., Pinidiyaarachchi, A., and Wählby, C., "Robust signal detection in 3d fluorescence microscopy," *Cytometry Part A*, vol. 77A, no. 1, pp. 86–96, 2010.

Andersen, J. B., Sternberg, C., Poulsen, L. K., Bjorn, S. P., Givskov, M., and Molin, S., "New unstable variants of green fluorescent protein for studies of transient gene expression in bacteria." *Applied and environmental microbiology*, vol. 64, no. 6, pp. 2240–6, 1998.

Anscombe, F. J., "The transformation of poisson, binomial and negative-binomial data," *Biometrika*, vol. 35, no. 3/4, pp. 246–254, 1948.

Bakal, C., Aach, J., Church, G., and Perrimon, N., "Quantitative morphological signatures define local signaling networks regulating cell morphology." *Science*, vol. 316, no. 5832, pp. 1753–6, 2007.

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

Chan, T. F. and Vese, L. A., "Active contours without edges," *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 266–277, 2001.

Chen, H. and Chan, D. C., "Mitochondrial dynamics–fusion, fission, movement, and mitophagy–in neurodegenerative diseases." *Human molecular genetics*, vol. 18, no. R2, pp. R169–76, oct 2009.

Chen, S., Tran, S., Sigler, A., and Murphy, T. H., "Automated and quantitative image analysis of ischemic dendritic blebbing using in vivo 2-photon microscopy data," *Journal of Neuroscience Methods*, vol. 195, no. 2, pp. 222–231, 2011.

Chowdhury, S., Kandhavelu, M., Yli-Harja, O., and Ribeiro, A. S., "Cell segmentation by multi-resolution analysis and maximum likelihood estimation (MAMLE)." *BMC bioinformatics*, vol. 14 Suppl 1, no. Suppl 10, p. S8, 2013.

Cortes, C. and Vapnik, V., "Support-Vector Networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

Cover, T. M., "Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition," *IEEE Transactions on Electronic Computers*, pp. 326–334, 1965.

Dabov, K., Foi, A., and Egiazarian, K., "Video denoising by sparse 3D transform-domain collaborative filtering," *European Signal Processing Conference*, vol. 16, no. 8, pp. 145–149, 2007.

D'Ambrosio, M. V. and Vale, R. D., "A whole genome RNAi screen of Drosophila S2 cell spreading performed using automated computational image analysis." *The Journal of cell biology*, vol. 191, no. 3, pp. 471–8, 2010.

Debray, F.-G., Lambert, M., and Mitchell, G. a., "Disorders of mitochondrial function." *Current opinion in pediatrics*, vol. 20, pp. 471–482, 2008.

Eleyan, A., Özkaramanli, H., and Demirel, H., "Complex Wavelet Transform-Based Face Recognition," *EURASIP Journal on Advances in Signal Processing*, vol. 2008, pp. 1–14, 2008.

Ferrari, S. and Cribari-Neto, F., "Beta Regression for Modelling Rates and Proportions," *Journal of Applied Statistics*, vol. 31, no. 7, pp. 799–815, 2004.

Fischer, F., Hamann, A., and Osiewacz, H. D., "Mitochondrial quality control: An integrated network of pathways," *Trends in Biochemical Sciences*, vol. 37, no. 7, pp. 284–292, 2012.

Fisher, R. A., "The use of multiple measurements in taxonomic problems," *Annals of Eugenics*, vol. 7, no. 2, pp. 179–188, 1936.

Fornasiero, E. F. and Opazo, F., "Super-resolution imaging for cell biologists: Concepts, applications, current challenges and developments." *BioEssays : news and reviews in molecular, cellular and developmental biology*, vol. 37, no. 4, pp. 436–51, 2015.

Frangi, A. F., Niessen, W. J., Vincken, K. L., and Viergever, M. a., "Multiscale vessel enhancement filtering," *Medical Image Computing and Computer-Assisted Invervention - MICCAI'98. Lecture Notes in Computer Science, vol 1496*, vol. 1496, pp. 130–137, 1998.

Frigault, M. M., Lacoste, J., Swift, J. L., and Brown, C. M., "Live-cell microscopy - tips and tools." *Journal of cell science*, vol. 122, no. Pt 6, pp. 753–767, 2009.

Geman, S. and Geman, D., "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images." *IEEE transactions on pattern analysis and machine intelligence*, vol. 6, no. 6, pp. 721–41, jun 1984.

Gerencser, A. A. and Nicholls, D. G., "Measurement of instantaneous velocity vectors of organelle transport: mitochondrial transport and bioenergetics in hippocampal neurons." *Biophysical journal*, vol. 95, no. 6, pp. 3079–99, 2008.

Golding, I., Paulsson, J., Zawilski, S. M., and Cox, E. C., "Real-time kinetics of gene activity in individual bacteria." *Cell*, vol. 123, no. 6, pp. 1025–36, dec 2005.

Gonzalez, R. C. and Woods, R. E., *Digital image processing (3rd edition).* Upper saddle river, New Jersey: Prentice-Hall, Inc., 2006.

Guo, Z. and Hall, R. W., "Parallel thinning with two-subiteration algorithms," *Communications of the ACM*, vol. 32, no. 3, pp. 359–373, 1989.

Haralick, R. M., Shanmugam, K., and Dinstein, I., "Textural Features for Image Classification," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 3, no. 6, pp. 610–621, nov 1973.

Hastie, T., Tibshirani, R., and Friedman, J., *The Elements of Statistical Learning*, ser. Springer Series in Statistics. New York, NY: Springer New York, 2009.

Hoerl, A. E. and Kennard, R. W., "Ridge Regression: Biased Estimation for Nonorthogonal Problems," *Technometrics*, vol. 12, no. 1, pp. 55–67, 1970.

Horn, B. K. and Schunck, B. G., "Determining optical flow," *Artificial Intelligence*, vol. 17, no. 1-3, pp. 185–203, 1981.

J. A. Nelder, R. W. M. W., "Generalized linear models," *Journal of the Royal Statistical Society. Series A (General)*, vol. 135, no. 3, pp. 370–384, 1972.

Jaqaman, K., Loerke, D., and Mettlen, M., "Robust single-particle tracking in live-cell time-lapse sequences," *Nature methods*, vol. 5, no. 8, pp. 695–702, 2008.

Ji, S., Yang, M., Yu, K., and Xu, W., "3D convolutional neural networks for human action recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 221–31, 2013.

Jonker, R. and Volgenant, A., "A shortest augmenting path algorithm for dense and sparse linear assignment problems," *Computing*, vol. 38, no. 4, pp. 325–340, 1987.

Kass, M., Witkin, A., and Terzopoulos, D., "Snakes: Active contour models," *International journal of computer vision*, vol. 331, pp. 321–331, 1988.

Katajisto, P., Döhla, J., Chaffer, C. L., Pentinmikko, N., Marjanovic, N., Iqbal, S., Zoncu, R., Chen, W., Weinberg, R. A., and Sabatini, D. M., "Asymmetric apportioning of aged mitochondria between daughter cells is required for stemness," *Science*, vol. 348, no. 6232, pp. 340–343, 2015.

Koopman, W. J. H., Visch, H.-J., Smeitink, J. A. M., and Willems, P. H. G. M., "Simultaneous quantitative measurement and automated analysis of mitochondrial morphology, mass, potential, and motility in living human skin fibroblasts," *Cytometry Part A*, vol. 69A, no. 1, pp. 1–12, 2006.

Krizhevsky, A., Sulskever, I., and Hinton, G. E., "ImageNet classification with deep convolutional neural networks," *Advances in Neural Information and Processing Systems (NIPS)*, pp. 1–9, 2012.

Lecellier, F., Fadili, J., Jehan-Besson, S., Aubert, G., Revenu, M., and Saloux, E., "Region-based active contours with exponential family observations," *Journal of Mathematical Imaging and Vision*, vol. 36, no. 1, pp. 28–45, 2009.

Lehmussola, A., Ruusuvuori, P., Selinummi, J., Huttunen, H., and Yli-Harja, O., "Computational framework for simulating fluorescence microscope images with cell populations," *IEEE Transactions on Medical Imaging*, vol. 26, no. 7, pp. 1010–1016, 2007.

Leung, T. and Malik, J., "Representing and recognizing the visual appearence of materials using three-dimensional textons," *International Journal of Computer Vision*, vol. 43, no. 1, pp. 29–44, 1999.

Li, Y., Chen, H., Rohde, G. K., Yao, C., and Cheng, L., "Texton analysis for mass classification in mammograms," *Pattern Recognition Letters*, vol. 52, pp. 87–93, 2015.

Lin, Y., Lin, C., Tsai, Y., and Ku, T., "A spectral graph theoretic approach to quantification and calibration of collective morphological differences in cell images," *Bioinformatics*, vol. 26, no. 12, pp. i29–i37, 2010.

Lindén, M., Ćurić, V., Boucharin, A., Fange, D., and Elf, J., "Simulated single molecule microscopy with SMeagol," *Bioinformatics*, p. btw109, 2016.

Liu, R. R. and Murphy, T. H., "Reversible cyclosporin A-sensitive mitochondrial depolarization occurs within minutes of stroke onset in mouse somatosensory cortex in vivo: a two-photon imaging study." *The Journal of biological chemistry*, vol. 284, no. 52, pp. 36 109–17, 2009.

Lucas, B. D. and Kanade, T., "An iterative image registration technique with an application to stereo vision," in *Proceedings of the 1981 DARPA Image Understanding Workshop*, vol. 130, 1981, pp. 121–130.

Magidson, V. and Khodjakov, A., "Circumventing photodamage in live-cell microscopy," *Methods in Cell Biology*, vol. 114, pp. 545–560, 2013.

Mäkitalo, M. and Foi, A., "Optimal inversion of the generalized anscombe transformation for poisson-gaussian noise," *IEEE Transactions on Image Processing*, vol. 22, no. 1, pp. 91–103, Jan 2013.

Manjunath, B. and Ma, W., "Texture features for browsing and retrieval of image data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 837–842, 1996.

McClatchey, P. M., Keller, A. C., Bouchard, R., Knaub, L. A., and Reusch, J. E. B., "Fully automated software for quantitative measurements of mitochondrial morphology." *Mitochondrion*, vol. 26, pp. 58–71, 2015.

Mingqiang, Y., Kidiyo, K., and Joseph, R., "A Survey of Shape Feature Extraction Techniques," in *Pattern Recognition Techniques, Technology and Applications*. InTech, 2008, pp. 43–90.

Misgeld, T., Kerschensteiner, M., Bareyre, F. M., Burgess, R. W., and Lichtman, J. W., "Imaging axonal transport of mitochondria in vivo." *Nature methods*, vol. 4, no. 7, pp. 559–561, 2007.

Mortensen, K. I., Churchman, L. S., Spudich, J. A., and Flyvbjerg, H., "Optimized localization analysis for single-molecule tracking and super-resolution microscopy," *Nature Methods*, vol. 7, no. 5, pp. 377–381, may 2010.

Mumford, D. and Shah, J. M., "Optimal approximations by piecewise smooth functions and associated variational problems," *Comm. Pure Appl. Math.*, vol. 42, pp. 577–685, 1989.

Murphy, R. F., "Building cell models and simulations from microscope images," *Methods*, vol. 96, pp. 33 – 39, 2016, high-throughput Imaging.

Murtagh, F., Starck, J.-L., and Bijaoui, A., "Image restoration with noise suppression using a multiresolution support." *Astronomy and Astrophysics, Suppl. Ser*, vol. 112, p. 179, Jul. 1995.

Newberg, J. Y., Li, J., Rao, A., Pontén, F., Uhlén, M., Lundberg, E., and Murphy, R. F., "Automated analysis of human protein atlas immunofluorescence images," *Proceedings - 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, ISBI 2009*, pp. 1023–1026, 2009.

Nikolaisen, J., Nilsson, L. I. H., Pettersen, I. K. N., Willems, P. H. G. M., Lorens, J. B., Koopman, W. J. H., and Tronstad, K. J., "Automated Quantification and Integrative Analysis of 2D and 3D Mitochondrial Shape and Network Properties," *PLoS ONE*, vol. 9, no. 7, p. e101365, jul 2014.

Norman, K., "Techniques: Intravital microscopy – a method for investigating disseminated intravascular coagulation?" *Trends in Pharmacological Sciences*, vol. 26, no. 6, pp. 327–332, 2005.

Obara, B., Fricker, M., and Grau, V., "Local phase approaches to extract biomedical networks," in *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2012, pp. 1796–1799.

Obara, B., Grau, V., and Fricker, M. D., "A bioimage informatics approach to automatically extract complex fungal networks," *Bioinformatics*, vol. 28, no. 18, pp. 2374–2381, 2012.

Olivo-Marin, J. C., "Extraction of spots in biological images using multiscale products," *Pattern Recognition*, vol. 35, no. 9, pp. 1989–1996, 2002.

Otsu, N., "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.

Paul, G., Cardinale, J., and Sbalzarini, I. F., "Coupling image restoration and segmentation: A generalized linear model/bregman perspective," *International Journal of Computer Vision*, vol. 104, no. 1, pp. 69–93, 2013.

Payne, B. A. I. and Chinnery, P. F., "Mitochondrial dysfunction in aging: Much progress but many unresolved questions," *Biochimica et Biophysica Acta - Bioenergetics*, vol. 1847, no. 11, pp. 1347–1353, 2015.

Peng, J.-Y., Lin, C.-C., and Hsu, C.-N., "Adaptive Image Enhancement for Fluorescence Microscopy," in *2010 International Conference on Technologies and Applications of Artificial Intelligence*, no. 1. IEEE, nov 2010, pp. 9–16.

Peng, J.-Y., Lin, C.-C., Chen, Y.-J., Kao, L.-S., Liu, Y.-C., Chou, C.-C., Huang, Y.-H., Chang, F.-R., Wu, Y.-C., Tsai, Y.-S., and Hsu, C.-N., "Automatic morphological subtyping reveals new roles of caspases in mitochondrial dynamics." *PLoS computational biology*, vol. 7, no. 10, p. e1002212, oct 2011.

Rafelski, S. M., "Mitochondrial network morphology: building an integrative, geometrical view." *BMC biology*, vol. 11, no. 1, p. 71, 2013.

Reid, D., "An algorithm for tracking multiple targets," *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.

Reis, Y., Bernardo-Faura, M., Richter, D., Wolf, T., Brors, B., Hamacher-Brady, A., Eils, R., and Brady, N. R., "Multi-parametric analysis and modeling of relationships between mitochondrial morphology and apoptosis," *PLoS ONE*, vol. 7, no. 1, pp. 1–19, 01 2012.

Rizk, A., Paul, G., Incardona, P., Bugarski, M., Mansouri, M., Niemann, A., Ziegler, U., Berger, P., and Sbalzarini, I. F., "Segmentation and quantification of subcellular structures in fluorescence microscopy images using Squassh," *Nature Protocols*, vol. 9, no. 3, pp. 586–596, feb 2014.

Ronneberger, O., Fischer, P., and Brox, T., "U-Net: Convolutional networks for biomedical image segmentation," ser. Lecture Notes in Computer Science, Navab, N., Hornegger, J., Wells, W. M., and Frangi, A. F., Eds. Cham: Springer International Publishing, 2015, vol. 9350, pp. 234–241.

Rosenblatt, F., "The perceptron: A probabilistic model for information storage and organization in the brain." *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1958.

Ruusuvuori, P., Aijö, T., Chowdhury, S., Garmendia-Torres, C., Selinummi, J., Birbaumer, M., Dudley, A. M., Pelkmans, L., and Yli-Harja, O., "Evaluation of methods for detection of fluorescence labeled subcellular objects in microscope images." *BMC bioinformatics*, vol. 11, p. 248, 2010.

Saunter, C. D., Perng, M. D., Love, G. D., and Quinlan, R. A., "Stochastically determined directed movement explains the dominant small-scale mitochondrial movements within non-neuronal tissue culture cells," *FEBS Letters*, vol. 583, no. 8, pp. 1267–1273, 2009.

Sbalzarini, I. F. and Koumoutsakos, P., "Feature point tracking and trajectory analysis for video imaging in cell biology." *Journal of structural biology*, vol. 151, no. 2, pp. 182–95, 2005.

Schölkopf, B., Smola, A. J., Williamson, R. C., and Bartlett, P. L., "New support vector algorithms," *Neural Comput.*, vol. 12, no. 5, pp. 1207–1245, May 2000.

Sedarat, F., Lin, E., Moore, E. D., and Tibbits, G. F., "Deconvolution of confocal images of dihydropyridine and ryanodine receptors in developing cardiomyocytes," *Journal of Applied Physiology*, vol. 97, no. 3, pp. 1098–1103, 2004.

Selesnick, I. W., Baraniuk, R. G., and Kingsbury, N. G., "The dual-tree complex wavelet transform," *Signal Processing Magazine*, vol. 22, no. 6, pp. 123–151, 2005.

Sezgin, M. and Sankur, B., "Survey over image thresholding techniques and quantitative performance evaluation," *Journal of Electronic Imaging*, vol. 13, no. 1, p. 146, 2004.

Shariff, A., Kangas, J., Coelho, L. P., Quinn, S., and Murphy, R. F., "Automated image analysis for high-content screening and analysis." *Journal of biomolecular screening*, vol. 15, no. 7, pp. 726–34, 2010.

Shaw, P. J., "Comparison of Widefield Deconvolution and Confocal Microscopy for Three-Dimensional Imaging," in *Handbook Of Biological Confocal Microscopy*, Pawley, J., Ed. New York City, United States: Springer US, 2006, pp. 453–467.

Sheng, Z.-H. and Cai, Q., "Mitochondrial transport in neurons: impact on synaptic homeostasis and neurodegeneration." *Nature reviews Neuroscience*, vol. 13, no. 2, pp. 77–93, 2012.

Smal, I., Grigoriev, I., Akhmanova, A., Niessen, W. J., and Meijering, E., "Microtubule Dynamics Analysis Using Kymographs and Variable-Rate Particle Filters," *IEEE Transactions on Image Processing*, vol. 19, no. 7, pp. 1861–1876, 2010.

Smith, M. B., Li, H., Shen, T., Huang, X., Yusuf, E., and Vavylonis, D., "Segmentation and tracking of cytoskeletal filaments using open active contours." *Cytoskeleton*, vol. 67, no. 11, pp. 693–705, 2010.

Sommer, C., Straehle, C., Köthe, U., and Hamprecht, F. A., "Ilastik: Interactive learning and segmentation toolkit," in *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2011, pp. 230–233.

Steger, C., "An unbiased detector of curvilinear structures," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 2, pp. 113–125, 1998.

Stephens, D. J., "Light Microscopy Techniques for Live Cell Imaging," *Science*, vol. 300, no. 5616, pp. 82–86, 2003.

Suen, D.-F., Norris, K. L., and Youle, R. J., "Mitochondrial dynamics and apoptosis," *Genes & Development*, vol. 22, no. 12, pp. 1577–1590, jun 2008.

Tibshirani, R., "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.

Trifunovic, A., Wredenberg, A., Falkenberg, M., Spelbrink, J. N., Rovio, A. T., Bruder, C. E., Bohlooly-Y, M., Gidlöf, S., Oldfors, A., Wibom, R., Törnell, J., Jacobs, H. T., and Larsson, N.-G., "Premature ageing in mice expressing defective mitochondrial DNA polymerase," *Nature*, vol. 429, no. 6990, pp. 417–423, 2004.

Ulman, V., "Improving Accuracy of Optical Flow of Heeger's Original Method on Biomedical Images," 2010, pp. 263–273.

Ulman, V., Orémuš, Z., and Svoboda, D., "TRAgen: A Tool for Generation of Synthetic Time-Lapse Image Sequences of Living Cells," in *Image Analysis and Processing - ICIAP 2015, Pt I*, 2015, vol. 9279, pp. 623–634.

Vapnik, V. N., "An overview of statistical learning theory," *Trans. Neur. Netw.*, vol. 10, no. 5, pp. 988–999, 1999.

Vapnik, V. N. and Lerner, A. Y., "Recognition of patterns with help of generalized portraits," *Avtomat. i Telemekh*, vol. 24, no. 6, pp. 774–780, 1963.

Viana, M. P., Lim, S., and Rafelski, S. M., "Quantifying mitochondrial content in living cells," in *Methods in Cell Biology*. Elsevier Ltd, 2015, vol. 125, pp. 77 – 93.

Vliet, L. V., Sudar, D., and Young, I., "Digital fluorescence imaging using cooled charge-coupled device array cameras," in *Cell Biology, Second Edition*, Celis, J. E., Ed. New York: Academic Press, 1998, vol. III, no. Castleman 1996, pp. 109 – 120.

Vowinckel, J., Hartl, J., Butler, R., and Ralser, M., "MitoLoc: A method for the simultaneous quantification mitochondrial network morphology and membrane potential in single cells," *Mitochondrion*, vol. 24, pp. 77 – 86, 2015.

Wäldchen, S., Lehmann, J., Klein, T., van de Linde, S., and Sauer, M., "Light-induced cell damage in live-cell super-resolution microscopy," *Scientific Reports*, vol. 5, p. 15348, 2015.

Walter, T., Held, M., Neumann, B., Hériché, J.-K., Conrad, C., Pepperkok, R., and Ellenberg, J., "Automatic identification and clustering of chromosome phenotypes in a genome wide RNAi screen by time-lapse imaging." *Journal of structural biology*, vol. 170, no. 1, pp. 1 – 9, 2010.

Walther, P., Höhn, K., and Krisp, H., "What is the true size of the mitochondrial intermembrane space? a study using high-pressure freezing and stem tomography," in *Proceedings of the First Joint Meeting of the 9th Multinational Congress on Microscopy 2009 and of the 7th Dreiländertagung für Elektronenmikroskopie 2009, Vol. 2: Life Sciences*, 2009.

Westermann, B., "Mitochondrial fusion and fission in cell life and death," *Nature Reviews Molecular Cell Biology*, vol. 11, no. 12, pp. 872 – 884, 2010.

Wu, Z., Kunz, T. H., and Betke, M., "Efficient track linking methods for track graphs using network-flow and set-cover techniques," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2011, pp. 1185 – 1192.

Xiao, X., Geyer, V. F., Bowne-Anderson, H., Howard, J., and Sbalzarini, I. F., "Automatic optimal filament segmentation with sub-pixel accuracy using generalized linear models and b-spline level-sets," *Medical image analysis*, vol. 32, pp. 157 – 172, 2016.

Xu, T., Li, H., Shen, T., Ojkic, N., Vavylonis, D., and Huang, X., "Extraction and analysis of actin networks based on open active contour models," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro, 2011*, Chicago, IL, 2011, pp. 1334 – 1340.

Ylikallio, E. and Suomalainen, A., "Mechanisms of mitochondrial diseases." *Annals of medicine*, vol. 44, no. 1, pp. 41 – 59, 2012.

Zhong, Q., Busetto, A. G., Fededa, J. P., Buhmann, J. M., and Gerlich, D. W., "Unsupervised modeling of cell morphology dynamics for time-lapse microscopy." *Nature methods*, vol. 9, no. 7, pp. 711 – 3, 2012.

Zou, H. and Hastie, T., "Regularization and variable selection via the elastic net," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301 – 320, apr 2005.

Zujovic, J., Pappas, T., and Neuhoff, D., "Structural texture similarity metrics for image analysis and retrieval," *IEEE Transactions on Image Processing*, vol. 22, no. 7, pp. 2545 – 2558, 2013.

# Publications

# Publication I

Eero Lihavainen, Jarno Mäkelä, Johannes N. Spelbrink, Andre S. Ribeiro, "Mytoe: Automatic analysis of mitochondrial dynamics," *Bioinformatics*, vol 28, no. 7, pp. 1050 – 1051 Feb. 2012.

BIOINFORMATICS APPLICATIONS NOTE

Vol. 28 no. 7 2012, pages 1050–1051
doi:10.1093/bioinformatics/bts073

Systems biology

Advance Access publication February 9, 2012

# Mytoe: automatic analysis of mitochondrial dynamics

Eero Lihavainen[1,*], Jarno Mäkelä[1], Johannes N. Spelbrink[2,3] and Andre S. Ribeiro[1,*]

[1]Department of Signal Processing, Tampere University of Technology, [2]Institute of Biomedical Technology, University of Tampere, FI-33520 Tampere, Finland, [3]Department of Pediatrics, Institute for Genetic and Metabolic Disease, Radboud University Nijmegen Medical Centre, P.O. Box 9101, 6500 HB Nijmegen, The Netherlands

Associate Editor: Martin Bishop

**ABSTRACT**

**Summary:** We present Mytoe, a tool for analyzing mitochondrial morphology and dynamics from fluorescence microscope images. The tool provides automated quantitative analysis of mitochondrial motion by optical flow estimation and of morphology by segmentation of individual branches of the network-like structure of the organelles. Mytoe quantifies several features of individual branches, such as length, tortuosity and speed, and of the macroscopic structure, such as mitochondrial area and degree of clustering. We validate the methods and apply them to the analysis of sequences of images of U2OS human cells with fluorescently labeled mitochondria.

**Availability:** Source code, Windows software and Manual available at http://www.cs.tut.fi/%7Esanchesr/mito

**Supplementary information:** Supplementary data are available at *Bioinformatics* online.

**Contact:** eero.lihavainen@tut.fi; andre.ribeiro@tut.fi

Received on December 5, 2011; revised on January 17, 2012; accepted on February 6, 2012

## 1 INTRODUCTION

Mitochondria play a key role in several cellular processes, from energy production to apoptosis and ageing. Malfunctioning in mitochondrial processes has been associated to several diseases (Westermann, 2010). Due to mitochondrial fusion and fission, the organelles form a highly dynamic structure that can change from fragmented to filamentous (Koopman *et al.*, 2005). A better understanding of the dynamics of this structure and its relationship with complex cellular processes may provide much insight on mitochondrial functioning and their role on the well-being of cells.

Recent works addressed the problem of segmenting mitochondria from fluorescence microscope images. (Koopman *et al.*, 2005) presented a segmentation method along with descriptors for mitochondrial properties. Another method uses 3D imaging (Song *et al.*, 2008). Mitochondrial motility has been assessed by the degree of colocalization between successive images and from the differences of subsequent image pairs (Koopman *et al.*, 2005; Yi *et al.*, 2004). These methods detect motion but do not yield data on directionality. Other methods to study motility include optical flow (OF) estimation (Gerencser *et al.*, 2008) and distance transform (Beraud *et al.*, 2009). Tracking methods for motility analysis exist, but rely heavily on human interaction (Saunter *et al.*, 2009; Silberberg *et al.*, 2008).

