



University
of Glasgow

Olesker, Daniel (2023) *Methods for fast volumetric fluorescence imaging based on 2D-3D volume reconstruction*. PhD thesis

<http://theses.gla.ac.uk/83581/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

**Methods for fast volumetric fluorescence
imaging based on 2D-3D volume
reconstruction**

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

Submitted in fulfillment of the requirements for the degree of *Doctor of
Philosophy*

Daniel Olesker

May 4, 2023

Abstract

Fluorescence microscopy offers unrivalled insight into the inner workings of dynamic biological processes across many temporal and spatial scales, with innovations such as light-sheet and confocal microscopy extending this imaging ability into all three spatial dimensions. When performed with these methods, however, 3D imaging lacks the high temporal resolution of 2D imaging. This is because acquisition of a volume requires the sequential acquisition of many individual 2D planes, with volumetric acquisition times often spanning multiple seconds. This means that such methods are unsuited to the 3D imaging of dynamic processes that occur over shorter timescales.

Ideally, it would be possible to perform full 3D imaging in a single snapshot, in a similar fashion to 2D imaging. Yet, while snapshot volume-imaging methods do exist, they typically require a significant compromise in spatial resolution. We see, therefore, that with existing methods there is a trade-off: for 3D fluorescence imaging we must choose between techniques that provide high-resolution in space, but slow, plane-by-plane 3D imaging; or those that provide fast 3D imaging but at the cost of spatial resolution.

In this thesis we introduce imaging techniques that overcome these limitations. Through the development of an image processing pipeline that offers volumetric reconstruction from 2D projection images, we demonstrate how 3D fluorescence imaging may be performed while maintaining high resolution in both time and space. We apply the volume reconstruction framework to both simulated and experimental images acquired via two distinct projection imaging modalities, demonstrating the versatility of the technique. Additionally, we explore the benefits provided by the volumetric reconstruction pipeline in the context of single-molecule localisation microscopy. Because of the high temporal resolution offered by the volumetric imaging techniques introduced in this thesis, the methods are applicable to a wide range of dynamic biological samples and have the potential to offer new insight into dynamic biological processes. To the best of the author's knowledge, the snapshot volume imaging methods introduced here are the first such methods that maintain the resolution of the parent imaging system.

Publications

Olesker, D., Harvey, A. R., & Taylor, J. M. (2022). Snapshot volumetric imaging with engineered point-spread functions. *Optics Express*, 30(19), 33490-33501.

Olesker, D., Handley, M., Harvey, A. R., & Taylor, J. M. (in preparation). Volume reconstruction for single-molecule localisation microscopy.

Conference Proceedings

Olesker, D., Harvey, A. R., & Taylor, J. M. (2021, July). Video-Rate Volumetric Imaging of Extended Structures with Engineered PSFs. In *Computational Optical Sensing and Imaging* (pp. CTu4B-1). Optical Society of America.

Olesker, D., Chen, B., Manton, J. D., Fiolka, R., Harvey, A. R., & Taylor, J. M. (2022, July). Volumetric Reconstruction from 2D Projection Data for Fast 3D Imaging. In *Computational Optical Sensing and Imaging* (pp. CTu4F-2). Optica Publishing Group.

Declaration

I certify that the thesis presented here for examination for a PhD degree of the University of Glasgow is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it) and that the thesis has not been edited by a third party beyond what is permitted by the University's PGR Code of Practice.

The copyright of this thesis rests with the author. No quotation from it is permitted without full acknowledgement.

I declare that the thesis does not include work forming part of a thesis presented successfully for another degree.

I declare that this thesis has been produced in accordance with the University of Glasgow's Code of Good Practice in Research.

I acknowledge that if any issues are raised regarding good research practice based on review of the thesis, the examination may be postponed pending the outcome of any investigation of the issues.

Signed:

Date:

Acknowledgements

I am extremely grateful to my PhD supervisors, Dr Jonathan M. Taylor and Professor Andrew R. Harvey, for their continuous guidance, support and supervision throughout the course of my PhD studies, especially during the Covid-19 lockdowns that introduced unexpected challenges to the project. Thank you for being brilliant PhD supervisors.

I would also like to express my gratitude to current and former members of the Imaging Concepts Group, it has been an honour working with you. In particular I would like to thank Dr Yongzhuang Zhou and Dr Chas Nelson for their guidance and help during the early stages of my PhD, and to Michael Handley for interesting discussions along the way. Further thanks go to the various collaborators with whom I've worked during my PhD, in particular Dr James D. Manton for hosting me in his Cambridge lab during a research visit.

Finally, I would like to thank my friends and family who have supported me throughout my PhD studies, in particular my partner Sophia for her continued support while I wrote this thesis.

Contents

1	Introduction	1
1.1	Thesis organisation	3
1.2	Background information	4
1.2.1	Principle of fluorescence microscopy	4
1.2.2	Microscope and image characteristics	5
1.2.3	The microscope optical transfer function	8
1.2.4	PSF simulation	11
1.2.5	Image convolution and deconvolution	12
1.3	Literature review	13
1.3.1	Epi-fluorescence focal stacks	14
1.3.2	Confocal microscopy	14
1.3.3	Two-photon scanning microscopy	15
1.3.4	Light-sheet fluorescence microscopy	16
1.3.5	Bi/multi-plane imaging	18
1.3.6	Light-field microscopy	18
1.3.7	PSF engineering	19
1.4	Chapter 1 summary	24
2	2D-3D Reconstruction	26
2.1	Snapshot volume imaging	27
2.2	RL deconvolution	28
2.3	Image formation	31
2.4	Volume reconstruction	32
2.5	Multi-view RL	33
2.6	Edge artefacts	34
2.6.1	Edge correction	36
2.7	Regularisation	37
2.8	Chapter 2 summary	38

3	3D-EPM	39
3.1	Cubic PSF	40
3.1.1	The cubic OTF	40
3.2	Encoding depth information	43
3.3	PSF phase retrieval	44
3.3.1	Estimating the pupil function	46
3.3.2	Low pass filtering in the Zernike basis	49
3.4	Registration	50
3.5	Snapshot volume imaging	52
3.5.1	Simulated point source imaging	53
3.5.2	Experimental validation	55
3.5.3	Extended structure	56
3.5.4	Simulated filaments	58
3.5.5	Lens-tissue	58
3.6	Dynamic Samples	60
3.6.1	Simulated filament growth	61
3.6.2	Chloroplast motion in <i>Egeria densa</i>	61
3.7	Performance and limitations	62
3.7.1	Resolution	62
3.7.2	Sample sparsity	64
3.7.3	The influence of differential defocus	67
3.7.4	Convergence and regularisation	67
3.8	Chapter 3 summary	68
4	3D-EPM localisation	70
4.1	Point localisation	70
4.2	3D-SMLM	72
4.3	3D-EPM localisation	73
4.4	Unbiased estimator	75
4.5	CRLB.	76
4.6	Simulation parameters	78
4.6.1	Performance metrics	79
4.6.2	Twin-Airy comparison	80
4.6.3	Image formation model	81

4.6.4	Iterations used for 3D-EPM localisation	84
4.7	Results	85
4.7.1	Emitter density	85
4.7.2	Signal-to-noise ratio	87
4.7.3	3D-EPM localisation is not limited to a single PSF	89
4.7.4	3D-EPM localisation performs well close to the edges of the imaging domain	91
4.8	Discussion	92
4.9	Simulated microtubules	94
4.10	Chapter 4 summary	94
5	Other applications	96
5.1	Projection imaging	97
5.2	Multi-angle imaging	98
5.2.1	Multi-angle projection image formation	100
5.2.2	Multi-angle projection reconstruction	102
5.2.3	Simulated HeLa cell organelles	105
5.2.4	Number of projections	106
5.3	OPM Multi-angle imaging	109
5.3.1	Geometry of OPM	110
5.3.2	Multi-angle OPM reconstruction	113
5.3.3	Live-cell multi-angle projection reconstruction	114
5.3.4	Discussion	115
5.4	Tomographic 3D-EPM	118
5.5	Chapter 5 summary	121
6	Conclusion and future Work	124
	Contributions	128
	Bibliography	128

List of Figures

1.1	Motion during light-sheet acquisition introduces significant artefacts. . .	2
1.2	Schematic of a fluorescence microscope.	5
1.4	Simulated imaging of point-sources with decreasing SNR	8
1.5	Diffraction limited PSF and MTF.	9
1.6	Simulated PSF (a) and OTF (b) of the widefield microscope.	12
1.7	OTFs for a widefield and confocal microscope	15
1.8	Light sheet fluorescence microscopy.	17
1.9	Schematic of a light-field microscope.	19
1.10	Pupil-engineering is typically implemented through the addition of a $4f$ lens relay behind the native image plane.	20
1.11	Phasemask and corresponding PSFs over an axial range of $\pm 15 \mu\text{m}$. . .	22
2.1	Normalised zx projection of a 0.8NA widefield PSF.	28
2.2	RL deconvolution improves image contrast.	31
2.3	Schematic of the model for projection image formation	32
2.4	Simulation demonstrating how edge artefacts can degrade the decon- volved image	35
3.1	Cubic PSF with increasing α	42
3.2	Differential-Airy PSF	45
3.3	Phase retrieval summary flow-chart.	47
3.4	Retrieval of the pupil phase of an experimentally-acquired 0.5NA cubic PSF.	48
3.5	Refinement of the recovered pupil function through Zernike decomposi- tion and low-pass filtering.	51
3.6	Improvement in the 3D PSF through phase-retrieval and Zernike filtering.	52
3.7	Registration process for 3D-EPM	53
3.8	Simulated 0.5NA snapshot volumetric imaging of a single point-source.	54

3.9	Experimental demonstration of 3D-EPM: imaging fluorescent beads . . .	55
3.10	Fluorescent bead reconstruction analysis	57
3.11	Simulated 0.5NA 3D-EPM imaging of a network of filaments inside a 102 × 102 × 50 μm volume.	59
3.12	Experimental 3D-EPM imaging of a sample of lens-tissue.	60
3.13	Simulated 3D-EPM imaging of filament growth	62
3.14	3D-EPM imaging of chloroplasts undergoing cytoplasmic streaming in a sample of <i>Egeria Densa</i>	63
3.15	Convergence of a single point source	65
3.16	3D-EPM resolution	66
3.17	Plot of convergence with varying differential defocus	67
4.1	Schematic of localisation with 3D-EPM	74
4.2	Reconstruction-localisation as an unbiased estimator	75
4.3	CRLB	76
4.4	Schematic of the point-matching approach for twin-Airy point localisation.	82
4.5	3D-EPM localisation performs well with increasing emitter density.	85
4.6	Localisation performance with varying SNR	88
4.7	3D-EPM localisation identifies many false positives at extremely low signal-to-noise ratio.	89
4.8	3D-EPM localisation with different PSFs.	90
4.9	Jaccard index as points are allowed closer to the edges of the image domain.	91
4.10	3D-EPM localisation takes longer to implement than the twin-Airy point- matching approach	93
4.11	Simulated microtubule imaging with 3D-EPM localisation.	95
5.1	Schematic of a multi-angle projection microscope	98
5.2	Geometry of optically-induced shear	100
5.3	Equivalence of shearing and rotation after projection and scaling.	102
5.4	Reconstruction of simulated multi-angle projection imaging of HeLa mi- tochondria	107
5.5	Ghost artefacts are introduced when reconstructing multi-view projec- tion data with only two projections.	108
5.6	Reconstruction performance with increasing numbers of projection angles.	109
5.7	Chromatin sample is marginally too dense to reconstruct well from multi-angle projection data	110

5.8	Schematic of an oblique plane microscope	111
5.9	Schematic illustrating the geometry of an OPM system	112
5.10	Multi-angle reconstruction of zebrafish vasculature	114
5.11	Multi-angle live-cell reconstruction.	115
5.12	Schematic of proposed tomographic EPM setup	119
5.13	Proof of concept results for tomographic EPM	123
6.1	Experimental setup built to implement tomographic 3D-EPM	126

List of Tables

3.1	Data acquisition and calibration pipeline for differential-Airy 3D-EPM	53
3.2	Acquisition parameters for each experiment detailed in chapter 3. . . .	69
4.1	EMCCD camera parameters used in simulation.	84

Acronyms

3D-EPM 3D engineered point-spread function microscopy.

CKM complimentary kernel matching.

CPM cubic phase mask.

CRLB Cramer-Rao lower bound.

DH double-helix.

DoF depth of field.

EDoF extended depth of field.

EMCCD electron-multiplying charge-coupled device.

FoV field of view.

LFM light-field microscopy.

LSFM light-sheet fluorescence microscopy.

MIP maximum-intensity projection.

MLA microlens array.

MTF modulation transfer function.

NA numerical aperture.

OPM oblique-plane microscopy.

OTF optical transfer function.

PSF point-spread function.

RL Richardson-Lucy.

RMSE root mean-square error.

SLM spatial light modulator.

SMLM single-molecule localisation microscopy.

SR super-resolution.

Chapter 1

Introduction

Ever since its inception at the end of the 16th century, optical microscopy has played a fundamentally important role in scientific discovery. In its infancy, the compound microscope piqued the interest of large parts of the scientific community through its ability to make visible the otherwise unobservable microscopic detail of everyday items, captured famously in Robert Hooke's 1665 publication of *Micrographia* [1]. In the ensuing years optical microscopy continued to enable key scientific discoveries, including those of cells and bacteria. The 19th century brought key technical and theoretical advancements: Joseph J. Lister's innovation in lens combination saw vast improvements in image quality, while Ernest Abbe linked the resolving power of a microscope to the wavelength of the light being observed. Mass production began, which continued into the 20th century and enabled the wider uptake of optical microscopy across the sciences. Since then, it has been adopted across a broad array of fields, including material science, pharmaceutical research and medicine; and remains perhaps the most widely used tool in the life sciences to the present day.

Several types of optical microscopy exist, each exploiting different optical phenomena to address specific challenges. Fluorescence microscopy, the focus of this thesis, has emerged as an especially powerful technique because of its ability to isolate specific structures of interest through targeting distinct fluorescent dyes or proteins, offering unrivalled insight into cellular structure and function. As such, it has become a staple tool in bio-imaging with applications spanning plant and animal biology across scales ranging from the single molecule through to whole organisms.

That said, fluorescence microscopy remains an imperfect practice and continues to be advanced through current and ongoing research. This is because the performance of fluorescence microscopy is fundamentally limited by the influence of several distinct but important phenomena. The ideal fluorescence imaging system would possess a number of key characteristics. For instance, it would acquire images that feature excellent contrast while providing a high spatial resolution across a large field of view (FoV). It would image many fluorophores simultaneously while maintaining a high optical throughput. Importantly, it would acquire images over a large depth range and in all




Figure 1.1: Motion during light-sheet acquisition introduces significant artefacts. Image shows a xz maximum-intensity projection of a 3D light-sheet image of a beating zebrafish heart. The motion of the heart during normal light-sheet acquisition causes significant artefacts and seriously degrades the image. Image acquired on a light-sheet system, courtesy of Dr Jonathan M. Taylor. Scale-bar 25 μm .

three spatial dimensions, with contributions from only in-focus light so as to provide optical sectioning. And finally, it would acquire images quickly in order to properly capture the dynamics of rapidly-moving scenes.

However, each of these characteristics is, in practice, restricted by one or more limiting factor. For example, acquisition speed is often limited by the typically low photon budget in fluorescence imaging experiments, or the requirement to scan during acquisition. The achievable spatial resolution is governed by the laws of diffraction, which state that the resolution is both proportional to the emission wavelength and inversely proportional to the numerical aperture (NA) of the objective lens. The spectral range that can be imaged simultaneously is restricted by the available laser lines, fluorophores and fluorescence filters.

This thesis focuses on the challenges associated with performing 3D imaging of highly dynamic samples. These challenges arise from the fact that conventional fluorescence microscopy is limited to the fast acquisition of 2D images, with volumes being imaged by repeating this process while moving the sample relative to the illumination in order to image at different depths. However, this process breaks down if the sample is undergoing motion during acquisition, since the acquired volume is not a true reflection of the sample at any single point in time. When imaging rapidly-moving scenes, this can lead to two issues in the acquired data. Firstly, the temporal sampling interval between volumes acquired with scan-based methods mean that the dynamics of the sample may not be properly captured. Secondly, motion artefacts may be introduced into acquired volumes resulting from slow acquisition speeds, possibly leading to erroneous sample measurements been made. An example of this type of motion artefact is displayed in figure 1.1.

Faster volumetric acquisition speed in a z -scanned system may somewhat solve

this problem for slow-moving samples, providing the Nyquist sampling criterion is met in the time domain as well as space. However, for highly dynamic scenes where motion occurs over much shorter timescales, for instance the beating hearts of small organisms or the flow of blood, conventional scan-based acquisition methods are no longer suitable: indeed, *any* temporal delay between imaging different regions of the sample is likely to introduce this type of artefact. For these purposes, scan-free, ideally snapshot volume-acquisition methods are desirable, that image the entire sample depth at the same time-point. While some scan-free volume imaging methods do exist, e.g. light-field microscopy (LFM) [2], they typically require a significant compromise in spatial resolution compared to that offered by the parent imaging system. This raises the question: is it possible to acquire a 3D image in a single snapshot while maintaining high spatial resolution? By combining various optical techniques with custom image processing frameworks, several imaging pipelines are proposed in this thesis that offer fast volumetric imaging while maintaining high spatial resolution. Chapters 2 and 3 focus on the development of a snapshot volume imaging technique that combines point-spread function (PSF) engineering methods with a custom image deconvolution framework. In chapters 4 and 5, we then demonstrate how this framework may be applied to various other image acquisition methods, illustrating its versatility in the reconstruction of volumes from few projection views.