*To whom correspondence should be addressed.

We present an easy-to-use software, Mytoe, for automated study of mitochondrial structural dynamics from temporal confocal images. It has a simple graphical user interface and requires only a few parameters from the user. Mytoe includes a novel segmentation method which identifies individual branches of the organelles' structure by thresholding and by morphological image processing. In contrast to previous methods, this allows quantifying properties of both single branches and the macroscopic structure formed by the mitochondria. OF is used for motion analysis. The output can be examined in Mytoe with branch-level data being visualized by color-coding individual branches based on the data, and motion vectors being visualized as quiver plots. The results are saved as MATLAB MAT-files and comma-separated values. We first describe the methods in Mytoe and results of the validation procedure. Finally, we use Mytoe to analyze time-lapse sequences of U2OS human osteosarcoma cells.

## 2 METHODS, APPLICABILITY AND VALIDATION

Mytoe segments the nucleus and cell membrane from fluorescence microscope images. The outlines can be manually drawn if the automatic segmentation is unsatisfactory. The mitochondria are segmented in two steps. The first is similar to the method of (Koopman *et al.*, 2005) but uses morphological top-hat (TH) (Soille, 2003) instead of linear filtering. Each image of mitochondria is first denoised with a median filter, followed by TH, which enhances the separation of mitochondria from the background. The result is median filtered to remove noise enhanced by the TH. Finally, the denoised image is contrast stretched and binarized by Otsu's method (Otsu, 1979) to produce a mask. The next step, developed here for our aims, extracts individual branches of the network by applying to the mask a two-iteration thinning procedure (Guo *et al.*, 1989) and locating the branch points of the resulting skeleton.

Motion is analyzed by OF estimation, using a pyramidal implementation (Bouguet, 2000) of the method proposed in (Lucas *et al.*, 1981). OF yields the displacements of objects between each pair of subsequent frames in a movie, enabling the calculation of velocities. Results using this method are provided in the Supplementary Material. If desired, motion can also be analyzed by measuring the colocalization of mitochondria between successive images as proposed in (Koopman *et al.*, 2005).

Mytoe extracts various statistics. From individual branches it computes properties such as thickness, length and orientation. From the OF, the average speed and the average direction of each branch are calculated. In addition, it extracts cell-level quantities such as number of branches and total mitochondrial area. The list of extracted quantities is shown in Table 1. The methods of computation of each feature and their practical implementation are described in Supplementary Material.

We apply Mytoe to confocal images of U2OS cells. Cells were transfected with a vector expressing mitoDsRED2, an red fluorescent protein targeting the mitochondrial matrix. The nuclei and cell membranes are visualized by

**Table 1.** Quantitative features from Mytoe

| | |
|---|---|
| Length | Distance to cell membrane |
| Thickness | Orientation relative to *x*-axis |
| Tortuosity | Orientation relative to centroid |
| Intensity | Number of branches (C) |
| Speed | Number of connected components (C) |
| Direction | Total mitochondrial area (C) |
| Wiggle ratio (Gerencser *et al.*, 2008) | Degree of clustering (C) |
| Distance to centroid | Colocalization (C) |
| Distance to nucleus membrane | |

Cell-level features marked with (C).

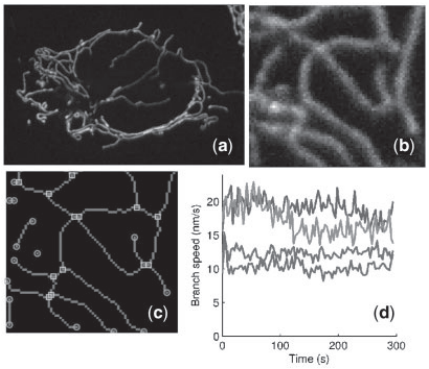


**Fig. 1.** Original image (**a**) and cropped region (**b**). (**c**) The skeleton (grey), branch points (squares) and end points (circles) found from the mask. (**d**) Time series of mean mitochondrial speed in four cells.

Hoechst 33342 and WGA647 fluorescent dyes, respectively. The images were acquired with a Nikon Eclipse Ti with 100× Apo, a Wallac-Perkin Elmer Ultraview spinning-disk confocal system, Andor EMCCD camera and an autofocus system. Each cell was imaged every 3 s for 10 min in 2D. An example image is shown in Figure 1a. In Figures 1b and 1c, we show a region of the same image, and the structure segmented from that region, respectively.

The quality of microscope imaging is degraded by noise and saturation, thus we added noise and increasingly saturated the images (Supplementary Material). We applied Mytoe and observed how the features changed with increasing levels of degradation. Visual inspection showed that the methods are robust to levels of degradation much higher than in real images. In Supplementary Material we present the results of the analysis of structure and motion and of the features in Table 1 for a set of images. As an example, Figure 1d shows time series of mean branch speeds of mitochondria of four cells, revealing that they differ widely, likely due to differing cell cycle stage.

## 3 DISCUSSION

Mytoe provides an easy means to study mitochondrial structure and kinetics by automated image analysis of temporal images of fluorescently labeled mitochondria. The structure analysis was found robust to noise and insensitive to saturation. Thus, the method can be used to obtain reliable measurements from saturated images, provided that the mitochondrial branches are separable.

So far, we tested Mytoe on one cell type. Its efficiency may depend on the shapes of cell type and mitochondria. Here we showed that it is accurate enough to, for example, distinguish the kinetics of the mitochondria structure from one cell to another.

In the future, this framework can be extended by including additional descriptors and adding other shape analysis techniques that can be chosen as a function of the cell type. Additionally, we aim to further develop Mytoe so as to automatically detect abnormalities in the mitochondria structure, due to chemicals (such as in drug screening) or disease.

## ACKNOWLEDGEMENT

We sincerely thank Howard T. Jacobs for advices throughout the project and support.

*Funding*: Work supported by Academy of Finland (A.S.R.) and Finnish Funding Agency for Technology and Innovation (E.L., J.M.). J.N.S. is supported by the Academy of Finland (CoE funding), the Tampere University Hospital Medical Research Fund (9J119, 9K126 and 9L097) and the Netherlands Organization for Scientific Research (NWO: VICI grant 865.10.004).

*Conflict of Interest*: none declared.

## REFERENCES

Beraud,N. *et al.* (2009) Mitochondrial dynamics in heart cells: very low amplitude high frequency fluctuations in adult cardiomyocytes and flow motion in non beating HL-1 cells. *J. Bioenerg. Biomembr.*, **41**, 195–214.

Bouguet,J.-Y. (2000) Pyramidal implementation of the Lucas-Kanade feature tracker. *Technical Report*, Intel Corporation, Microprocessor Research Labs.

Gerencser,A. *et al.* (2008) Measurement of instantaneous velocity vectors of organelle transport: mitochondrial transport and bioenergetics in hippocampal neurons. *Biophys. J.*, **95**, 3079–3099.

Guo,Z. *et al.* (1989) Parallel thinning with two-subiteration algorithms. *Commun. ACM*, **32**, 359–373.

Koopman,W. *et al.* (2005) Simultaneous quantitative measurement and automated analysis of mitochondrial morphology, mass, potential, and motility in living human skin fibroblasts. *Cytometry A*, **69A**, 1–12.

Lucas,B. *et al.* (1981) An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI '81)*, Vancouver, B.C., Canada, pp. 674–679.

Otsu,N. (1979) A threshold selection method from gray level histograms. *IEEE Trans. Syst. Man Cybern.*, **9**, 62–66.

Saunter,C. *et al.* (2009) Stochastically determined directed movement explains the dominant small-scale mitochondrial movements within non-neuronal tissue culture cells. *FEBS Lett.*, **583**, 1267–1273.

Silberberg,Y.R. *et al.* (2008) Tracking displacements of intracellular organelles in response to nanomechanical forces. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium.* Paris, France, pp. 1335–1338.

Soille,P. (2003) *Morphological Image Analysis: Principles and Applications*. 2nd edn. Springer New York, Inc., Secaucus, NJ, USA.

Song,W. *et al.* (2008) Assessing mitochondrial morphology and dynamics using fluorescence wide-field microscopy and 3D image processing. *Methods*, **46**, 295–303.

Westermann,B. (2010) Mitochondrial fusion and fission in cell life and death. *Nat. Rev. Mol. Cell Biol.*, **11**, 872–884.

Yi,M. *et al.* (2004) Control of mitochondrial motility and distribution by the calcium signal: a homeostatic circuit. *J. Cell Biol.*, **167**, 661–72.

# Mytoe

Supplementary material

Eero Lihavainen, Jarno Mäkelä, Johannes N. Spelbrink, Andre S. Ribeiro

December 2, 2011

In this supplement we describe the methods used in Mytoe. Also, we present the results of the validation of the methods and of the application of Mytoe to movies of U2OS cells.

## 1 Image analysis background

### 1.1 Pixel neighbourhood and connectivity

A *neighbourhood* of a pixel $p = (x, y)$ in an image $I$ is a set of pixels close to $p$, including $p$ itself. Two often encountered neighbourhoods are the *4-neighbourhood*

$$I_4 = \begin{bmatrix} \cdot & p_1 & \cdot \\ p_4 & p & p_2 \\ \cdot & p_3 & \cdot \end{bmatrix}$$

and the *8-neighbourhood*

$$I_8 = \begin{bmatrix} p_1 & p_2 & p_3 \\ p_8 & p & p_4 \\ p_7 & p_6 & p_5 \end{bmatrix}.$$

The pixels $p_1, p_2 \ldots$ are referred to as *neighbours* of $p$.

Two pixels $p$ and $q$ are *n-connected* if one belongs to the n-neighbourhood of the other. For example, $p$ and $q$ are 4-*connected*, if $q = p_i$ for some $p_i$ in the 4-neighbourhood of $p$. A set $S$ of pixels is *n-connected*, if, for any $p, q \in S$ there exists a sequence $\{x_1, x_2, \ldots, x_m\}$ of pixels in $S$ such that $p$ is $n$-connected to $x_1$, $x_m$ is $n$-connected to $q$, and $x_i$ is $n$-connected to $x_{i+1}$ for $i < m$. Such $S$ is also referred to as a *connected component*. Connected component *labeling* refers to assigning a unique gray value to each connected component in a binary image. The process yields a *label image* $L$, where $L(x, y)$ has value $k$ if $(x, y)$ belongs to component $k$.

### 1.2 Distance transform

Distance transform is a means to calculate distances between objects. It is a representation of a digital image in which the value of each pixel is the distance from that pixel to a nearest *obstacle pixel*. In binary images, a '1' may be considered an obstacle pixel. In that case

the distance transform gives the shortest distance of each pixel to an object. Formally, the distance transform of image $I(x,y)$ is defined as

$$\mathscr{D}\{I\}(x,y) = \min_{(i,j)\in F} d((x,y),(i,j)), \tag{1}$$

where $F$ is the set of obstacle pixels, and $d(p_1,p_2)$ is a distance function. The distance function $d(p_1,p_2)$ can be any valid metric, but the most useful is the Euclidean distance:

$$d(p_1,p_2) = ||p_1 - p_2||. \tag{2}$$

### 1.3 Thresholding

Thresholding refers to a class of methods where an image is segmented by comparing its intensity values to a *threshold value*. Thresholding an image $I(x,y)$ produces a binary image $B(x,y)$ of the same size as $I(x,y)$, where $B(i,j) = 1$ if the pixel $I(i,j)$ belongs to the foreground, otherwise $B(i,j) = 0$. The threshold value $T$ is selected such that $B(i,j) = 1$ when $I(i,j) > T$ and $B(i,j) = 0$ if $I(i,j) \leq T$. The binary image $B$ is referred to as a *mask*, and *masking* an image $I$ with $B$ refers to the point-wise multiplication of $I$ and $B$. Masking thus produces an image which is zero where $B$ is zero, and is equal to $I$ where $B$ is one.

In automatic image analysis, the threshold $T$ has to be selected automatically. A popular method for automatic threshold selection is Otsu's method. The algorithm searches a threshold value which best separates the intensities of the two classes, the foreground, $C_0$, and the background, $C_1$. Let $I$ be an image represented with values $[1,2,..,L]$, and denote the number of pixels in $I$ with value $i$ with $n_i$, and let $N = n_1 + n_2 + ... + n_L$ be the total number of pixels in $I$. Based on discriminant analysis, the method selects the threshold value $T$ to maximize the *between-class variance* $\sigma_B^2 = \omega_0\omega_1(\mu_1 - \mu_0)^2$, where $\omega_i$ and $\mu_i$ denote the probability of occurrence and mean level of class $i$, respectively, and are calculated as

$$\omega_0 = \sum_{i=1}^{T} p_i, \tag{3}$$

$$\omega_1 = \sum_{i=T+1}^{L} p_i, \tag{4}$$

$$\mu_0 = \sum_{i=1}^{T} \frac{i p_i}{\omega_0} \tag{5}$$

and

$$\mu_1 = \sum_{i=T+1}^{L} \frac{i p_i}{\omega_1}, \tag{6}$$

where $p_i = n_i/N$ specifies the normalized histogram of the image. [1] Due to its assumptions, Otsu's method cannot be expected to work well when the intensity histogram of the image is not bimodal.

We use the following method for computing the areas of objects: For a binary image $I$ with one connected component, we assign a value $n$ for each $4\times4$-neighborhood in $I$, i.e., for each

block of 4 pixels, as follows: if the block contains 0, 1, 3 or 4 ones, the respective values for $n$ are $0, 1, 7/2$ and 4. For blocks with exactly two pixels, $n = 2$, if the pixels are horizontally or vertically adjacent, otherwise, $n = 3$. The total area of the connected component is obtained by summing the values of $n$ and dividing the result by 4, because each pixel contributes to 4 blocks. [2]

## 1.4 Morphological image analysis

### 1.4.1 Morphological operations

The two basic operations in mathematical morphology are *erosion* and *dilation*. Both operate on a binary image $I$ and a *structuring element* (SE) $S$, which is a binary image smaller than $I$, and the images are interpreted as sets. The coordinates of the SE are defined such that a point, usually the center, of the SE is taken to be the origin $(0,0)$, and the SE is said to be at point $(i, j)$ in the image when its origin coincides with the point $(i, j)$. Some common SEs are shown in Figure 1. Erosion is defined as

$$I \ominus S = \bigcap_{s \in S} I_{-s} \tag{7}$$

and dilation

$$I \oplus S = \bigcup_{s \in S} I_{-s}, \tag{8}$$

where $I_{-s}$ denotes a translation of $I$ by the vector $-s$, i.e. $I_{-s}(x) = I(x + s)$. Thus, erosion is the set of points of $I$ where the SE fits inside the image, whereas dilation is the set of points where the SE touches the image (Figure 2). With erosion, one can remove objects where the SE does not fit – for example, objects that are too small, and dilation can be used to connect objects that are close to each other, smooth object boundaries, etc. [3]

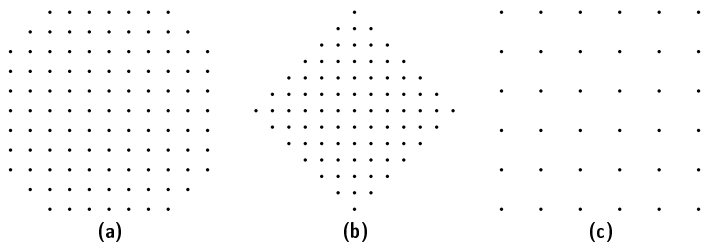


**Figure 1**: Some common structuring elements. The dots represent the pixels of the SE. a) A disk-shaped SE. b) A diamond-shaped SE. c) A square SE.

When using erosion to remove small objects, all of the structures in the image will shrink, which is, in general, irreversible. Morphological *closing* dilates the image after erosion, thus correcting for the shrinkage caused by erosion. It is defined as

$$I \bullet S = (I \ominus S) \oplus S \tag{9}$$

for an image $I$ and a SE $S$. Its dual operation, *opening*, erodes the image after dilation, and is obtained by

$$I \circ S = (I \oplus S) \ominus S. \tag{10}$$

[3]

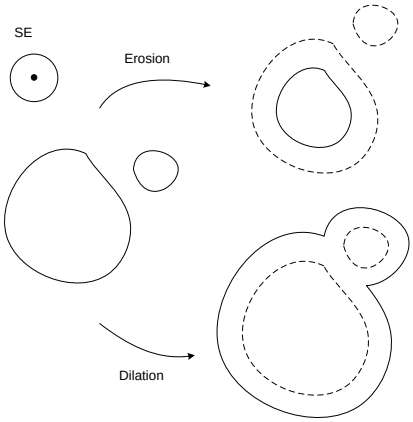


**Figure 2**: Operation of morphological erosion and dilation with a disk-shaped structuring element (SE). The original image, which is shown also using dashed lines, has two separate objects. Erosion removes the smaller of them because the SE does not fit inside it. On the other hand, dilation joins the two objects together.

### 1.4.2 Grayscale morphology

Mathematical morphology generalizes to grayscale digital images by defining erosion and dilation as

$$I \ominus S = \min_{s \in S} I_{-s} \tag{11}$$

and

$$I \oplus S = \max_{s \in S} I_{-s}, \tag{12}$$

respectively. Opening and closing are defined for grayscale images, as they are for binary images, in terms of erosion and dilation. In grayscale morphology, *top-hat filters* can be used to enhance objects of interest by first removing them with an opening or closing, and then recovering them by subtracting the opening or closing from the original image. The *white top-hat* is defined as

$$WTH(I,S) = I - (I \circ S), \tag{13}$$

[3] that is, the difference of $I$ and its closing.

### 1.4.3 Thinning

Morphological *thinning* is a process where foreground pixels are removed from a binary image so as to find a *medial line*, as demonstrated in Figure 3. Such a procedure is also called *skeletonization*, and the medial line is referred to as a *topological skeleton*. A desired property for a skeletonization algorithm is that it is connectivity preserving, i.e. objects are not disconnected, previously disjoint objects are not connected, and objects are not deleted in the process [4].

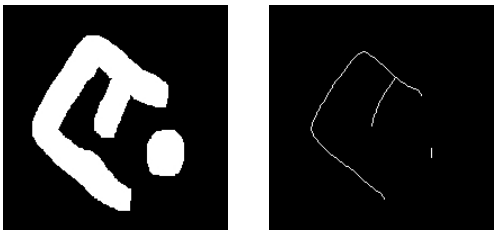


**Figure 3**: Morphological skeleton (b) obtained by thinning the original image (a).

The following parallel algorithm for thinning was presented in [4]. Let $F$ be the set of foreground points in the binary image $I$. The algorithm deletes points from $F$ in two subiterations, until no point can be deleted. The resulting set of points is a thin skeleton. Let $C(p)$ denote the number of distinct 8-connected components in the 8-neighbourhood of $p$, and let us define $N(p) = \min\{N_1(p), N_2(p)\}$, where

$$N_1(p) = (p_1 \vee p_2) + (p_3 \vee p_4) + (p_5 \vee p_6) + (p_7 \vee p_8) \tag{14}$$

and

$$N_2(p) = (p_2 \vee p_3) + (p_4 \vee p_5) + (p_6 \vee p_7) + (p_8 \vee p_1), \tag{15}$$

where $\vee$ is the logical or-operator, and the binary values 1 and 0 are interpreted as true and false, respectively.

## 1.5 Optical flow

Horn and Schunk [5] define optical flow as "the distribution of apparent velocities of movement of brightness patterns in an image". The problem of determining the optical flow between two images can be formulated as finding the displacements $\delta x$ and $\delta y$ in

$$I(x, y, t) = I(x + \delta x, y + \delta y, t + \delta t), \tag{16}$$

where $I$ is a *spatiotemporal image* with spatial coordinates $x$ and $y$, and a temporal coordinate $t$ [5]. Thus, optical flow is essentially the movement of individual pixels, and Eq. 16

**Algorithm 1** Parallel thinning with two subiterations

$i \leftarrow 0$
$F_1 = F$
**repeat**
  $i \leftarrow i + 1$
  **for all** $p \in F_i$ **do**
    **if** $C(p) = 1$ and $2 \leq N(p) \leq 3$ **then**
      **if** $i$ is odd and $(p_2 \vee p_3 \vee \neg p_5) \vee p_4 = 0$ **then**
        $F_i \leftarrow F_{i-1} \cap \{p\}$
      **else if** $i$ is even and $(p_6 \vee p_7 \vee \neg p_1) \wedge p_8 = 0$ **then**
        $F_i \leftarrow F_{i-1} \cap \{p\}$
      **end if**
    **end if**
  **end for**
**until** $F_i = F_{i-1}$

includes the *brightness constancy* assumption, that is, the intensity of the displaced pixel is assumed not to change.

Differential optical flow methods approach the problem by applying to Eq. 16 the first-order Taylor approximation

$$I(x+\delta x, y+\delta y, t+\delta t) \approx I(x,y,t) + \frac{\partial I}{\partial x}\delta x + \frac{\partial I}{\partial y}\delta y + \frac{\partial I}{\partial t}\delta t,$$

giving

$$\frac{\partial I}{\partial x}\delta x + \frac{\partial I}{\partial y}\delta y + \frac{\partial I}{\partial t}\delta t = 0$$

or

$$\frac{\partial I}{\partial x}\frac{\delta x}{\delta t} + \frac{\partial I}{\partial y}\frac{\delta y}{\delta t} + \frac{\partial I}{\partial t}\frac{\delta t}{\delta t} = 0,$$

which can be written as

$$I_x u + I_y v = -I_t, \tag{17}$$

where $I_x = \frac{\partial I}{\partial x}$, $I_y = \frac{\partial I}{\partial x}$, $I_t = \frac{\partial I}{\partial t}$, and $u = \frac{\delta x}{\delta t}$ and $v = \frac{\delta y}{\delta t}$ are the $x$ and $y$ components of the displacement velocity, i.e. the optical flow. As there are two unknowns, $u$ and $v$, Eq. 17 is underdetermined – this is known as the *aperture problem*. Additional constraints need thus to be introduced to compute the flow. [5] The two main approaches to differential optical flow estimation are the global approach proposed originally by Horn and Schunk [5], and the local approach of Lucas and Kanade [6].

### 1.5.1 The Lucas-Kanade method

The Lucas-Kanade method (LK) assumes that the optical flow is constant inside some neighbourhood $R$ of pixel $(x, y)$ and computes a least-squares estimate for $(u, v)$ in $R$, i.e. the $(u, v)$

that minimizes

$$E = \sum_{(x,y)\in R} (I_x u + I_y v + I_t)^2.$$

Setting $\frac{\partial E}{\partial u} = \frac{\partial E}{\partial v} = 0$ gives

$$\begin{bmatrix} 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \sum_{(x,y)\in R} 2I_y(I_x u + I_y v + I_t) \\ \sum_{(x,y)\in R} 2I_x(I_x u + I_y v + I_t) \end{bmatrix}$$

$$= \begin{bmatrix} 2\sum_{(x,y)\in R} (I_y I_x u + {I_y}^2 v + I_y I_t) \\ 2\sum_{(x,y)\in R} ({I_x}^2 u + I_x I_y v + I_y I_t) \end{bmatrix}$$

$$\iff \begin{bmatrix} \sum_{(x,y)\in R} ({I_x}^2) & \sum_{(x,y)\in R} (I_x I_y) \\ \sum_{(x,y)\in R} (I_x I_y) & \sum_{(x,y)\in R} ({I_y}^2) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} -\sum_{(x,y)\in R} (I_x I_t) \\ -\sum_{(x,y)\in R} (I_y I_t) \end{bmatrix}.$$

Thus,

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \sum_{(x,y)\in R} ({I_x}^2) & \sum_{(x,y)\in R} (I_x I_y) \\ \sum_{(x,y)\in R} (I_x I_y) & \sum_{(x,y)\in R} ({I_y}^2) \end{bmatrix}^{-1} \begin{bmatrix} -\sum_{(x,y)\in R} (I_x I_t) \\ -\sum_{(x,y)\in R} (I_y I_t) \end{bmatrix}, \tag{18}$$

given that the inverse matrix exists. This is not the case if $R$ is a region of constant intensity, as the gradients $I_x$ and $I_y$ vanish: in such regions, the optical flow is not defined.

To further refine the displacement estimates, a Newton-Raphson iteration can be applied. That is, on each iteration a LK estimation step is performed. On iteration $i$, the temporal gradient is computed as

$$I_t(x, y) = I_1(x, y) - I_2(x + u, y + v), \tag{19}$$

and the estimation result $(u_i, v_i)$ of the iteration is added to the total result by setting

$$(u, v) = (u + u_i, v + v_i). \tag{20}$$

The iteration is initialized such that $u = v = 0$, and the final result of the iteration is $(u, v)$. [6]

The problem with Eq. 19 is that the image $I_2$ may be indexed in noninteger points. This is a problem of *warping* the image, that is, moving each pixel of an image to another position based on a mapping. In practice, this warp is implemented by interpolation: the pixel value at $(x + u, y + v)$ can be interpolated using the values of the pixels nearest to the point with non-integer coordinates. Using bilinear interpolation, the warping of an image $I$ based on displacements $(u, v)$ works as follows. Denoting the nearest integer coordinates $A = \lfloor x + u \rfloor$, $B = \lfloor y + v \rfloor$, $C = \lceil x + u \rceil$, and $D = \lceil y + v \rceil$,

$$I_1(x + u, y + v) = (1 - \alpha)(1 - \beta) I_1(A, B) + \alpha(1 - \beta) I_1(C, B) + (1 - \alpha)\beta I_1(A, D) + \alpha\beta I_1(C, D),$$

where $\alpha = x - A$ and $\beta = y - B$.

### 1.5.2 Pyramidal Lucas-Kanade algorithm

The Lucas-Kanade method works only with small displacements, thus the motion estimation may fail if the objects in the frames move fast. The *Pyramidal Lucas-Kanade* (PLK) algorithm implements LK using a coarse-to-fine strategy: optical flow is first estimated at the lowest level of a Gaussian approximation pyramid, and the resulting flow field is used as an initial guess for the estimation at the second lowest level. This process is continued until the estimate on the highest level is obtained.

The algorithm starts by creating truncated Gaussian approximation pyramids for $I_1$ and $I_2$, with $P$ levels. First, using the lowest level images, a normal *LK* estimation is performed. The displacement vectors $(u^{l-1}, v^{l-1})$ from level $l-1$ are propagated to level $l$ by setting the initial guess

$$(g_x^l, g_y^l) = (2g_x^{l-1} + 2u^{l-1}, 2g_y^{l-1} + 2v^{l-1}). \tag{21}$$

The time derivative $I_t$ at level $l$ is then calculated as

$$I_t^l = I_1^l - I_2^l(x + g_x^l, y + g_y^l), \tag{22}$$

after which Eq. 18 can be applied to obtain the displacement $(u^l, v^l)$. As in the original formulation, the estimate at each level can be refined iteratively, and the warping in Eq. 22 can be done by interpolation. The final optical flow is obtained as

$$(u, v) = (u^L + g_x^L, v^L + g_y^L). \tag{23}$$

A detailed description of the algorithm can be found in [7].

## 2 Materials and methods

Here we present the image processing pipeline in Mytoe to analyze a sequence of confocal microscope images of mitochondria. Cells and images were prepared as described in the main document.

An overview of the pipeline is shown in Figure 4.

### 2.0.3 Image preprocessing

To cope with the noise in the images we use median filtering with a neighbourhood of size $5 \times 5$. The size was selected such that the filter is large enough to remove noise spikes and smooth out the noise inside the mitochondria, but not too large to blur the edges of mitochondria. In our images, the mitochondria are approximately five pixels wide, thus the width of the filter corresponds to the width of the mitochondria. After median filtering, the image intensities are scaled between 0 and 1 by

$$I_N = \frac{I_M - \min(I_M)}{\max(I_M - \min(I_M))}, \tag{24}$$

to make further processing independent of the gray scale.

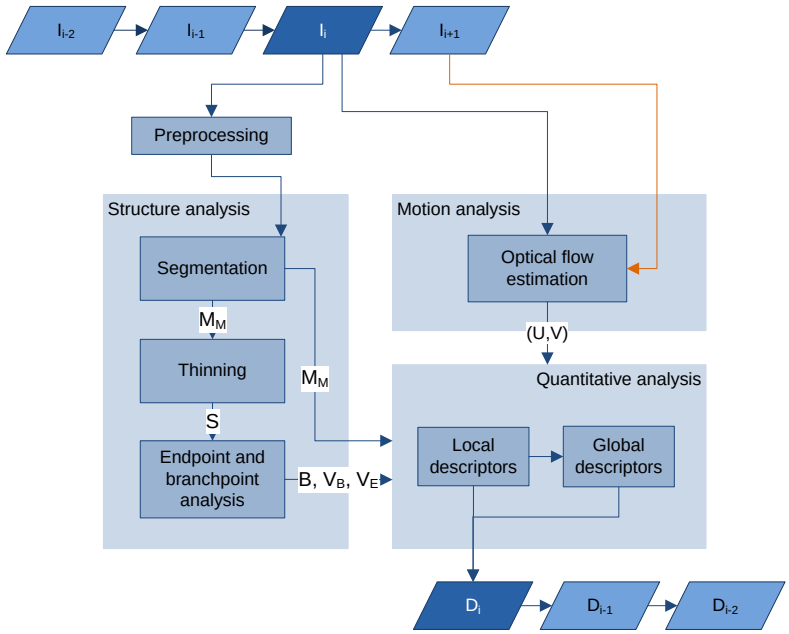


**Figure 4**: Overview of the image analysis system. $I_i$ is the $i$th image from the beginning of the sequence, and $D_i$ is the quantitative data extracted from it. The orange line emphasizes that while processing an image, the next image in the sequence is used for the optical flow estimation.

### 2.0.4 Segmentation of mitochondria

Our method for binarizing $I_M$ is similar to what was proposed in [8], but with some differences which should improve its robustness to noise. To remove the effect of varying background fluorescence we apply a white top-hat filter to $I_M$ to obtain $I_{th}$. The choice of SE for the morphological filter is guided by prior knowledge about the objects of interest: the SE should be selected to be large enough such that it does not fit inside the mitochondria. We found that a disk-shaped SE with radius 5 works well for our images. The morphological top-hat filter does not enhance noise pixels as much as the linear high-pass filter used in [8]. Because it may still enhance noisy objects in the background, we filter $I_{th}$ with a $5 \times 5$ median filter. The contrast of the median filtered result is then stretched by saturating 1% of the lowest and highest intensities, which further removes possible background noise. Finally, Otsu's method is used to binarize the contrast adjusted image, yielding the mitochondria mask $M_M$. Figure 5 shows images of all the intermediate steps in the segmentation process.

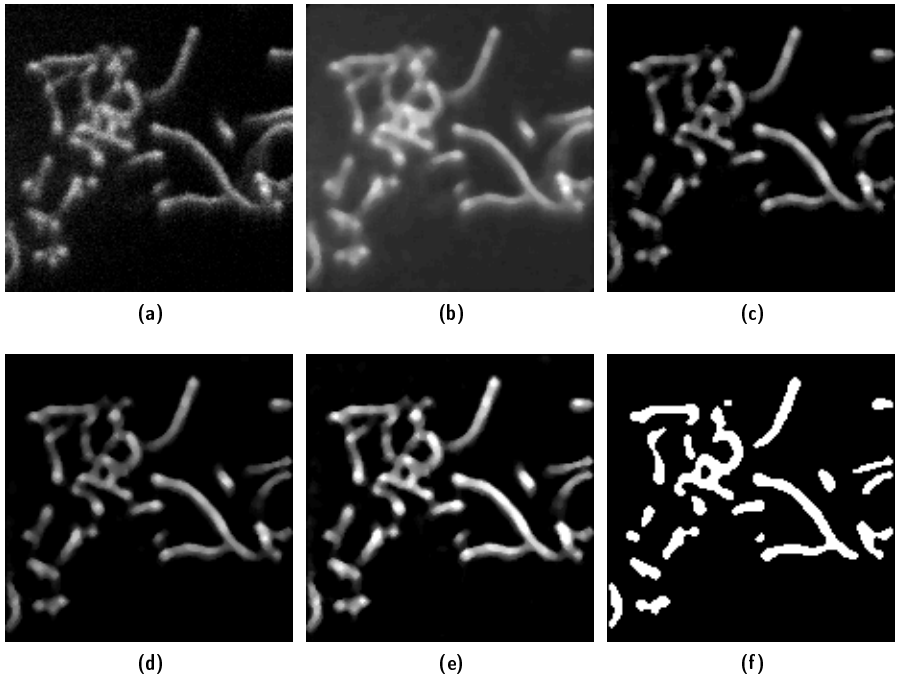


**Figure 5**: The process of mitochondria segmentation. The original image in (a) goes through the following image processing steps, the outputs of which are shown: (b) median filtering, (c) top-hat filtering, d) median filtering, e) contrast adjustment, and finally, f) thresholding.

#### 2.0.5 Detecting individual branches

We use Algorithm 1 to find the morphological skeleton $S$ of the mask $M_M$. To identify the individual branches, we need to find the end and branch points in the skeleton. A point $p \in S$ is considered an end point, if the zeros in its 8-neighbourhood are 4-connected. On the other hand, if there are more than three ones in the 8-neighbourhood of $p$, it is considered a branch point. We denote the set of end points $V_E$ and the set of branch points $V_B$. Individual branches are the 8-connected components in $B = S \setminus V_B$. The individual branches are identified by labeling $B$, yielding the label image $L_B$. The connected components of mitochondria are identified by labeling $S$, which yields the label image $L_S$, and relating the pixels of each branch to the corresponding pixels in $L_S$. An example of a skeleton along with its end and branch points is shown in Figure 6.

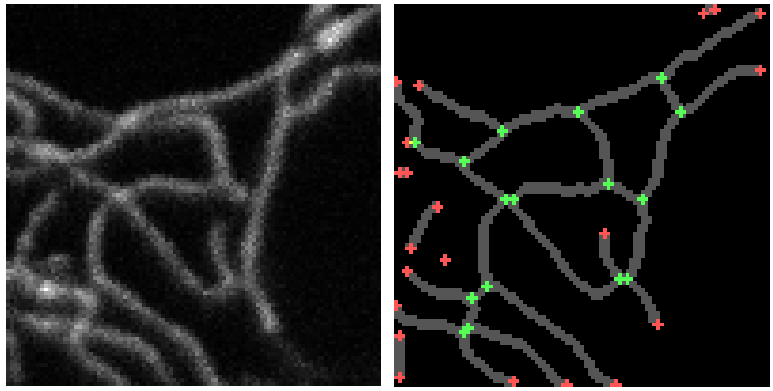

(a) (b)

**Figure 6**: a) Original image. b) The skeleton (gray), branch points (green) and end points (red). The skeleton as well as the branch and end points have been dilated for visualization purposes: the actual skeleton is one pixel wide, and each branch and end point consists of one pixel.

## 2.1 Motion analysis

### 2.1.1 Optical flow estimation

Between a mitochondria image $I$ and the next image in the sequence, $I_2$, we compute the optical flow to obtain the displacement $(U, V)$. $U$ and $V$ are matrices of the same size as $I$, and $(U(x,y), V(x,y))$ is the displacement vector starting from pixel $(x,y)$. We implement PLK, because it is computationally inexpensive, and local optical flow methods have been previously applied successfully to analysis of fluorescence microscope images of organelles [9] [10].

Prior to the motion estimation, both $I$ and $I_2$ are smoothed using a Gaussian kernel of size $5 \times 5$ with $\sigma = 1$. In both methods, we use a weighted version of Eq. 18, that is, the sums are replaced with weighted sums, with Gaussian weights:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} G*({I_x}^2) & G*(I_x I_y) \\ G*(I_x I_y) & G*({I_y}^2) \end{bmatrix}^{-1} \begin{bmatrix} -G*(I_x I_t) \\ -G*(I_y I_t) \end{bmatrix}, \tag{25}$$

where $G$ is a Gaussian kernel. This applies more importance to pixels closer to the center pixel in a neighbourhood. [11] We choose $G$ to have size $13 \times 13$ and $\sigma = 3$, as it produces the best results for our images. The spatial derivatives $I_x$ and $I_y$ are computed by convolving $I$ with the operators $[-0.5, 0, 0.5]$ and $[-0.5, 0, 0.5]^T$. The time derivative $I_t$ is computed simply as the difference $I_2 - I$. Iteration was not used in either of the methods, as it did not improve the results significantly. The Gaussian image pyramids are constructed using a filter with size $5 \times 5$ and $\sigma = 1$. We found $P = 3$ to be a sufficient number of levels, as increasing it would not significantly improve the estimation result. For warping, bilinear interpolation is used.

### 2.1.2 Semi-synthetic data for validation

Evaluation of the performance of optical flow estimation on real-world data is often hard because of the lack of ground-truth data. In some cases, computer-generated synthetic and semi-synthetic data can be used. Synthetic approaches generate images according to a model – for example, by drawing lines mimicking the shape of organelles and their movement, as was done in [9] and [10]. The most challenging aspect of a fully synthetic approach is producing realistic data. Semi-synthetic approaches, on the other hand, take real images and construct new images from them by transforming them with computer-generated flow fields [12].

Our method for generating semi-synthetic data for optical flow estimation works as follows. For a mitochondria image $I_1$ of size $N \times M$, and its mask $M_M$, the method outputs a synthetic ground-truth flow field and an image $I_2$ obtained by transforming $I_1$ based on the flow field. One can use the data for evaluation by estimating the optical flow between $(I_1, I_2)$, and comparing the result with the ground-truth flow.

We generate a synthetic flow field by creating two $N \times M$ matrices, $U_S$ and $V_S$, of normally distributed random numbers with zero mean and unit variance. $U_S$ and $V_S$ are interpreted as the $x$- and $y$-components of the synthetic flow field, respectively. $U_S$ and $V_S$ are smoothed with a Gaussian filter of size $60 \times 60$ with $\sigma = 30$, after which they are normalized, i.e. the value of each pixel $(x, y)$ is divided by the magnitude of the corresponding flow vector $(U_S(x, y), V_S(x, y))$. Next, $U_S$ and $V_S$ are masked with a version of $M_M$ dilated with a disk-shaped SE of radius 10, such that the motion will only be applied to the mitochondria and nearby pixels. The result is then smoothed with a Gaussian filter of size $15 \times 15$ with $\sigma = 7$, so as to get rid of the discontinuities caused by masking. Finally, the flow fields are scaled by multiplying by a magnitude factor $m$.

The synthetic flow has to be applied to $I_1$ to obtain $I_2$. Since $U_S$ and $V_S$ give a displacement for each pixel in $I_1$, we can put

$$I_2(x, y) = I_1(x + U(x, y), y + V(x, y)). \tag{26}$$

$I_2$ can thus be obtained by warping, for which we use bilinear interpolation. An example of a resulting image pair and the corresponding ground-truth flow field with $m = 2$ is shown in Figure 7.

### 2.1.3 Segmentation of the cytoplasm and of the nucleus

The nucleus and the cytoplasm are segmented separately but by the same method. We describe the method when applied to segment the nucleus.

We first preprocess $I_M$ by smoothing it with a $50 \times 50$ Gaussian filter with $\sigma = 10$. We use a large kernel to reduce the intensity variation inside the nucleus. The filtered image is thresholded using Otsu's method, which yields a mask $M$. We mask $M$ with $M_P$ to remove possible nuclei from other cells in the image. Because $M$ may also contain incorrectly detected objects caused by fluorescence from outside the nucleus, we obtain the nucleus mask $M_N$ as the object in $M$ with the largest area. The steps of the segmentation process are illustrated in Figure 8.

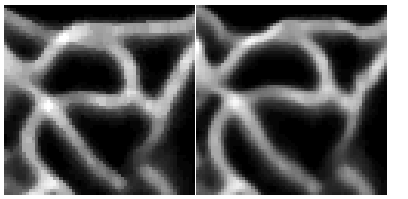

(a) (b)

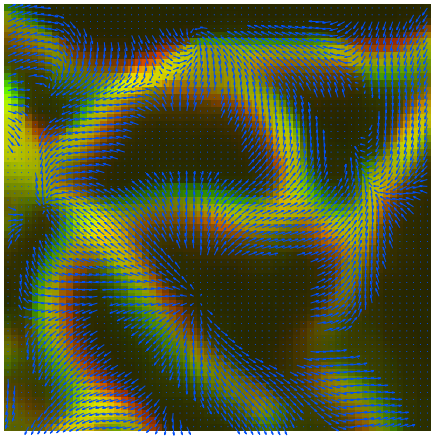

(c)

**Figure 7**: Semi-synthetic ground-truth data for testing optical flow estimation, with $m = 2.0$. a) Original image $I_1$. b) Image obtained by warping $I_1$ using the synthetic optical flow. c) The synthetic optical flow (arrows) plotted on top of an overlay of $I_1$ (red) and $I_2$ (green).

## 2.2 Quantitative analysis

From the mitochondria mask $M_M$, the decomposed structure $(B, V_E, V_B)$ and the estimated optical flow $(U, V)$ it is possible to compute various quantities, which can then be used to characterize mitochondrial properties using statistical analysis. Next, such quantities are considered and methods for measuring them are presented. Table 1 lists all the measured features. In the following, we denote the set of pixels of a branch as $F_B$.

### 2.2.1 Branch-level features

The branch length $l_B$ is approximated by the area of the branch. We compute the thickness, or width, of a branch, $w_B(p)$ at each $p \in F_B$ by using the distance transform: the thickness at point $p$ is two times the shortest distance from the point to the border of the mitochondrion,

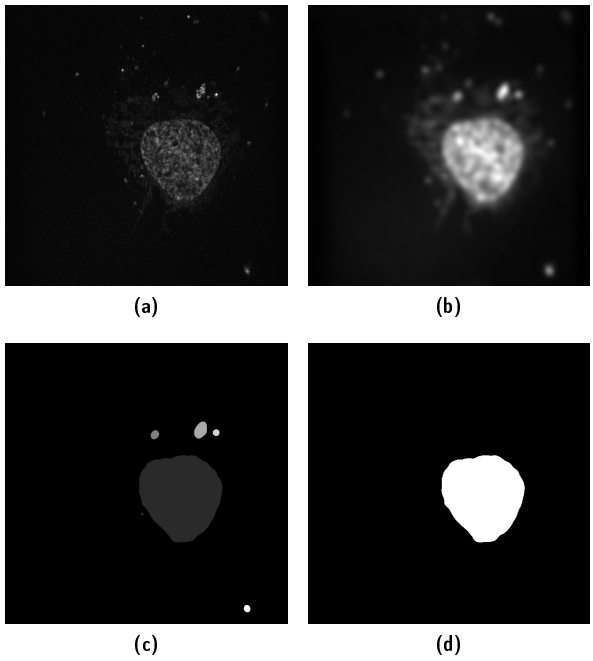


**Figure 8**: Nucleus segmentation. a) Original image. b) Blurred image. c) Label image of thresholding result, where different shades of gray indicate different labels. d) Segmented nucleus, obtained as the largest object in the label image.

thus we assign $w_B(p)$ the value of $2\mathscr{D}\{M_B\}(p)$. $M_B = \complement M_M$ is the background mask, where $\complement$ denotes set complementation. Since the background, instead of the mitochondria, is considered an object in this distance transform, the values at points $p$ will represent the shortest distance to the background, that is, to the border. To convert $l_B$ and $w_B$ to actual distances, both are multiplied by $w$.

The *tortuosity* $T$ of a branch is a measure of how curved the branch is. A simple tortuosity measure is the *arc-chord ratio*, defined as

$$T = \frac{l_B}{d(e_1, e_2)} \tag{27}$$

where $e_1$ and $e_2$ are the endpoints of the branch, and $d(e_1, e_2)$ is their Euclidean distance. For a straight line $T = 1$, and for a curved line $T > 1$. In [13], a similar measure was suggested for studying tortuosity in retinal vasculature, which is related to several diseases.

We measure branch orientation, $O_C$, relative to the nucleus centroid: let $x$ be the line segment connecting the endpoints of a branch, and let $y$ be the line passing through the nucleus center and the midpoint of $x$. We define $O$ as the angle between $y$ and $x$. We also measure the orientation relative to the x-axis of the image, $O$, as the angle between $x$ and the

| Feature | Symbol | Type |
|---|---|---|
| Length | $l$ | B |
| Thickness | $w_B$ | B |
| Tortuosity | $T$ | B |
| Intensity | $I$ | B |
| Speed | $s$ | B |
| Direction | $d$ | B |
| Wiggle Ratio [10] | $W$ | B |
| Distance to centroid | $d_C$ | B |
| Distance to nucleus membrane | $d_N$ | B |
| Distance to plasma membrane | $d_P$ | B |
| Orientation relative to x-axis | $O$ | B |
| Orientation relative to centroid | $O_C$ | B |
| Number of branches | $n_B$ | C |
| Number of connected components | $n_C$ | C |
| Total mitochondrial area | $A_M$ | C |
| Degree of clustering | $L$ | C |

**Table 1**: Features extracted by Mytoe. B and C denote branch-level and cell-level measurements, respectively.

x-axis.

Intracellular distances from the mitochondria to an object can be measured using the distance transform. Specifically, given a mask $M$ of the object, the distance $d_{obj}$ from a pixel $p \in F_B$ to the object is given by $\mathscr{D}\{M\}(p)$. If the object is selected to be the plasma membrane, the mitochondria are inside the mask $M_P$, thus the mask used is $\complement M_P$, and we denote the distance $d_P$. If the object is the nucleus boundary, the mask is $M_N$, and the distance is denoted $d_N$. In the case of $d_N$, the mitochondria inside the nucleus mask are given the value $-\mathscr{D}\{\complement M\}(p)$, that is, the distance to the nucleus membrane with sign $-$. Finally, we compute the distance to the nucleus centroid, $d_C$, where the mask consists of only a single pixel, the centroid of $N_M$.

Knowing the image acquisition time interval and the pixel distance in the images, we can convert the displacements $(U,V)$ to velocities as

$$\begin{aligned} v_x &= \frac{U \cdot w}{\Delta t} \text{ and} \\ v_y &= \frac{V \cdot w}{\Delta t}. \end{aligned} \tag{28}$$

Given $\mathbf{v} = [v_x, v_y]^T$ we can compute the speed $s$ and direction $d$ of each branch by averaging the speed and direction of each pixel in the branch. As a measure of the directedness of mitochondrial movement, a quantity termed *Wiggle Ratio*, $W$, was proposed in [10]. It is based on the notion that when a mitochondrion is moving directionally, the mean velocity

vector over the $N$ measured points in a mitochondrion,

$$\overline{\mathbf{v}} = \sum_{i=1}^{N} \mathbf{v}_i, \tag{29}$$

has a magnitude equal to the mean speed

$$\overline{s} = \sum_{i=1}^{N} ||\mathbf{v}_i||. \tag{30}$$

Thus, the ratio $W = \dfrac{\overline{s}}{||\overline{\mathbf{v}}||} \approx 1$ in case of directional movement. On the other hand, if different parts of the mitochondrion are moving in different directions, the magnitude of $\overline{\mathbf{v}}$ will be smaller than that of $\overline{s}$, and therefore $W > 1$.

#### 2.2.2 Cell-level features

The cell and mitochondrial areas, $A_P$ and $A_M$, respectively, are obtained by computing the areas of their masks and converted to $\mu m^2$ by multiplying by $w^2$. By enumerating the individual branches we get the branch count, $N_B$. The number of connected components, $N_C$, is obtained as the number of connected components in $M_M$.

To define a measure of how clustered the mitochondria are, we utilize the *Ripley's L* function. Ripley's L can be used to measure the randomness of a spatial point process, that is, a stochastic process whose samples can be interpreted as locations of events. It is a normalized version of *Ripley's K*, which is defined as

$$\hat{K}(s) = \frac{1}{\lambda n} \sum_i \sum_{j \neq i} I(d(i,j) < s), \tag{31}$$

where $\lambda$ is the average density of points, $n$ is the number of points, $I(x)$ is the identity function, $d(x,y)$ is a distance metric, and $s$ is the scale parameter. The value of $\lambda$ can be estimated as $A/n$, where $A$ is the area where the points are measured. For a homogeneous process, $\hat{K}$ equals approximately $\pi s^2$, which motivates the use of a variance-normalized version

$$\hat{L}(s) = \sqrt{\frac{\hat{K}(s)}{\pi}}, \tag{32}$$

which is Ripley's L. [14] In our analysis, we consider two separate point sets, the mitochondria endpoints $V_E$ and the branchpoints $V_B$.

## 3 Validation of the methods

### 3.1 Robustness of the structure analysis

In microscope imaging, among the most common factors that degrade the quality of the images are noise and saturation. Saturation may occur, for instance, when the intensity of the

light signal from the specimen increases during the acquisition of a time-lapse, and the detector reaches the ceiling of its dynamic range. We model the effect of these degradation sources by adding noise to the images and saturating them.

The accuracy of the segmentation of mitochondria was assessed qualitatively, by visually comparing the obtained masks to the mitochondria images, and it was observed to be working well in all of the sequences. However, as we only use a single optical section per frame, and the mitochondria are not necessarily oriented exactly parallel to the focal plane, the method may not correctly segment a mitochondrion, or a part of it, which is out of focus. The method is not able to distinguish overlapping mitochondria from those that are connected, which is also a limitation of the two-dimensional imaging method.

The degradation of a preprocessed and segmented mitochondria image $I$ was performed as follows. We modeled the image noise with Poisson noise as follows. The noisy image $I_d$ was generated by setting, for each $(i, j)$, $I_d(i, j)$ to be a random number from a Poisson distribution with the parameter $I(i, j)$. The level of noise was varied by multiplying the pixel values of $I$ prior to the degradation by a factor of $1/a$, where $a$ is a real number greater than 1, and scaling back $I_d$ by multiplying it by $a$. As a measure of the noise level, we used the signal-to-noise ratio (SNR) definition from [15]:

$$\mathrm{SNR} = \frac{I_o - I_b}{\sigma}, \tag{33}$$

where $I_o$ and $I_b$ are the mean object and background intensities, respectively, and $\sigma$ is the standard deviation of the noise. As was done in [15], we estimate $\sigma$ as the RMS error over the object,

$$\sigma_{RMS} = \sqrt{\sum_{(x,y)\in M} (I(x, y) - I_d(x, y))^2}, \tag{34}$$

where $M$ is the mitochondria mask interpreted as a set of points. To model the saturation, we set $I_S(i, j) = (1 - \alpha) \max I(x, y)$ for all $I(i, j) > (1 - \alpha) \max I(x, y)$, where the maximum is taken over all pixels of the image, with $\alpha = 0, 0.01, 0.05, 0.1, 0.5, 0.75, 0.8$.

To test the effect of degradation on the quantification of mitochondrial parameters, we computed the number of branches, mean branch thickness and mean branch length for each degraded image. The first image from each sequence was used to create the degraded images. From Figure 9 it can be seen that the lowest SNR for which the method produces results consistent with those obtained from the better quality images, is between 3 and 7. When $SNR \leq 3$, the segmentation method fails, and noise pixels are detected as branches, which results in a high branch count. When $3 \leq SNR \leq 7$, the noise affects the thickness and length measurements, but does not cause dramatic changes in the branch counts. Compared to the previously presented method in [8], although a quantitative comparison is difficult because of different imaging systems, it seems that our method is more robust to noise as it performs well without image averaging.

It is visible from Figure 9 that the method is fairly robust to saturation when $\alpha < 0.4$. The performance under saturation suggests that reliable structure measurements may be obtained from images that are slightly degraded by saturation. This result is not unexpected, because saturation of the mitochondria pixels will actually remove intensity variation from

the mitochondria, and the segmentation performance will only be affected when nearby mitochondrial objects blend together due to saturation. The implication of this is that one does not necessarily have to exclude slightly saturated images when using the method to study mitochondrial structure.