1.1 Thesis organisation

This thesis is organised as follows. The current chapter contains a brief theoretical background to concepts that are central to the work presented in this thesis, including optical transfer functions (OTFs), image deconvolution and PSFs. This is followed by a review of relevant literature and related optical methodologies. In chapter 2, we develop a general mathematical framework that forms the basis of the image processing used in the imaging methods introduced later in this thesis. This mathematical framework, based on Richardson-Lucy (RL) deconvolution [3, 4], enables full 3D volumetric reconstruction from 2D projection images, providing certain conditions are met. In chapter 3, we introduce 3D engineered point-spread function microscopy (3D-EPM), a technique that combines this mathematical framework with PSF engineering to enable high-resolution snapshot 3D fluorescent imaging. Chapter 4 demonstrates the application of 3D-EPM to the specific challenge of 3D point localisation in fluorescence microscopy, where we see its application introduces significant advantages over existing approaches. Finally, chapter 5 demonstrates the application of the mathematical framework developed in chapter 2 to different imaging modalities including multi-angle projection imaging [5], enabling 3D reconstruction from few projection images, as well as a tomographic implementation of 3D-EPM, enabling snapshot 3D imaging with a isotropic 3D resolution.

1.2 Background information

In this section we introduce some theoretical background material to topics that are relevant to this thesis. We begin with the working principles of fluorescence microscopy, before discussing how a microscope may be characterised in terms of its OTF or PSF. We also discuss some theoretical performance limits of an optical microscope, describe image formation in terms of an optical convolution, before finally introducing image deconvolution, an image processing technique that underpins the methods developed in later chapters of this thesis.

1.2.1 Principle of fluorescence microscopy

We start our discussion with an introduction to the type of imaging that is the focus of this thesis: fluorescence microscopy. This type of microscopy exploits the principle of fluorescence, a process where certain fluorescent molecules, known as fluorophores, emit light of a characteristic spectrum upon excitation with incident light, also of a characteristic spectrum or wavelength. The fluorescence process can be described in two stages: excitation and emission. Excitation occurs when an incident photon is absorbed by the fluorophore, elevating it to an excited electron state. The fluorophore remains in an excited state for a short period of time, known as its excited lifetime, after which time a photon is emitted, returning the fluorophore to its ground energy state. Due to energy dissipation that occurs during the excited state, the emitted photon is of longer wavelength than the excitation photon¹. Fluorescence microscopy exploits the characteristic nature of the emission and excitation spectra of fluorophores. By labelling a certain structure within the sample with a known fluorophore, that structure is imaged by illuminating the sample with light of the required excitation wavelength and detecting fluorescence at the corresponding emission wavelength through use of appropriate spectral filters. Therefore, fluorescence microscopy offers molecular specificity and is able to target specific sample components and isolate them from other structures within the sample, where other microscopy methods might suffer from poor contrast. This is especially useful in, for instance, the study of individual organelles within cells. Fluorescence microscopy may also be used to image samples expressing proteins that fluoresce, for instance green fluorescent protein [6], or the auto-fluorescence emitted by samples under certain irradiation wavelengths.

Fluorescence microscopes come in a number of different varieties. Figure 1.2 displays a schematic of a simple widefield epi-fluorescence microscope. Here, the fluorescence excitation light is delivered to the sample through the same objective lens used for imaging. Other types of fluorescence microscopes, e.g. light-sheet microscopes, decouple the illumination and detection, delivering the illumination via a second objective lens. In the epi-illumination configuration displayed in figure 1.2, the sample

¹Assuming single photon excitation.




Figure 1.2: Schematic of a fluorescence microscope.

is illuminated with collimated light, formed by focusing the illumination laser beam onto the back pupil plane of the objective lens. The desired excitation wavelength is passed by the excitation filter, before the beam is reflected towards the objective lens with a dichroic mirror. Fluorescence (green) is excited at the sample and collected by the same objective lens. The fluorescence passes the dichroic filter and the target wavelength is filtered with an emission filter. The tube lens then focuses the detected fluorescence to an image plane, where an image may be recorded with a camera.

1.2.2 Microscope and image characteristics

The characteristics of an image recorded by such an optical microscope depend on a number of factors. The purpose of any microscope is to produce a magnified image of the sample in question. The magnification of the image depends primarily on the choice of objective lens, as well as the other lenses in the imaging path. The magnification alone, however, does not completely describe the imaging system. Other important features include the resolution and contrast achieved by the microscope. Here we explore some of these characteristics and their influence on imaging performance in fluorescence microscopy.

Resolution

The resolution of a microscope, or any imaging system, is defined as the minimum distance two points must be separated in object space for them to be resolved as distinct points in the image. Again, the resolution is determined primarily by the choice of objective lens used for imaging, or more specifically, its numerical aperture (NA). The NA of an objective lens indicates the range of angles it can accept and is defined as:

$$\text{NA} = n \sin \theta \quad (1.1)$$

where n is the refractive index of the medium between the sample and the objective

lens (known as the immersion medium), and θ is the half-angle of the cone of light that the objective lens can accept. The resolution of a perfectly-aligned microscope was found in 1873 by Ernest Abbe to be [7]:

$$r_{x,y} = \frac{\lambda}{2\text{NA}} \quad (1.2)$$

Therefore, to resolve smaller sample features, higher NA lenses are required. Increasing the NA above 1 requires use of an immersion medium with a refractive index greater than 1, such as water, glycerol or oil. Similarly, the Abbe definition of resolution in the axial direction is given by:

$$r_z = \frac{2\lambda}{\text{NA}^2} \quad (1.3)$$

The Abbe resolution limits given in equations 1.2 and 1.3 were defined in terms of the size of the Airy disk formed when imaging a single point-emitter (discussed further in section 1.2.3). Rayleigh later refined the definition of resolution to require that two points are separated such that the central, or principal diffraction maximum in the Airy disk from one point-source coincides with the first diffraction minimum of the Airy disk from the other point source in order for them to be resolved. The Rayleigh definition for lateral resolution is given by:

$$r_{x,y} = \frac{0.61\lambda}{\text{NA}} \quad (1.4)$$

Both definitions of resolution are common in microscopy and differ only by a constant, reflecting different interpretations of what separation is needed for two points to be resolved. The key point to note, however, is that lateral resolution is inversely proportional to the objective NA, whereas axial resolution is inversely proportional to NA^2 .

Sampling

To achieve a resolution approaching the theoretical maximum determined by equation 1.2, in addition to the microscope being perfectly aligned, the image formed at the image plane must also be sampled correctly. In modern fluorescence microscopes, images are recorded on scientific cameras which feature arrays of discrete pixels of certain dimensions. This means the continuous image signal is sampled in a discrete fashion and therefore the size of the pixels must be considered. For instance, if the theoretical resolution of a microscope is $1\ \mu\text{m}$, but the pixel dimensions are greater than this value, the theoretical resolution will not be achieved in practice, since two points separated by the theoretical minimum resolvable distance will be imaged in the same or neighbouring pixels. The Nyquist-Shannon sampling criterion requires the sampling interval to be less than half the period of highest frequency passed by the microscope




Figure 1.3: Simulated imaging of point-sources where the image is Nyquist-sampled (left) and undersampled (right). Grid lines are visible (when zoomed in to the PDF version) to emphasize the pixel dimensions. Scale-bar is $1\ \mu\text{m}$.

for the signal, or resolution, to be preserved [8]. If the image is undersampled, i.e. the pixel dimensions exceed $\frac{r}{2}$, aliasing of high-frequency features may occur, degrading image quality, an example of which is illustrated in figure 1.3.

Noise

The above sampling criterion places an upper limit on the pixel dimensions for optimal imaging performance. The magnification of the microscope should therefore be chosen such that this criterion is met. However, the pixel dimensions should not be decreased far beneath this limit for two reasons. Firstly, the detector has a finite number of pixels and therefore finite dimensions. Since the NA places an upper limit on the resolution of the imaging system, no resolution is gained by increasing the sampling rate. Therefore, to maximise the dimensions of the field imaged without sacrificing resolution, the pixel dimensions should be as close to the sampling requirement as possible. The second factor is the influence of noise on acquired images. Scientific cameras are imperfect measuring devices that convert the number of photons detected by individual pixels within a given exposure time to an electronic charge. At the end of the exposure, the accumulated charge in each pixel is quantified during a process known as readout, and converted to a final pixel value. However, as well as quantifying the signal photons detected at each pixel, images are also susceptible to multiple sources of noise, the exact nature of which depends on the type of camera used. There are generally two primary sources of camera-induced noise present in images: read noise and dark current noise. Read noise is the accumulation of charge generated by all electronic components during readout. In charge-coupled device cameras read noise is typically characterised by a Gaussian distribution [9]. Dark current noise, on the other hand, results from thermally-generated electrons that are independent of the incident light. Dark current noise is reduced through sensor cooling and is often negligible for high-performance cameras. Other sources of noise include clock-induced charge and pattern noise. In addition to the camera-induced noise, the stochastic nature of fluorescence emission and the discrete nature of photon detection mean that fluorescence images follow a Poisson noise profile. This type of noise is commonly referred to as photon shot noise and is quantified by the square root of the measured signal [10]. The influence of shot noise




Figure 1.4: Simulated imaging of point-sources with decreasing signal-to-noise ratio (l-r). Scale-bar is $2\ \mu\text{m}$.

is therefore reduced at higher incident photon fluxes, and is a sample-dependent noise, independent of the choice of camera. Shot noise is a key consideration in fluorescence microscopy where there is typically a limited photon budget, owing to the sensitivity of biological samples to harmful laser excitation light.

The combination of the camera noise and the sample noise mean that an acquired fluorescence image is a noise-corrupted estimate to the true underlying fluorophore distribution. Noise acts to degrade image contrast, introducing a grainy appearance into images that worsens with decreasing levels of signal photons. This is illustrated in figure 1.4, which shows fluorescence images of a sample of point-sources being simulated at decreasing signal levels, where the mean and standard deviation of the readout noise distribution is kept constant. Additionally, noise can inhibit the performance of many of the image processing methods discussed in this thesis. The level of signal above the noise in the image is quantified by the signal-to-noise ratio and should be kept as high as possible for optimal image quality and image processing results. Fluorescence imaging is typically subject to low levels of signal photons, owing to the sensitivity of biological samples to excitation light and the bleaching of fluorophores. Therefore, to maintain the highest-possible signal-to-noise ratio, the pixel dimensions should be kept as large as possible while still sampling the image adequately, so that the signal photons are spread between the fewest possible pixels.

1.2.3 The microscope optical transfer function

In the above section we saw that the resolution limit of a perfectly aligned microscope can be calculated from theoretical predictions. The ability to characterise an optical system is important in understanding its theoretical performance limits, as well as in the design of new imaging systems. One such method for characterising an imaging system in the spatial domain is through measurement of its response to a single point-source of light. The response of a system to a single point-source is known as its PSF. Its Fourier transform, the OTF, equivalently characterises the imaging system in the spatial frequency domain [11]. In this section we describe the OTF, the PSF, and a closely related quantity, the modulation transfer function (MTF). It is worth noting that the following discussion is limited to incoherent imaging systems, such as




Figure 1.5: Diffraction limited PSF and MTF.

a fluorescence microscope.

Point-spread function

The PSF is a function that describes the impulse response of an imaging system in the spatial domain [11]. The theoretical 3D PSF represents the noise-free image that a perfect detector would record of a single sub-diffraction point source being translated through the focus along the optical axis. In this thesis, the optical axis is assigned the coordinate z . When a single point-source at the focal plane of a fluorescence microscope is imaged, a portion of the emitted light is collected by the objective lens and focused to an image at the image plane. For a typical, well-aligned microscope, the in-focus 2D PSF takes the form of an Airy disk, the size of which is determined by the NA of the objective lens. A 3D PSF estimate may be obtained by stacking sequentially acquired images as the point-source is translated along the optical axis. A 3D image formed in this fashion is herein referred to as a z -stack. As the point source is moved away from the focal plane, the size of the Airy-disk increases and its peak intensity drops. A PSF recorded in this fashion corresponds to the intensity PSF, which is the squared magnitude of the complex-valued amplitude PSF. The amplitude-PSF contains both a phase term and a magnitude term. Often in microscopy, the phase terms are disregarded, with the system instead being described in terms of its intensity-PSF (the squared magnitude of the amplitude term of the PSF), or the magnitude of the OTF, the MTF. To avoid confusion in this thesis, unless otherwise noted, PSF is taken to mean the intensity-PSF (i.e. $|\text{amplitude-PSF}|^2$), not the complex-valued function.

Optical transfer function

The OTF is a similar function that characterises the imaging system in the spatial frequency domain [11]. Similar to how the intensity PSF is the squared-magnitude of the complex-valued amplitude PSF, the intensity OTF is calculated as the auto-correlation of the complex-valued amplitude OTF [12]. Again, in this thesis, OTF is taken to mean the intensity-OTF. The OTF shares a Fourier transform relationship with the PSF:

$$PSF(x, y, z) = \left| \iiint OTF(k_x, k_y, k_z) e^{i(k_x x + k_y y + k_z z)} d_{k_x} d_{k_y} d_{k_z} \right|^2 \quad (1.5)$$

Here, x, y, z represent the 3D spatial coordinates, and k_x, k_y, k_z represent the corresponding coordinates in the frequency domain. The Fourier relationship between the OTF and the PSF means that the information content is equivalent between the two functions. Because both functions lack phase information, the OTF and PSF are considered incomplete descriptions of an optical system, however are sufficient analysis tools in many cases. Since the OTF is calculated as the Fourier transform of the PSF, it is a complex-valued function, whose absolute value, or modulation transfer function (MTF), describes the systems ability to transfer contrast from object-space to image-space as a function of spatial frequency, whereas the complex argument describes translation as a function of spatial frequency. Often, the complex argument is disregarded, with the MTF being used to characterise the imaging system. In many cases, for instance a diffraction-limited system, the OTF is real and therefore equivalent to the MTF [11].

Modulation transfer function

The MTF describes how well different spatial frequencies are modulated by the imaging system. At the focus, low spatial frequencies are well modulated, representing good information transfer through the imaging system, while those above a certain frequency are not modulated at all. This is known as the cutoff frequency f_c and is given by:

$$f_c = 2 \frac{NA}{\lambda} \quad (1.6)$$

The cutoff frequency is observed to be the inverse of the lateral resolution r defined in equation 1.2. Away from the focus, however, the MTF drops much faster and contains zeros at lower spatial frequencies, corresponding to information loss (figure. 1.5). The depth range over which the MTF remains useful, and indeed the cutoff frequency, is limited by the NA of the microscope objective, and for the highest NA lenses becomes the diffraction limit. As shall be seen later in this chapter, consideration of the system MTF is essential in PSF design and engineering.

The problem of the missing cone

Inspection of the 3D OTF for a widefield microscope reveals a toroidal shape with missing regions of frequencies above and below the focal plane, see figure 1.6. The highest spatial frequency that can be passed by the imaging system is limited by the objective NA. However, the limited acceptance angle of the objective lens means that lower spatial frequencies are not as well modulated by the imaging system away from the focal plane as they are at its focus. This phenomenon gives rise to the toroidal shape

of the 3D OTF, with the frequency voids above and below the focal plane known as the missing cones, indicated by the dashed red lines in figure 1.6b. The missing cones are responsible for the relatively poor axial resolution of widefield systems and the presence of significant out-of-focus light in 2D images. The ideal microscope, therefore, would collect fluorescence from only a small depth-region close to the native focal plane, rejecting out-of-focus light and effectively filling the missing cone. A microscope that is capable of rejecting out-of-focus light is said to achieve optical sectioning, since it images only a narrow depth-region of the sample at any one time. Optical sectioning is the key to acquiring high-quality 3D images, which are otherwise severely degraded by out-of-focus light and suffer severe resolution anisotropy. As shall be seen later in this chapter there are a number of ways to achieve optical sectioning, including confining the excitation light to only a narrow depth region close to the focal plane, or through use of pinholes to reject out-of-focus light.

1.2.4 PSF simulation

As well as estimating the PSF experimentally, the ability to simulate a PSF is important in the design and comparison of optical systems. Simulating a PSF requires mathematical description, or model, of the PSF. Numerous PSF models exist: in this thesis, all PSF simulation is performed using the PSF model described in refs. [12, 13]. This model describes how, for monochromatic light, the amplitude OTF is non-zero only at the cap of a spherical shell, with a radius k determined by the ratio of the refractive index of the microscope immersion medium n to the wavelength of the emitted light, λ . Therefore, by expressing the axial Fourier coordinate k_z as a function of the lateral coordinates k_x, k_y [13]:

$$k_z(k_x, k_y) = \sqrt{\left(2\pi\frac{n}{\lambda}\right)^2 - \left(k_x^2 + k_y^2\right)} \quad (1.7)$$

the Fourier relationship in (1.5) can instead be expressed as the 2D integral [13]:

$$PSF(x, y, z) = \left| \int \int P(k_x, k_y) e^{i2\pi(k_x x + k_y y)} e^{i2\pi k_z(k_x, k_y)z} dk_x dk_y \right|^2 \quad (1.8)$$

over the 2D pupil function $P(k_x, k_y)$, which is non-zero over a domain determined by the acceptance angle of the system (equivalently NA). The pupil function is a complex-valued function that describes both the phase and magnitude of the wavefront produced by a single point-source in the sample. Because it contains both phase and magnitude information, the pupil function is a complete description of the imaging system. Hanser et al. note that it is necessary to express k_z in the form given by equation 1.7 because while low-NA systems have an OTF that can be approximated as flat, at high NA the surface takes on a spherical shape [13]. The second exponential term in this description




Figure 1.6: Simulated PSF (a) and OTF (b) of the widefield microscope.

of the PSF serves as a defocus term acting to shift the focus of the pupil function along the optical axis by a distance z [12, 13].

As shall be seen in chapter 3, this description of the PSF is especially useful, as the PSF is calculated through the Fourier transform of a simple multiplication, between a complex-valued pupil function $P(k_x, k_y)$ and a defocus term.