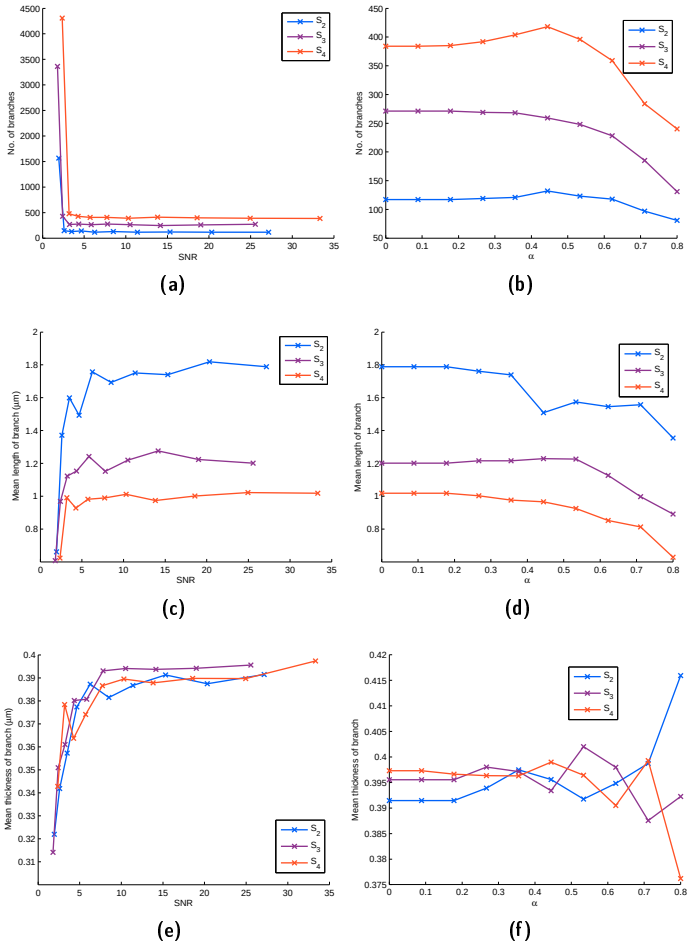


**Figure 9**: Robustness of the structure analysis to degradation in three sequences, $S_2$, $S_3$ and $S_4$. From each sequence, a frame was picked and degraded with varying levels of noise and saturation. In (a), (c) and (e) the robustness to noise is assessed by computing the number of branches, mean branch length and mean branch thickness, respectively. In (b), (d) and (f) the same descriptors were computed and the level of saturation was varied.

| Sequence | $100\% \cdot R_{max}$ | $100\% \cdot R_{mean}$ |
|---|---|---|
| 1 | 103.96% | 100.47% |
| 2 | 101.29% | 100.01% |
| 3 | 101.61% | 100.00% |
| 4 | 104.94% | 99.99% |
| 5 | 102.90% | 100.00% |
| 6 | 116.39% | 100.00% |
| 7 | 102.33% | 100.00% |
| 8 | 102.10% | 99.97% |

**Table 2**: Maximum and average ratios of median intensities between pairs of subsequent frames, converted to percentages.

### 3.2 Motion analysis

Changes in the intensity of the light emitted by the specimen will decrease the accuracy of optical flow estimation, because the brightness constancy assumption will be violated [16]. We confirmed visually that the RFP fluorescence of individual mitochondria was homogeneous between frames in each sequence. Global variation in intensity was assessed by comparing the median and mean intensities of the mitochondria pixels, obtained by segmenting them, in each pair of subsequent frames. For each sequence $S$ we computed the ratios

$$R_{max} = \max_i \frac{\tilde{M}_i}{\tilde{M}_{i+1}} \tag{35}$$

and

$$R_{mean} = \frac{1}{N} \sum_{i=1}^{N} \frac{\tilde{M}_i}{\tilde{M}_{i+1}}, \tag{36}$$

where $\tilde{}$ denotes the sample median, $M_i$ is the set of mitochondria pixels in the $i$th frame of $S$, and $N$ is the number of frames in $S$. $R_{max}$ represents the maximum change in median intensity from one frame to the next, while $R_{mean}$ is the average intensity change between subsequent frames. The results are listed in Table 2. The average changes are less than 1% in each sequence. The maximum changes are between 1% and 5%, except in $S_6$, which has a 16% maximum increase in median intensity between frames 45 and 46 – this was confirmed to be due to a temporary change in focus in frame 45. From Figure 10, which shows the median intensities for all frames in every sequence, it is visible that there is no significant decrease in the fluorescence intensity during the sequences, except for $S_1$, in which the intensity drop was due to mitochondria moving out of focus.

To assess the performance of optical flow estimation, we generated five synthetic flows for the first frame of each sequence, with flow magnitudes 1.0, 2.0, 3.0, 4.0 and 5.0. Each pair of images $(I_1, I_2)$ corresponding to a flow was degraded with varying levels of Poisson noise, as was done for the structure analysis. We compared PLK with one pyramid level, i.e. the original LK, and PLK with 3 pyramid levels.

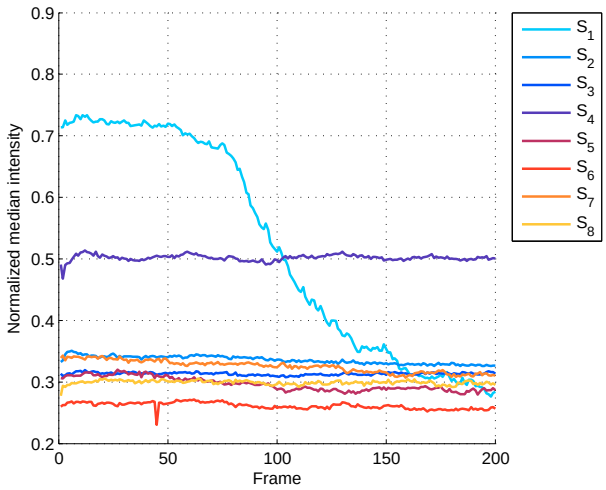


**Figure 10**: Median mitochondrial intensity of the frames of each imaged sequence. The intensities were normalized by dividing by $2^{14}$, which is the maximum of the dynamic range of the imaging system. The notable intensity decrease in $S_1$ is due to defocusing.

For measuring the performance of optical flow algorithms we adopted the angular error (AE) and magnitude of differences error (MDE) from [17], as well as the *magnitude error*

$$ME = \left| \frac{||\mathbf{c}|| - ||\mathbf{e}||}{||\mathbf{c}||} \right|, \tag{37}$$

which measures the difference in the magnitudes of the true and estimated vectors.

For each flow estimate, we averaged the error metric values inside the mitochondria mask $M$, to obtain the average angular error (AAE), average magnitude error (AME) and average magnitude of difference error (AMDE). Figure 11 shows the measured AAE, AME and AMDE for varying noise and flow magnitude, for one synthetic flow. The results were similar for all test flows.

The performance of PLK with 3 pyramid levels is constantly better compared to LK, but both methods are sensitive to noise, and become less accurate when magnitude is increased. An AME of 0.3 reports that, on average, the difference between the magnitude of a true flow vector and the magnitude of the estimated flow is 0.3 times the magnitude of the true flow vector. While the measured errors are fairly high, it should be noted that the mean is affected by outliers. These are, for example, flow estimates with a large magnitude or a direction opposite to that of the ground truth. In practice the error can be reduced by applying outlier removal. It is also worth noting that the semi-synthetic images generated by our method tend to become less realistic with increasing flow magnitude, and the error measurements consequently become less reliable.

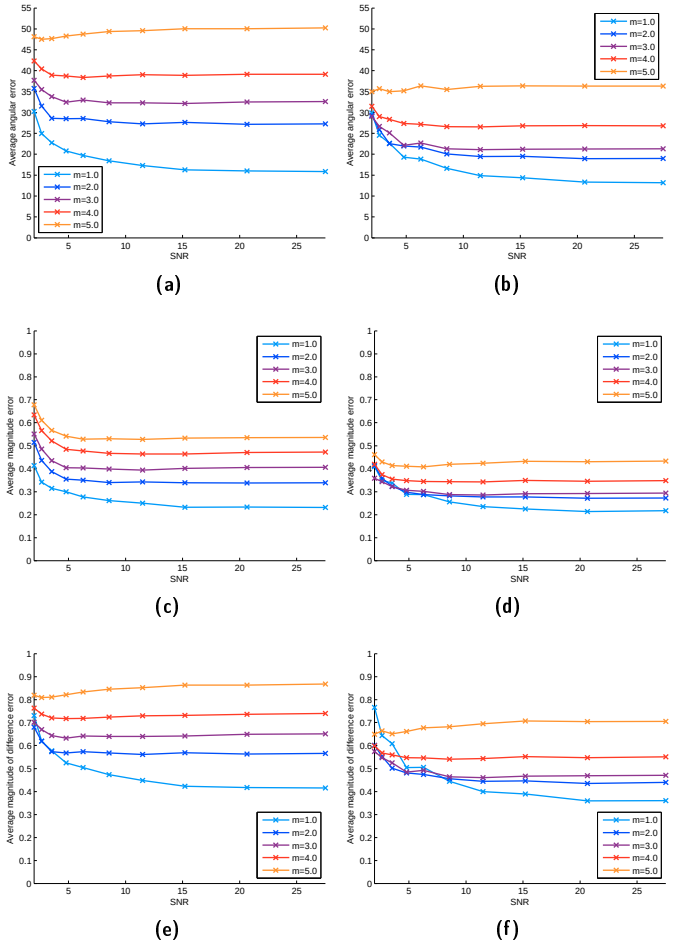


**Figure 11**: Effect of noise and flow magnitude on optical flow estimation. (a), (c) and (e) show the AAE, AME and AMDE for the Lucas-Kanade method, respectively, whereas (b), (d) and (f) are the corresponding results for the Pyramidal Lucas-Kanade method.

The results indicate that the method produces consistent results with SNRs higher than 10, and is reasonably accurate when the flow magnitudes do not exceed 3 pixels per frame.

## 4 Application of Mytoe to U2OS cells

We imaged sequences of eight cells referred to as $S_1,\ldots S_8$. Figure 12 shows example frames of the sequences $S_3$ and $S_6$.

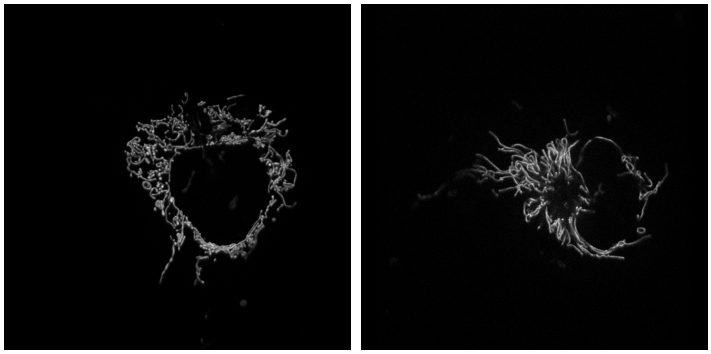


**Figure 12:** Example frames of the image sequences (a) $S_3$ and (b) $S_6$.

A scatter plot of branch thickness and length in sequences $S_2$–$S_5$ (Figure 13) shows differences in the size of the branches over time and between cells. Each point in the plot corresponds to a single frame of one cell in a sequence. The variation between points in the same cell is due to inaccuracies in the segmentation and movements of the organelles that can change length and thickness. Defocusing of a part of a long mitochondrial branch could also cause the branch to be segmented as two shorter branches, thus decreasing the average measured length.

Time series of the average tortuosity $T$ in three cells (Figure 14) shows how this quantity differs between cells.

Figure 15 shows the distributions of branch distances to the nucleus membrane, to the cell membrane, and to the nucleus centroid, obtained from a single cell image. Figure 15 a) suggests that the mitochondrial structure is slightly denser near the nucleus. On the other hand, Figure 15 b) does not show such packing near the cell membrane. This example shows that the intracellular distribution of mitochondria can be represented in many ways, and the choice should depend on the application. For example, the distance to the nucleus, or another organelle, may be an appropriate measure, if one wishes to compare differences between mitochondria near that organelle and those further away from it.

Figure 16 shows histograms of branch orientations $O$ and $O_C$ for the first image in $S_3$. The peaky histogram in Figure 16 (a) shows that there are preferred orientations. On the other hand, Figure 16 (b) shows a uniform distribution of orientation relative to the centroid.

Time series of the mean branch speed in some of the sequences are shown in Figure 17, along with speed distributions. The mitochondria in $S_1$ and $S_8$ move at speeds between 0.05–0.07 $\mu m/s$, while those in $S_2$ and $S_4$ move slower, at $0.03-0.05$ $\mu m/s$. It was shown in [10]

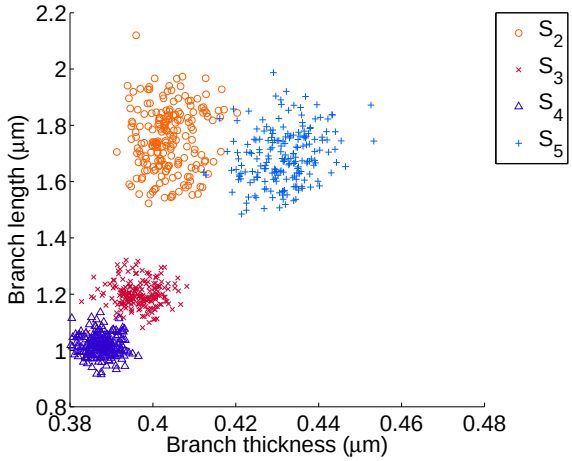


**Figure 13**: Mean branch thickness versus mean branch length for different cells. For a sequence, each data point corresponds to a frame in the sequence.

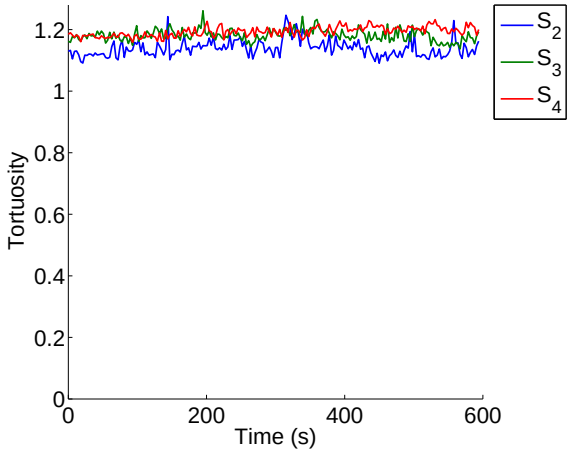


**Figure 14**: Mean $T$ in three cells.

that, in hippocampal neurons, the speed of mitochondria is inversely correlated with their size. In $S_2$ the mitochondria are long, which may explain their slower speed. However, in $S_4$, the mitochondria are shorter compared to $S_1$ and $S_8$, but they move slower. There are, of course, many factors that cause differences in the speeds of mitochondria, and Figure 17 merely demonstrates how accurately those differences can be measured with our method. When measuring the speed at a certain moment, one should measure the speeds between a

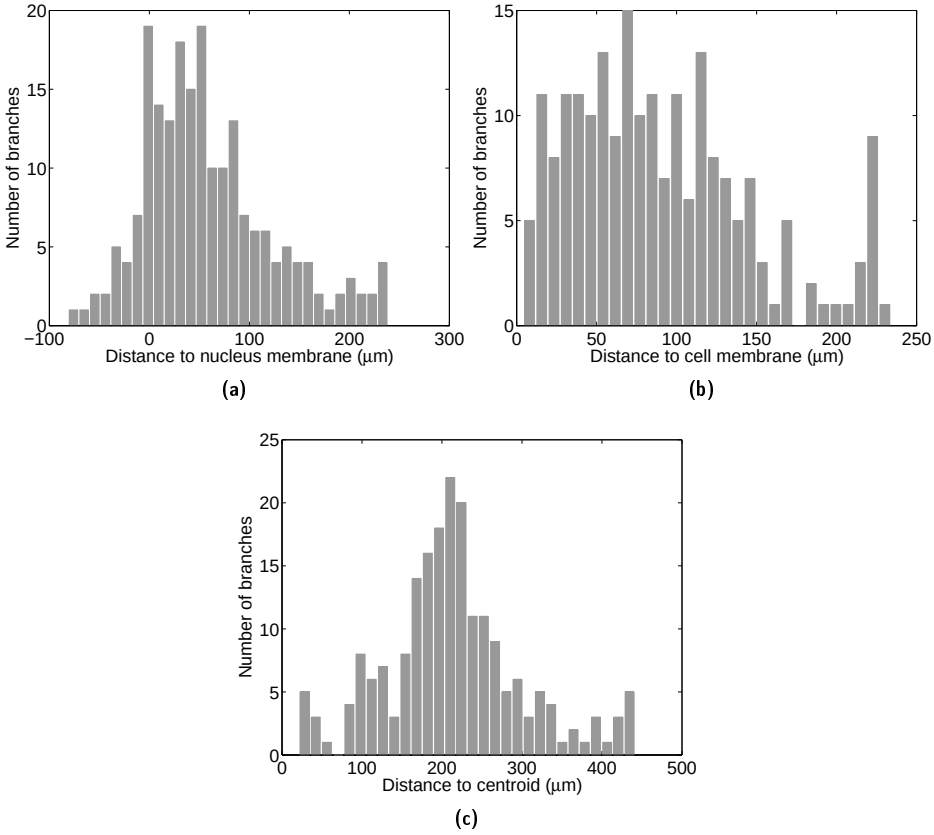


**Figure 15**: (a) Histogram of branch distances to the nucleus membrane in the first frame of $S_1$. (b) Histogram of branch distances to the cell membrane in the same frame. (c) Histogram of branch distances to the nucleus centroid.

few image pairs and average them, in order to reduce the noise.

Time series of average $W$ are shown in Figure 18. The results suggest that the movement of mitochondria in $S_8$ is more directed than that of $S_1$.

Time series of $n_B, n_C$ and $A_M$ for different cells can be seen in Figure 19. While there is fluctuation in the data, in part due to inaccuracies in the segmentation, it is evident that during the time series the mitochondria are not changing in numbers or size.

The plots in Figure 20 show the measurements of degree of clustering of the mitochondria endpoints and branchpoints for different cells. As $\hat{L}(s)$ equals $s$ if the points are homogeneously distributed, a line that is close to zero in the $s - \hat{L}(s)$ plot indicates that the corresponding points are not clustered. A negative value of $s - \hat{L}(s)$ suggests that the points are

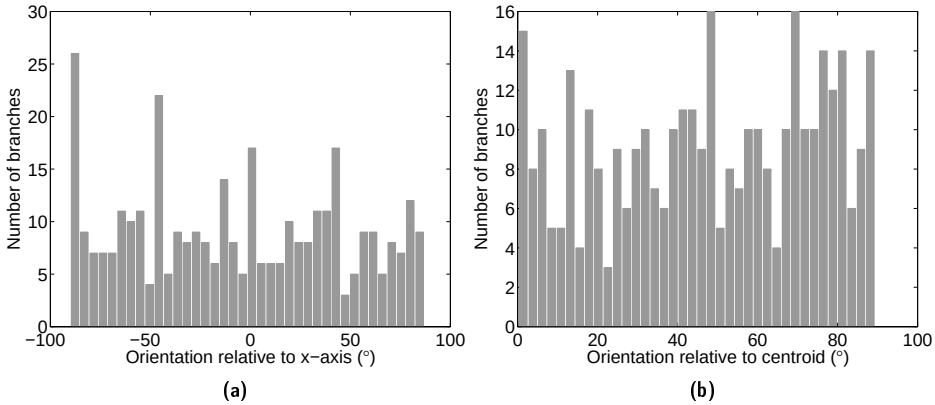


**Figure 16**: Branch orientations (a) relative to x-axis ($O$), (b) and relative to cell centroid ($O_C$) for $S_3$.

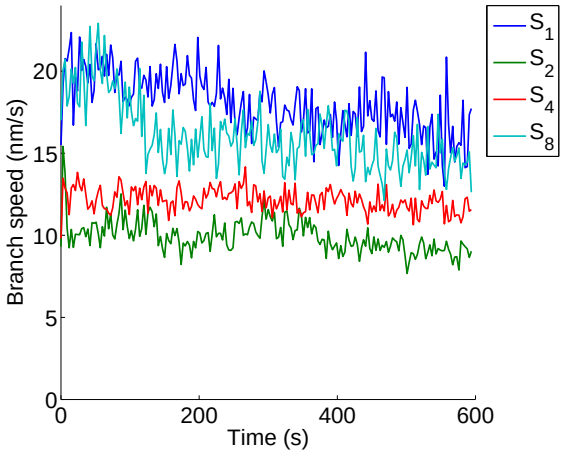


**Figure 17**: Time series of mean branch speeds reveal differences in mitochondrial motility.

clustered at scale $s$, and a positive value would indicate spatial regularity. Based on these measurements, both the branch points and end points $S_3$ show increased clustering at increasing scales, while in $S_4$, the branch points show similar clustering but the end points do not.

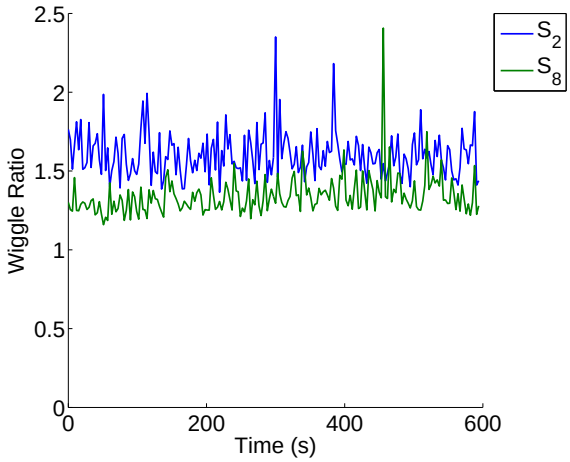


**Figure 18**: Time series of mean Wiggle Ratio for two cells.

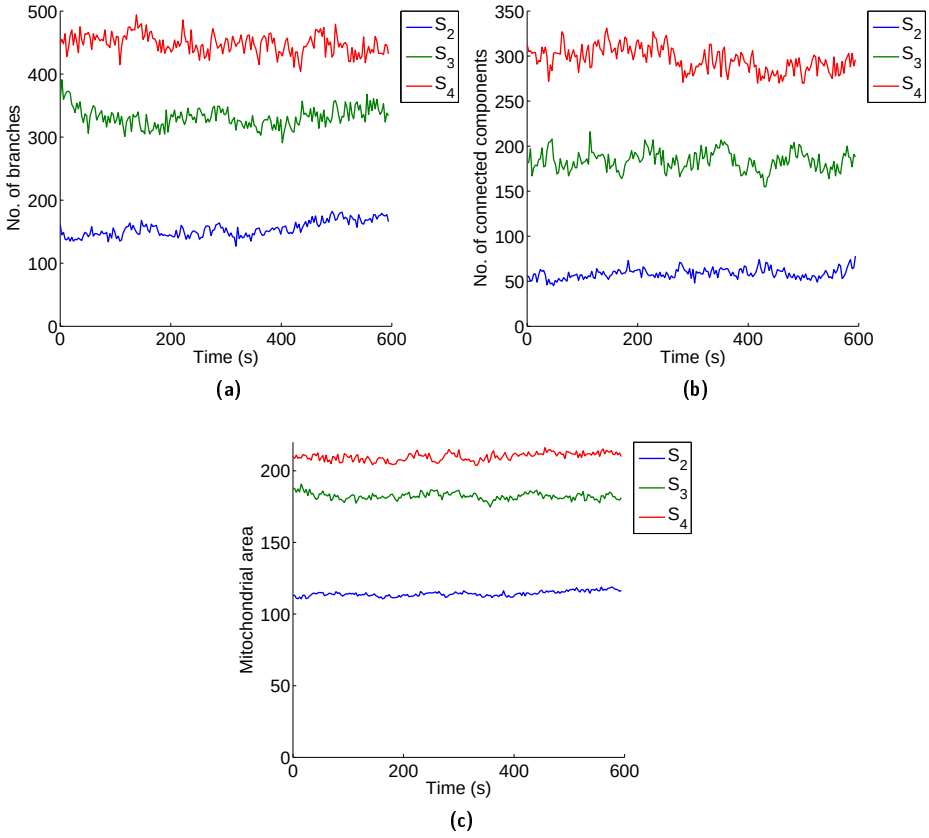


**Figure 19**: (a) $n_B$, (b), $n_C$ and (c), $A_M$ for different cells.

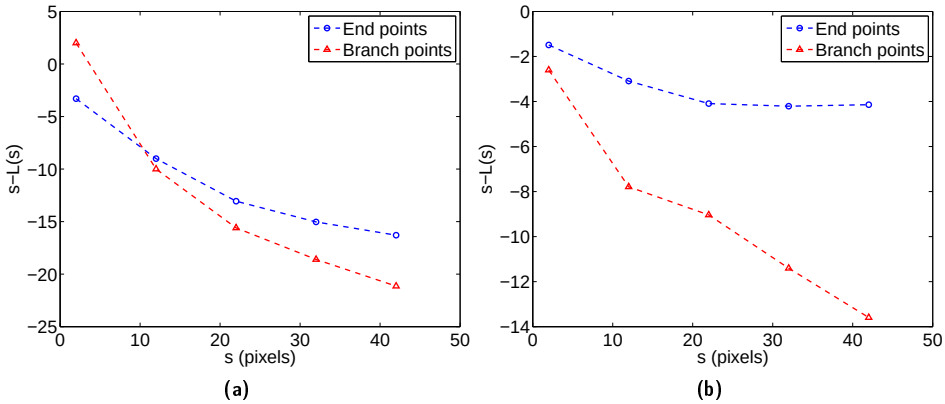


**Figure 20**: Clustering of branch points and end points of mitochondria in the first images of a) $S_3$ and b) $S_4$. Notice the different scales.

## References


[1] N. Otsu, "A threshold selection method from gray level histograms," *IEEE Trans. Systems, Man and Cybernetics*, vol. 9, pp. 62–66, Mar. 1979.

[2] W. K. Pratt, *Digital Image Processing*, 2nd ed. New York, NY, USA: John Wiley & Sons Inc., 1991.

[3] P. Soille, *Morphological Image Analysis: Principles and Applications*, 2nd ed. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2003.

[4] Z. Guo *et al.*, "Parallel thinning with two-subiteration algorithms," *Communications of the ACM*, vol. 32, no. 3, pp. 359–373, March 1989.

[5] B. Horn *et al.*, "Determining optical flow," *Artificial Intelligence*, vol. 17, no. 1-3, pp. 185–203, August 1981.

[6] B. Lucas *et al.*, "An iterative image registration technique with an application to stereo vision," in *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI '81)*, April 1981, pp. 674–679.

[7] J.-Y. Bouguet, "Pyramidal implementation of the Lucas-Kanade feature tracker," Tech. Rep., 2000.

[8] W. Koopman *et al.*, "Simultaneous quantitative measurement and automated analysis of mitochondrial morphology, mass, potential, and motility in living human skin fibroblasts," *Cytometry Part A*, vol. 69A, no. 1, pp. 1–12, 2005.

[9] D. Uttenweiler *et al.*, "Motion determination in actin filament fluorescence images with a spatio-temporal orientation analysis method," *Biophysical Journal*, vol. 78, no. 5, pp. 2709–2715, 2000.

[10] A. Gerencser *et al.*, "Measurement of instantaneous velocity vectors of organelle transport: Mitochondrial transport and bioenergetics in hippocampal neurons," *Biophysical Journal*, vol. 95, no. 6, pp. 3079–3099, 2008.

[11] A. Bruhn, J. Weickert, and C. Schnörr, "Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods," *International Journal of Computer Vision*, vol. 61, pp. 211–231, 2005.

[12] V. Ulman *et al.*, "Pseudo-real image sequence generator for optical flow computations," in *Image Analysis*, ser. Lecture Notes in Computer Science, B. Ersbøll and K. Pedersen, Eds. Springer Berlin / Heidelberg, 2007, vol. 4522, pp. 976–985.

[13] W. E. Hart, M. Goldbaum, B. Côté, P. Kube, and M. R. Nelson, in *Proceedings of the AMIA Annual Fall Symposium*, 1997, pp. 459–463.

[14] P. M. Dixon, *Ripley's K Function*. Chichester: John Wiley & Sons, Ltd, 2006.

[15] M. K. Cheezum, W. F. Walker, and W. H. Guilford, "Quantitative comparison of algorithms for tracking single fluorescent particles," *Biophys. J.*, vol. 81, no. 4, pp. 2378–2388, 2001.

[16] H. Haussecker *et al.*, "Computing optical flow with physical models of brightness variation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, pp. 661–673, June 2001.

[17] B. McCane, K. Novins, D. Crannitch, and B. Galvin, "On benchmarking optical flow," *Computer Vision and Image Understanding*, vol. 84, no. 1, pp. 126–143, Oct. 2001.

# Publication II

Eero Lihavainen, Jarno Mäkelä, Johannes N. Spelbrink, Andre S. Ribeiro, "Detecting and tracking the tips of fluorescently labeled mitochondria in U2OS cells," *Lecture Notes in Computer Science*, 9280:363 – 372, Aug. 2015

# Detecting and Tracking the Tips of Fluorescently Labeled Mitochondria in U2OS Cells

Eero Lihavainen[1(✉)], Jarno Mäkelä[1], Johannes N. Spelbrink[2], and Andre S. Ribeiro[1]

[1] Tampere University of Technology, Tampere, Finland
{eero.lihavainen,jarno.makela,andre.ribeiro}@tut.fi
[2] Nijmegen Centre for Mitochondrial Disorders, Radboud University Medical Center, Nijmegen, The Netherlands
hans.spelbrink@radboudumc.nl

**Abstract.** We present a method for automatically detecting the tips of fluorescently labeled mitochondria. The method is based on a Random Forest classifier, which is trained on small patches extracted from confocal microscope images of U2OS human osteosarcoma cells. We then adopt a particle tracking framework for tracking the detected tips, and quantify the tracking accuracy on simulated data. Finally, from images of U2OS cells, we quantify changes in mitochondrial mobility in response to the disassembly of microtubules via treatment with Nocodazole. The results show that our approach provides efficient tracking of the tips of mitochondria, and that it enables the detection of disease-associated changes in mitochondrial motility.

**Keywords:** Mitochondria · Detection · Tracking · Image analysis

## 1 Introduction

Mitochondria are involved in many cellular processes, and their dysfunctions have been linked to several diseases. In particular, abnormal mitochondrial dynamics such as an increased rate of fission, have been reported in the case of neurodegenerative diseases (see [2]).

In order to better understand the underlying mechanisms behind abnormal mitochondrial dynamics, it is necessary to analyze time-lapse image data from a large number of cells. So far, studies have relied on qualitative descriptions of mitochondrial movement [3] and manual image analysis [11], which limit the amount of data that can be analyzed. For more detailed studies, e.g. focusing on how interactions may affect the mobility, automatic image analysis methods are needed.

Previously described methods for automatic quantification of mitochondrial motion have mostly been restricted to measuring instantaneous velocity distributions using e.g. Optical Flow estimation [9] among other techniques [1]. Such methods yield no information about long-term dynamics of individual mitochondria. For tracking individual mitochondria, Silberberg et al. [12] applied a


V. Murino and E. Puppo (Eds.): ICIAP 2015, Part II, LNCS 9280, pp. 363–372, 2015.
DOI: 10.1007/978-3-319-23234-8_34

particle tracking method, consisting of the detection of mitochondria and a subsequent tracking step. A limitation of their detection method is that it assumes that the mitochondria appear globular, which is not true in general, as mitochondria often exhibit elongated and networked morphologies.

As mitochondria are similar in appearance to other elongated cellular structures, such as cytoskeletal filaments, when imaged with a fluorescence microscope, methods for tracking such filaments should be applicable to tracking mitochondria as well. For tracking cytoskeletal filaments, active contour tracking methods have been used [13]. These methods have the disadvantage of requiring the adjustment of several, non-intuitive parameters. In addition, methods for tracking the tips of microtubules have recently been proposed [5,6], but they rely on an initial manual detection of the tips.

In this work, we present a novel, automatic approach for detecting the tips of mitochondria, and apply the tracking framework of [8] to track the detected tips. Our detection method is based on supervised learning, namely a Random Forest classifier. Previous methods for automatically detecting the tips of microtubules [10] or mitochondria [9] have relied on segmentation, by applying filters that enhance curvilinear structures, binarizing the filtered image via thresholding, and extracting the tips from the morphological skeleton of the binarized image. Often such approaches will either over- or undersegment parts of the mitochondrial structure, which leads to false positives and misses. In contrast, directly detecting the tips should lead to a more robust method.

Here, we present the method and its validation using synthetic data. We also present a comparison of the method to a segmentation-based approach, and find it more reliable. Finally, we demonstrate the applicability of the method to experimental data, by measuring changes in mitochondrial motility caused by treatment with Nocodazole [14,15].

## 2 Materials and Methods

### 2.1 Image Acquisition

We transfected U2OS cells with a vector expressing mitoDsRED2, a red fluorescent protein targeting the mitochondrial matrix. The nuclei were labeled with the Hoechst 33342 fluorescent dye. The images were acquired with a Nikon Eclipse Ti-E with 100x Apo, a Wallac-Perkin Elmer Ultraview spinning-disk confocal system, Andor EMCCD camera, and a Nikon PFS autofocus system.

Prior to imaging, the cells were treated with $5\mu g/ml$ Nocodazole. We then selected four cells to be imaged. At 0, 30, 60, 90, 120 and 180 minutes after the application of Nocodazole, we imaged one optical slice of each cell every 3 seconds, for 10 minutes. This resulted in 6 movies of 10 minutes for each cell.

### 2.2 Training and Test Data

We selected one representative image of a cell not affected by Nocodazole, in which the sizes and appearances of mitochondria varied widely. In this image,

we manually marked 50 points at the tips of mitochondria, and extracted square patches of size $9 \times 9$ around them; we will refer to these as *positive* patches. Next, we extracted 1000 patches at random points, to serve as *negative* examples. This random selection was justified, because less than 1% of the points in the image contain tips of mitochondria. Finally, we manually marked 50 points at non-tip locations that shared visual features with tips: points along the mitochondria filaments, borders between two mitochondria, and curved edges of mitochondria. Figure 1 shows examples of each of these three subsets of the training data.

Although we estimated that 50 patches suffice to cover most of the variation in the appearance of the tips of mitochondria, as well as of the non-tip regions, any particular appearance may be represented in few orientations. In order to make our detector invariant to orientation, we augmented the data set with transformed versions of each manually selected patch. In particular, we applied each of the symmetries of a square to the patch: First, we mirrored the patch horizontally. Second, we rotated both the original patch and its mirrored version by 90, 180 and 270 degrees. This amounted to 7 new patches for each old patch. The reason we selected these transformations instead of, e.g., rotations of arbitrary angles, is that they require no interpolation, and thus do not introduce artifacts to the patch.

After extracting the training data, we further manually extracted 25 positive and 25 negative patches to serve as test data.

### 2.3 Detecting the Tips of Mitochondria Using a Random Forest Classifier

Our detection method works by classifying each sub-patch of the image in a sliding window, using a binary Random Forest (RF) classifier. As features for the classifier, we use the pixel values of the patches, read in column-major order, and normalized to zero mean and unit variance, in order to achieve invariance to intensity scaling. That is, for an image patch $P_k = \{p_{ij}\}$, where $i \in 1..9$ and $j \in 1..9$ are the row and column indices, respectively, the corresponding unnormalized feature vector is

$$x_k = [p_{11}, p_{21}, \ldots, p_{91}, p_{12}, p_{22}, \ldots, p_{19}, p_{29}, \ldots p_{99}], \tag{1}$$

and the final, normalized feature vector is obtained as

$$y_k = \frac{x_k - \langle x_k \rangle}{\mathrm{std}(x - \langle x_k \rangle)}, \tag{2}$$

where $\langle \cdot \rangle$ denotes the sample mean, *std* denotes the sample standard deviation.

This detection procedure results in a binary image, where each connected component (CC) corresponds to one tip of a mitochondrion. However, the image will also contain some CCs that are false positive detections. We verified by visual inspection that these are typically small, approximately $1 - 2$ pixels in size. Noting that there is some uncertainty to the exact location of a manually marked tip, true positive CCs should contain more pixels in the tip region.

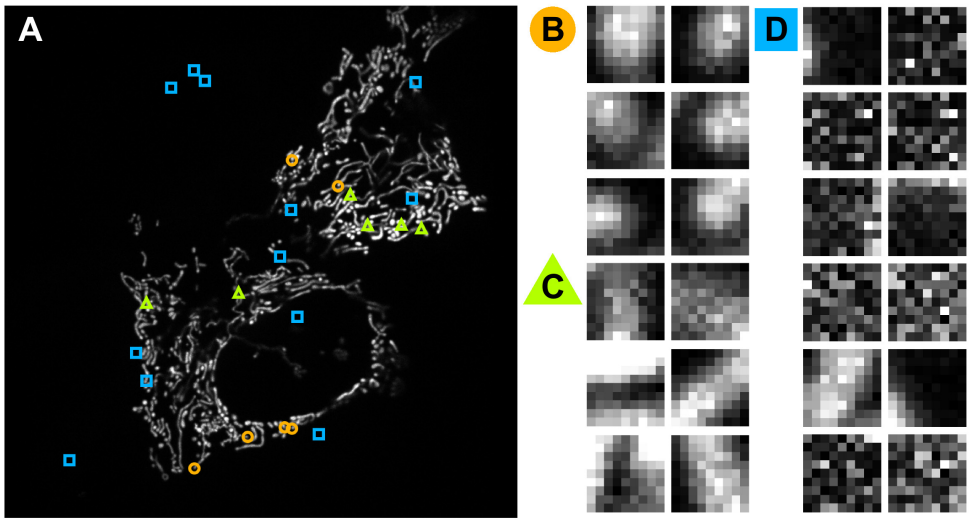


**Fig. 1.** Panel A shows the image from which the training data was extracted. The markers denote different types of training samples: manually picked positive samples (orange circles), manually picked negative samples (green triangles) and randomly selected negative samples (blue squares). The corresponding patches are visualized in panels B, C and D, where the colors and shapes correspond to the training sample types in panel A.

This is evident from the example classification results shown in Figure 2. Thus, as a post-processing step, we remove connected components that contain less than 3 pixels.

From the remaining CCs, we compute the centroids, which we use as the estimates of the tip locations. Examples of the final detection results are shown in Figure 2C.

### 2.4 Tracking

For tracking, we adopted the framework of [8], and used the authors' publicly available MATLAB implementation. In short, the method constructs tracks for detected objects in two steps: First, in each pair of subsequent frames, the objects are linked by solving a Linear Assignment Problem (LAP). If an object disappears temporarily, this procedure results in track segments instead of a complete track. Thus, as a second step, another LAP is solved to link track segments from the first step. A more detailed description of the method can be found in [8].

### 2.5 Generation of Synthetic Image Data

In order to test the performance of the method quantitatively, we needed movies for which the ground-truth location of each tip is known in each frame. To this end, we generated movies with simulated mitochondria. The advantages of using

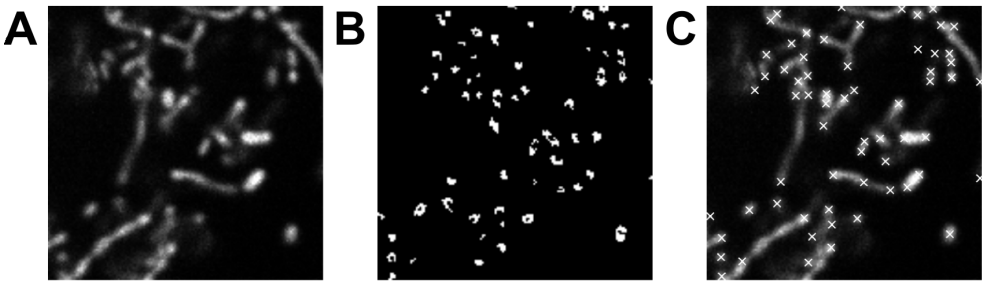


**Fig. 2.** A. Cropped region of a test image. B. Classification result. C. Tip locations (white crosses) detected as centroids of the connected components in B.

simulated images over manually-analyzed real images are that the ground-truth locations are free from human error, and that all parameters of the simulation can be varied to produce diverse test data.

To generate these synthetic movies, we modeled mitochondria as cubic splines, and subjected them to both brownian-like "wiggling" motion, and translation by applying a single random displacement vector to all control points of a mitochondrion. To simulate the wiggling, we first generated a random displacement vector for each control point. We then averaged the displacement vectors of nearby control points, in order to restrict the movements of the mitochondrial filaments to a realistic level of rigidity. To simulate the flat morphology of the U2OS cells, the distances moved by the mitochondria in the z-dimension were on average 1% of the distances they moved in the x- and y-dimensions.

The imaging process was modeled as follows: first, we rasterized the splines. To the resulting three-dimensional binary image, we added Poisson noise in order to generate variability in intensity inside the mitochondria filaments. Next, we convolved the noisy image with a Gaussian approximation of a fluorescence microscope point-spread-function. Finally, we added Gaussian noise to the image to simulate noise from the imaging system. Figure 3 shows example frames from one synthetic movie, as well as the locations of the true and detected tips.

## 3 Results

### 3.1 Performance of Tip Detection

For the RF, we trained 50 trees, and the number of features per split was selected to be 8. The latter parameter was selected using cross validation, by maximizing the area-under-ROC-curve. On the test data set, the classifier correctly classified 20/25(80%) of the positive examples, and 24/25(96%) of the negative examples.

We also compared the detection performance to that of a previous approach [9] based on segmentation. Briefly, the alternative method, here referred to as *SEG*, finds tips in 3 steps: first, it applies a median filter and a morphological

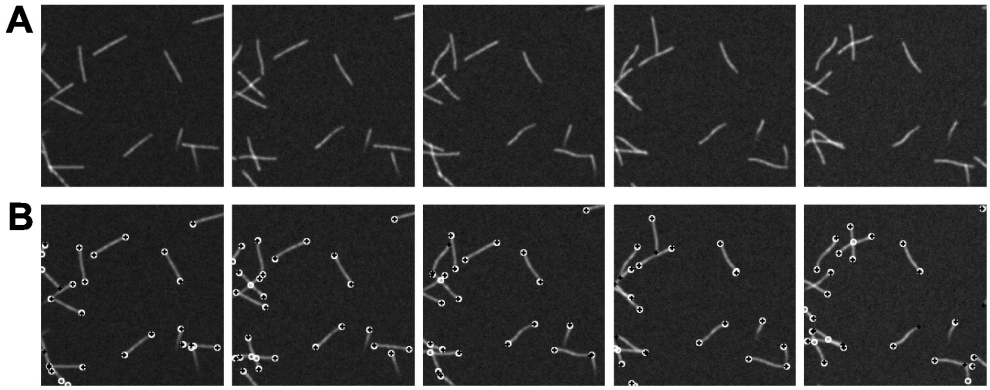


**Fig. 3.** A. Frames 0, 10, 20, 30 and 40 from a synthetic movie. B. The same frames with ground truth points (black crosses) and detected tips (white circles) overlaid.

top-hat filter to the image, to enhance the mitochondria structures. Next, it segments the image via Otsu's thresholding. Finally, it applies morphological thinning to find the morphological skeleton, from which the tips can be uniquely identified.

The methods were compared by generating 11 synthetic images, each containing approximately 20 tips, and comparing the locations of the detected tips to the ground truth tip locations. Specifically, the detected tips were paired with the ground truth tips via a Linear Assignment Problem (LAP); pairings were only made when the distance between the points was less than X pixels – points lacking such a real pair were paired with dummy points in the LAP. From the pairings, we obtained false positives as unpaired detections, true positives as paired detections, and false negatives as unpaired ground truth points. From these, we computed the true positive rate (TPR) and positive predictive value (PPV), and the F1-score defined as the harmonic mean of PPV and TPR. The results for both methods are shown in Table 1. The proposed method had only a slightly smaller TPR compared to SEG, but a significantly larger PPV. Consequently, the proposed method also yielded a larger F1-score.

**Table 1.** Detection performance of proposed method and alternative method (SEG) on synthetic images

| | Proposed | SEG |
|---|---|---|
| TPR | 0.91 | 0.92 |
| PPV | 0.89 | 0.79 |
| F1 | 0.90 | 0.85 |

### 3.2 Accuracy of Tracking

For validating the method, we generated 11 synthetic movies of 101 frames each, and stored the locations of the mitochondrial tips, to serve as the ground truth data. We quantified the tracking accuracy on the synthetic data using a similar approach to [4]: we paired each ground truth track with a *hypothesis track* generated by the tracking method. This was done via solving a Linear Assignment Problem (LAP), where the cost $c$ of each pairing is defined as

$$c(g,h) = \sum_{i \in G \cap H} \min(||g_i - h_i||_2, d_{max}) + |G \ominus H| \cdot d_{max} \tag{3}$$

Here, we used the following definitions: $|\cdot|$ denotes set cardinality, and $\ominus$ is the symmetric difference between sets. We defined $d_{max} = 2$ pixels to be the maximum distance between a ground truth point and a hypothesis point at which we can consider a point detected. $G$ and $H$ are sets that contain the indices of the movie frames in which a ground truth track and the paired hypothesis track are present (the former corresponding to frames where the object is visible, i.e. in focus), $g$ and $h$ are ground truth and hypothesis tracks, and the track points in frame $i$ are denoted by $g_i$ and $h_i$.

The effect of the latter summand in Eq. 3 is to increment $c$ by $d_{max}$ for all ground truth points that do not have a matching hypothesis point (misses), and all the hypothesis points that do not have a matching ground truth point (false detections). In the solution of the LAP, a pairing was not allowed between tracks that had no points less than $d_{max}$ pixels apart; unpaired tracks were handled by assigning them to dummy elements in the LAP.

From the paired tracks, we wanted to answer the following questions: How likely is a hypothesis track to be paired with a ground truth track? Also, how likely is a true object track to be detected by the tracking method? To address these questions, we defined the following error measures, which were calculated over all the tracks in the data set:

$$E_1 = 1 - \frac{\#\{\text{hypothesis tracks paired with a ground truth track}\}}{\#\{\text{hypothesis tracks}\}} \tag{4}$$

$$E_2 = 1 - \frac{\#\{\text{ground truth tracks paired with a hypothesis track}\}}{\#\{\text{ground truth tracks}\}}. \tag{5}$$

In addition, we wanted to quantify the quality of the detected tracks. To this end, we asked: in a correctly detected true track, how likely is a point to be missed or assigned to the wrong object? For this, we calculated the following:

$$E_3 = \frac{\sum_{i \in G} \mathcal{I}(i \notin H)}{|G|} \tag{6}$$

$$E_4 = \frac{|H \setminus G| + \sum_{i \in G \cap H} \mathcal{I}(||g_i - h_i||_2 > d_{max})}{|H|}. \tag{7}$$

Here, $\mathcal{I}$ denotes the indicator function. The quantity $E_3$ measures the fraction of ground truth points that were missed, and $E_4$ measures the fraction of hypothesis points that were assigned to a wrong object.

The results from the experiment were as follows: $E_1$, 0.09; $E_2$, 0.29; $E_3$, 0.4 and $E_4$, 0.09. Thus, 9% of the hypothesis tracks were not paired; 29% of the true tracks were not detected; 40% of the ground truth points were missed; and 9% of the hypothesis points were assigned to a wrong object. We confirmed by visual inspection that the high values for $E_2$ and $E_3$ were largely due to failures of the detection method. Another source of missed points was that some track segments were not linked by the tracking method.

### 3.3 Quantification of Mitochondrial Motion in U2OS Cells

After validating the method on synthetic data, we tested whether it can quantify changes in mitochondrial movement caused by Nocodazole, which causes the depolymerization of microtubules (see e.g. [14]). Since intracellular transport of mitochondria occurs, at least to some extent, along microtubules (see [7]), the application of Nocodazole should result in reduced mobility of mitochondria. Such a reduction in mobility has been observed in various cell types [14,15]. We can quantify this effect with the proposed method.

To this end, we computed the mean speed for each track as the mean displacement between movie frames. Figure 4 shows the distributions of the mean speeds in each cell, for the first and the last movie, and the medians of these distributions in all movies. Between the first and the last movie, the median speed has decreased by 14–44%, with the difference being statistically significant for each cell ($P < 0.01$, Wilcoxon rank-sum test, N=416–755).

## 4 Discussion

With the ongoing development of live, single-cell time-lapse imaging techniques, the objective analysis of the kinetics of mitochondria is expected to play a significant role in the detection of mitochondria-related diseases, among other. This will require the use of image analysis tools capable of automatic detection and tracking mitochondria.

We demonstrated that our method can detect the tips of fluorescently labeled mitochondria with reasonable accuracy. In addition, we showed that the direct detection of tips using the new method results in fewer false positives than a segmentation-based approach, as indicated by the higher PPV.

In our tests with synthetic data, we found that the detection and tracking approach proposed here tends to result in incomplete tracks, as well as completely missed tracks. This is in part due to some tips being missed by the detection method, and also due to the tracking method sometimes failing to link track segments. Still, as long as the application does not require complete tracks, or unique tracks for each object, the method produces reliable results in this regard.

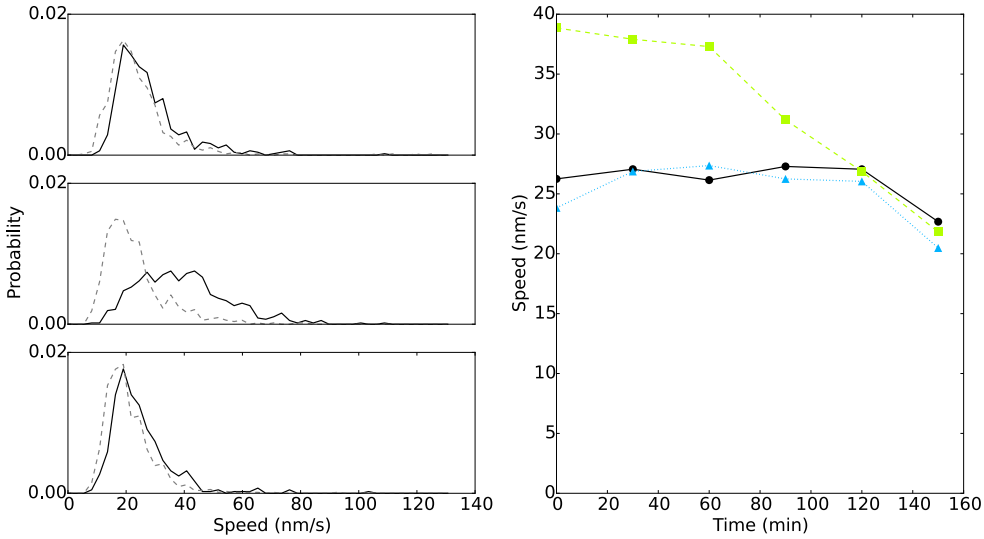


**Fig. 4.** Left: normalized histograms of mean speeds for cells 1-3 (top to bottom). The solid and dashed lines correspond to the movies captured 0 minutes and 180 minutes after application of Nocodazole, respectively. Right: medians of mean speeds for cells 1 (circles), 2 (squares) and 3 (triangles); each data point corresponds to one movie.

The results also showed that, for a correctly identified track, although many points will be missed (relating to the aforementioned issues), the tracks consist mainly of correctly identified points. This, along with the previous result, suggests that the method's results can be trusted.

Finally, by applying the method to images of U2OS cells treated with Nocodazole, we detected a decrease in mitochondrial motility in all cells; this result is consistent with previous studies [14,15].

For a more complete characterization of mitochondrial dynamics, it would be useful to keep track of not only the tips, but the whole mitochondrial filaments. This should be feasible by, for example, coupling the present method with an active contour segmentation method.

Finally, we expect our approach to be applicable to tracking other subcellular structures with similar shapes, such as microtubules and other cytoskeletal filaments, as well.