1.2.5 Image convolution and deconvolution

Images acquired via fluorescence imaging systems follow a linear image formation process. That is, the image of many points is equivalent to the summed images of each individual point, owing to the temporal incoherence of fluorescence emission and resulting lack of interference effects. Additionally, a widefield fluorescence microscope can be approximated as spatially shift-invariant, meaning the image of an object does not change as the object moves across the FoV. With these two assumptions, and the fact that the image of a point source is the PSF, image formation can thus be modelled as the summed intensity resulting from the shifted and intensity-scaled images of many instances of the PSF. Mathematically, this is equivalent to a convolution operation, and when extended to all 3 spatial dimensions, formation of a 3D image $I(x, y, z)$ of a 3D object $O(x, y, z)$ can be modelled as the 3D convolution given by:

$$I(x, y, z) = O(x, y, z) \circledast PSF(x, y, z) \quad (1.9)$$

By convolution theory, this can be equivalently represented as a multiplication in the Fourier domain:

$$\mathcal{F}\{I(x, y, z)\} = \mathcal{F}\{PSF(x, y, z)\} \times \mathcal{F}\{O(x, y, z)\} \quad (1.10)$$

where \mathcal{F} denotes a Fourier transform. Equation 1.10 clearly illustrates how information loss can occur during imaging: if the MTF contains zeros then all information at the corresponding spatial frequencies is irretrievably lost. Convolution is inherently a blurring operation: the image of a point-source is blurred by the optical convolution and appears as the PSF. Therefore, fluorescence images are a blurred estimate of the

true underlying fluorophore distribution at the sample. This raises the question: given the image formation model in equation 1.9, and the fact that the microscope PSF can be measured, can this blur be reversed during image processing to reveal the true underlying sample?

Image deconvolution

The most obvious method to attempt this blur removal is through simple inversion of the Fourier-domain multiplication in equation 1.10. If the Fourier transform of the image can be calculated as the product of the respective Fourier transforms of the PSF and object, then it follows that the object should be restored through the inverse operation. However, as seen earlier in this chapter, MTFs often contain zeros, or very small values close to the cutoff frequency, causing amplification of these frequency components during image restoration. This is problematic, since fluorescence imaging experiments suffer the characteristic noise degradation discussed in section 1.2, with numerous contributions from background and readout sources as well as the stochastic nature of fluorescence emission. Therefore, if noise contributions dominate the spatial frequencies that are close to zero in the MTF, noise amplification occurs during inverse filtering, seriously degrading image reconstruction.

Instead, many image restoration frameworks exist that are designed to not amplify noise in this fashion. Collectively these are known as deconvolution algorithms and can be separated into two general categories: those that are linear, and those that are non-linear. Linear deconvolution algorithms search for a solution that is a linear combination of the Fourier components of the recorded image whereas non-linear algorithms do not share this constraint. Linear algorithms are generally quicker to implement than their non-linear alternatives and include the popular Wiener filter, H_{wiener} , defined as:

$$H_{wiener} = \frac{\mathcal{F}\{PSF(x, y, z)\}^*}{|\mathcal{F}\{PSF(x, y, z)\}|^2 + k} \quad (1.11)$$

where the asterisk denotes the complex conjugate and k is a small constant, typically the noise-to-signal rate of the image data. Non-linear algorithms include the RL algorithm [3, 4], which is perhaps the most widely-used image restoration method in fluorescence microscopy. The RL algorithm forms the basis of much of the image processing in this thesis, and is discussed further in chapter 2.

1.3 Literature review

In this section we take a closer look at existing fluorescence microscopy methods that are able to acquire 3D images. We start with techniques that require scanning to image a volume, before moving on to techniques that do not share this requirement. For each technique, we comment on the resolution they provide in both the temporal and spatial

domains, as well as their optical sectioning abilities and other relevant strengths and weaknesses.

Scan-based volumetric imaging

We begin by considering volumetric imaging techniques that require the scanning of either the sample or illumination to acquire a 3D image. These methods form 3D images by stacking sequentially-acquired 2D images in the course of a z -stack. While such methods are generally able to provide high resolution in the spatial domain, the requirement to scan during acquisition fundamentally limits the resolution offered in the temporal domain. Nonetheless, these methods are widely used in bio-imaging and each provide their own benefits and drawbacks.

1.3.1 Epi-fluorescence focal stacks

The simplest method to image a volume via fluorescence microscopy is to sweep the sample through the focus while imaging on a widefield fluorescence microscope, where the entire sample is illuminated simultaneously via the same objective lens used for imaging with, for instance, a laser beam that is collimated at the sample. Imaging may either be performed during the course of a single camera exposure-window, leading to a z -projection of the sample volume [14, 15], or acquiring one image per z -step, yielding a 3D image. The time needed to acquire a 3D image is determined by the number of z -slices the user wishes to acquire and the exposure-time needed per slice, whereas the spatial resolution of images is determined by the choice of objective lens. However, as we discussed in the previous section, the epi-fluorescence microscope does not inherently reject out-of-focus light and therefore does not achieve optical sectioning, meaning acquired volumes suffer blurring from out-of-focus light, reducing contrast and 3D image quality.

1.3.2 Confocal microscopy

Confocal microscopy [16] is perhaps the most popular method for 2D and 3D fluorescence imaging in the study of cell biology. Instead of illuminating the sample with a collimated laser beam, as in widefield epi-fluorescence microscopy, the confocal microscope illuminates only a small region of the sample with a laser beam that is focused to a diffraction-limited spot. The fluorescence excited within this small region is then imaged through a conjugate pinhole aperture placed at the image plane. The pinhole acts to minimise detection of out-of-focus fluorescence, meaning the confocal microscope achieves optical sectioning, resulting in high-contrast optically-sectioned images. The illumination is then scanned across the sample in a raster fashion to construct either a 2D or 3D image. The combined confocal illumination and detection mean the

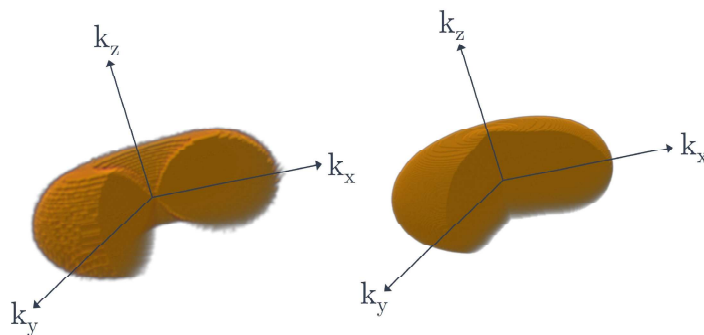


Figure 1.7: OTFs for a widefield (left) and confocal (right) microscope. Both OTFs are shown with a quadrant removed so that their 3D shape is apparent.

confocal PSF is more compact than its widefield equivalent, and the missing cone in the OTF is largely filled [17]. A visualisation of the confocal and widefield OTF is provided in figure 1.7. The temporal resolution of the confocal microscope is limited by both the sensitivity of biological samples to the illumination intensity (which determines the acceptable excitation intensity) and the requirement to scan the illumination relative to the sample during acquisition. Vast improvement in temporal resolution is realised via parallel-acquisition modalities such as spinning-disk [18] or line-scanning confocal microscopy [19], however scanning is still required in the axial direction to achieve 3D imaging.

1.3.3 Two-photon scanning microscopy

Two-photon fluorescence microscopy is another technique that provides optical sectioning and is similar to confocal microscopy in that the sample is most commonly illuminated with focused laser light [20]. In conventional fluorescence microscopy techniques, fluorescence is generated via single photon excitation, where a single excitation photon generates a single emission photon of slightly greater wavelength. Both photons typically fall in the visible region of the spectrum. However, single photon excitation with visible light cannot penetrate far into thick tissue samples, so is not suited to deep-tissue imaging. Two-photon microscopy is based on the principle of two-photon excitation, where two almost-simultaneous excitation photons generate a single emission photon [21]. Through energy conservation arguments, the two incident photons must be of longer wavelength than that needed for single-photon excitation (and the emission wavelength); typically two-photon excitation is in the near infra-red, allowing far deeper penetration into tissue than single-photon microscopy. Two-photon excitation is a non-linear process²: fluorescence emission only occurs within a region very close to the focus of the illumination spot, meaning the PSF is more compact than the widefield microscope and optical sectioning may be achieved without a confocal pinhole present, allowing for higher optical throughput, albeit at a reduced lateral resolution

²Two photon excitation is proportional to the square of the illumination intensity.

compared with a comparable confocal microscope because of the longer emission wavelength [22]. The non-linearity is also beneficial in light-sensitive experiments since the sample undergoes less photobleaching during acquisition compared with single-photon excitation as fluorescence is excited within a smaller region. However while widefield two-photon excitation has been demonstrated [23, 24], two-photon microscopy is more commonly implemented as a point-scanning technique that requires the illumination spot to be scanned across the sample in a raster fashion, limiting the temporal resolution and rendering it unsuitable for imaging highly dynamic scenes in 3D.

Multi-photon imaging

Extending the concepts of two-photon imaging, fluorescence microscopy techniques have been developed that exploit three-photon [25, 26] and even four-photon [27] excitation. Three-photon excitation requires a longer excitation wavelength to excite the same fluorophore compared with single or two-photon excitation. This means that scattering is reduced and deeper penetration into turbid or scattering media, e.g. biological tissue, is facilitated [28]. Further benefits of three-photon microscopy compared with two-photon microscopy include lower background excitation and better optical sectioning [29] owing to an axially-smaller excitation spot [28].

1.3.4 Light-sheet fluorescence microscopy

Light-sheet fluorescence microscopy (LSFM) is another 3D imaging technique that provides optically-sectioned volumes [30]. Here, the illumination is delivered in the form of a light-sheet, formed either by the rapid unidirectional scanning of a laser beam [31, 32] or by focusing a beam along one axis via a cylindrical lens (figure. 1.8). Volumes are then constructed by sequentially imaging xy planes while scanning the sample relative to the light-sheet. Traditional LSFM implementations delivered the light-sheet illumination via a second objective in an orthogonal orientation compared with the imaging objective. While decoupling the excitation and imaging optical paths is beneficial in enabling optical sectioning, the requirement for the launch objective to be in the same plane as the sample and orthogonal to the imaging objective means that the application of conventional LSFM to samples prepared by some common means, for instance those mounted on glass microscope slides, is more challenging [33]. The orthogonal geometry also limits the objective lenses that can be used for imaging, largely restricting the resolution of conventional light-sheet imaging to the cellular level. Additionally, there exists a trade-off between axial resolution and FoV: a higher axial resolution requires a thinner, more tightly-focused light-sheet, which in turn features a shorter Rayleigh-range and therefore limits the FoV. Various methods have been proposed to extend the FoV including exciting the sample with Airy [34] or Bessel beams [35] (the latter of which benefits from combination with two-photon excitation [36]). While LSFM enables rapid acquisition of xy planes, the temporal resolution is




Figure 1.8: Light sheet fluorescence microscopy.

still limited by the requirement to scan the sample or illumination in order to acquire a 3D image.

Oblique Plane Microscopy (OPM)

Oblique-plane microscopy OPM has emerged as an especially powerful fluorescence imaging modality that shares the advantages of light-sheet imaging while enabling samples mounted by conventional means to be imaged [33]. In OPM, a tilted light-sheet is delivered through the same objective used for imaging, with additional corrective optics employed in the imaging path that allow the tilted plane to be imaged. The plane imaged by the primary objective is relayed to the back of a second objective, which reforms the tilted image in a remote-focused volume at its focus. This plane is then imaged onto a camera via a tertiary objective, tilted at the angle of the light-sheet. Volumes are acquired by scanning the light-sheet or sample relative to each other. Computational post-processing is then applied to acquired volumes to remove the shear introduced from the tilt of the light-sheet. A limitation of OPM is that some lateral resolution is compromised owing to the tilt of the tertiary objective, as high-angle rays are not captured. This can be mitigated, however, through use of refractive-index mismatching at the interface between the secondary and tertiary objectives and custom-designed tertiary objectives [37, 38, 39].

Scan-free volumetric imaging

We now consider volumetric imaging methods that do not require scanning during acquisition. The requirement to scan in order to acquire a 3D image via confocal or light-sheet microscopy fundamentally limits the temporal resolution they can provide, making them unsuitable for capturing 3D dynamics that occur over millisecond timescales. For such applications, scan-free volumetric-imaging techniques are desirable.

1.3.5 Bi/multi-plane imaging

Multiplane imaging is perhaps the simplest way of acquiring volumetric information without scanning during acquisition. Here, multiple focal planes at different depths within the sample volume are imaged simultaneously. This may be performed using beamsplitters, which project images of the different focal planes either onto different cameras [40], or onto the same large camera sensor [41, 42]. Alternatively, diffractive optical elements can be used to separate the different images onto different regions of the camera sensor [43], which have been demonstrated to image up to 25 focal planes simultaneously. Multi-plane imaging does not inherently provide optical sectioning, since the planes that are imaged in-focus at one sensor/sensor region contribute out-of-focus light to all others, although this can be suppressed post-acquisition with image deconvolution [42]. Additionally, the axial separation between imaged focal planes typically far exceeds the lateral pixel dimensions, potentially leading to multi-plane focal stacks that are undersampled in the axial direction and feature highly anisotropic voxel dimensions.

1.3.6 Light-field microscopy

Light-field microscopy (LFM) is a scan-free volumetric imaging technique that permits the acquisition of a 3D image within a single snapshot [2, 44]. By inserting a microlens array (MLA) into the imaging path of a widefield microscope, the light-field microscope samples angular as well as lateral information about the observed scene. Originally, plenoptic treatments of the acquired data allowed computational refocussing of the observed scene to be performed post-acquisition, which, when stacked and deconvolved, yielded a 3D image [2]. More recently, a wave optics model of the light-field microscope was developed, permitting volumetric reconstruction from the acquired 2D light-field data at increased resolution compared to the original methods [45]. This work demonstrated how RL deconvolution can be adapted to map between a volume and a 2D image in a similar context to the volume-reconstruction framework developed in chapter 2 of this thesis. However, LFM requires a fundamental sacrifice to spatial resolution, owing to the way the light field is sampled by the MLA [45]. This has mostly limited the application of LFM to imaging sample structure at the cellular level or above. While sub-cellular LFM has been demonstrated, it was limited to use with extremely high-NA objective lenses and sample structure confined to a small depth region close to the coverslip [46]. Additionally, reconstructed LFM volumes suffer from characteristic reconstruction artefacts, owing to a non-uniform sampling pattern across the sample depth, with a particularly information-lacking region at the native focus [45]. Efforts have been made to suppress/avoid these artefacts, through e.g. imaging on only one side of the focal plane, combining with wavefront coding [47], and compressive sensing methods. Despite these drawbacks, LFM has found use in a broad array of biological applications, including imaging neuronal activity [48], calcium imaging [49],




Figure 1.9: Schematic of a conventional light-field microscope (not to scale) [44]. A microlens array (MLA) is placed at the native image plane. The detector (not shown) images the back focal plane of the MLA. Schematic illustrates a ray-optics depiction of a single point-source centered in the FoV at the native object plane. F_T and F_{MLA} are the focal lengths of the tube lens and MLA respectively.

and imaging medaka heart beat and blood flow [50].

Fourier-plane LFM

Conventional LFM placed the MLA at the native image plane of the microscope [2]. This leads to a PSF that is spatially variant, meaning the convolutional model of image formation given by equation 1.9 cannot be applied, and therefore LFM deconvolution is extremely computationally burdensome. Recently, Fourier LFM was introduced, where the MLA was instead placed at a Fourier plane [51]. This yields a PSF that is spatially invariant, simplifying the volume reconstruction process, as well as significantly reducing reconstruction artefacts. Fourier LFM also offers improved resolution compared with conventional LFM, although there is still a significant loss of resolution compared with that provided by the parent imaging system.

1.3.7 PSF engineering

PSF engineering, or pupil plane engineering, is another technique that has found popularity in fast 3D imaging. Here, additional optical elements are introduced into the imaging path of the microscope such that its PSF becomes altered to better suit the needs of a particular experiment. Originally termed wavefront coding, the technique was first introduced in [52], where the authors demonstrated that by introducing a cubic phase term at the exit pupil of an incoherent imaging system, the PSF and MTF became less sensitive to defocus, effectively providing an extension to the depth of field (DoF) (the depth range over which an imaging system can form an in-focus image, measured in object-space) with no reduction to the NA. However, the PSF of a wavefront coded imaging system is typically larger and of a different form than that of an unaltered system (figure 1.11), meaning acquired images must be restored via computational post-processing (typically deconvolution). This combination of optical and computational techniques is termed hybrid imaging [52]. PSF engineering techniques, as well as computational post-processing, underpin the work presented in this thesis. Therefore, it is worth exploring the field a little more in-depth than the




Figure 1.10: Pupil-engineering is typically implemented through the addition of a $4f$ lens relay behind the native image plane.

above-mentioned microscopy methods. The following sections provide a brief theoretical background to the techniques and explore its application in various experimental settings.

Engineering the pupil

The exit pupil represents an ideal location to adjust the characteristics of an imaging system because of the Fourier relationship it shares with the image plane, see equation 1.8. Additionally, the PSF model discussed in section 1.2.3 demonstrates how a defocused PSF can be modelled as the the Fourier transform of a 2D pupil function multiplied by a defocus term [13]. Therefore, the 3D properties of a system with a given pupil function $P(k_x, k_y)$ can be easily predicted, providing scope to design a PSF based on particular experimental requirements. However, the exit pupil of most microscope objective lenses (especially infinity-corrected lenses) is located inside the lens housing and therefore not easily accessible. Therefore, in PSF engineering techniques the exit pupil is typically accessed through the addition of a $4f$ relay behind the primary image plane, see figure 1.10. The PSF is then altered through placement of some phase-modulating device, e.g a glass phase mask, a digital micro-mirror device or a spatial light modulator (SLM), at the re-imaged pupil. Both the phase and amplitude can be modulated at the pupil (the pupil-function is complex-valued); glass masks are typically designed to modulate only phase, whereas SLMs may also serve as an amplitude mask.

Cubic phase-mask

Dowski and Cathey [52] first introduced pupil-engineering in their seminal 1995 work which demonstrated the resulting MTF of a system with a cubic phase mask (CPM) at the exit pupil is free of zeros across an extended DoF. Such a system therefore exhibits less information loss than an equivalent system with no phase-mask present. The cubic phase mask corresponds to a phase-only pupil function given by

$$\phi(k_x, k_y) = \alpha(k_x^3 + k_y^3) \quad (1.12)$$

where ϕ denotes the pupil phase and α is a constant that determines the strength of

the CPM. A visualisation of the cubic phase function, as well as other phase functions is provided in figure 1.11. The resulting PSF of such a system exhibits a distinct “L” shape, with a prominent lobe structure and strong asymmetry, see figure 1.11b. Extended depth of field EDoF imaging using the cubic PSF was first demonstrated experimentally in [53] in a macroscopic setting.