## References

1. Beraud, N., Pelloux, S., Usson, Y., Kuznetsov, A.V., Ronot, X., Tourneur, Y., Saks, V.: Mitochondrial dynamics in heart cells: very low amplitude high frequency fluctuations in adult cardiomyocytes and flow motion in non beating HL-1 cells. Journal of bioenergetics and biomembranes **41**(2), 195–214 (2009)
2. Chen, H., Chan, D.C.: Mitochondrial dynamics-fusion, fission, movement, and mitophagy-in neurodegenerative diseases. Human Molecular Genetics **18**(R2), R169–76 (2009)

3. Chen, H., Detmer, S., Ewald, A.J., Griffin, E.E., Fraser, S.E., Chan, D.C.: Mitofusins Mfn1 and Mfn2 coordinately regulate mitochondrial fusion and are essential for embryonic development. The Journal of cell biology **160**(2), 189–200 (2003)
4. Chenouard, N., Smal, I., Chaumont, F.D.: Objective comparison of particle tracking methods. Nature Methods **11**(3), 281–289 (2014)
5. Demchouk, A.O., Gardner, M.K., Odde, D.J.: Microtubule Tip Tracking and Tip Structures at the Nanometer Scale Using Digital Fluorescence Microscopy. Cellular and molecular bioengineering **4**(2), 192–204 (2011)
6. Hadjidemetriou, S., Toomre, D., Duncan, J.: Motion tracking of the outer tips of microtubules. Medical image analysis **12**(6), 689–702 (2008)
7. Hales, K.G.: The machinery of mitochondrial fusion, division, and distribution, and emerging connections to apoptosis. Mitochondrion **4**(4), 285–308 (2004)
8. Jaqaman, K., Loerke, D., Mettlen, M.: Robust single-particle tracking in live-cell time-lapse sequences. Nature methods **5**(8), 695–702 (2008)
9. Lihavainen, E., Mäkelä, J., Spelbrink, J.N., Ribeiro, A.S.: Mytoe: automatic analysis of mitochondrial dynamics. Bioinformatics **28**(7), 1050–1 (2012)
10. Saban, M., Altinok, A., Peck, A., Kenney, C., Feinstein, S., Wilson, L., Rose, K., Manjunath, B.S.: Automated tracking and modeling of microtubule dynamics. In: 3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, pp. 1032–1035 (2006)
11. Saunter, C.D., Perng, M.D., Love, G.D., Quinlan, R.A.: Stochastically determined directed movement explains the dominant small-scale mitochondrial movements within non-neuronal tissue culture cells. FEBS letters **583**(8), 1267–73 (2009)
12. Silberberg, Y.R., Pelling, A.E., Yakubov, G.E., Crum, W.R., Hawkes, D.J., Horton, M.A.: Tracking displacements of intracellular organelles in response to nanomechanical forces. In: 2008 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 1335–1338, May 2008
13. Smith, M.B., Li, H., Shen, T., Huang, X., Yusuf, E., Vavylonis, D.: Segmentation and tracking of cytoskeletal filaments using open active contours. Cytoskeleton **67**(11), 693–705 (2010)
14. Steinberg, G., Schliwa, M.: Organelle movements in the wild type and wall-less fz; sg; os-1 mutants of Neurospora crassa are mediated by cytoplasmic microtubules **564**, 555–564 (1993)
15. Yi, M., Weaver, D., Hajnóczky, G.: Control of mitochondrial motility and distribution by the calcium signal: a homeostatic circuit. The Journal of cell biology **167**(4), 661–72 (2004)

# Publication III

Eero Lihavainen, Mikhail Kislin, Dmytro Toptunov, Leonard Khiroug, Andre S. Ribeiro, "Automatic quantification of mitochondrial fragmentation from two-photon microscope images of mouse brain tissue," *Journal of Microscopy*, vol 260, no. 3, pp. 338 – 351 Dec. 2015.

Journal of Microscopy

*Journal of Microscopy, Vol. 260, Issue 3 2015, pp. 338–351* doi: 10.1111/jmi.12301
*Received 13 April 2015; accepted 9 July 2015*

# Automatic quantification of mitochondrial fragmentation from two-photon microscope images of mouse brain tissue

E. LIHAVAINEN*, M. KISLIN†, D. TOPTUNOV‡, L. KHIROUG† & A.S. RIBEIRO*
*Laboratory of Biosystem Dynamics, Department of Signal Processing, Tampere University of Technology, Tampere, Finland
†Neuroscience Center, University of Helsinki, Helsinki, Finland
‡Neurotar Ltd., Helsinki, Finland

**Key words.** Beta regression, image analysis, intravital imaging, mitochondrial fragmentation, mitochondrial morphology, two-photon microscopy.

## Summary

The morphology of mitochondria can inform about their functional state and, thus, about cell vitality. For example, fragmentation of the mitochondrial network is associated with many diseases. Recent advances in neuronal imaging have enabled the observation of mitochondria in live brains for long periods of time, enabling the study of their dynamics in animal models of diseases. To aid these studies, we developed an automatic method, based on supervised learning, for quantifying the degree of mitochondrial fragmentation in tissue images acquired via two-photon microscopy from transgenic mice, which exclusively express Enhanced cyan fluorescent protein (ECFP) under Thy1 promoter, targeted to the mitochondrial matrix in subpopulations of neurons. We tested the method on images prior to and after cardiac arrest, and found it to be sensitive to significant changes in mitochondrial morphology because of the arrest. We conclude that the method is useful in detecting morphological abnormalities in mitochondria and, likely, in other subcellular structures as well.

## Introduction

Mitochondria play a key role in the functioning of neurons, regulating energy consumption, ionic homeostasis, $Ca^{2+}$ and redox signalling, among other processes (Mattson et al., 2008). In axons and dendrites, mitochondria function within a dynamic reticulum that is modified by growth, fission, fusion, traffic and autophagy. Proper mitochondrial dynamics in neurons is required for normal synapse density and activity-dependent synapse formation (Li et al., 2004). Disturbance in mitochondrial physiology has been linked, directly or indirectly, to a variety of pathogenic processes in the central nervous system. Studies on Alzheimer's, Huntington's and Parkinson's diseases identified mitochondrial dysfunction as a responsible factor in neurodegeneration (Schon & Przedborski, 2011). The existence of disease-related mutations in genes encoding proteins associated with mitochondrial dynamics, such as Mfn2 and OPA1, further highlight the role of mitochondrial dynamics in neuronal health (Knott & Perkins, 2008).

There are several reports of morphological abnormalities in neuronal mitochondria in animal and *in vitro* disease models. Evidence suggests a bidirectional relationship between mitochondrial morphology and function (Picard et al., 2013). Some diseases alter the mitochondrial morphology and dynamics (see Lin & Beal, 2006, for a review). By contrast, an imbalance in mitochondrial fusion and fission leads to excessive mitochondrial elongation or fragmentation and can trigger neurodegenerative diseases (see, e.g. Knott & Perkins, 2008). Thus, understanding how mitochondrial dynamics and fusion–fission mechanisms are related to the shape of the mitochondrial network may be central in the development of novel treatments for a broad spectrum of neurological disorders. We believe that the combination of animal disease models with the analysis of mitochondrial morphology in living brains will assist in this effort.

In this paper, we are concerned with the quantification of the degree of mitochondrial fragmentation. The problem of automatically quantifying fragmentation of cellular organelles and structures has been tackled by some authors. Recently, Chen et al. (2011) published a technique for quantifying morphological changes in dendrites. This problem bears similarity to ours, as the visual appearance of the dendrites changes from curvilinear to fragmented. Their method relies on comparing an image to a reference image, where the dendrites are assumed to be intact, that is, to exhibit only elongated morphologies. As mitochondrial morphology, including their degree of fragmentation, varies even between healthy cells, this method is not directly applicable to our problem. Furthermore, the reliance on a reference image restricts the applicability of the method to cases where such an image is available.

Correspondence to: Eero Lihavainen, Laboratory of Biosystem Dynamics, Department of Signal Processing, Tampere University of Technology, P.O. Box 527, FI 33101 Tampere, Finland, Tel: +35-8408280531; e-mail: eero.lihavainen@tut.fi

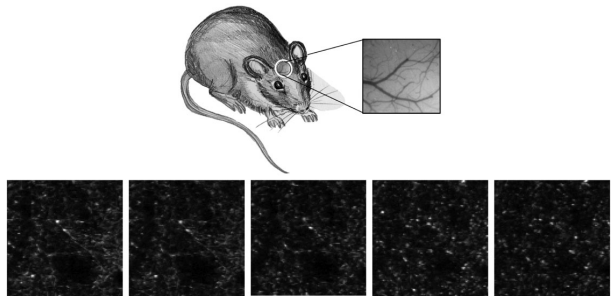

**Fig. 1.** Mitochondria undergoing fragmentation in mouse brain tissue. Top: The microscopy was performed in anaesthetized mice through a cranial window (see 'Materials and Methods' for details). Bottom: From left to right, this sequence of two-photon fluorescence microscope images exemplifies the progression of mitochondrial fragmentation over time, following cardiac arrest (see 'Materials and Methods' for details). The images are taken at intervals of 66 s.

Mitochondrial fragmentation results in an increase in the number of disconnected mitochondria, with shorter average lengths than before; indeed, a large number of short mitochondria has been proposed as a criterion of fragmentation (Farrand et al., 2013). To measure the lengths, the organelles need to be segmented. Although methods exist for the segmentation of mitochondria from images of tissue culture cells (Peng et al., 2011; Lihavainen et al., 2012), these were found to not always be reliable to analyse the images of mitochondria in cortical tissue here used as source of information. In particular, the resolution and signal-to-noise ratio in these images (when acquired without optical zooming and at non-damaging laser exposure conditions) are too low for accurate segmentation and, consequently, accurate length measurements. These features were limited to allow studying the fragmentation of mitochondria in a large region, while minimizing photodamage.

We propose a method for the absolute quantification of the degree of mitochondrial fragmentation from two-photon microscope images of mouse brain tissue, where mitochondria are fluorescently labelled. Figure 1 shows a sequence of images illustrating the process of fragmentation of initially intact mitochondria. Noting from experience that this task is relatively easy for expert human observers, we use a supervised learning approach. Our goal is to train a model to predict a *fragmentation score* from a local, square image region herein referred to as a *patch*. Because we are estimating a continuous quantity, our problem is one of regression. Thus, we formulate the problem as a Beta regression model (Ferrari & Cribari-Neto, 2004), which we train using example patches scored by several experts. The continuous response variable is restricted to the interval (0,1), representing the range from completely intact to completely fragmented mitochondria. Our approach is therefore somewhat similar to that of (Lin et al., 2010), where the authors used k-nearest neighbour regression.

As explanatory variables for the regression model, we use features extracted from the patch. First, we extract several morphometric features based on the segmented mitochondria. In addition, to cope with the limited accuracy of the segmentation, we quantify alternative features that inform on the degree of fragmentation. Namely, we use features that describe image texture because images with different levels of mitochondrial fragmentation exhibit different textures. In addition, to cope with the low resolution and signal-to-noise ratio of the tissue images, our method uses a Support Vector Machine classifier (Cortes & Vapnik, 1995) to automatically exclude locations where the mitochondrial phenotype is unclear.

Here, we first describe our approach to quantifying mitochondrial fragmentation, and evaluate its performance. Next, we compare the performance of our method to previous ones [namely, the one proposed in (Lin et al., 2010), which also does not require use of reference images]. In addition, we test our method on synthetic data. We then present results from the application of our method to experimental data, namely, on two-photon microscope images of fluorescently labelled mitochondria in cortical tissues, prior to and following cardiac arrest (CA). As mitochondria have been shown to undergo fragmentation during ischemia *in vitro*, using simulated ischemia in HL-1 cells (Brady et al., 2006), and *in vivo*, in rat kidneys (Hall et al., 2013), we use it to test the performance and applicability of the method. The results show that the method is able to distinguish, quantitatively, the levels of fragmentation between animals and in individuals, prior to and after CA.

## Materials and methods

### *Animal procedures and in vivo two-photon microscopy*

Recently, transgenic (Tg) mouse lines with mitochondria-targeted fluorescent proteins (FPs) have been developed (Misgeld et al., 2007; Shitara et al., 2010). These have several properties that ensure a robust fluorescence signal, making them useful for studies of mitochondria in disease models: First, the expression of FPs is stable, providing the constitutive

localization of mitochondria. Second, to the best of our knowledge, mitochondrial remodelling or fragmentation does not lead to redistribution of the FPs from the mitochondrial matrix to the cytosol. Finally, the FPs are resistant to photobleaching and have low pH sensitivity, which permits reproducible long-term analysis of mitochondrial dynamics *in vivo* over extended periods of time. These mice can be used in inducible animal models of diseases, or crossed with chronic neurodegenerative genetic models to evaluate the role and importance of mitochondria in cell or tissue degeneration. A particularly attractive technique for studying mitochondria in these lines is intravital two-photon microscopy, which provides repetitive access to mitochondria in their native environment (Liu & Murphy, 2009; Bilsland et al., 2010; Nikić et al., 2011), and makes it possible to quantify mitochondrial fragmentation in three-dimensional (3D) volumes.

For the above reasons, in this study, we used nine Thy1-mito-CFP (Misgeld et al., 2007) 2–5-month-old mice. All the procedures were performed according to the guidance for animal care (EU directive 2010/63/EU). Animal license (ESAVI/2857/04.10.03/2012) was obtained from the local authority (ELÄINKOELAUTAKUNTA-ELLA). The mice were bred in the certified University's animal facility and kept in group cages on a 12-h light/12-h dark cycle with food and water *ad libitum*. For chronic cranial window implantation and imaging sessions, the mice were anaesthetized (i.p.) with a mixture of ketamine (80 mg $kg^{-1}$) and xylazine (10 mg $kg^{-1}$). Depth of anaesthesia was evaluated by hind-paw pinching, and to maintain adequate anaesthesia during the procedure, an additional one-quarter of the dose was administered when needed. During surgical operations, imaging sessions and the first hour after recovery from anaesthesia, the core temperature was maintained at 37.0°C using a heating pad.

Implantation of the cranial window was performed according to published protocols (Ricardo & Carlos, 2008; Holtmaat et al., 2009) with minor modifications. Briefly, the skull was gently drilled over the somatosensory cortex to form a window 3–3.5 mm in diameter. The skull was then covered with sterile cortical buffer (125 mM NaCl, 5 mM KCl, 10 mM glucose, 10 mM HEPES, 2 mM $CaCl_2$, 2 mM $MgSO_4$). Cranial bone was carefully removed, and a 5 mm diameter #1.5 glass cover slip was placed over the cranial window and sealed to the skull using polyacrylic glue. To stabilize the head during the imaging sessions, a custom made metal holder was glued to the top of the glass coverslip and then hardened with dental cement. After the operation, the mice were placed in a warmed cage with easily accessible food and water. When the mice fully recovered from anaesthesia they were returned to the group-housed cage and allowed to recover for at least 1 month. During this recovery period, postoperational damage and inflammation is subsiding and the window regains its transparency.

The *in vivo* two-photon microscopy was performed on FV1000MPE (Olympus Corporation, Tokyo, Japan) equipped with the 1.05NA 25× water immersion objective. Animals were placed under the microscope by attaching the metal holder to the custom built frame. Cyan FP (CFP) was excited with Mai Tai DeepSee laser (Spectra-Physics, Santa Clara, CA, USA) tuned to 860 nm and fluorescence was collected using 420–500 and 515–560 nm band pass filters.

To produce complete global cerebral ischemia, we induced CA via KCl infusion utilizing a tail vein catheter (Braintree Scientific, Inc., Braintree, MA, USA). To monitor the morphology of neuronal mitochondria a z-stack with an imaging field comprising 500 × 500 $\mu$m in the *xy* axes (0.497 $\mu$m pixel size) and with spacing of 1 $\mu$m was recorded before and after CA. To access time course of mitochondria fragmentation in 1 cortical layer, typically within 30–70 $\mu$m of the pial surface, a small stack of images with the same imaging field and 3 $\mu$m z-spacing was acquired every 10 seconds, during 10 min. A total of 50 frames of the same z-plane were taken per animal CA was induced during imaging by intravenous injection of 300 $\mu$l 3M KCl, and the control animals received 0.1M Phosphate buffered saline (PBS). Heart rates were monitored using a pulse oximeter (MouseOx, Starr Life Sciences, Oakmont, PA, USA) with a mouse thigh sensor to confirm CA, typically within 0.5–3 min after KCl injection.

### *Extraction of the training data*

To train the regression model, we selected 1000 random locations from all the image data, and extracted 103 × 103 patches around each location; here, we will refer to these 1000 patches as *exemplars*. The patch size was set such that there are enough mitochondria in the patch to allow an expert to perform their assessment, while still keeping the patch as small as possible to provide localized quantification.

Five experts independently scored each patch using a 21-point scale ranging from 0 to 100 with intervals of 5 (see Materials and Methods for details on the scoring procedure). Patches that either contained no mitochondria or were too noisy to allow an assessment of the mitochondrial morphology were not given a score. The scoring from one expert was excluded, because it was weakly correlated (Pearson correlation coefficient less than 0.6) with the scorings of a majority of the other experts, indicating an incongruent interpretation of the task.

Before the manual scoring of patches, the experts were familiarized with 12 large-scale z-stacks and 200 random patches from control animals and animals subject to CA. They were then given the following guidelines for scoring:

- At least 25% of the image should contain perceptible signal from mitochondria to assign a score.
- Based on the ratio of globular to continuous (tubular) structures of mitochondria, the image should be allocated to the score range 0–100 (in intervals of 5). It was then suggested that, usually, if the tissue looks healthy (similar to control)

**Table 1.** Pairwise biases, computed as mean differences, of all expert scorings (E1–E4), computed from the test data

| | E1 | E2 | E3 | E4 |
|---|---|---|---|---|
| E1 | 0 | 0.07 | 0.00 | −0.06 |
| E2 | | 0 | −0.05 | −0.05 |
| E3 | | | 0 | 0.06 |
| E4 | | | | 0 |

the score should be below 20; if the image contains at least one tubular mitochondrion, the score should be below 50; and if there are no visible tubular mitochondria and less than half of patch contain fragmented mitochondria, the score should be below 70. As such, scores above 70 should be assigned only when at least half of the image contains swollen or globular mitochondria (similar to the images after CA).

As the ground truth score for each patch, we considered the mean of the scores given by the four remaining experts. For this to be valid, we needed to make sure that there was no significant bias in the scorings of any of the experts. To study this, we computed the mean differences between each pair of scorings (Table 1). Given that these pairwise biases were small, and that each expert utilized the whole range of scores, we found no need for bias corrections. When more than two experts gave no score for a patch, that patch was marked as *rejected*, and was not used for training the regression model. There were 483 rejected patches in total, and thus 517 accepted patches. Examples of patches with different scores and rejected patches are shown in Figure 2.

For training the regression model, we split the 517 accepted exemplars to training and test sets, their respective sizes being 362 and 155. To ensure that all the score values were represented in similar frequencies in both sets, we compared the sets with a two-sample Kolmogorov–Smirnov test; we repeated the random splitting until the test returned a $p$-value smaller than $10^{-4}$, which indicated that the distributions were sufficiently similar.

As the Beta regression model assumes the response variable to lie in the open interval (0,1), we had to convert the training and test scores to satisfy this assumption. To this end, we divided the scores by 100, resulting in values in the closed interval [0,1]. We then further mapped these normalized scores $y$ to (0,1) by setting $y' = (y(N-1)+0.5)/N$, where $N$ was the number of samples, as suggested in Smithson and Verkuilen (2006).

## Image analysis

All methods were implemented in the MATLAB language. Software code implementing the method, along with pretrained models, is available at http://www.cs.tut.fi/~lihavain/phenometer. Figure 3 illustrates the components of the method.

### *Extraction of morphometric features*

We defined a set of 13 simple morphometric features (Table S1) that, intuitively, should be descriptive of mitochondrial fragmentation. These were derived from various shape properties such as lengths and widths, which can be measured following the segmentation of mitochondria.

We segment mitochondria using a method modified from (Lihavainen et al., 2012). First, a Difference-of-Gaussians (DOG) filter is applied to the image. This has the effect of making the mitochondria brighter while removing low-frequency intensity fluctuations. We threshold the resulting image, considering pixels with value above zero to be mitochondria and all others to be background. This results in a binary mask inside which the mitochondria are segmented, but that may contain some false-positive objects as well; we apply a morphological erosion to remove the smallest, noise-induced objects, and also to shrink the masks to better match the sizes of the mitochondria. To further remove possible nonmitochondrion objects from the mask, we discard each connected component whose mean intensity, computed from the original image, is lower than the mean + 1 standard deviation of the background intensity, computed from the original image at pixels not containing mitochondria. Finally, we remove all objects smaller than three pixels or larger than 3000 pixels in area, as they do not fit the *a priori* expectations of the size of a mitochondrion, and likely originate from segmentation errors.

From each connected component in the segmented image, we compute the lengths of the major and minor axes of an ellipse fitted to the shape of the component [using the 'regionprops' function from MATLAB Image Processing Toolbox (MATLAB and Image Processing Toolbox Release 2012b, 2012)], and the area of the morphological skeleton of the component, which approximates the length of the mitochondrion. The ellipse fitting and the morphological skeleton are illustrated in panels B and C of Figure 4, respectively. The full list of morphometric features used is: mean major axis length; mean minor axis length; minimum, mean and maximum of the reciprocal aspect ratios, that is the ratios of minor axis lengths to major axis lengths; the minimum, maximum and the 1st, 2nd and 3rd quartiles of the skeleton areas; mean and variance of the skeleton areas; and, finally, a weighted mean of the major axis lengths, where the weights are the major axis lengths themselves. The rationale for the last feature is that the longer the mitochondria are, the fewer will fit into the image, and thus longer mitochondria should have a larger weight.

### *Extraction of texture features*

To complement the morphometric features, we additionally computed a set of features describing image texture, consisting of summary statistics of responses from processing the image with a filter bank (Zujovic et al., 2013). As differences in illumination or in the fluorescent signal emitted by different

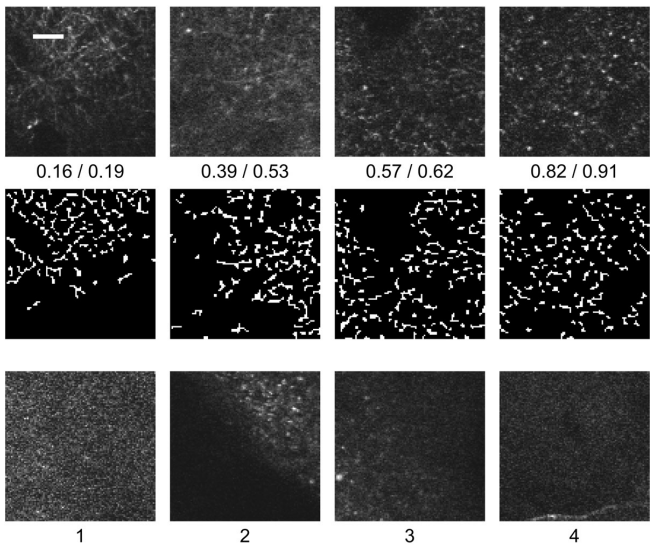


**Fig. 2.** Examples of training data and segmentation results. (B) The top row shows four examples of accepted patches, and their respective scores (expert average/score output by our method). The middle row shows the corresponding segmentations, used for computing the morphometric features. The bottom row shows examples of typical low-quality patches that were rejected: numbered from left to right, these represent (1) lack of visible mitochondria/low signal-to-noise ratio, (2) clear occlusion of a large portion of the patch by a vessel and (3–4) too few visible mitochondria for making a decision. The scale bar corresponds to 10 $\mu$m. The image brightness was adjusted for better visualization.

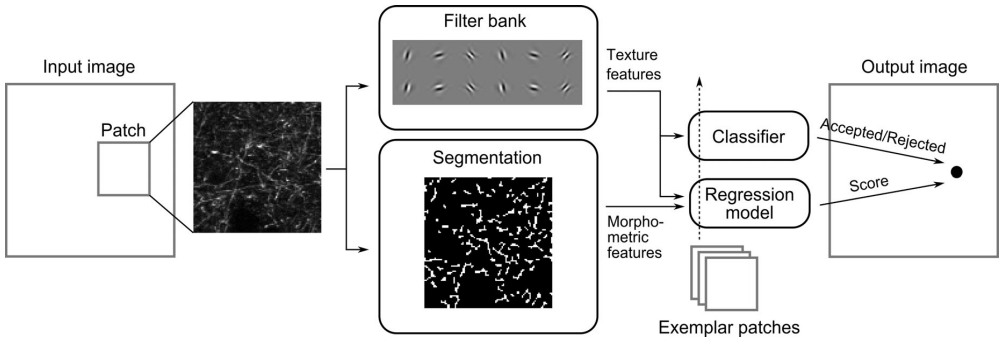


**Fig. 3.** Overview of the method. The fragmentation score is calculated for individual patches cropped from an image. Extraction of features from a patch is performed in two steps. To extract texture features, the patch is processed by a multiscale wavelet filter bank (impulse responses of the 12 filters at one scale are shown). Summary statistics of the filter responses are used as features. Morphometric features are extracted from segmented mitochondria by measuring their shape properties. The features extracted from a patch are then fed to the classifier and the regression model, previously trained using the exemplar patches scored by experts (training represented by the dashed arrow). Note that the classifier uses only the texture features. The regression model outputs a score in the range (0,1), whereas the classifier outputs a decision of whether the patch is accepted or not, i.e. whether the obtained fragmentation score can be considered reliable or not.

regions of the cell tissue would affect the values of the texture features, and therefore the quantification of the degree of fragmentation of the mitochondria, we normalized the intensities of each exemplar image to zero mean and unit variance prior to texture feature extraction. This makes the method invariant to such differences.