While the cubic PSF does retain a compact shape over an EDoF, it also occupies a larger lateral footprint than the equivalent in-focus unaltered PSF, which acts to blur acquired images by the convolution relationship given in equation 1.9. This means that raw cubic-encoded³ images appear blurred and feature repetition-like artefacts owing to its lobe-structure. Through image deconvolution, however, this blur can be successfully reversed throughout the EDoF to restore a high-quality image, owing to a lack of zeros in the cubic MTF. Originally performed using a simple inverse filter [53], this is more commonly implemented with Wiener deconvolution. [54]. However, while the appearance of the cubic PSF remains largely unchanged over the EDoF, it does exhibit a lateral translation across the FoV, owing to a term in the cubic OTF that is linear with defocus [52]. Additionally, a phase-shift term in the cubic OTF leads to image artefacts being introduced if image restoration is performed using only a single deconvolution kernel, manifesting as replication artefacts in restored images [55]. The cubic OTF is explored further in chapter 3. While attempts have been made to improve deconvolution algorithms to find an optimal kernel to reduce image artefacts [56], performing image restoration with only a single deconvolution kernel is fundamentally unsuitable for cubic-encoded images because convolution with a single kernel is an incomplete description of image formation. Despite these problems, however, the CPM has found use in many different imaging applications, including in conjunction with other imaging modalities such as light-field imaging [47] and light-sheet microscopy [57]. For the same reasons the CPM has found popularity in imaging optics, it has also found applications in the illumination path of light-sheet microscopy, offering a larger FoV and lower peak irradiance than conventional Gaussian light-sheets [34].

Complimentary kernel matching

Complimentary kernel matching (CKM) was introduced as a method to avoid the replication artefacts introduced during image restoration [58]. By imaging the scene simultaneously with a pair of PSFs that respond differently to defocus, an optimal deconvolution kernel for local image regions could be found by varying the kernel defocus until artefacts in both deconvolved image regions were minimised. As well as restoring an artefact-free image, this method also benefited from providing a depth-map of small patches within the image; extracting latent depth-information contained in the acquired images through their disparity. Because the method sub-divides the images into smaller patches, it is optimal for scenes containing smooth surfaces, however does

³Raw images acquired on a system exhibiting a cubic PSF are said to be “cubic-encoded”.




Figure 1.11: Phasemask and corresponding PSFs over an axial range of $\pm 15 \mu\text{m}$. Left column shows the phase function for (a) the unaltered, diffraction-limited PSF; (b) the cubic PSF; (c) the Twin-Airy PSF and (d) the COSA psf. Center 3 columns show the PSFs at $-15 \mu\text{m}$, $0 \mu\text{m}$ and $15 \mu\text{m}$. Right column displays a maximum intensity projection of each PSF over the full $30 \mu\text{m}$ range, colour-coded for depth.

not work as well for more sparse scenes that feature large depth discontinuities within the smaller image regions. Additionally, since the method provides only a 2D depth-map rather than a full 3D volume, it can not be described as a volumetric imaging technique. However, the method for encoding depth information into snapshot images embodied by CKM forms the basis of the snapshot volumetric imaging technique presented in chapter 3 of this thesis. By applying the mathematical framework developed in the following chapter to images encoded in this way, we demonstrate how full 3D volume reconstruction is possible by exploiting the depth-information encoded into PSF-engineered images.

Point localisation

Point localisation and single molecule localisation microscopy (SMLM) applications are perhaps the most common use of PSF engineering techniques in the life sciences. Here, the spatial coordinates of individual molecules or point-emitters are determined from the acquired 2D PSF-engineered images. In SMLM, localisation precisions of

beneath 50 nm are possible [59]. Hence, SMLM is said to achieve super-resolution. Often, SMLM is used alongside fluorophores that undergo photo-switching in techniques such as (fluorescence) photo-activated localisation microscopy (fPALM) [60, 61] and stochastic optical reconstruction microscopy (STORM) [62]. Here, fluorophores randomly switch between an “on” state and an “off” state. Consequently, only a sparse subset of the fluorophores labelling a sample are emitting fluorescence photons at any given time-point. This helps ensure sparsity in images, allowing the coordinates of spatially-separated PSFs to be calculated. In 2D, this is often performed by fitting e.g. Gaussian profiles to PSFs. However, performing SMLM in only two dimensions is often insufficient, since most biological samples of interest exhibit 3D structure and dynamics. It is therefore desirable to also obtain the axial position of fluorophores in addition to the coordinates in x and y .

Determining the axial coordinate of fluorophores introduces significant additional challenges to SMLM experiments for a number of reasons. Firstly, super-resolution (SR) imaging of this kind typically requires extremely high NA lenses, with an associated DoF of just a few-hundred nanometers either side of the focal plane. Additionally, the diffraction-limited PSF is approximately symmetric about the focal plane and contains very little information about depth, meaning it is not possible to extract the axial coordinate of emitters. PSF engineering has been implemented as a tool to solve both these challenges and enable 3D SMLM. Early implementations exploited astigmatic PSFs, formed by inserting a cylindrical lens into the imaging path, yielding an asymmetric PSF whose shape is indicative of depth [63]. However, while the astigmatic PSF remains compact, it does not inherently extend the system DoF, meaning the depth range over which 3D-SMLM could be performed was limited to approximately 600 nm. More recently, several other PSFs have been designed that simultaneously extend the DoF while encoding depth information, thus permitting 3D-SMLM to be performed over an extended depth range. The earliest example of such a PSF was the double-helix (DH), that exhibited a pair of lobes whose orientation is indicative of depth [64]. The DH PSF effectively doubled the available DoF to approximately 2 μm while maintaining strong localisation performance. Since then, several other PSFs have been proposed which further extend the DoF. These include the Tetrapod [65], Cropped-oblique secondary astigmatism (COSA) [66], Airy/self-bending PSFs [67] and the twin-Airy PSF [68].

Performing 3D localisation from PSF-engineered image data presents significant additional challenges compared with performing 2D point-localisation. This is because (i) a Gaussian profile may no longer be an adequate description of the PSF, and (ii) the image domain lateral footprint of engineered PSFs is typically far greater than a diffraction-limited PSF, meaning signal-to-noise ratio in images is typically far worse, and samples must be labelled less densely to avoid overlapping PSFs in images. This is explored further in chapter 4 of this thesis, where the volume reconstruction methodology developed in the following chapter is applied to SMLM, where we see it introduce

significant improvements to localisation performance.

Extended depth-of-field and projection imaging

The final method for obtaining 3D fluorescence information that we consider in this chapter is through EDoF or projection imaging, where an extended depth-region within the sample is projected into a single 2D acquisition. This type of imaging may be implemented by extending the DoF with, for instance, PSF engineering where no axial scanning is required to image throughout an extended axial range [52]. The CKM method discussed above falls into this category [58]. Alternatively, projection images may be acquired by rapidly sweeping the focus of the imaging system axially through the sample during the course of a single camera exposure-window. Such focus-sweeping may be implemented with, for instance, a tunable lens in the imaging and/or illumination paths [69, 70], or with deformable mirrors used to rapidly adjust the focus of the imaging system [71]. While projection imaging methods do not enable fast 3D imaging, they do permit 2D imaging through entire sample depths within the course of a single camera exposure-window, offering video-rate imaging speeds. However, images acquired in this way contain no information about the position of sample features along the projection axis. Recently, multi-angle projection imaging was introduced, where it was demonstrated that by incorporating a scan-unit into the imaging path of any microscope capable of optical-sectioning, the angle at which the sample is projected could be varied [5]. This enables real-time viewing of projections of the sample from varying angular perspectives, as well as allowing a projection-series to be acquired where the sample is optically rotated during an acquisition. While it remains true that a single projection contains no information about the position of sample features along the projection-axis, we shall see in chapter 5 that two or more projections do capture 3D information about the sample, providing scope to perform a full 3D reconstruction of the sample volume.

1.4 Chapter 1 summary

In this chapter we've discussed image formation on a fluorescence microscope and introduced topics that are central to the methods developed in the remainder of this thesis, including PSF engineering and image deconvolution. We then discussed methods that are available to acquire 3D fluorescence images. Broadly, these techniques fall into two categories: those that require some form of scanning to capture 3D information, and those that do not. Scan-based methods such as light-sheet and confocal imaging offer high spatial resolution, yet the temporal resolution is fundamentally limited by the requirement to scan. Conversely, scan-free volumetric imaging methods such as LFM offer fast volume-imaging, at the expense of spatial resolution. In the remainder of this thesis we develop imaging methods that overcome this trade-off: offering high

resolution in both the spatial and temporal domains by performing full 3D volume reconstruction from rapidly-acquired projection images, obtained via various imaging modalities. We start in the next chapter by modelling projection image formation and developing a mathematical framework that allows us to perform 3D reconstruction from such projection images, providing certain criteria are satisfied.

Chapter 2

2D-3D volume reconstruction

In the previous chapter we discussed a number of methods for acquiring 3D image data via fluorescence microscopy. Confocal microscopy remains the most popular method for volumetric imaging in the life sciences owing to its resolution, versatility, and ease-of-use, in-part due to the large number of commercial systems available. Yet, its slow acquisition speed fundamentally prevents its application to dynamic scenes. Light-sheet microscopy offers some improvement by acquiring entire xy planes simultaneously, however scanning is still required along the axial direction, again preventing its application to highly dynamic scenes. Conversely, LFM enables 3D imaging of fast-moving samples through its snapshot volumetric imaging capability. However, this capability is provided at the expense of significant spatial resolution loss, limiting the type of sample it can be applied to, as well as characteristic artefacts being introduced into reconstructed volumes [45].

The goal of this thesis is to overcome some of these limitations by developing an imaging pipeline that enables high-resolution volume reconstruction from 2D projection data, with the reconstruction process extracting depth information that is encoded into the acquired 2D projection images. In doing so, this permits fast volumetric imaging by requiring only a few projection images, rather than an entire z -stack or raster scan to image a 3D sample. We will see in later chapters that if these projection images are acquired simultaneously, this enables high-resolution snapshot volumetric imaging. In the current chapter, we develop a general mathematical framework that enables 3D reconstruction from 2D projection images. This framework underpins the volume-imaging techniques that are presented in the rest of this thesis. In chapters 3-5 we then apply this framework to different imaging scenarios including SMLM, PSF-engineered microscopy [72] and multi-angle projection imaging [5]. We begin this chapter by outlining the challenges associated with imaging an entire volume in a snapshot that we must solve in order to achieve high-resolution snapshot volumetric imaging.

2.1 Snapshot volume imaging

Acquiring a 3D fluorescence image of an entire volume in a single snapshot introduces two key challenges. Firstly, the depth-range over which a microscope can form an in-focus image, known as its DoF¹, will typically be far smaller than the thickness of the sample of interest. The DoF is defined as the depth, measured in object space, between the nearest and farthest planes that simultaneously appear sharply in focus [10] and is calculated as:

$$DoF = n \left(\frac{\lambda}{NA^2} + \frac{e}{M \cdot NA} \right) \quad (2.1)$$

where M is the lateral magnification of the system and e is the pixel size of a given detector placed at the image plane [73]. Any fluorescence contributing to images that originates from outside this region will be imaged out-of-focus. The diffraction-limited DoF of a microscope is primarily determined by the objective lens used for imaging and shares an inverse square relationship with the NA, meaning that as the NA of the lens is increased, the associated DoF becomes smaller. For instance, the high-NA lenses commonly used in live-cell imaging have an associated DoF of just a few hundreds of nanometers. There is, therefore, a trade-off between spatial resolution and depth-of-field. We saw in section 1.2 how significant contribution to images from out-of-focus light will degrade image quality, particularly when the sample is imaged with conventional widefield illumination. A system that enables snapshot volume imaging, therefore, should extend the DoF so that entire samples can be imaged in-focus. Ideally, this should avoid a reduction in the system NA so that spatial resolution and optical throughput aren't compromised.

The second challenge associated with imaging an entire volume in a single snapshot is more fundamental than the first. Data acquired in a single snapshot (or camera exposure-window) is, at best, going to be 2D. The data that we wish to reconstruct, however, is 3D. This means that the depth of image features must be extracted during the computational reconstruction process to return to a 3D volume. For this to be possible, depth information must first be present within the 2D snapshot data. Images acquired on widefield systems do not meet this requirement: the PSF of an unaberrated widefield microscope is approximately symmetric about the focal plane, as illustrated in figure 2.1, and does not change significantly within the DoF, meaning that depth information is not effectively encoded into images. Consequently, the calculation of depth from snapshot images acquired with standard widefield microscopes is not possible.

We see, therefore, that to effectively perform snapshot volumetric imaging, additional steps must be taken to (i) extend the DoF of our microscope, and (ii) effectively encode depth information into acquired images. We saw in chapter 1 how one method

¹when measured in object-space. The equivalent quantity measured in image space is the depth-of-focus.

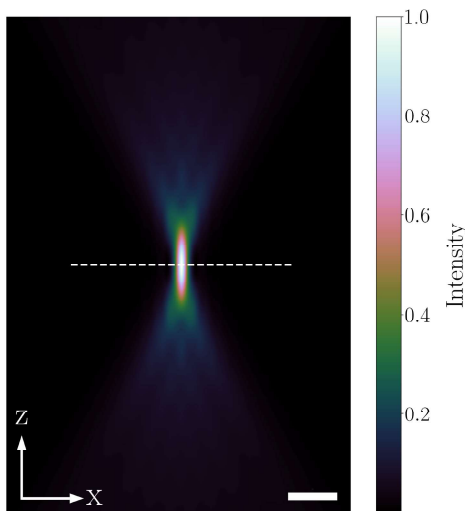


Figure 2.1: zx projection of a simulated 0.8NA widefield PSF, showing approximate axial symmetry about the focal plane, indicated by the dashed white line. Scale-bar is $2\ \mu\text{m}$.

to solve both of our challenges was through use of PSF engineering techniques. For example, CKM encodes depth information across an EDoF into snapshot images by imaging simultaneously with a pair of engineered PSFs that respond differently to defocus [58]. However, the image processing employed in CKM yielded an artefact-free 2D image and localised 2D depth-map, not a 3D volume. Since we wish to reconstruct an entire 3D volume from the acquired 2D data, a new image processing framework is required.

In the remainder of this chapter, a general methodology for full 3D volume reconstruction from 2D projection data is developed. At this point, we make no assumptions regarding *how* the required projection images are acquired; we simply state that both challenges outlined in this section have been addressed, i.e. our chosen imaging modality effectively encodes depth information into the acquired projection data across the necessary DoF. We shall see in later chapters that the volume-reconstruction pipeline developed here is not limited to any particular imaging scenario: it represents a general result, applicable to data acquired via *any* projection imaging modality, providing our two challenges are addressed. In chapters 3 and 4 we employ it alongside PSF-engineering; whereas in chapter 5 we explore its use alongside other imaging modalities and geometries. The volume-reconstruction framework is based on iterative RL deconvolution, and that is where our discussion starts.

2.2 Richardson-Lucy deconvolution

As we saw in chapter 1, the RL algorithm is a non-linear image restoration process, derived independently by William H. Richardson in 1970 [3] and Leon B. Lucy in 1974 [4] through Bayesian statistical considerations, treating the image and PSF as probability density functions. Since the volume reconstruction pipeline developed in

this chapter is based on RL deconvolution, it is worth summarising the derivation here. For simplicity, we constrain our initial discussion to one dimension.

We start with the fact that fluorescence imaging is characterised by Poisson noise statistics. That is to say, if we have a perfect noise-free detector, then the probability p of a pixel detecting a given number of photons n is given by

$$p(n) = \frac{e^{-\lambda} \lambda^n}{n!} \quad (2.2)$$

where $!$ is the factorial operation and λ is the expected (mean) number of photons. This is because fluorescence emission occurs stochastically, and fluorescence detection is a discrete process. The probability of recording the photon count observed across an entire image $I(x)$ is the product of these probabilities for every pixel in the image:

$$p(I(x)) = \prod_x \frac{e^{-\lambda(x)} \lambda(x)^{I(x)}}{I(x)!} \quad (2.3)$$

Here, x is the pixel index and $\lambda(x)$ is the expected value for each pixel (i.e. the underlying object). In Bayesian statistical terms, the probability of recording an image $I(x)$ given an underlying object $O(x)$ is expressed $p(I(x) | O(x))$ and is known as the likelihood function. The posteriori probability, $p(O(x) | I(x))$, is the probability a given object is the true underlying object given (i) the recorded (image) data $I(x)$, and (ii) the prior model (in our case, the Poisson distribution in equation 2.3). To solve the inverse imaging problem, we wish to estimate the most probable underlying object $\hat{O}(x)$ (the maximum likelihood solution), such that $I(x)$ is explained by our prior model. To achieve this, we see by equation 2.3 that we need to maximise the likelihood function $p(I(x))$. We observe that maximising the likelihood function (equation 2.3) is equivalent to finding the maximum of its natural logarithm:

$$\begin{aligned} \ln p(I(x)) &= \sum_x \left(\ln \left(e^{-\lambda(x)} \lambda(x)^{I(x)} \right) - \ln (I(x)!) \right) \\ &= \sum_x \left(-\lambda(x) + I(x) \ln (\lambda(x)) - \ln (I(x)!) \right) \end{aligned} \quad (2.4)$$

where \sum denotes discrete summation. As we saw in section 1, if our system PSF is spatially invariant, our image formation can be modelled as a convolution operation between the true underlying object $O(x)$ and the PSF. Therefore, the expected distribution $\lambda(x)$ can be written with this convolution, and our log-likelihood becomes

$$\ln \sum_x \left(- (O(x) \otimes PSF(x)) + I(x) \ln (O(x) \otimes PSF(x)) \right) \quad (2.5)$$

where the last constant term in Equation 2.4 has been dropped since it is constant and therefore does not affect the position of a maximum. At this point we note that maximising the log-likelihood is equivalent to minimising the negative-log likelihood.

The negative likelihood, which is a convex function, then becomes our loss function, and is conventionally given the symbol J . In solving the inverse problem, therefore, we wish to find the object $O(x)$ that minimises:

$$J(O(x)) = \sum_x \left(O(x) \otimes PSF(x) - I(x) \ln (O(x) \otimes PSF(x)) \right) \quad (2.6)$$

To find the minimum of this function we calculate its gradient and set to zero:

$$\nabla J(O(x)) = \left[1 - \frac{I(x)}{O(x) \otimes PSF(x)} \right] \otimes PSF(-x) = 0 \quad (2.7)$$

A full calculation of this derivative can be found in [74]. Assuming the PSF is normalised to unity, the RL deconvolution algorithm attempts to solve for $O(x)$ in a multiplicative fashion through iterative updates to an estimate of the object $\hat{O}(x)$:

$$\hat{O}^{n+1}(x) = \hat{O}^n(x) \cdot \frac{I(x)}{\hat{O}^n(x) \otimes PSF(x)} \otimes PSF(-x) \quad (2.8)$$

where n is the iteration index. The RL implementation given in Equation (2.8) is applicable only when the PSF is spatially invariant. A more general description, applicable to any PSF is given by

$$\hat{\mathbf{O}}^{n+1} = H^T \frac{\mathbf{I}}{H \hat{\mathbf{O}}^n} \cdot \hat{\mathbf{O}}^n \quad (2.9)$$

where \mathbf{I} is a column vector of the recorded image pixel values, H is a transfer matrix, describing the mapping between image space and object space; $\hat{\mathbf{O}}$ is a column vector of estimate at iteration n of the true underlying object, and the superscript T denotes the matrix transpose. The matrix operation $H\mathbf{O}$ is known as the forward operation, mapping the true object to the recorded image. Conversely, the backward operation maps image space back to object space, and is described by the matrix H^T . For a system that is spatially invariant (for example, many microscopy platforms), the transfer matrix H takes on a diagonal-constant (Toeplitz) form, the effect of which is to perform a discrete convolution [75]. This means that equation 2.9 can be written instead with the operations described by convolutions, as given in equation 2.8.

Because RL deconvolution improves likelihood given a Poisson noise model, it has become ubiquitous in the processing of fluorescence microscopy data. However, we see from equation 2.8 that the convolutional implementation of RL deconvolution preserves dimensionality between the image and the object, and is therefore not immediately applicable to our problem of reconstructing a 3D volume from 2D image data. In the following sections we adapt our model for image formation to reflect projection image formation across an EDoF, mapping between the 3D sample and 2D projection images. We will then see how the RL deconvolution algorithm may be modified to reflect this




Figure 2.2: RL deconvolution improves contrast and removes blurring in fluorescence images. (a) Simulated 2D 20X 0.3NA fluorescence image of sparse point-sources, subject to both Poisson and Gaussian noise. (b) Restored image after 25 RL deconvolution iterations. Insets show the same regions in both panels, indicated by the cyan boxes. Plots show normalised traces along the dashed lines in (a) and (b). The restored image shows increased contrast between the point-sources and background, as well as reduction in the size of the points. Scale-bar is $10\ \mu\text{m}$. The display ranges for both images are set to their respective minimum and maximum values.

model for image formation, permitting the reconstruction of a 3D volume from our 2D projection images.

2.3 Model for image formation

In chapter 1, we saw how the formation of a microscope image, providing the system PSF does not vary across the FoV, can be modelled as the convolution between the object and the PSF. This is also seen in equation 2.8, where the forward operation at iteration n is the convolution between the PSF and the current estimate to the underlying object \hat{O}^n . Generally, a 2D fluorescence image recorded at the image plane of such a system by a scientific camera can be modelled as:

$$I(x, y) = \text{Pois}\left\{O(x, y) \otimes \text{PSF}(x, y)\right\} + N(x, y) \quad (2.10)$$

where Pois represents the Poission distribution describing fluorescence emission and $N(x, y)$ is a noise term reflecting multiple sources including background fluorescence, dark current and readout noise, as discussed in chapter 1. This equation is easily ex-




Figure 2.3: Projection image formation can be modelled as the sum of 2D convolutions between the PSF and object throughout the DoF.

tended into 3D for modelling volumetric image formation in modalities such as confocal or light-sheet imaging. However, since the dimensionality of the acquired image and object are the same in this description, it does not adequately describe the acquisition of the type of EDoF image under consideration in this chapter. Instead, our model must describe the projection of the 3D sample onto the 2D snapshot images. Again, assuming spatial invariance of the PSF, this EDoF image formation can be modelled as the incoherent summation of 2D convolutions between xy planes in the 3D object and 3D PSF at each discrete axial position throughout the DoF:

$$I(x, y) = \text{Pois} \left\{ \sum_k O(x, y, z = k) \otimes \text{PSF}(x, y, z = k) \right\} + N(x, y) \quad (2.11)$$

Figure 2.3 illustrates a visualisation of this model of image formation. Equation 2.3 describes how the 3D sample object is projected onto the 2D camera sensor during snapshot volume acquisition with an EDoF. The noise term describes all sources of noise present in acquired images, including those discussed in chapter 1 (readout and dark current noise), as well as any background fluorescence that may originate from outwith the DoF.

2.4 Volume reconstruction

Now we have described how the image of a 3D sample is formed onto a 2D camera, we must consider how to invert this model in order to reconstruct the 3D sample and achieve snapshot volume imaging. With this model of image formation, the standard convolutional implementation of the RL algorithm is not applicable, since the dimensionality of the object and image are different. Instead, we revert to the the matrix description in equation 2.9. Generally, the model for image formation serves as the forward operator in RL deconvolution and is described by the matrix equation $\mathbf{I} = \mathbf{H}\mathbf{O}$. In our case, the matrix H has a well-defined sub-structure:

$$\mathbf{I} = [H_1, H_2, \dots, H_N]\mathbf{O} \quad (2.12)$$

where N is the number of discrete z planes we wish to reconstruct. Here, each sub-matrix H_n has a Toeplitz structure, the effect of which is to perform a 2D convolution. This equation is an equivalent matrix representation of the model of image formation given by equation 2.11, without the noise term present. The inverse operation, or backward operator, must map the 2D snapshot images back to the 3D sample volume. We see in equation 2.9 that in the matrix description of RL deconvolution, the backward operator is given by H^T . Indeed, that is also true in our case, and we therefore see that, in a similar fashion to the forward operation, the backward operation may be implemented as a series of 2D convolutions. Therefore, our adapted RL deconvolution scheme that reconstructs a 3D volume from 2D projection data is given by:

$$\hat{O}^{n+1}(x, y, z = k) = \hat{O}^n(x, y, z = k) \left[\frac{I(x, y)}{\sum_k \{PSF(x, y, z = k) \otimes \hat{O}^n(x, y, z = k)\}} \otimes PSF(-x, -y, z = k) \right] \quad (2.13)$$

This equation demonstrates how a 3D sample volume may be reconstructed from snapshot image data and a 3D PSF, with the model of image formation described by equation 2.11.

2.5 Multi-view volume reconstruction

The next piece of our puzzle is to incorporate the information content from multiple images into our deconvolution pipeline. We saw in the above sections, that in image formation modalities such as that employed by CKM, two independent imaging channels with PSFs that respond differently to defocus are required to encode depth information into the snapshot images. Extracting this information, therefore, requires the deconvolution to consider the information content from both snapshot images. RL deconvolution has been shown to enable the information content from multiple views to be combined together to improve the overall quality of the reconstructed object [76, 77]. Examples of this application of RL deconvolution include improving the spatial resolution of light-sheet microscopy [78, 79] and combining orthogonal views in LFM [50]. Generally, the inclusion of multiple perspective views improves deconvolution performance by reducing the ill-posedness of the inverse problem. However, in the case of, for instance, CKM, the multiple imaging channels are essential to encode the information we need into images: without the two PSFs that respond differently to defocus, depth information is not present in our data and it is not possible to reconstruct a volume. Therefore, it is crucial that our deconvolution considers the information content of all imaging channels simultaneously. This is achieved through a joint deconvolution scheme, where the back-projections from all imaging channels are averaged

each iteration:

$$\hat{O}^{n+1}(x, y, z = k) = \hat{O}^n(x, y, z = k) \frac{1}{M} \sum_m \left[\frac{I_m(x, y)}{\sum_k \{PSF_m(x, y, z = k) \otimes \hat{O}^n(x, y, z = k)\}} \otimes PSF_m(-x, -y, z = k) \right] \quad (2.14)$$

where m is the index of the imaging channel and M is the total number of imaging channels present. For ease, if we define a back-projection $B_m^n(x, y, z)$ term for the m^{th} imaging channel at iteration n to be:

$$B_m^n(x, y, z = k) = \frac{I_m(x, y)}{\sum_k \{PSF_m(x, y, z = k) \otimes \hat{O}^n(x, y, z = k)\}} \otimes PSF_m(-x, -y, z = k) \quad (2.15)$$

then our volume-reconstruction algorithm can be expressed as:

$$\hat{O}^{n+1}(x, y, z = k) = \hat{O}^n(x, y, z = k) \cdot \frac{1}{M} \sum_m B_m^n(x, y, z = k) \quad (2.16)$$

2.6 Edge artefacts

A key problem faced generally in deconvolution is the introduction of artefacts close to edges of the reconstructed images. This type of artefact can be attributed to the fact that convolution is inherently a blurring operation and causes the image of a point to spread out. In fluorescence imaging, this leads to two specific issues that must be considered by a deconvolution framework to prevent edge artefacts seriously degrading the restored image: (i) some light emitted by regions of the sample close to the edges of the imaging domain may exceed the camera FoV and therefore not be present in images, and (ii) similarly, sample features located outwith the imaging domain may emit light that is recorded within the FoV. This leads to large discontinuities at the image boundaries. These factors become particularly problematic when imaging with PSFs that exhibit a lateral footprint far larger than the diffraction-limited PSF, such as the cubic PSF.

One strategy to reduce these types of artefact is to pad the acquired image data with zeros prior to deconvolution [80]. This allows the image-restoration procedure to reconstruct a larger domain than the acquired image, which can then be cropped down to the required size post-deconvolution. This assumption of abrupt boundary conditions, however, is a poor description of true imaging conditions and can lead to poor deconvolution performance in these image regions. Different boundary conditions that can instead be imposed are reflective or periodic, where the image outside the FoV




Figure 2.4: Simulation demonstrating how edge artefacts can degrade the deconvolved image. (a) Ground-truth image. (b) cubic-encoded image, formed by convolution of the ground-truth with a 2D $\alpha = 3$ 0.5NA cubic PSF. (c) Solution after 25 deconvolution iterations. (d) Solution after 25 iterations, after zero-padding the image. (e) Solution after 25 iterations, with reflective boundary conditions. (f) Solution after 25 iterations, with the edge correction implemented during deconvolution. Panes above each plot show zoomed-in versions of the regions bounded by the red boxes. Scale-bar is 20 μm .

is assumed to be either a mirrored or exact copy of the recorded data [81]. However, while such boundary conditions do remove discontinuities at the image edges, they are also a poor description of the true conditions and can again lead to the introduction of artefacts. Fundamentally, the image outwith the FoV is unknown and therefore the assumption of any boundary conditions should be avoided.

An alternative method for the reduction of edge artefacts was introduced in [82] that makes no assumptions about conditions of the image outside the recorded image domain. Instead, the RL algorithm is free to choose what happens at the boundary according to the image data and the deconvolution kernel. Originally proposed for the deconvolution of astronomical images, the method makes no prior assumptions about the sample or image outside the FoV. For the 2D-3D volume reconstruction algorithm employed in this thesis, a similar method to reduce edge-artefacts is applied through modification to the joint deconvolution scheme given by equation 2.14. This modification is discussed below.

2.6.1 Edge correction

The modification made to reduce edge effects is designed to deal with the two specific scenarios mentioned above and takes the form of a corrective mask applied every iteration. The corrective mask is formed as follows:

Firstly, the raw 2D images are zero-padded in x and y by a small number of pixels, typically in the range 10-25. Then, for each imaging channel m , a binary mask G is created that reflects the information content of the now-extended imaging domain: everywhere we have information, the mask is set to one; elsewhere it is zero. For a raw image with N_x and N_y pixels in the x and y directions respectively, this mask is defined as:

$$G^m(i, j) = \begin{cases} 1, & \text{if } |i| \leq \frac{N_x}{2}, |j| \leq \frac{N_y}{2} \\ 0, & \text{otherwise} \end{cases} \quad (2.17)$$

where i and j are the pixel coordinates in x and y respectively, for a coordinate system with the origin positioned at the centre of the image. This binary mask G is then back-projected according to the back-projection operation defined in equation 2.14 to form a 3D mask, with the same lateral dimensions as G and an axial dimension the same as the volume we wish to reconstruct. The inverse of this mask is the corrective weighting C_m that is applied every iteration during deconvolution:

$$C_m(x, y, z = k) = \frac{1}{G_m(x, y) \otimes PSF_m(x, y, z = k)} \quad (2.18)$$

The effect of this mask is to suppress the influence of discontinuities at the image boundaries during deconvolution, allowing the deconvolution to reconstruct structure outside the domain originally captured by the FoV of the camera. This region will, however, be incomplete, since the information content in the acquired images about this region is only the light that has leaked into the imaging domain. Therefore, the final reconstructed volume is cropped in xy to match the domain that was originally imaged. The effect of the edge correction modification discussed in this section is illustrated in figure 2.4. Here, a cubic-encoded image is deconvolved with different methods of handling the artefacts introduced at the edges of the image domain. Significant artefacts at the edge of the reconstructed image are introduced with both an unmodified RL deconvolution (figure 2.4c) and a scheme where the image is zero-padded prior to deconvolution (figure 2.4d). While the artefacts are reduced when reflective boundary conditions are assumed (figure 2.4e), the greatest improvement is seen in figure 2.4f, where the edge modification discussed in this section is applied and artefacts at the edges of the reconstructed image are mitigated almost entirely.

2.7 Noise amplification and regularisation

RL deconvolution is a particularly attractive image restoration process in fluorescence imaging for a number of reasons. For instance, inspection of equation 2.8 reveals that there is an implicit non-negativity constraint “built-in” to the algorithm, providing the input data does not contain any negative values. This is unlike, for instance, Wiener deconvolution, which can introduce non-physical negative values into the restored image. Other advantages include conservation of flux (providing the PSF is normalised), and its convergence to the maximum likelihood solution for a Poisson description of the data. However, there are a number of drawbacks to the algorithm that must also be considered. Firstly, the computational burden of RL deconvolution is significantly greater than other image restoration techniques such as Wiener or inverse filtering. While this may be somewhat mitigated through use of graphics-processing units for parallel computation, it remains a major drawback for the processing of large multi-dimensional datasets. Another significant drawback of RL deconvolution is the phenomenon of noise amplification. This occurs when the algorithm fits too closely to noise fluctuations in the raw data. In order to replicate small pixel-to-pixel noise fluctuations in the observed data via the forward operation, the deconvolution increases the flux in the brightest pixels, while the surrounding pixels fall closer to zero as the solution iterates. This is explained by considering the fact that convolution is a blurring operation: to reproduce a small noise spike in the observed data, a larger spike in the restored image is required. The effect worsens as the solution iterates further: the bright pixels become brighter and, because of flux conservation, the surrounding pixels fall closer to zero. This concept is particularly problematic for regions of extended structure in a noisy image and can lead to a speckled appearance in the restored image that is unrepresentative of any true sample structure, limiting how far the solution can be left to iterate.

Several methods exist to attempt to reduce the effect of noise amplification during RL deconvolution. The simplest involve attempting to prevent noise amplification by removing noise fluctuations from the raw data, through median or Gaussian filtering the raw images prior to deconvolution [83]. Other methods involve filtering the restored images to remove the amplified noise [84]. However, neither of these solutions are optimal: filtering the raw data removes physical structure from images as well as noise, whereas Gaussian filtering the deconvolved images causes resolution loss. Instead, modifications to the deconvolution scheme itself may be implemented, commonly derived from additional terms in the optimisation problem, in the process of regularisation.

Regularisation

Regularisation, in the context of RL deconvolution, refers to modifying the algorithm in some way to avoid the introduction of undesirable effects or artefacts into the solution, for instance those caused by noise amplification. Generally, regularisation is informed

by the incorporation of some prior knowledge into the optimisation problem, meaning the deconvolution process becomes one of *maximum a posteriori* estimation. In this thesis, the total variation regularisation term derived in [74] is applied during RL deconvolution to suppress noise amplification. Total variation regularisation benefits from preserving edges in images while suppressing noise amplification, unlike other methods such as those based on Tikhonov-Miller regularisation, which are known to smooth out edges [85, 86]. The total variation regularisation term, applied at the n^{th} iteration R_{TV}^n , is given by [87]:

$$R_{TV}^n(x, y, z) = \frac{1}{1 - \lambda_{TV} \operatorname{div} \left(\frac{\nabla \hat{O}^n(x, y, z)}{|\nabla \hat{O}^n(x, y, z)|} \right)} \quad (2.19)$$

where div denotes divergence, ∇ is the vector differential operator and λ_{TV} is a small constant. The final deconvolution algorithm, with the back-projection term B , the edge-corrective term C and regularisation term R_{TV} , then becomes:

$$\hat{O}^{n+1}(x, y, z = k) = \hat{O}^n(x, y, z = k) \cdot R_{TV}^n(x, y, z) \cdot \frac{1}{M} \sum_m B_m(x, y, z = k) \cdot C_m(x, y, z) \quad (2.20)$$

2.8 Chapter 2 summary

In this chapter we developed a method for the full 3D volume reconstruction from one or more projection images of the sample in question. We saw how the RL deconvolution algorithm could be derived by maximising the likelihood function for data governed by Poisson statistics. Furthermore, when the model for image formation was amended to map between the 3D sample and 2D projection data, we saw how the RL algorithm could also be amended to reconstruct the 3D sample from one or more 2D projection images. Additionally, we introduced terms to regularise the deconvolution to limit the effect of noise amplification and to avoid the introduction of edge artefacts arising from fluorescent structure positioned in close proximity to the edge of the imaging domain. However, up to this point, we made no assumptions about how our projection data is acquired: we simply stated that our chosen projection method effectively encodes depth information across an EDoF. In the following chapter, we see how the volume reconstruction framework developed in this chapter may be applied to PSF-engineered projection images that satisfy this criteria, enabling volume-reconstruction from cubic-encoded images acquired in a single snapshot.