To extract texture features from a patch, it is processed with a filter bank of six directional filters at three different scales. For the filter bank, we chose to use the Dual-Tree Complex Wavelet Transform (DT-CWT) because of its computational efficiency (Selesnick et al., 2005); it provides an alternative for the commonly used Gabor filters (Eleyan et al., 2008). The

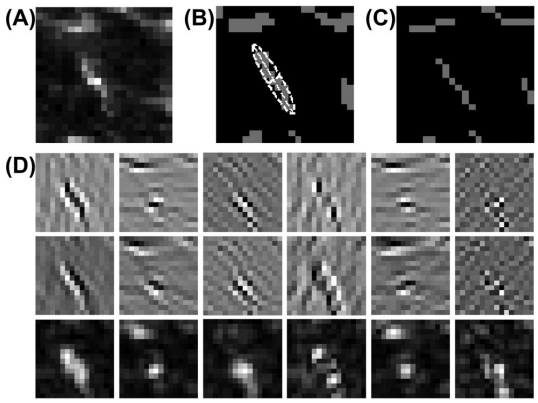


**Fig. 4.** Feature extraction. (A) A cropped region showing a few mitochondrial filaments. (B) For extracting morphometric features, an ellipse (white dashed line) is fit to the shape of a segmented mitochondrion (shown in gray). The lengths of the major and minor axes of the ellipse are used to estimate the length and width of the mitochondrion. (C) Morphological skeleton of the segmentation result (gray), from which the areas of connected components are used to estimate the lengths of the mitochondrial filaments. (D) Output of the DT-CWT filter bank with one scale for the region in panel A. The first two rows show the real and imaginary parts of the complex filter responses, and the third row shows their magnitudes. The columns correspond to filters of different angles. The magnitudes of the filter responses have been rescaled for better visualization.

transform is implemented as two parallel filter banks, whose outputs are interpreted as the real and imaginary parts of a complex-valued filter bank. In effect, the transform filters the image at each scale with six complex filters that respond to oriented structures. For computing the texture features, we use only the magnitudes of the complex filter responses corresponding to each complex wavelet (see Fig. 4D for an example of the filter responses);there are thus 20 filter responses, or subbands: 18 from the directional filters (six directions for each three scales), and two low-pass subbands. From the subbands, we extract several summary statistics, which have been found to perform well as texture features (Zujovic et al., 2013): mean and variance of each subband, mean horizontal and vertical autocorrelations for each subband, cross-band correlations for each pair of subbands having the same orientation but different scale, and cross-band correlations for each pair of distinct subbands at the same scale. As the subbands are required to be the same size for computing the cross-band correlations, we use an undecimated version of the DT-CWT, obtained in a manner analogous to the undecimated Discrete Wavelet Transform (Holschneider et al., 1990). A full list of the features is presented in Table S1.

Finally, we attempted to include in the model a common texture feature set, namely, Haralick's gray-level co-occurrence matrix features (Haralick et al., 1973), but this did not improve the results. Thus, we left these out of the final model.

*Classifier for detecting regions of interest*

We train a binary classifier to detect regions of interest (ROIs), that is parts of the image where mitochondria are visibly present. For this, we use a Support Vector Machine with a radial basis function kernel. The two hyperparameters of the Support Vector Machine – box coefficient and radial basis function bandwidth – are optimized using a grid search, where both parameters take 25 distinct logarithmically spaced values between $10^{-5}$ and $10^{5}$. The best hyperparameters are chosen as the ones resulting in the highest precision among those yielding a recall of at least 0.9.

*Feature selection*

Before training the final Beta regression model, for the purposes of feature selection, we assume a linear model on a transformed version of the response. We first map the scores from the interval (0,1) to the set of real numbers via the probit function $g$, which is the inverse of the Gaussian cumulative distribution function. We assume that the transformed score $g(y)$ of a patch depends on the feature vector computed from the patch via a linear model. We fit this model using Elastic Net (EN) (Zou & Hastie, 2005), with mixing parameter $\alpha = 0.5$, and with 5-fold cross-validation to estimate the prediction error. Using this fitting method allows the automatic selection of relevant features while including several correlated features instead of choosing just one; this is helpful in interpreting the resulting coefficients. From the set of increasingly complex models fit by EN, following a common heuristic (Hastie et al., 2009), we select the simplest model $\boldsymbol{\beta}$ whose cross-validation error is no more than 1 standard error higher than the minimum cross-validation error. Given a feature vector $\boldsymbol{x}$, the predicted score is then given by $\hat{y} = g^{-1}(\boldsymbol{x}^{\mathrm{T}}\boldsymbol{\beta})$.

*Regression model for quantifying mitochondrial fragmentation*

Using the set of automatically selected features, we apply Beta regression (Ferrari & Cribari-Neto, 2004) to find the final model. In this framework, the score is modelled as having a Beta distribution, whose domain is the open interval (0,1), and which can be characterized by a location parameter – the mean $\mu$ – and a dispersion parameter $\phi$ (Ferrari & Cribari-Neto, 2004). The mean of the distribution is assumed to depend linearly on the features $\boldsymbol{x}$ via the link function: $y \sim Beta(g^{-1}(\boldsymbol{x}^{\mathrm{T}}\boldsymbol{\beta}), \phi)$, and the parameters $\boldsymbol{\beta}$ and $\phi$ are estimated by numerically maximizing the likelihood function, which was derived in (Ferrari & Cribari-Neto, 2004). As in the feature selection, we use the probit link function. The advantage of this model over alternatives, e.g. the ordinary least squares with probit-transformed response that we use in the feature selection step, is that it accounts for non-constant variance.

In general, it is not necessary to evaluate the model at each pixel in an image because patches centred at adjacent pixels

overlap heavily and therefore yield similar scores. In addition, since the features are averaged inside the patch, the patch size limits the frequency of local changes that can be observed, and too dense sampling will result in artefactual high-frequency changes. In our experiments, we set the distance between horizontally and vertically adjacent windows at 20 pixels, corresponding to 9.94 $\mu$m in the sample. This value was chosen because it was the smallest one for which the result did not appear to contain artefacts, therefore yielding the highest resolution at which the local changes were informative.

### *Computational efficiency*

On a laptop computer with a 2.50 GHz dual-core processor, the processing of one of our microscopy images took 2 min on average. For our analyses, we used a computational grid at Tampere University of Technology. By distributing the processing of individual images to different nodes, 20-fold speedups were achieved.

## Generation of synthetic image sequences

As the data for training and validation of the method was acquired with one imaging system with fixed settings, it was uncertain if the method would work following changes in the properties of the imaging system; thus, we studied the robustness of the quantification to changes in properties such as resolution and noise. In particular, we investigated how much the score varied as a function of the resolution and the noise level of the imaging system, by applying the method to synthetic time-lapse image sequences of mitochondria, which were generated by simulating the dynamics of the organelles, and the effects of the imaging system.

To generate synthetic image sequences, we initialized N mitochondria as 3-D cubic splines, such that at least one end point of each was contained in a volume of size $200\times200\times15$ pix (here, "pix" represents the distance between two consecutive pixels in the x-y-plane), and the initial shape of the mitochondria was linear. In our simulations, N was selected randomly from the range 51—151, and each mitochondrion was approximately 150 pix in length. At each instant, we resampled the splines such that the distance between consecutive control points was approximately constant at 0.5 pix.

To simulate the diffusive motion of a mitochondrion, we randomly displaced each control point by adding a zero-mean Gaussian random number to each coordinate. We then applied a running mean filter of size 20, corresponding to a length of approximately 5 pix, to the coordinates of the spline, in order to simulate the rigidity of the organelle. Finally, each mitochondrion was displaced by a small "drift", by generating one zero-mean Gaussian random vector and adding it to each control point.

Fission was simulated at each time step by splitting each mitochondrion into two with a probability of $p = 0.1$, and choosing the split point randomly, such that both resulting mitochondria consisted of at least two control points. In the first few time steps, we simulated only the motion of mitochondria without fission, in order for the mitochondria to deviate from their initial, linear shapes prior to fission events. These first steps were not used.

For simulating imaging, we created an image of zeros with size $200\times200\times15$ pix, whose coordinates corresponded to those of the volume where the mitochondria were initialized. Each spline was rasterized by densely resampling it, rounding the coordinates of each control point towards the nearest integer, and setting a pixel to the value 1 if it coincided with any of the resulting points. This resulted in mitochondria approximately one pixel wide. Next, to simulate both background noise and uneven distribution of fluorescence, we replaced the value of each pixel with a Poisson random variate, with means of 1 and 100 for background and mitochondria pixels, respectively. We then filtered the image with the 3-D anisotropic Gaussian filter used to approximate the point spread function. Also, we added Gaussian noise to the result, and set all negative values to zero.

Finally, from the centre of each simulated z-stack, we cropped a $103\times103\times1$ pix patch to be the final patch used for the analysis.

## Results

### *Validation of the ROI classifier*

For training the ROI classifier, we extracted texture features from all exemplars, and split the feature vectors randomly into training, validation and test sets of sizes 400, 400 and 200, respectively. Using the training set, we then trained several Support Vector Machine classifiers with different parameter combinations, and selected as the final classifier the one that had optimal performance on the validation set, as described in Materials and Methods. The performance of the final classifier was estimated using the test set; the classifier achieved a precision of 0.92 and a recall of 0.90.

### *Validation of the regression model*

First, we wanted to analyse the importance of the individual feature sets, segmentation-based morphometric features (MF) and texture features (TF). For this purpose, we trained three Beta regression models: one with only MF, one with only TF and one with both (MF + TF). Before training each model, we applied automatic feature selection to reduce the number of features used. This was done to prevent overfitting and also to find a relevant set of features, making the interpretation of the model easier. In the case of MF + TF, from the total of 160 features, 40 were selected by the feature selection procedure. Of these, 10 were morphometric and 30 were texture features. The full list of selected features is in Table S1.

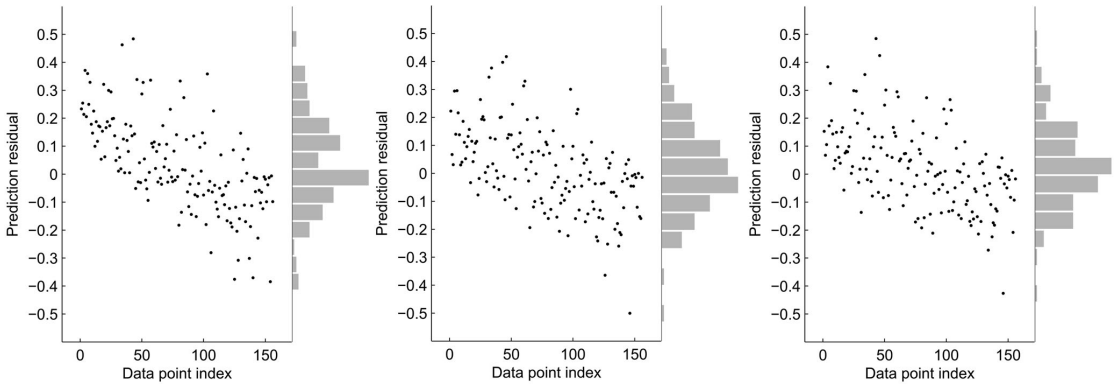


**Fig. 5.** Prediction residuals. The scatterplots show the predicted scores for the models with only morphometric features (left), only texture features (middle) and both sets of features (right). Each point in the scatterplot is the difference between a ground truth score from the test dataset and the corresponding predicted score. The *x*-axis values are the indices of the ground truth values ordered from low to high. In each case, the residuals exhibit a decreasing trend, which shows that the scores are biased in the extreme ends of the range. Also, in all cases, the histogram of residuals on the right side of the scatterplot spreads almost symmetrically around zero, indicating that there is no significant overall bias.

To summarize the feature selection results, the reciprocal aspect ratio features had positive coefficients, and the features based on the area of the morphological skeleton had negative coefficients. In other words, as expected, the fragmentation score is negatively correlated with the length of the mitochondrial filaments. The coefficients of the texture features show the following patterns: for a small number of filters, the mean filter response was selected with a positive coefficient, and for some filters the variance of the filter response was selected with a negative coefficient. The vertical and horizontal autocorrelations of the filter responses had both positive and negative coefficients. Finally, the cross-band correlations had negative coefficients for many pairs of bands of the same direction but different scale, whereas most of the correlations between filters of different directions had positive coefficients.

We next trained the models with the selected features, and assessed the performance of each model on the test dataset. The fraction of variance in the scores that is explained by a model was estimated by a pseudo-$R^2$ coefficient (Ferrari & Cribari-Neto, 2004). As a measure of the prediction accuracy, we computed the mean absolute error (MAE) between the ground truth scores and the scores predicted by the model. The respective values for the pseudo-$R^2$ and the MAE were: 0.592 and 0.130 for MF, 0.668 and 0.118 for TF and 0.683 and 0.113 for MF + TF. Thus, for further studies, we selected MF + TF as the model to be employed. The prediction residuals for the test samples for each of the different models are shown in Figure 5. The figure reveals a slight bias in the prediction for all models. The predicted score for the nonfragmented mitochondria is overestimated, and the predicted score for the fragmented mitochondria is underestimated; however, this bias is least apparent in the case of MT + TF.

**Table 2.** Pairwise mean absolute differences of all expert scorings (E1–E4) and the model prediction, computed from the test data

| | E1 | E2 | E3 | E4 | Model |
|---|---|---|---|---|---|
| E1 | 0 | 0.17 | 0.11 | 0.14 | 0.11 |
| E2 | | 0 | 0.21 | 0.24 | 0.20 |
| E3 | | | 0 | 0.23 | 0.13 |
| E4 | | | | 0 | 0.20 |
| Model | | | | | 0 |

Next, we compared the method to that proposed in (Lin et al., 2010) for quantifying mitochondrial fragmentation from images of single cells. Their method uses a different set of features, and the fragmentation is quantified using weighted 6-nearest neighbour regression. To perform the comparison, we trained our model with their data, which consisted of 392 training images along with expert scores in the range 0 (intact)–1 (fragmented), and a 42-sample test dataset, for which their method's output was also given. As their images differ from ours in that they are more blurry (because of using epifluorescence instead of two-photon microscopy) and differ in image resolution, we preprocessed them by applying an unsharp mask filter (to remove background fluorescence), and resizing to 0.32 times the original width and height (such that the thickness of the mitochondria corresponded approximately to that in our images). We then padded the resulting image to size 103 × 103 by mirroring it along the edges. On their test dataset, our method performed better as it reached an MAE of 0.073 and a pseudo-$R^2$ of 0.869, while using their method the values were 0.103 and 0.748, respectively.

We also tested how our method would work with images of lower resolution than the ones used here. For that, we

reduced the resolution of the exemplar patches by decimating the images. We then repeated the training procedure on the decimated images. We observed that the performance was not significantly worse provided that the image size (in number of rows and columns) equalled at least 0.7 times the original size (pixel distance of 0.497 $\mu$m); as such, a pixel distance of 0.71 $\mu$m can be considered to be an upper limit for the method proposed.

Finally, to determine if the results of our method can be considered a reliable replacement for expert scores, we determined the MAE between the model predictions and the scorings of the individual experts for the test data. The MAEs for all expert-expert pairs and expert-model pairs are shown in Table 2. It is apparent that the agreement of the model with any expert, as measured by MAE, is close to the best agreement between experts.

### *Analysis of the method using synthetic images*

For studying the performance of the method with synthetic data, we generated 10 image sequences with randomized initial numbers and lengths of mitochondria, and applied the imaging system model with each parameter combination ($\sigma_{\mathrm{PSF}}$, $\sigma_{\mathrm{noise}}$). Examples of different resolutions are shown in Figure 6, and of different noise levels in Figure 7. The results are visualized in Figure 8. From the latter figure (top and bottom rows), it is visible that a resolution too low or too high (i.e., not of similar resolution to the microscopy images) will cause the score to be underestimated when the true degree of fragmentation is high. The effect of noise is similar: strong noise will cause the score to be underestimated. Nevertheless, in the range of $\sigma_{\mathrm{PSF}} = 0.35$–$0.52$ $\mu$m and $\sigma_{\mathrm{noise}} = 1$–$5$, the results are good: the mean scores at the last frame range from 0.716 ($\sigma_{\mathrm{PSF}} = 0.43$ $\mu$m and $\sigma_{\mathrm{noise}} = 5$) to 0.845 ($\sigma_{\mathrm{PSF}} = 0.35$ $\mu$m and $\sigma_{\mathrm{noise}} = 3$). Also, in the first time moments, where the mitochondria are longer than in real images, the method gives scores very close to zero.

### *Mitochondrial fragmentation following CA*

Next, as an application, we used the method to quantify the degree of mitochondrial fragmentation in mouse brain tissue during ischemia, resulting from induced CA. In particular, we assessed how the scores differ before and after CA, and how the scores depend on the distance to *pia mater*. To this end, we imaged a cranial window of 7 anaesthetized Thy1-mito-CFP mice using a two-photon microscope. A z-stack of 121–441 slices with spacing of 1 $\mu$m was recorded at least 20 min before inducing CA, and another one starting 3–5 min after CA. In all images, we evaluated the score in a subset of the image pixels, corresponding to a sampling grid of $47 \times 47$ points, as described in the Materials and Methods. An example pair of z-stacks before and after CA is visualized in Movie S1.

As an example of the results, in Figure 9A we show the scores of one individual, as heatmaps, at five different depth levels, prior to and after CA. Below these, we show the normalized histograms of scores of the pre-CA (blue) and the post-CA (red) scores, for each depth level, as calculated from all animals by pooling all scores from the same depth for each z-stack. From the heatmaps, close to the pia mater surface, the scores are lower before CA than after CA. However, as the distance from the pia mater surface increases, these scores increase and, from the score histograms, become very similar between pre-CA and post-CA. This is also visible in Figure 9B, which shows the average score as a function of the distance to pia mater for the pre- and post-CA images. Near the pia mater, the pre- and post-CA scores differ widely from one another, and in both cases the average scores change little between 0 and 45 $\mu$m. However, as the depth increases, the difference between the pre-CA and post-CA scores decreases, and after 100 $\mu$m, the difference between the mean pre-CA score and the mean post-CA score becomes not statistically significant (*t*-test, $p < 0.01$).

Next, to obtain a more detailed view of the dynamics of the process of fragmentation induced by CA (Movie S2), we performed an experiment using small z-stacks (three slices, with spacing of 3 $\mu$m) that were recorded every $11.0 \pm 0.2$ s for almost 10 min. CA was induced in two animals, whereas three animals were imaged without inducing CA, as a control. As before, in each image, we evaluated the score in a subset of the image pixels (sampling grid of $47 \times 47$ points). Three example images from this experiment are shown in Figure 10A. Figure 10B, as an example, shows the quantified progression of mitochondrial fragmentation in one animal.

We calculated the median score for each image of the z-stack, for each animal. In Figure 10C, we show the mean and the standard deviation of these scores over the z-planes for all time moments, each animal being represented by one line. From the figure, the average score in an image was between 0.2 and 0.3 before CA and, after a delay of approximately 200–300 s, it increased to nearly 0.8 in both animals, during a rapid period of fragmentation of nearly 100 s. Figure 10D shows the results from the control data: as expected, the scores remain approximately constant throughout the time series.

From these measurements, it is possible to establish a quantitative, repeatable method for assessing when the mitochondria start their rapid fragmentation because both average score curves exhibit a sigmoidal shape with a slowly increasing linear trend (Figure 10C). For this, we fitted to each time series the sum of a logistic function, a line, and a constant. We then calculated the threshold point from the sigmoidal part of the function using the following definition (see, e.g. McDowall & Dampney, 2006). The threshold is the point after which the function's value increased by 5% between the lower and upper plateau. These two points are indicated as a circle in each curve in the figure and suffice to show that there is significant variability of responses between individuals.

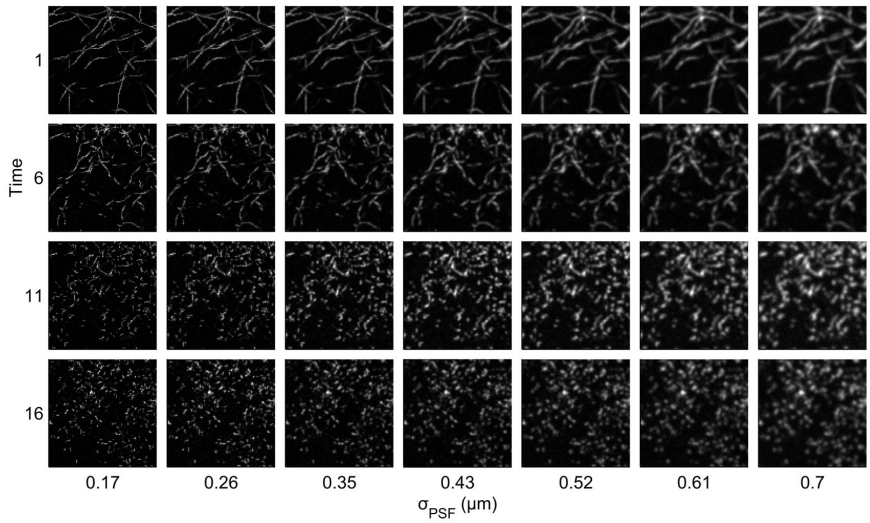


**Fig. 6.** Examples of synthetic patches with varying resolution. Each column shows patches at four time points of the same synthetic image sequence, at varying resolutions. For this figure, the noise was kept constant at $\sigma = 1$.

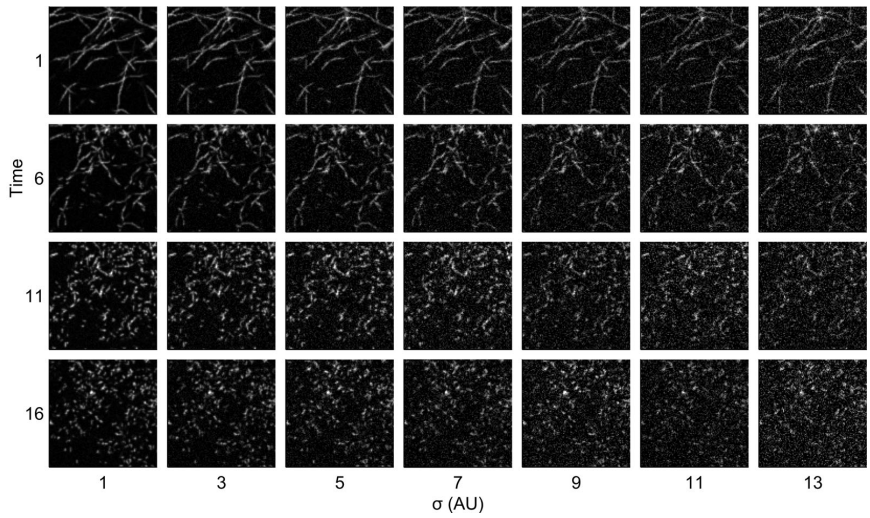


**Fig. 7.** Examples of synthetic patches with varying noise level. Each column shows patches at four time points of the same synthetic image sequence, at varying levels of noise. For this figure, the resolution was kept at $\sigma_{PSF} = 2$.

## Discussion

We proposed a method based on supervised learning for absolute quantification of the degree of fragmentation in mitochondria from *in vivo* two-photon fluorescence microscope images of cell tissues. We tested it on images of neuronal cells of mouse somatosensory cortex. The training process itself revealed the need for such a method, as all five experts had small differences in their scorings, and one of them diverged so significantly from the other four that we did not use his scores in the training.

We demonstrated that using only the morphometric features based on a segmentation of the mitochondria resulted in a rather poor predictive performance, compared to using only texture features. This can be attributed to the fact that, because of image noise and low levels of the fluorescence signal,

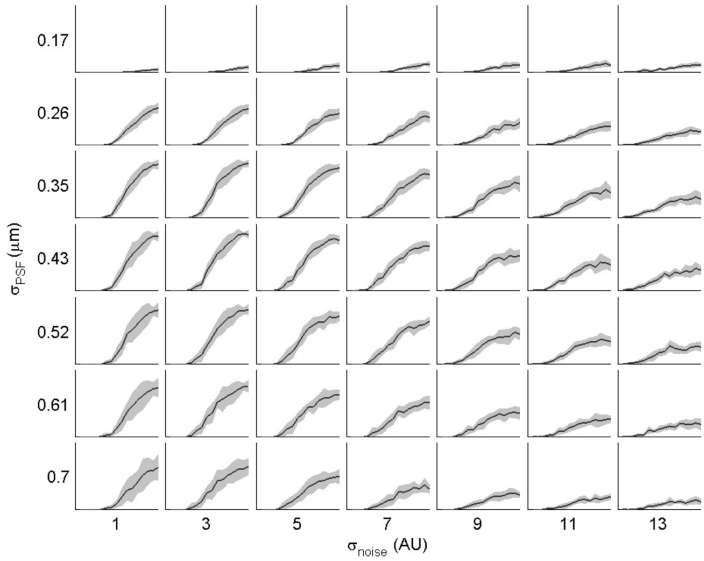


**Fig. 8.** Scores for synthetic time-lapse image sequences. Each subplot shows the mean score (black line) and its standard deviation (shading) over 10 simulated patches; time is on the *x*-axis and moves from left to right (there are 16 time points), and the score is on the y-axis. The bottom and the top of the *y*-axis in each subplot correspond to 0 and 1, respectively. Each column corresponds to one noise level, and each row corresponds to one resolution.

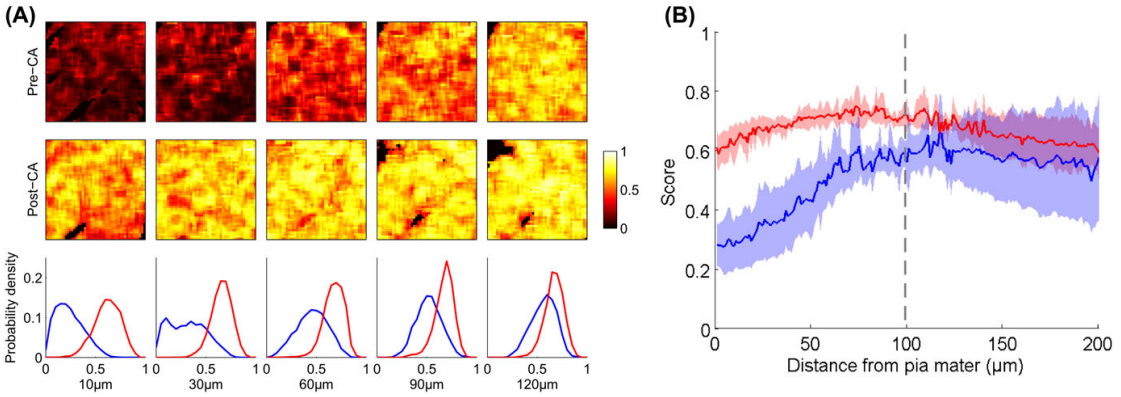


**Fig. 9.** Mitochondrial fragmentation after cardiac arrest. (A) Pre-CA (top row) and post-CA (middle row) scores from one animal for five different depths. The bottom row shows the normalized histograms of the pre-CA (blue) and the post-CA (red) scores, calculated from all animals by pooling all scores from the same depth for each z-stack. The contiguous black regions are non-ROI pixels, as determined by the classifier. (B) Mitochondrial fragmentation scores pre-CA (blue) and post-CA (red). The solid lines represent the mean scores, and the shading represents the one standard deviation intervals (mean $\pm$ SD). Both scores increase as the distance to the pial surface increases, until 100 $\mu$m. The dashed line at that point denotes the depth after which the difference between the two mean scores is no longer statistically significant (*t*-test with significance level 0.01).

the segmentation results contain false positive 'objects', which introduce error to the measurements of mitochondrial morphology. The best performance was obtained by combining the two sets of features, which indicates that the morphometric features carry information that is not encoded by the texture features.