Chapter 3

3D engineered PSF microscopy

In chapter 1 we saw a number of different methods for acquiring volumetric images via fluorescence microscopy. While confocal, light-sheet and multi-photon scanning imaging methods provide high-resolution optically-sectioned images, their acquisition speeds are fundamentally limited by the requirement to scan the sample relative to the illumination, or vice versa. This makes them unsuitable for the study of biomechanics where the sample of interest undergoes motion during acquisition. LFM, on the other hand, permits volumetric imaging from image data acquired in a single snapshot, therefore providing the temporal resolution required for the study of highly dynamic samples. However, the associated compromise in spatial resolution limits the application of LFM to the cellular level, and the characteristic reconstruction artefacts can degrade reconstructed volumes.

In chapter 2, we developed a mathematical framework for the computational reconstruction of 3D volumes from an arbitrary number of 2D projection images. The two main criteria for such a projection-imaging system are that (i) the DoF is extended such that the entire sample can be imaged in-focus, and (ii) depth information is encoded into the acquired projection images. Light-field images might, at first, seem like ideal candidates for this volume-reconstruction process. LFM achieves EDoF imaging and effectively encodes depth through its parallel acquisition of angular and spatial information. However, the light-field PSF (in its original implementation) is not spatially invariant, meaning the convolutional image formation model in equation 1.9 can not be applied. Fourier-LFM solves the issue of the spatially-variant PSF, however the associated resolution-loss is still problematic. In this thesis we aim to develop volume-imaging techniques that provide high resolution in both time and in space. Therefore, the ideal imaging system would not require any decrease in spatial resolution in order to provide snapshot volumetric imaging.

PSF-engineering seems, on the surface at least, a better candidate for tackling our problem. While PSF-engineering can successfully encode depth information across an EDoF, the introduction of a phase-mask at the exit pupil does not automatically compromise the resolution of the imaging system. Indeed, its adoption in SMLM proves

its capability of effectively addressing these two challenges. However, the uptake of the technique in volumetric imaging has been more limited: instead it is more-often used as a tool to enhance the performance of other imaging modalities, e.g. light-sheet imaging, two-photon imaging or LFM [57, 88, 47], with a few notable exceptions. Zammit *et al.* argued that this slow uptake is due to the artefacts introduced during image restoration with e.g. Wiener deconvolution [58]. Yet, by varying the deconvolution kernel, the same authors demonstrated artefact-free 2D EDoF image restoration from images in which depth information was encoded with their technique of CKM.

In this chapter, we build on the concept of CKM and introduce a snapshot volumetric imaging technique which we name 3D engineered point-spread function microscopy (3D-EPM). We demonstrate how full 3D volumetric reconstruction is enabled by combining PSF-engineered imaging methods with the mathematical framework developed in chapter 2. By acquiring multiple images simultaneously with PSFs that respond differently to defocus, we show how snapshot volumetric imaging can be achieved, while maintaining high resolution in both time and space. We implement this across a range of samples, and demonstrate the benefit of being able to perform snapshot volume imaging by imaging a sample that undergoes continuous motion during acquisition. We begin our discussion by exploring the CKM method in more detail, and determine why the cubic PSF in particular is well-suited to this imaging modality.

3.1 Cubic PSF

For a given PSF to be useful for our purposes, it must extend the DoF while enabling depth information to be encoded into images. To extend the DoF, the MTF must contain no nulls away from the focus, thus enabling more complete information transfer through the imaging system. We saw in chapter 1 how a cubic PSF is formed through the introduction of a phase function $p(x, y) = \alpha(x^3 + y^3)$ at the exit pupil, where x and y are the pupil coordinates. This leads to a distinctly “L” shaped PSF that exhibits a lobe structure, see figure 1.11. In this section, we take a closer look at the OTF of a system with a cubic phase mask at the exit pupil, and examine how it provides an extension to the DoF. In particular, we focus on two key aspects: (i) the insensitivity of the MTF to defocus, and (ii) the lateral translation of the PSF.

3.1.1 The cubic OTF

For simplicity, we constrain the analysis of the cubic OTF to one-dimension. For low numerical-aperture systems, the phase delay introduced by defocus is given by

$$P_\psi(x) = \begin{cases} \exp[i\psi x^2], & |x| \leq 1 \\ \text{Undefined}, & \text{otherwise} \end{cases} \quad (3.1)$$

where ψ is the defocus parameter and the pupil radius is normalised to 1. The calculation of ψ is unimportant to our discussion, however it is worth noting that with no defocus, i.e. when $\psi = 0$, the defocus phase term introduces no phase delay at the pupil. With a cubic phase mask $P_c(x) = \exp i\alpha x^3$ present at the pupil of strength α , the pupil phase becomes the multiplication of these terms:

$$P(x) = \begin{cases} \exp [i\psi x^2] \exp [i\alpha x^3], & |x| \leq 1 \\ \text{Undefined,} & \text{otherwise} \end{cases} \quad (3.2)$$

We saw in chapter 1 how the PSF is calculated as the Fourier transform of the pupil-function, and that the OTF is defined as the Fourier transform of the PSF. Therefore, by auto-correlation theorem, the defocused OTF may be calculated as the auto-correlation of the pupil phase

$$OTF(u, \psi) = \int_{-1+\frac{|u|}{2}}^{1-\frac{|u|}{2}} \exp \left[i\psi \left(x + \frac{u}{2} \right)^2 \right] \exp \left[i\alpha \left(x + \frac{u}{2} \right)^3 \right] \exp \left[-i\psi \left(x - \frac{u}{2} \right)^2 \right] \exp \left[-i\alpha \left(x - \frac{u}{2} \right)^3 \right] dx \quad (3.3)$$

where u denotes spatial frequency. Expanding and simplifying the exponents containing α , this integral becomes:

$$OTF(u, \psi) = \int_{-1+\frac{|u|}{2}}^{1-\frac{|u|}{2}} \exp [i\alpha 3ux^2] \exp [i\alpha \frac{u^3}{4}] \exp \left[i\psi \left(x + \frac{u}{2} \right)^2 \right] \exp \left[-i\psi \left(x - \frac{u}{2} \right)^2 \right] dx \quad (3.4)$$

Repeating this process for the exponents containing ψ terms, and taking the constant term outside the integral, this leaves us with:

$$OTF(u, \psi) = \exp \left[i\alpha \frac{u^3}{4} \right] \int_{-1+\frac{|u|}{2}}^{1-\frac{|u|}{2}} \exp [i\psi 2ux] \exp [i\alpha 3ux^2] dx \quad (3.5)$$

which is a general expression for the OTF of a cubic-encoded imaging system [52].

Insensitivity to misfocus

Dowski and Cathey [52] and others used the stationary-phase approximation, assuming large α , to approximate this integral as:

$$OTF(u, \psi) = \sqrt{\frac{\pi}{12|\alpha u|}} \exp \left[\frac{i\alpha u^3}{4} \right] \exp \left[\frac{-i\psi^2 u}{3\alpha} \right], u \neq 0. \quad (3.6)$$

This description of the OTF contains two phase terms: one that is linear with spatial frequency; the other being cubic. Importantly, the magnitude of the OTF (the MTF)




Figure 3.1: Normalised $z = 0 \mu\text{m}$ cubic PSFs at phase-mask strengths of $\alpha = 3$ (a), $\alpha = 6$ (b) and $\alpha = 9$ (c). (d) shows a plot of peak signal with α , which decreases as the strength of the phase-mask increases.

does not depend on ψ : in this approximation, defocus only affects the phase of the cubic OTF. It is this feature of the MTF that extends the DoF of our imaging system and solves the first of our two challenges.

Defocus-dependent translation

The second phase term in equation 3.6 is a function of defocus and is linear with spatial frequency u . This was observed in [52] to cause a lateral shift of the PSF with defocus. For a separable rectangular aperture, this lateral shift was shown in [89] to equal

$$s(\psi) = \frac{\psi^2}{6\pi\alpha}, \quad (3.7)$$

i.e. the lateral shift exhibits a quadratic relationship with defocus, but may be reduced by increasing the phase-mask strength α . The authors of [52] suggested that the lateral PSF-shift $s(\psi)$ is mitigated by use of a sufficiently high mask strength, and becomes negligible for large α . However, it is not always practicable to use a high- α phase-mask, especially in fluorescence experiments where there is a limited photon budget. This is because with increasing phase-mask strength, the lateral footprint of the PSF becomes larger, meaning the signal is spread across a greater number of pixels. This is demonstrated in figure 3.1a-c, where the $z = 0 \mu\text{m}$ xy slices of a 0.5NA cubic PSF with increasing phase-mask strength are displayed. Figure 3.1d displays a plot of the

peak signal level of the normalised $z = 0 \mu\text{m}$ PSFs as the strength of the phase-mask increases, which at a value of $\alpha = 2$ falls to approximately 15% of the peak signal level of the $\alpha = 0$ (no phase-mask) PSF. It was shown in [55] that the characteristic artefacts in Wiener-restored cubic-encoded images resulted from strong phase modulation on the OTF resulting from the the phase terms in equation 3.6. Additionally, the noise amplification phenomena discussed in section 1.2.5 becomes more severe with increasing phase-mask strength as the signal-to-noise ratio worsens.

Since the translation of the PSF is dependent only on defocus and phase-mask strength (and can be measured experimentally), it could, in principle, encode depth information into images. If, for instance, the lateral coordinates of a point are known, then the coordinates of the PSF in an image could be used to determine the point's axial position. However, generally this isn't possible for two reasons. Firstly, the lateral positions of points are generally unknown. Secondly, from equation 3.7 we see that the depth-dependent translation is symmetric about the focus, i.e $s(\psi) = s(-\psi)$, and therefore, unless the imaging domain is confined to only one side of the focal plane, there is a sign ambiguity in the calculated z coordinate.

Strength of the cubic mask

We see from equation 3.6 that the strength of the phase-mask, α , influences both the phase and magnitude of the PSF. Generally, a stronger phase-mask provides a larger depth-of-field while reducing the defocus-dependent translation. It follows, therefore, that a stronger phase-mask is better for imaging across a larger DoF. However, as we saw in figure 3.1d, the peak signal drops with increasing phase-mask strength as a greater number of lobes become apparent and the lateral extent of the PSF becomes larger. In photon-limited experiments, such as many fluorescence imaging applications, for instance, the choice of phase-mask strength becomes important: there exists a trade-off between signal-to-noise ratio and the extension to the DoF.

3.2 Encoding depth information

The above sections demonstrate how the insensitivity of the cubic OTF to defocus provides an extension to the DoF. While the second of our challenges, encoding depth information into images, could in principle be helped by the defocus-dependent lateral translation of the PSF, generally this alone does not solve the problem. In this section we take a closer look at the technique of CKM, and how imaging instead with two cubic PSFs can more-effectively encode depth information into images.

The general principle of CKM is centered on an imaging system that records two images simultaneously, where the PSF corresponding to each imaging channel responds differently to defocus. If the defocus-dependent response of each imaging channel is known (through measurement of calibration 3D PSF z -stacks), the system effectively

encodes depth information into the 2D images. The authors of [58] demonstrated this principle initially with the cubic PSF, where the defocus-dependent response between the two imaging channels was altered by means of rotating the CPM between cameras. The method is explained by considering the following scenario. If a point source, which is situated at some unknown distance z from the native focal plane, is imaged by such an imaging system, the two images will appear as single cubic PSFs, translated from the true xy coordinates by an amount determined by equation 3.7. If image restoration, performed with a Wiener filter, is then implemented, the point in the the restored images will be translated from its true position if the defocus of the recovery kernel does not match the true defocus of the point in the object. The correct value for defocus, therefore, can be found by varying the defocus of the deconvolution kernels until the translation of the point between the two restored images is minimised.

While the example above considered only a single point-source, the technique is generalizable to more extended structures and surfaces, varying the defocus of the deconvolution kernels until the minimum disparity between restored images is found. The method of CKM then repeats this process on small patches across the entire FoV, yielding a map of locally-averaged defocus of the observed scene. Additionally, the correctly defocused kernels are used to deconvolve each patch, which, when stitched together, leave an artifact-free EDoF image.

While the disparity generated through rotation of the cubic phase mask was used to initially demonstrate CKM, the authors noted that the method was not limited to this implementation. Indeed, CKM may be implemented with *any* pair of point-spread functions that respond differently to defocus, providing the DoF is adequate and the MTF contains no nulls, ensuring maximal information transfer. In the application of Airy-CKM in Ref. [54], the disparity was generated with a differential defocus between the two cubic PSFs, generated by means of a custom prism beamsplitter placed in the imaging path that directed the two images onto the same camera sensor. This method is advantageous over the rotated CPM implementation since it requires only one phase-mask in the imaging path for snapshot imaging. For the majority of the experimental results presented in this chapter, the disparity was generated in a similar fashion by means of differential defocus, however a conventional beam-splitter cube and two separate relay lenses were used to instead focus the images onto two separate cameras, see figure 3.2. In the remainder of this thesis, the overall PSF of such a system is referred to as the differential-Airy PSF. Imaging onto two cameras instead of a single camera is advantageous since it allows full use of the available FoV, allowing larger samples to be imaged in a snapshot.

3.3 PSF phase retrieval

High quality deconvolution performance depends on having good-quality PSF measurement. While several works report that the RL algorithm is robust against minor

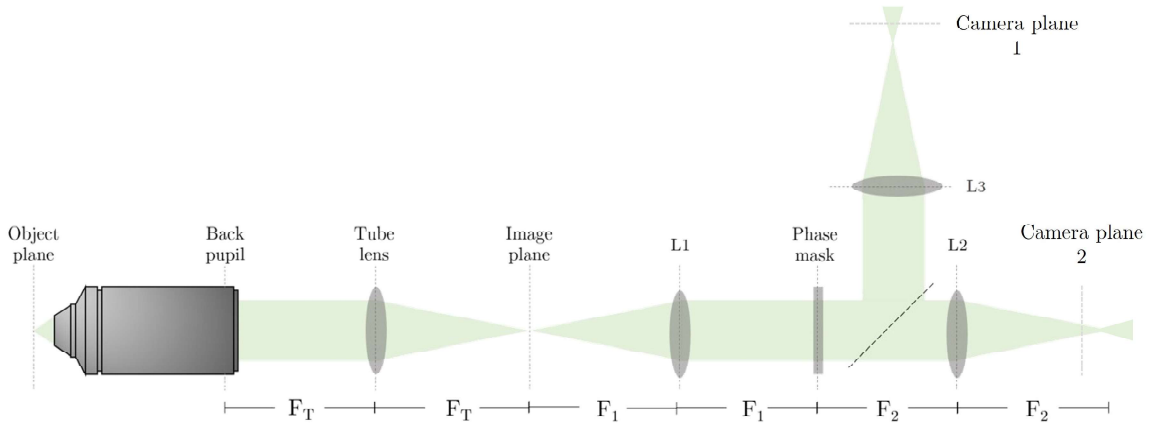


Figure 3.2: Generation of the differential-Airy PSF used for encoding depth information across an extended depth of field. A $4f$ relay is placed behind the native image plane of any microscope, with a cubic phase mask positioned at the re-imaged exit pupil. A 50:50 beam-splitter then directs the light towards two relay lenses which focus an image onto two separate camera sensors. The positions of the two cameras are adjusted with respect to the two final relay lenses (L2 and L3) to introduce a differential defocus between images.

imperfections in the PSF [84], this argument is mostly limited to astronomical imaging where the PSF can be approximated as an Airy disk. For EDoF imaging with, for instance, the cubic PSF, it has been shown that use of an incorrect deconvolution kernel introduces significant artefacts to the restored image [55]. In microscopy, PSF estimates are either obtained from theoretical predictions, i.e. simulation, or from experimental measurement. Generally, experimental PSFs are preferred, since they capture the influence of any aberrations present, that would be missed in simulated PSFs.

Acquiring an experimental estimate to the PSF is normally achieved by imaging a sub-diffraction fluorescent bead. 3D PSFs are acquired by repeating this process as the bead is translated in z . Since the image recorded using such methods is not directly the PSF, rather the convolution of the system PSF with the shape of the bead, a good PSF estimate requires the size of the bead to be beneath the diffraction limit of the microscope, ideally as small as possible. For the highest NA objective lenses, this means the bead must have a diameter beneath 200 nm, which limits the volume of fluorescent dye present and acquired images become corrupted by noise. The presence of noise in a PSF used for deconvolution also significantly inhibits deconvolution performance, introducing artefacts into restored images. There therefore exists a trade-off between the signal-to-noise ratio of the bead image and the quality of the corresponding PSF estimate: a larger bead will yield a better signal-to-noise ratio but provide a poorer estimate to the PSF. The signal-to-noise ratio may be improved through averaging the images of many beads, or through background signal removal with image processing (for instance, a simple threshold to remove background noise), however these methods may introduce non-physical structure into the PSF.

Another issue that must be considered when measuring 3D PSFs is the bead undergoing motion during acquisition. Inspection of large z -stacks of a single bead reveal small translations, caused by Brownian or other random motion, as well as intensity fluctuations between slices, see figure 3.4. This problem is especially noticeable in imaging modalities that require the bead to be suspended in e.g. agarose gel during acquisition, instead of being affixed to a coverslip. This type of artefact in PSF calibration data resulting from motion and intensity fluctuations also inhibits deconvolution performance and therefore should be minimised in the calibration PSF z -stacks.

The alternative to using experimentally-acquired PSFs for deconvolution is to use simulated PSFs, derived from a theoretical prediction of what the PSF should look like based on a particular PSF model. However, theoretical PSFs do not capture any system aberrations that are present, and there is therefore often a large discrepancy between the simulated PSF and an experimentally acquired estimate. Simulated PSFs do, however, benefit from being free of noise, which is advantageous for deconvolution. Therefore, in this work, a combination of experimental and theoretical approaches are employed to obtain a PSF estimate that both captures system aberrations while being free of noise, maximising deconvolution performance. This involves employing phase-retrieval methods to estimate the pupil function, capturing both the added phase-mask as well as system aberrations, before filtering and re-simulating the PSF.

3.3.1 Estimating the pupil function

We saw from equation 1.8 how a the complex-valued PSF (i.e. the amplitude PSF) may be calculated through the 2D Fourier Transform of a complex-valued pupil function. The pupil function describes both the phase and magnitude of the wavefront at the exit pupil produced by a point source in the sample. For a defocused PSF, the pupil function contains a defocus phase term given by equation 1.7 that depends only on the physical aperture. Since the pupil function shares a Fourier transform relationship with the complex-valued amplitude PSF, it follows that that it completely describes the imaging system, including system aberrations and the phase introduced through additional phase-masks. Therefore, if the pupil function for a given imaging system is known, or can be calculated, the system PSF can be simulated. However, calculating the pupil function is not trivial. This is because the PSF that is recorded experimentally is not the complex-valued amplitude PSF, but instead its squared magnitude, the real-valued intensity PSF, meaning phase information is lost.

Hanser et al. demonstrated, however, that the pupil function could be estimated via a modified Gerchberg-Saxton phase retrieval scheme [13, 90]. Originally designed for estimating aberrations in diffraction-limited systems, it is applied here to estimate the phase across the exit pupil of PSF-engineered systems. The Gerchberg-Saxton algorithm is an iterative procedure that estimates the phase from two related planes, e.g. the image and object plane. Hanser *et al* modified the scheme to incorporate the

information from a number of defocused planes to calculate the phase across the exit pupil (within the confines of the aperture), since the defocus phase is a known prior. The modified scheme also allowed the magnitude to vary across the pupil, since it was argued in [13] that the magnitude could not be assumed constant across the pupil. The flowchart in figure 3.3 provides a summary of the phase-retrieval process.

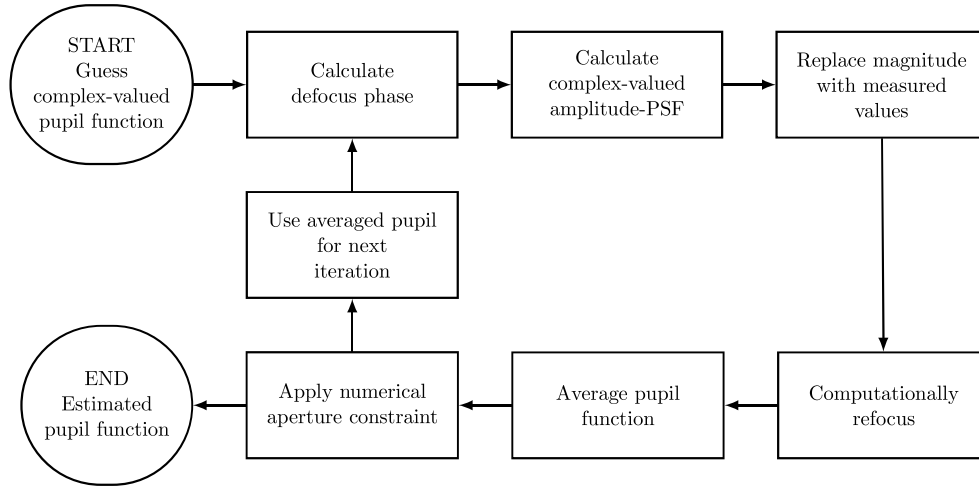


Figure 3.3: Flow chart summarising the phase retrieval algorithm.

The phase retrieval begins with an estimate to both the phase and magnitude of the pupil function that is unity across the pupil aperture and zero elsewhere:

$$P(k_x, k_y)^{n=0} = \begin{cases} 1, & \sqrt{k_x^2 + k_y^2} \leq r \\ 0, & \text{otherwise} \end{cases} \quad (3.8)$$

where k_x and k_y are the pupil coordinates, n is the iteration index and r is the radius of the pupil aperture. Next, the phase associated with the defocus of each slice in the reference PSF stack is calculated:

$$\Phi(k_x, k_y)_z = \exp \left[iz \sqrt{\left(\frac{2\pi n}{\lambda} \right)^2 - (k_x^2 + k_y^2)} \right], \quad (3.9)$$

where z is the distance in front of the native focal plane, n is the refractive index of the immersion medium and λ is the wavelength of the detected fluorescence emission. The term in the exponent in equation (3.9) is the expression for k_z given by equation (1.7), which takes into account the spherical shape of the complex-valued amplitude OTF surface, valid for any aperture [13]. At this point, the 2D Fourier transform of each defocused pupil function is taken:

$$PSF_A(x, y)_z^n = \mathcal{F}\{\Phi(k_x, k_y)_z P(k_x, k_y)^n\}. \quad (3.10)$$

This yields the (complex) amplitude-PSF that our system would exhibit, if our guessed pupil-function was correct. Next, the magnitude of the guessed amplitude PSF is




Figure 3.4: Retrieval of the pupil phase of an experimentally-acquired 0.5NA cubic PSF. (a) and (b) show xx and zy maximum-intensity projections through the raw PSF data. The PSF stack is corrupted by noise, motion artefacts and intensity fluctuations, all of which inhibit deconvolution performance. (c-e) and (f-h) display the recovered pupil phase and magnitude respectively after 1, 5 and 100 iterations of the phase-retrieval procedure. After 100 iterations the pupil phase resembles a cubic phase function, and the magnitude varies very little across the pupil as expected. The display range of each panel is set to the min-max values of each image.

replaced with the measured value:

$$|PSF_A(x, y)_z^n| = \sqrt{PSF_I(x, y)_z} \quad (3.11)$$

where PSF_I denotes the measured intensity-PSF. Note the square root of the intensity-PSF is taken to yield the magnitude of the system's amplitude PSF. This inverse Fourier transform of this amplitude-PSF is then calculated and the defocus removed to leave an estimate to the pupil function for each slice in our PSF:

$$P(x, y)_z^n = \mathcal{F}^{-1}\{PSF_A(x, y)_z^n\}\Phi(k_x, k_y)_{-z} \quad (3.12)$$

the final step of the phase-retrieval iteration is to average the pupil function across the defocus range and constrain it to the numerical aperture of the system, to leave us with a complex-valued 2D estimate to the pupil function:

$$P(k_x, k_y)^{n+1} = P(k_x, k_y)^{n=0} \frac{1}{Z} \sum_z P(k_x, k_y)_z^n \quad (3.13)$$

where Z is the total number of z planes in our calibration PSF stack. Figure 3.4 displays the convergence of the pupil function for an experimentally acquired 0.5NA cubic PSF that suffers from the motion and intensity-fluctuation artefacts detailed in the previous section. When phase-retrieval is applied, the cubic phase across the pupil resulting from the phase-mask is apparent after 100 phase-retrieval iterations.

3.3.2 Low pass filtering in the Zernike basis

While the pupil function can be found to a good approximation with the above method, the solution may suffer minor degradation as a result of noise or imperfections in the raw data, or slight miscalculations of defocus. In particular, it may exhibit a speckle-like appearance that becomes more apparent the further the phase-retrieval iterates. If this degraded pupil function is used directly to re-simulate the PSF, the reformed PSF is also likely to suffer degradation.

To avoid these artefacts, the pupil function may be refined before being used to re-simulate the system PSF. Using a method similar to that presented in [12], in this work the pupil-refinement was performed via a Zernike decomposition, before low-pass filtering in the Zernike basis. Zernike polynomials are orthogonal on the unit disk, and are commonly used to describe aberrations in optical systems at the pupil plane, in terms of deviation from a zero-mean. They are defined in radial coordinates (r, θ) in terms of two non-negative integers $m, n, m \leq n$. There are a number of conventions for ordering Zernike polynomials; the Noll convention defines them in terms of a single index i [91]:

$$Z_i(r, \theta) = \begin{cases} \sqrt{2(n+1)}R_n^m(r)G^m(\theta), & m \neq 0 \\ R_n^0(r), & m = 0 \end{cases} \quad (3.14)$$

with $R_n^m(r)$ and $G^m(\theta)$ being the radial and angular terms respectively defined as:

$$R_n^m(r) = \sum_{s=0}^{\frac{1}{2}(n-m)} \frac{(-1)^s (n-s)!}{s! \binom{n+m}{2}! \binom{n-m}{2}!} r^{n-2s} \quad (3.15)$$

$$G^m(\theta) = \begin{cases} \sin m\theta, & i \text{ odd} \\ \cos m\theta, & i \text{ even} \end{cases} \quad (3.16)$$

The conversion from the Noll index i to the m, n indices is complex and not given in this thesis since the m and n indices are unimportant to our discussion, but a list of the first 36 conversions and corresponding aberration names can be found in [92].

The refinement of the estimated pupil function implemented in this work is performed in 2 distinct stages. Firstly, both the magnitude and the phase of the pupil function are decomposed into the Zernike basis. This is achieved by performing a least squares fit to find the coefficients of the first k Zernike coefficients that minimises

$$|\hat{P}(k_x, k_y) - \sum_{i=1}^k C_i Z_i(r, \theta)| \quad (3.17)$$

where C_i is the coefficient for the i^{th} Zernike polynomial. In this thesis all pupil functions were decomposed into $k = 200$ modes. Once the coefficients have been found,

the phase and magnitude of the pupil-function are regenerated with only the lowest 20 Zernike modes, leaving a smooth pupil function that captures low-order system aberrations, without the high-frequency content left by the phase-retrieval procedure detailed in the previous section. Figure 3.5 displays the Zernike decomposition of the pupil function from figure 3.4, and figure 3.6 illustrates the improvement after re-simulating the PSF.

3.4 Registration

The final pre-processing step we consider before implementing the 3D-EPM volume reconstruction is the registration of image data. When recorded on two separate cameras, as shown in figure 3.2, the differential defocus introduced by the second relay lenses (L2 and L3) can introduce small scale differences between the two images. Additionally, optical misalignment may introduce small translations and/or rotations between images. Such differences between images can introduce significant artefacts into reconstructed volumes, since the deconvolution volume is inconsistent between imaging channels. To mitigate this type of artefact, an affine transform that maps the field recorded by the second camera to that recorded by the first is calculated prior to data acquisition. All subsequent calibration and image data is then registered by warping the data acquired on the second camera according to this affine transform via a third-order flux-conserving affine transform.

The affine transform is calculated from registration images acquired on both cameras prior to data acquisition. The sample used for registration is unimportant, providing the acquired images feature regions of high contrast. For all experimental data shown in this chapter, a sample of fluorescent beads, either embedded in an agarose medium or affixed to a coverslip, was used for registration. However, because of the differential defocus introduced between cameras, it is not possible to calculate from a simple snapshot, since different features will appear in-focus in each image. Instead, a maximum-intensity projection of a z -stack acquired over a depth-range larger than the differential defocus is used, such that the same in-focus features are present in both images. Additionally, the registration data should be acquired without the phase-mask present, so that no depth-dependent lateral translation is introduced that may lead to an incorrect affine transform being calculated.

From the registration data, features are identified and matched between images, from which the affine registration that best describes the matched features is calculated. Because some features may be erroneously matched, the affine transform is calculated via a random sample consensus (RANSAC) fitting approach. Due to the random nature of such a fitting approach, this process is repeated several times, with the affine transform that provides the lowest mean-square error (i.e. closest match) between images after warping being chosen as the correct transform and used to correct all subsequent calibration and image data. The feature-matching, image warping

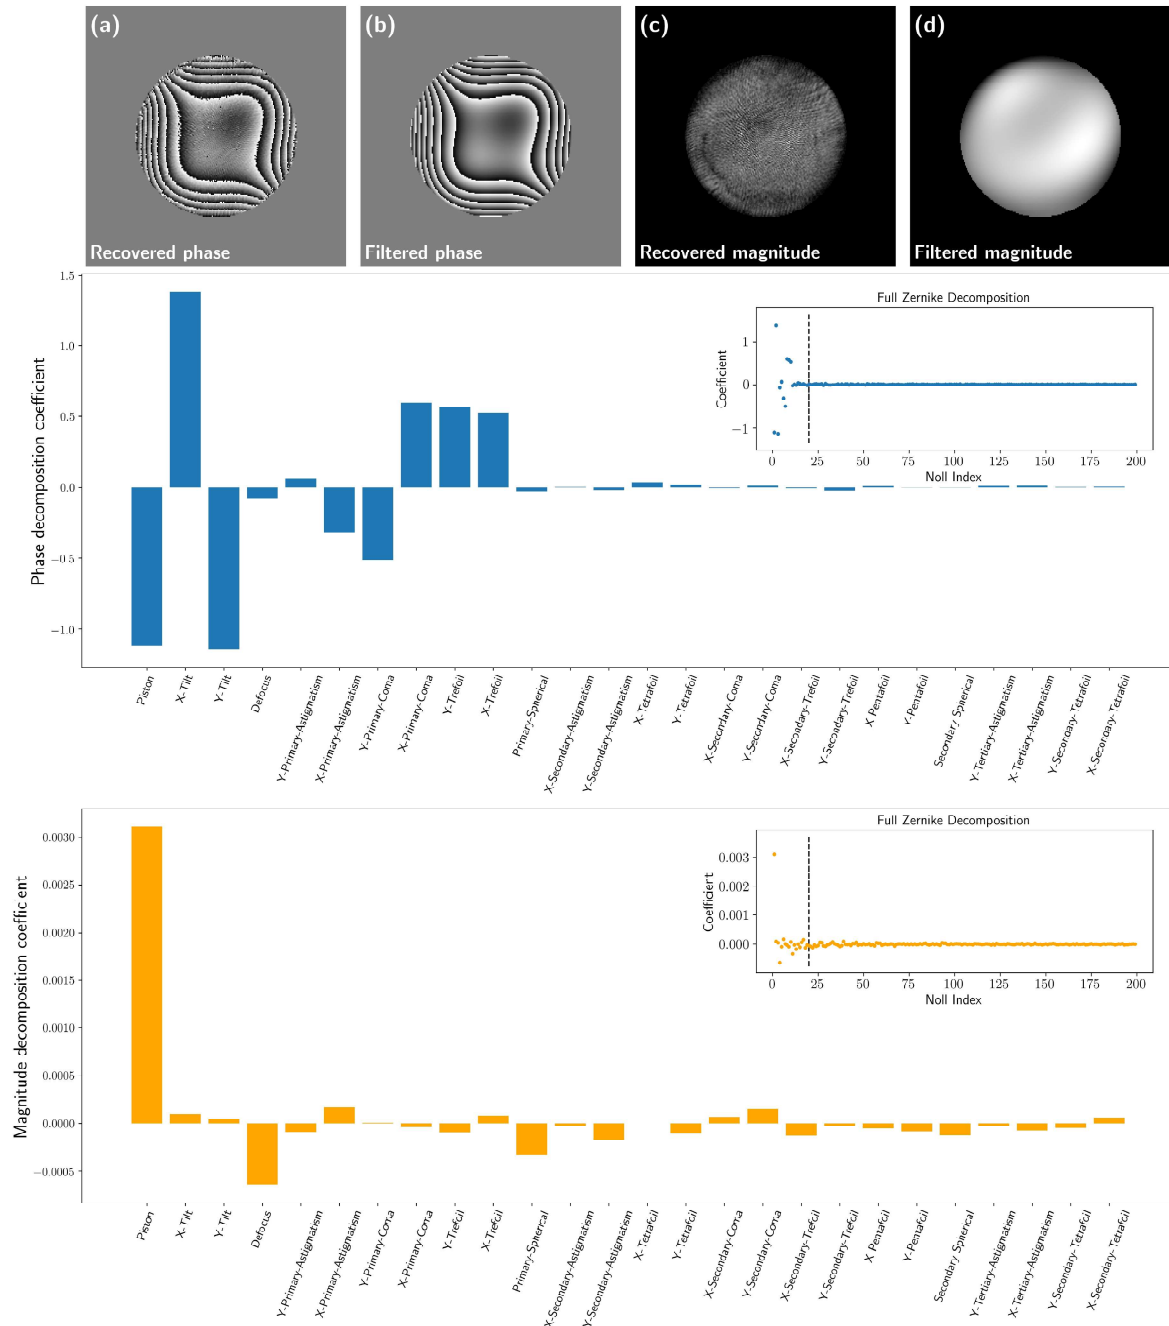


Figure 3.5: Refinement of the recovered pupil function through Zernike decomposition and low-pass filtering. Panels (a) and (b) display the pupil phase, before and after the pupil refinement procedure respectively. Panels (c) and (d) display the same but for the pupil magnitude. Note how (b) and (d) appear cleaner and less grainy than the unrefined results of the phase retrieval in (a) and (c). The blue bar chart shows the first 26 coefficients found by the least-squares Zernike decomposition of the pupil phase. The inset plot shows the full decomposition into 200 Zernike modes. The vertical dashed line displays the cut-off mode, above which the coefficients were ignored when regenerating the pupil function. The orange plot displays the same for the pupil magnitude. Note the difference in scale between the phase and magnitude: the magnitude varies far less across the pupil than the phase does. The display range of each panel in (a)-(d) is set to the min-max values of each image.