Most of the automatically selected features have a straightforward interpretation, namely, they reflect the transition of mitochondria from an elongated to a globular shape because of fragmentation. For example, mitochondrial lengths were negatively correlated with fragmentation. Also, the cross-correlation between subbands of different scales and same

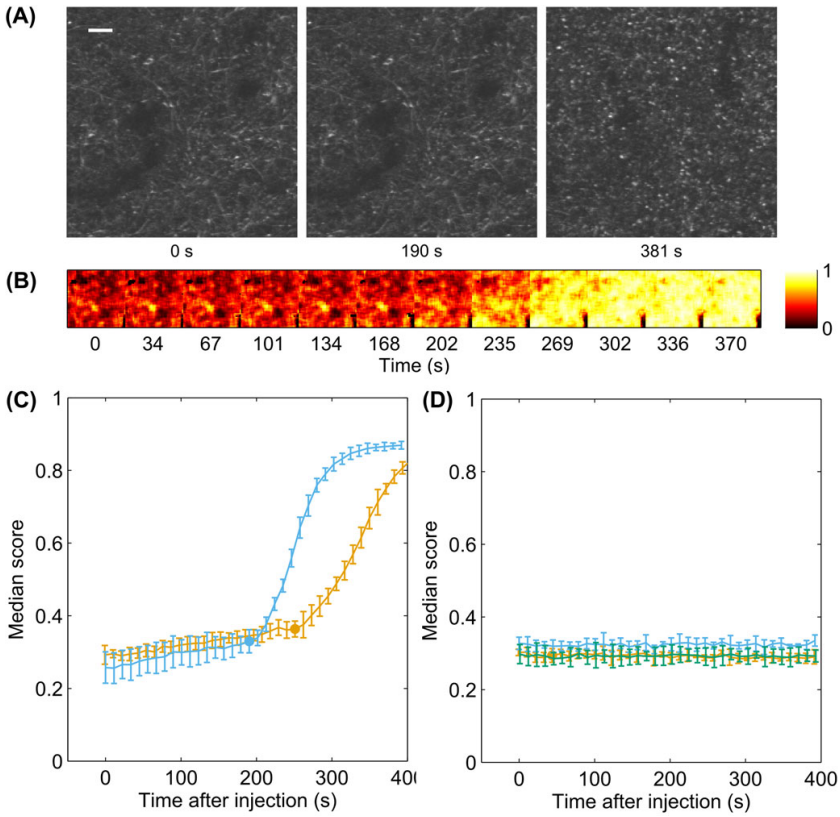


**Fig. 10.** Time series of mitochondrial fragmentation after cardiac arrest. (A) Microscope images showing mitochondrial fragmentation 0 s (left), 190 s (middle) and 381 s after cardiac arrest (images cropped and brightened for easier visualization). The scale bar corresponds to 10 $\mu$m. (B) Scores from the time series depicted as heatmaps. The contiguous black regions are non-ROI pixels, as determined by the classifier. (C) Time series of per-slice median scores from two animals for which CA was induced (yellow and blue represent different animals), with three z-slices per animal. The lines represent the mean, and the error bars represent the standard deviation. The circle on each line marks an estimated threshold moment, that is, a moment when the score starts to rapidly increase. (D) Time series of per-slice median scores from three control animals (different colours represent different animals), with three z-slices per animal.

direction is negatively correlated with the fragmentation score. This is because the orientation of structures like elongated mitochondria is scale invariant, whereas globular mitochondria do not possess a well-defined orientation. However, the interpretation of a few of the selected features is not as clear. For example, the horizontal and vertical autocorrelations of subbands had both positive and negative coefficients, although the mitochondria are randomly oriented in the patches. This inconsistency is likely because of the strong correlation between these features (Hastie et al., 2009). Nevertheless, future analysis making use of information on both mitochondria fragmentation as well as on their functioning should allow associating some of the features of the model with physiological changes in the mitochondrial network.

During the validation procedure, we found that a substantial fraction of the variance in the ground truth scores is not explained by the model. This can be because of the features being insufficient to characterize the changes in the images, as well as the model structure being too limited. Another reason is the differences in subjective evaluations of the experts. Importantly, the MAEs between individual expert scores (0.11–0.24) were not significantly smaller than those between the model's predictions and expert scores (0.11–0.20). As such, we can conclude that our method is (at least) as reliable as using an expert to perform the scoring. In this regard, it is worth noting that because the scores for individual patches have high variance, it is advisable to average scores from a larger region to better estimate the overall degree of fragmentation.

Our tests on synthetic data provided limits for the resolution and signal-to-noise ratio of the imaging system. In addition, these results show that the quantification is not dependent on any particular artefacts in the imaging system used to obtain the training images. Thus, in the limits provided, the method ought to be applicable to images of fragmenting mitochondria, as well as images of other similar structures.

Next, we tested our method's ability in detecting differences in the degree of fragmentation in mitochondria, prior to and after a CA. The results demonstrated that the method is sensitive to the changes that occur in mitochondrial morphology following CA. Furthermore, the results, in particular, the threshold concerning the start of the process of fragmentation, were in agreement with a previous assessment (Liu & Murphy, 2009). Interestingly, preceding the period of rapid fragmentation, the measured scores exhibited a slight increase in both animals tested; suggesting that CAs may have an almost immediate, although slight, effect on mitochondrial morphology. Further studies are needed to validate this hypothesis. Finally, we also conclude from these results that our method is of use to study any process causing mitochondria fragmentation. Not only it detects changes but also, it is able to quantify the degree of the changes.

At the moment, the method has one limitation. Namely, the image analysis performed by the method is based on 2D images, although it is possible to acquire 3D images using two-photon microscopy. This has the advantage of allowing more efficient computation. The main drawback is that a mitochondrion will appear as a small spot, if oriented normal to the image plane. However, as long as the mitochondria are not biased towards the axis normal to the image plane, the method will produce reliable results. In the application shown here, the method worked well for neuronal mitochondria closer than 45 $\mu$m to the pial surface, but the scores became considerably biased at longer distances, because of biases in the orientation of the neurites and in the neuronal subpopulation that expresses mitochondrial-targeted ECFP under the Thy1-promoter (Misgeld et al., 2007). Thus, this bias ought not to affect significantly the results from images taken sufficiently close to the pial surface. Extracting features from 3D patches along with further research on appropriate features and their efficient implementation should solve this present limitation.

Nevertheless, our method can automatically assess mitochondrial morphology, thus reducing human workload and increasing the objectiveness of the results as compared to visual inspection. Unlike previous methods for similar applications, such as Chen et al. (2011), our method provides an absolute quantification, in the sense that no comparison to a 'baseline' image of the same sample is needed.

As a side note, the observed quantitative differences in the degree of mitochondria fragmentation in cells following CA versus cells prior to this event, extracted from the neurological imaging, show that it is possible to use our methodology to assess the damage caused by CAs at a cellular level (Friberg et al., 2013). In addition, our data highlight mitochondria as potential drug target for post CA treatment. Finally, our method can be further used in non-terminal CA animal models to validate efficacy of drug candidate in pre-clinical trials.

Overall, the proposed method can be used to quantify the degree of mitochondrial fragmentation, which is a phenomenon known to occur following various stresses or in various pathologies. As such, the method can be used to detect the occurrence of these events and pathologies, as well as to assess the degree of their effects on the mitochondria. Another possible application, related to the case-study of the effects of ischemia following CA, would be to quantify the dependency of the degree of fragmentation on the duration of blood flow disruption.

Finally, we expect that our method will be useful in quantifying the fragmentation of not only mitochondria, but of other curvilinear intracellular structures as well, such as neurites (Chen et al., 2011), endoplasmic reticula (Kucharz et al., 2011), and actin filaments (Popp et al., 2007).

## Acknowledgements

This work was supported by Emil Aaltonen foundation (www.emilaaltonen.fi) [EL], the Academy of Finland (http://www.aka.fi/en-gb/a/) [LK, grant no. 126803 to ASR], the Centre of International Mobility of Finland (CIMO) (http://www.cimo.fi/) [MK], the Finnish Graduate School of Neuroscience (FGSN) (http://www.brain-mind.fi/) [MK], the Portuguese Foundation for Science and Technology (http://alfa.fct.mctes.pt/index.phtml.en) [grant no. PTDC/BBB-MET/1084/2012 to ASR], and the Finnish Funding Agency for Technology and Innovation (http://www.tekes.fi/en/) [grant no. 40284/08 to ASR].

The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

We are deeply thankful to Evgeny Pryazhnikov, Dmitry Molotkov and Ramil Afzalov for technical assistance and comments on an early draft of the paper, and to Dzhessi Rait for helping with the illustration of Figure 1.

## References

Bilsland, L.G., Sahai, E., Kelly, G., Golding, M., Greensmith, L. & Schiavo, G. (2010) Deficits in axonal transport precede ALS symptoms *in vivo*. *Proc. Natl. Acad. Sci. U. S. A.* **107**, 20523–20528.

Brady, N.R., Hamacher-Brady, A. & Gottlieb, R.A. (2006) Proapoptotic BCL-2 family members and mitochondrial dysfunction during ischemia/reperfusion injury, a study employing cardiac HL-1 cells and GFP biosensors. *Biochim. Biophys. Acta.* **1757**, 667–678.

Chen, S., Tran, S., Sigler, A. & Murphy, T.H. (2011) Automated and quantitative image analysis of ischemic dendritic blebbing using *in vivo* 2-photon microscopy data. *J. Neurosci. Methods.* **195**, 222–231.

Cortes, C. & Vapnik, V. (1995) Support-Vector Networks. *Mach. Learn.* **20**, 273–297.

Eleyan, A., Özkaramanli, H. & Demirel, H. (2008) Complex wavelet transform-based face recognition. *EURASIP J. Adv. Signal Process.* **2008**, 1–14. doi:10.1155/2008/185281.

Farrand, L., Kim, J.Y., Im-Aram, A., Suh, J.-Y., Lee, H.J. & Tsang, B.K. (2013) An improved quantitative approach for the assessment of mitochondrial fragmentation in chemoresistant ovarian cancer cells. *PLoS One*. **8**, e74008.

Ferrari, S. & Cribari-Neto, F. (2004) Beta regression for modelling rates and proportions. *J. Appl. Stat*. **31**, 799–815.

Friberg, H., Rundgren, M., Westhall, E., Nielsen, N. & Cronberg, T. (2013) Continuous evaluation of neurological prognosis after cardiac arrest. *Acta Anaesthesiol. Scand.* **57**, 6–15.

Hall, A.M., Rhodes, G.J., Sandoval, R.M., Corridon, P.R. & Molitoris, B.A. (2013) *In vivo* multiphoton imaging of mitochondrial structure and function during acute kidney injury. *Kidney Int*. **83**, 72–83.

Haralick, R.M., Shanmugam, K. & Dinstein, I. (1973) Textural features for image classification. *IEEE Trans. Syst. Man. Cybern. B. Cybern*. SMC-3, 610–621.

Hastie, T., Tibshirani, R. & Friedman, J. (2009) *The Elements of Statistical Learning, Springer Series in Statistics*. Springer New York, New York, NY. doi:10.1007/b94608

Holschneider, M., Kronland-Martinet, R., Morlet, J., Tchamitchian, P. (1990) A real-time algorithm for signal analysis with the help of the wavelet transform. *Wavelets, Time-Frequency Methods and Phase Space, Inverse Problems and Theoretical Imaging* (ed. by J.-M. Combes, A. Grossmann & P. Tchamitchian), pp. 286–297. Berlin Heidelberg, Springer.

Holtmaat, A., Bonhoeffer, T., Chow, D.K. *et al*. (2009). Long-term, high-resolution imaging in the mouse neocortex through a chronic cranial window. *Nat. Protoc.* **4**, 1128–1144.

Knott, A. & Perkins, G. (2008). Mitochondrial fragmentation in neurodegeneration. *Nat. Rev. Neurosci*. **9**, 505–518.

Kucharz, K., Wieloch, T. & Toresson, H. (2011). Rapid fragmentation of the endoplasmic reticulum in cortical neurons of the mouse brain in situ following cardiac arrest. *J. Cereb. Blood Flow Metab*. **31**, 1663–1667.

Li, Z., Okamoto, K.I., Hayashi, Y. & Sheng, M. (2004). The importance of dendritic mitochondria in the morphogenesis and plasticity of spines and synapses. *Cell*. **119**, 873–887.

Lihavainen, E., Mäkelä, J., Spelbrink, J.N. & Ribeiro, A.S. (2012). Mytoe: automatic analysis of mitochondrial dynamics. *Bioinformatics*. **28**, 1050–1051.

Lin, M.T. & Beal, M.F. (2006). Mitochondrial dysfunction and oxidative stress in neurodegenerative diseases. *Nature*. **443**, 787–795.

Lin, Y., Lin, C., Tsai, Y. & Ku, T. (2010). A spectral graph theoretic approach to quantification and calibration of collective morphological differences in cell images. *Bioinformatics*. **26**, i29–i37.

Liu, R.R. & Murphy, T.H. (2009). Reversible cyclosporin A-sensitive mitochondrial depolarization occurs within minutes of stroke onset in mouse somatosensory cortex *in vivo*: a two-photon imaging study. *J. Biol. Chem*. **284**, 36109–17.

MATLAB and Image Processing Toolbox Release 2012b. (2012) The MathWorks, Inc., Natick, Massachusetts, United States.

Mattson, M.P., Gleichmann, M., Cheng, A. (2008). Mitochondria in neuroplasticity and neurological disorders. *Neuron*. **60**, 748–766.

McDowall, L.M. & Dampney, R.A.L. (2006). Calculation of threshold and saturation points of sigmoidal baroreflex function curves. *Am. J. Physiol. Heart Circ. Physiol*. **291**, H2003–H2007.

Misgeld, T., Kerschensteiner, M., Bareyre, F.M., Burgess, R.W., Lichtman, J.W. (2007). Imaging axonal transport of mitochondria *in vivo*. *Nat. Methods* **4**, 559–561.

Nikić, I., Merkler, D., Sorbara, C., *et al*. (2011). A reversible form of axon damage in experimental autoimmune encephalomyelitis and multiple sclerosis. *Nat. Med.* **17**, 495–499.

Peng, J.-Y. *et al*. (2011). Automatic morphological subtyping reveals new roles of caspases in mitochondrial dynamics. *PLoS Comput. Biol*. 7. e1002212.

Picard, M., Shirihai, O.S., Gentil, B.J. & Burelle, Y. (2013). Mitochondrial morphology transitions and functions: implications for retrograde signaling? *Am. J. Physiol. Regul. Integr. Comp. Physiol*. **304**, R393–R406.

Popp, D., Yamamoto, A. & Maéda, Y. (2007). Crowded surfaces change annealing dynamics of actin filaments. *J. Mol. Biol*. **368**, 365–374.

Ricardo, M. & Carlos, P. (2008). A craniotomy surgery procedure for chronic brain imaging. *J. Vis. Exp*. doi:10.3791/680.

Schon, E.A. & Przedborski, S. (2011). Mitochondria: The next (neurode)generation. *Neuron* **70**, 1033–1053.

Selesnick, I.W., Baraniuk, R.G. & Kingsbury, N.G. (2005). The dual-tree complex wavelet transform. *Signal Process. Mag*. **22**, 123–151.

Shitara, H., Shimanuki, M., Hayashi, J.-I. & Yonekawa, H. (2010). Global imaging of mitochondrial morphology in tissues using transgenic mice expressing mitochondrially targeted enhanced green fluorescent protein. *Exp. Anim*. **59**, 99–103.

Smithson, M., Verkuilen, J. (2006). A better lemon squeezer? Maximum-likelihood regression with beta-distributed dependent variables. *Psychol. Methods*. **11**, 54–71.

Zou, H. & Hastie, T. (2005). Regularization and variable selection via the elastic net. *J.R. Stat. Soc. Ser. B (Stat. Methodol.)*. **67**, 301–320.

Zujovic, J., Pappas, T. & Neuhoff, D. (2013). Structural texture similarity metrics for image analysis and retrieval. *IEEE Trans. Image Process*. **22**, 2545–2558.

## Supporting Information

Additional Supporting information may be found in the online version of this article at the publisher's website:

**Movie S1.** A movie showing z-stack slices before and after cardiac arrest, in increasing depth starting from the pia mater.
**Movie S2.** A movie of mouse brain tissue following the onset of cardiac arrest. Left: sECFP signal. Right: scores output by our method.
**Table S1.** List of automatically selected features and the corresponding regression coefficients.

| **MORPHOMETRIC FEATURES** | | |
|---|---|---|
| **Feature** | **Coefficient in model (empty: not selected)** | **Description** |
| MajorAxis | | Major axis of an ellipse fit to the shape of the segmented mitochondrion |
| MinorAxis | | Minor axis of an ellipse fit to the shape of the segmented mitochondrion |
| AspectRatioMin | 0,5711 | Smallest reciprocal aspect ratio of segmented mitochondria |
| AspectRatioMean | 0,8687 | Mean reciprocal aspect ratio of segmented mitochondria |
| AspectRatioMax | | Maximum reciprocal aspect ratio of segmented mitochondria |
| SkeletonArea | 0,0158 | Mean area of morphological skeleton of segmented mitochondria |
| MinSize | | The minimum, the three quartiles and the maximum of the morphological skeleton areas |
| 1stQuantileSize | -0,0787 | |
| MedianSize | -0,0843 | |
| 2ndQuantileSize | -0,0369 | |
| MaxSize | -0,0016 | |
| WeightedMeanSize | -0,0625 | $\frac{\sum_{i=1}^{N} s_i^2}{\sum_{i=1}^{N} s_i}$, where $s_i$ is the area of the morphological skeleton of the $i$ th segmented mitochondrion |
| SizeVariance | | Variance of the area of the morphological skeleton |
| | | |
| **TEXTURE FEATURES** | | |
| **Feature (S: scale, D: direction)** | **Coefficient in model (empty: not selected)** | **Description** |
| Mean_S1D1 | | Mean filter response |
| Mean_S1D2 | | |
| Mean_S1D3 | | |
| Mean_S1D4 | | |
| Mean_S1D5 | | |
| Mean_S1D6 | | |
| Mean_S2D1 | | |
| Mean_S2D2 | | |
| Mean_S2D3 | | |
| Mean_S2D4 | | |
| Mean_S2D5 | 0,5414 | |

| | | |
|---|---|---|
| Mean_S2D6 | | |
| Mean_S3D1 | | |
| Mean_S3D2 | | |
| Mean_S3D3 | | |
| Mean_S3D4 | | |
| Mean_S3D5 | | |
| Mean_S3D6 | 0,0147 | |
| Mean_LP1 | 0,0206 | |
| Mean_LP2 | | |
| Variance_S1D1 | | Variance of filter response |
| Variance_S1D2 | | |
| Variance_S1D3 | -1,0453 | |
| Variance_S1D4 | | |
| Variance_S1D5 | | |
| Variance_S1D6 | -1,0879 | |
| Variance_S2D1 | | |
| Variance_S2D2 | | |
| Variance_S2D3 | | |
| Variance_S2D4 | | |
| Variance_S2D5 | | |
| Variance_S2D6 | | |
| Variance_S3D1 | | |
| Variance_S3D2 | | |
| Variance_S3D3 | | |
| Variance_S3D4 | -0,0111 | |
| Variance_S3D5 | | |
| Variance_S3D6 | | |
| Variance_LP1 | | |
| Variance_LP2 | | |
| AutoCorrVertical_S1D1 | -0,8035 | Horizontal and vertical autocorrelations of |
| AutoCorrHorizontal_S1D1 | | filter responses |
| AutoCorrVertical_S1D2 | 3,5622 | |
| AutoCorrHorizontal_S1D2 | -4,3529 | |
| AutoCorrVertical_S1D3 | | |
| AutoCorrHorizontal_S1D3 | | |
| AutoCorrVertical_S1D4 | | |
| AutoCorrHorizontal_S1D4 | | |
| AutoCorrVertical_S1D5 | | |
| AutoCorrHorizontal_S1D5 | | |
| AutoCorrVertical_S1D6 | | |
| AutoCorrHorizontal_S1D6 | | |
| AutoCorrVertical_S2D1 | 4,4285 | |
| AutoCorrHorizontal_S2D1 | | |
| AutoCorrVertical_S2D2 | | |
| AutoCorrHorizontal_S2D2 | | |
| AutoCorrVertical_S2D3 | | |
| AutoCorrHorizontal_S2D3 | | |

| | | |
|---|---|---|
| AutoCorrVertical_S2D4 | | |
| AutoCorrHorizontal_S2D4 | | |
| AutoCorrVertical_S2D5 | | |
| AutoCorrHorizontal_S2D5 | -3,4926 | |
| AutoCorrVertical_S2D6 | -2,4877 | |
| AutoCorrHorizontal_S2D6 | | |
| AutoCorrVertical_S3D1 | -3,1216 | |
| AutoCorrHorizontal_S3D1 | | |
| AutoCorrVertical_S3D2 | | |
| AutoCorrHorizontal_S3D2 | | |
| AutoCorrVertical_S3D3 | | |
| AutoCorrHorizontal_S3D3 | | |
| AutoCorrVertical_S3D4 | | |
| AutoCorrHorizontal_S3D4 | | |
| AutoCorrVertical_S3D5 | | |
| AutoCorrHorizontal_S3D5 | | |
| AutoCorrVertical_S3D6 | | |
| AutoCorrHorizontal_S3D6 | | |
| AutoCorrVertical_LP1 | 3,082 | |
| AutoCorrHorizontal_LP1 | | |
| AutoCorrVertical_LP2 | | |
| AutoCorrHorizontal_LP2 | | |
| CrossbandCorr_S1D1_S1D2 | 0,1367 | Correlations of responses from different sub-bands |
| CrossbandCorr_S1D1_S1D3 | | |
| CrossbandCorr_S1D1_S1D4 | | |
| CrossbandCorr_S1D1_S1D5 | 2,602 | |
| CrossbandCorr_S1D1_S2D1 | | |
| CrossbandCorr_S1D1_S3D1 | | |
| CrossbandCorr_S1D1_LP1 | | |
| CrossbandCorr_S1D2_S1D3 | | |
| CrossbandCorr_S1D2_S1D4 | 2,2631 | |
| CrossbandCorr_S1D2_S1D5 | | |
| CrossbandCorr_S1D2_S2D2 | | |
| CrossbandCorr_S1D2_S3D2 | | |
| CrossbandCorr_S1D2_LP2 | | |
| CrossbandCorr_S1D3_S1D4 | | |
| CrossbandCorr_S1D3_S1D5 | | |
| CrossbandCorr_S1D3_S2D3 | -2,0135 | |
| CrossbandCorr_S1D3_S3D3 | | |
| CrossbandCorr_S1D4_S1D5 | | |
| CrossbandCorr_S1D4_S2D4 | | |
| CrossbandCorr_S1D4_S3D4 | | |
| CrossbandCorr_S1D5_S2D5 | | |
| CrossbandCorr_S1D5_S3D5 | | |
| CrossbandCorr_S1D6_S2D1 | | |
| CrossbandCorr_S1D6_S2D2 | | |
| CrossbandCorr_S1D6_S2D3 | | |

| | |
|---|---|
| CrossbandCorr_S1D6_S2D4 | |
| CrossbandCorr_S1D6_S2D5 | |
| CrossbandCorr_S1D6_S2D6 | |
| CrossbandCorr_S1D6_S3D6 | |
| CrossbandCorr_S2D1_S2D2 | |
| CrossbandCorr_S2D1_S2D3 | 0,2063 |
| CrossbandCorr_S2D1_S2D4 | |
| CrossbandCorr_S2D1_S2D5 | -0,2547 |
| CrossbandCorr_S2D1_S3D1 | |
| CrossbandCorr_S2D1_LP1 | |
| CrossbandCorr_S2D2_S2D3 | 0,3464 |
| CrossbandCorr_S2D2_S2D4 | 0,0251 |
| CrossbandCorr_S2D2_S2D5 | 1,7758 |
| CrossbandCorr_S2D2_S3D2 | -0,8758 |
| CrossbandCorr_S2D2_LP2 | |
| CrossbandCorr_S2D3_S2D4 | |
| CrossbandCorr_S2D3_S2D5 | |
| CrossbandCorr_S2D3_S3D3 | -0,6356 |
| CrossbandCorr_S2D4_S2D5 | 0,9427 |
| CrossbandCorr_S2D4_S3D4 | -0,4543 |
| CrossbandCorr_S2D5_S3D5 | |
| CrossbandCorr_S2D6_S3D1 | |
| CrossbandCorr_S2D6_S3D2 | |
| CrossbandCorr_S2D6_S3D3 | 1,0658 |
| CrossbandCorr_S2D6_S3D4 | |
| CrossbandCorr_S2D6_S3D5 | |
| CrossbandCorr_S2D6_S3D6 | |
| CrossbandCorr_S3D1_S3D2 | |
| CrossbandCorr_S3D1_S3D3 | |
| CrossbandCorr_S3D1_S3D4 | |
| CrossbandCorr_S3D1_S3D5 | |
| CrossbandCorr_S3D1_LP1 | |
| CrossbandCorr_S3D2_S3D3 | |
| CrossbandCorr_S3D2_S3D4 | |
| CrossbandCorr_S3D2_S3D5 | |
| CrossbandCorr_S3D2_LP2 | |
| CrossbandCorr_S3D3_S3D4 | |
| CrossbandCorr_S3D3_S3D5 | |
| CrossbandCorr_S3D4_S3D5 | |
| CrossbandCorr_S3D6_LP1 | |
| CrossbandCorr_S3D6_LP2 | |
| CrossbandCorr_LP1_LP2 | |

**Model intercept:** -1,6418

Tampereen teknillinen yliopisto
PL 527
33101 Tampere

Tampere University of Technology
P.O.B. 527
FI-33101 Tampere, Finland

ISBN 978-952-15-3799-8
ISSN 1459-2045