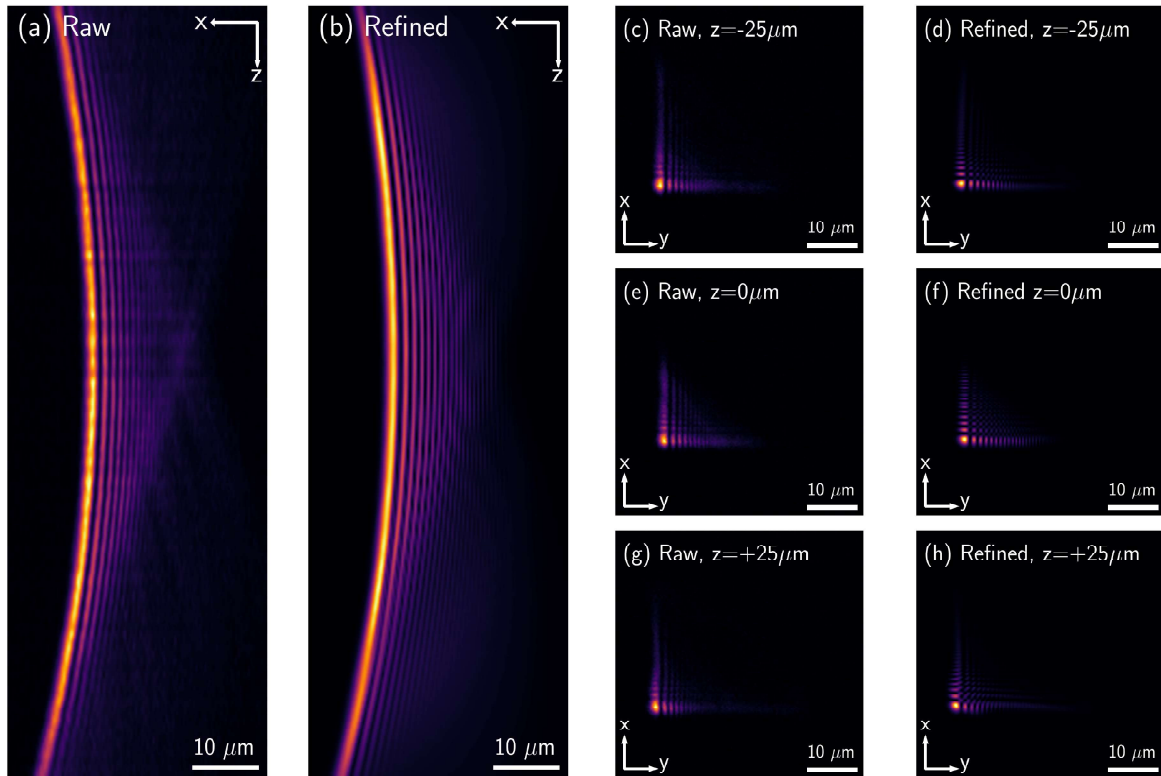


Figure 3.6: Improvement in the 3D PSF through phase-retrieval and Zernike filtering. Panel (a) displays a zx maximum-intensity projection through the raw, unrefined PSF stack. Panel (b) displays the same projection through the refined PSF after phase-retrieval and Zernike filtering. The refined PSF exhibits fewer artefacts and intensity fluctuations, as well as an improved contrast between lobes owing to better signal-to-noise ratio. Panels (c) and (d) display xy slices through the raw and refined PSFs respectively at a depth of $z = -25 \mu\text{m}$. The same slices are displayed in panels (e) and (f) at $z = 0 \mu\text{m}$, and in (g) and (h) at $z = +25 \mu\text{m}$. Note in particular the aberrated appearance of the raw PSF in (g) is also present in (h), demonstrating the refined PSF reflects system aberrations. Scale-bar is $10 \mu\text{m}$.

and random-sample consensus fitting for all data in this thesis was performed using functions from the scikit-image python library [93].

3.5 Snapshot volume imaging with the cubic PSF

The previous sections laid the foundations for performing snapshot volumetric imaging with the cubic PSF. We saw in chapter 2 how the RL deconvolution algorithm could be modified to reconstruct a volume from M snapshot 2D images, providing depth information is present. We saw in section 3.2 how depth information could be encoded with the cubic PSF when imaged onto two detectors with a differential defocus, and how experimentally measured PSFs may be refined for use with the 2D-3D volume reconstruction. In this section, it is demonstrated how these processes are combined to enable volumetric reconstruction from snapshot cubic-encoded images. We start by demonstrating the process with simulated imaging of a single point-source, before demonstrating the technique experimentally.




Figure 3.7: Registration process for 3D-EPM. Features are identified and matched between maximum-intensity projections of z -stacks acquired on each camera. Lines indicate matched features between images. Random sample consensus fitting is then used to calculate the affine transform that best describes the matched features. Plot was generated using the scikit-image python library [93].

Multi-View Reconstruction	
1.	Calculate registration between cameras
2.	Acquire experimental PSF z -stacks
3.	Process and refine PSFs with phase-retrieval and Zernike filtering
4.	Acquire snapshot data
5.	Calculate edge corrections
6.	Dark-field correct
7.	Deconvolve data & reconstruct sample volume

Table 3.1: Data acquisition and calibration pipeline for 3D-EPM.

3.5.1 Simulated point source imaging

Figure 3.8 displays the results of a simulation where a $58 \times 58 \times 58 \mu\text{m}$ (xyz) volume containing a single point-source was imaged via the CKM imaging method, before being reconstructed via the volume reconstruction method developed in the previous chapter. Figures 3.8a and 3.8b display the simulated images via the image formation model in equation (2.11). The imaging system was simulated to feature a 0.5NA objective with an $\alpha = 3$ CPM at the exit pupil. The image detectors were simulated to have a pixel size of $4.5 \mu\text{m}$, corresponding to typical pixel dimensions in a scientific charged-coupled device (CCD) camera, and the system had an overall magnification of 20 X (the object-space pixel size was therefore $0.225 \mu\text{m}$). The detectors were simulated to each detect 25000 signal counts from the point-source, and each image was subject to both Poisson noise and a small Gaussian background (mean 5 photons, standard deviation 2 photons). The detectors were simulated to have a baseline 100 counts per pixel to avoid negative readout, and the differential defocus was $\pm 7.5 \mu\text{m}$. The resulting disparity between the simulated raw camera images in figure 3.8a and 3.8b encodes the depth information necessary to reconstruct the volume. Figures 3.8c-e display maximum-intensity projection (MIP)s through the reconstructed volume along




Figure 3.8: Simulated 0.5NA snapshot volumetric imaging of a single point-source. Panels (a) and (b) display the simulated raw snapshot images. The disparity between images encodes the depth information necessary for volumetric deconvolution. Panels (c-e) display maximum-intensity projections along each axis after 1 reconstruction iteration. The same projections after 50 iterations are shown in (f-h) and 100 iterations in (i-k). Panel (l) shows a trace through the point along x after 100 iterations, and (m) shows a plot along z , where we see full-widths at half-maximum of $0.3\ \mu\text{m}$ and $4.3\ \mu\text{m}$ respectively.

each axis after the first deconvolution iteration. Inspection of the projections provides some insight into the workings of the reconstruction algorithm: the back-projections of the images along the curved 3D trajectory of the PSF-stacks is clearly visible. Where the PSF trajectories intersect is where sample structure is present; in this case the single point source. As the deconvolution iterates, the solution converges back to a single point source. Each iteration the estimate of the object is updated such that the forward projections of the multi-view deconvolution better matches the recorded images. Maximum-intensity projections through the reconstruction after one iteration are displayed in figure 3.8c-e; the same projections are shown after 50 iterations (f-h) and 100 iterations (i-k). After 100 deconvolution iterations the reconstruction has converged on a single point-source. Traces through the 100-iteration solution along x and z are displayed in figures 3.8l and 3.8m respectively, where we see full-width at half-maximum of $0.3\ \mu\text{m}$ in the x direction and $4.3\ \mu\text{m}$ along z . It is emphasised that the quantities quoted here do not directly correspond to the resolution of the reconstruction, however we note that the size of the point in z is significantly larger than that in xy , which is to be expected given the limited collection angle of the imaging system. The resolution of the reconstruction is explored further in section 3.7.1.




Figure 3.9: Experimental demonstration of 3D-EPM: snapshot volumetric imaging of a sample of $0.19\ \mu\text{m}$ diameter fluorescent beads suspended in agarose. Panels (a) and (b) display the raw, snapshot differential-Airy images. Panel (c) displays a zx maximum-intensity projection through the reconstructed volume after 200 deconvolution iterations. Panel (d) displays the same projection as (c) through a ground-truth light-sheet z -stack of the same sample. Panels (e) and (f) display the same as (c) and (d) but with zy projections. Scale-bar is $15\ \mu\text{m}$. Full experimental parameters are detailed in table 3.2.

3.5.2 Experimental validation

Next, we experimentally validate the snapshot volume imaging system. Figure 3.9 displays the results of an imaging experiment where a sample of $0.19\ \mu\text{m}$ diameter fluorescent beads (Bangs Laboratories) were imaged in a snapshot via the same method, using a custom-built microscope based around a 0.8NA water-dipping objective lens. The beads were suspended in an 0.5% agarose medium and mounted in fluorinated ethylene propylene tubing, before being placed in a water-immersion chamber for imaging. A $4f$ relay was placed behind the native image plane, with a beamsplitter placed before the second relay lenses, similar to the scheme in figure 3.2, which focused images with a small differential defocus onto two cameras. The cameras acquired snapshot images simultaneously with an exposure-length of $100\ \text{ms}$. The raw camera images are displayed in figures 3.9a and 3.9b. There is a high background noise evident owing to the thickness of the sample being illuminated. Prior to deconvolution, the calibration steps detailed in this chapter were implemented: an affine registration between the two cameras was calculated, and the PSFs were refined with the phase-retrieval and Zernike decomposition steps.

Figures 3.9c and 3.9e show MIPs through the reconstructed volume after 200 deconvolution iterations, where it is evident that each bead appears as a single point in the reconstructed 3D volume. A ground-truth light-sheet z -stack image of the same sample was then acquired via the same imaging path, with no phase-mask present. Figures 3.9d and 3.9f display zx and zy MIPs through the ground-truth image. The size of the beads appears larger in the light-sheet image than the reconstructed volume: this is because the reconstruction has undergone deconvolution whereas the light-sheet image has not. Nevertheless, the reconstruction is visually comparable to the light-

sheet image (the image fidelity is explored below). The light-sheet image comprises 70 xy -planes, therefore it is reasonable to state that the snapshot volumetric imaging displayed in this example represents a $70\times$ acquisition-time speedup than the conventional light-sheet imaging, although this speedup is greater when the scan-time is also considered.

To explore the similarity between the reconstructed volume and the ground-truth light-sheet image, the position of beads in each image was estimated and compared. The top row of figure 3.10 displays the localisation error of points across the reconstruction depth in x (blue), y (red) and z (black). We see that in z , the localisation error is evenly distributed around 0, suggesting the localisation on average recovers the correct z coordinate in the reconstruction, i.e. no systematic error is observed. The standard deviation of the z localisation error from this analysis was calculated to be $0.15\ \mu\text{m}$. In y , we see a very small systematic offset in the recovered coordinate of $0.08\ \mu\text{m}$. This systematic error is most likely the result of the calibration PSF not being completely centered in the volume: a small offset from the center would result in a shift of the deconvolved volume by an equal amount. This is explained by the sifting property of the convolution operation. In x , we see a trend in the localisation error across the sample depth. The localisation error is negative at one side of the volume depth, and positive towards the other, ranging from approximately $-0.2\ \mu\text{m}$ to $0.1\ \mu\text{m}$. This trend could be explained by a number of phenomena, most likely small optical misalignment: for instance a small lateral translation as well as axial during the PSF acquisition scanning, or a small rotation of the cameras with respect to the image planes. However, both of these possibilities could in principle be identified and corrected for with appropriate system calibration steps.

The bottom row of figure 3.10 shows plots of the recovered position of bead in the reconstruction (x : red, y : blue and z : black), against the position of the corresponding bead in the light-sheet image across the sample depth. The dashed diagonal line on each plot indicates perfect localisation: if a point lies completely on this line then the coordinate in each image matches. In x and y , the points all lie close to the line, however in z there is more of a spread, indicating a larger uncertainty in the measured coordinate. This is generally expected in point localisation experiments owing to the lower sensitivity of the system to changes in z position than changes in x and y .

3.5.3 Extended structure

The results presented in the previous section confirm that volume reconstructed by the snapshot volume imaging reconstruction pipeline closely match the ground-truth volume. However, the results presented so far are not uniquely obtained by this method: the coordinates of point-sources could also be acquired by 3D SMLM methods, albeit without reconstructing the 3D volume. Indeed, this topic is the subject of the next chapter, where we see that the method presented in this thesis does in fact introduce

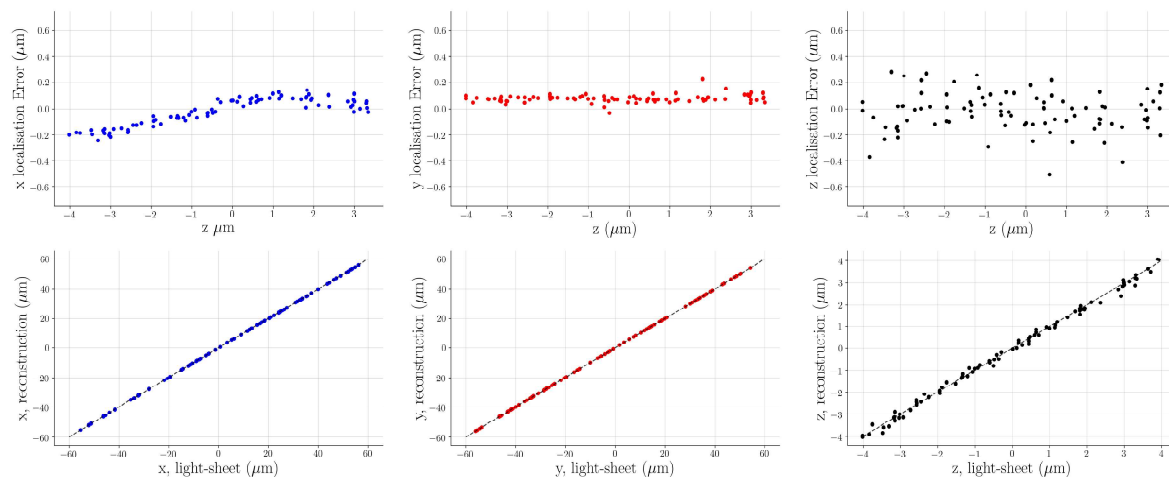


Figure 3.10: Fluorescent bead reconstruction analysis. Top row: plots localisation error in x (blue), y (red) and z (black) throughout the sample depth. Bottom row: plots of the estimated coordinates of each bead in the reconstructed volume against the coordinates of the matched bead in the reconstruction. Colours are the same as the top row.

significant advantages to point-localisation applications. However, many biological samples are not well-suited to SMLM imaging, since they exhibit extended fluorescent structure rather than discrete point-emitters. Therefore, the ability to image samples exhibiting regions of extended fluorescent emission as well as those featuring sparser point sources is crucial for a general microscopy technique.

In SMLM, the coordinates of point-sources are estimated directly from the acquired image data. This requires the points to be spatially and/or temporally separated such that individual PSFs can be identified. Substantial research effort has been directed at increasing the density of point-sources that can be imaged with SMLM, with examples including multi-fitting algorithms and deep-learning approaches. Deconvolution is often used as a tool in the processing of SMLM data to both increase image sparsity and aid localisation. However, SMLM methods fundamentally aim to identify the spatial coordinates of the imaged point-sources and therefore the application of these techniques to more general biological structure, which is often extended in nature, is not possible. Deconvolution, however, does not share this same constraint. The general aim of image deconvolution is not to identify the coordinates of point-sources, rather to improve image contrast through inversion of the optical blurring imposed by the imaging system. The type of structure being imaged, therefore, is less important. This makes the volume reconstruction method developed in this chapter, which is based on deconvolution, well-suited to the imaging of more general fluorescent structure. While there is still a modest requirement for sparsity, which is explored more later in this chapter, the 3D-EPM volume reconstruction may be readily applied to samples exhibiting more extended structure. In this section, the application of the 3D-EPM method to this type of sample is explored.

3.5.4 Simulated filaments

We begin with a simulated example of the application of 3D-EPM to sample exhibiting regions of extended fluorescent structure. Figure 3.11 depicts the results of a simulation where a network of computationally-generated filaments were simulated to be imaged in a snapshot with the differential-Airy PSF. The sample volume was then reconstructed with the volume-reconstruction pipeline. Figure 3.11a displays a 3D rendering of the ground-truth volume, which is colour-coded for depth. The trajectory of each filament was generated via 5th order polynomials with randomly generated coefficients, inside a discretised volume with isotropic voxel dimensions of $0.2\ \mu\text{m}$. The overall volume size was $102.4 \times 102.4 \times 50\ \mu\text{m}$ (xyz). Each filament was created by marking its trajectory (to the nearest voxel) with a value of one in the sample volume, which was otherwise zero. The filaments were given a 3D shape through convolution with a small 3D Gaussian kernel ($\sigma = 0.4\ \mu\text{m}$). Snapshot differential-Airy images were simulated via the image formation model in equation (2.11) with simulated 0.5NA $\alpha = 2$ cubic PSFs with a differential defocus of $\pm 10\ \mu\text{m}$. The images were subject to both a small Gaussian background noise, reflective of background and readout sources. This proof-of-concept simulation was performed with idealised high signal-to-noise imaging conditions: the peak signal-to-noise ratio of the snapshot images was approximately 250:1. The images were also subject to Poisson shot noise to reflect the noise statistics of fluorescence emission. Figs 3.11c-d; e-f, and g-h display zx and zy maximum-intensity projections through the reconstructed volume after 5, 50, and 150 iterations respectively. Due to the high signal-to-noise ratio in the simulated images, no regularisation term was required in the reconstruction. After 150 iterations, the reconstruction has converged: figure 3.11 displays a 3D rendering of the reconstructed volume after 150 iterations, which clearly resembles the ground-truth volume. Similar to the results displayed in figure 3.9, the reconstructed filaments are thicker axially than they are laterally. Again, this reflects the resolution of the imaging system, which is worse axially than laterally, and is determined by the numerical aperture. The resolution provided by 3D-EPM is explored further in section 3.7.1.

3.5.5 Lens-tissue

Next, the application of 3D-EPM to samples exhibiting extended fluorescent structure is demonstrated experimentally. Figure 3.12 displays the results of an experiment where a sample of lens-tissue was imaged and reconstructed via 3D-EPM. The lens-tissue was prepared by dipping in a solution of fluorescein salts before being left to dry, such that the fibrous structure would fluoresce under 488 nm laser illumination. Figures 3.12a and 3.12b display the raw, 2D cubic-encoded snapshot data, which was acquired on a custom-built upright microscope based around a 10X 0.3NA objective. While the sample, at this magnification, is still relatively sparse, it does exhibit extended regions of fluorescence emission. At a NA of 0.3, and with the CPM present, the entire sample

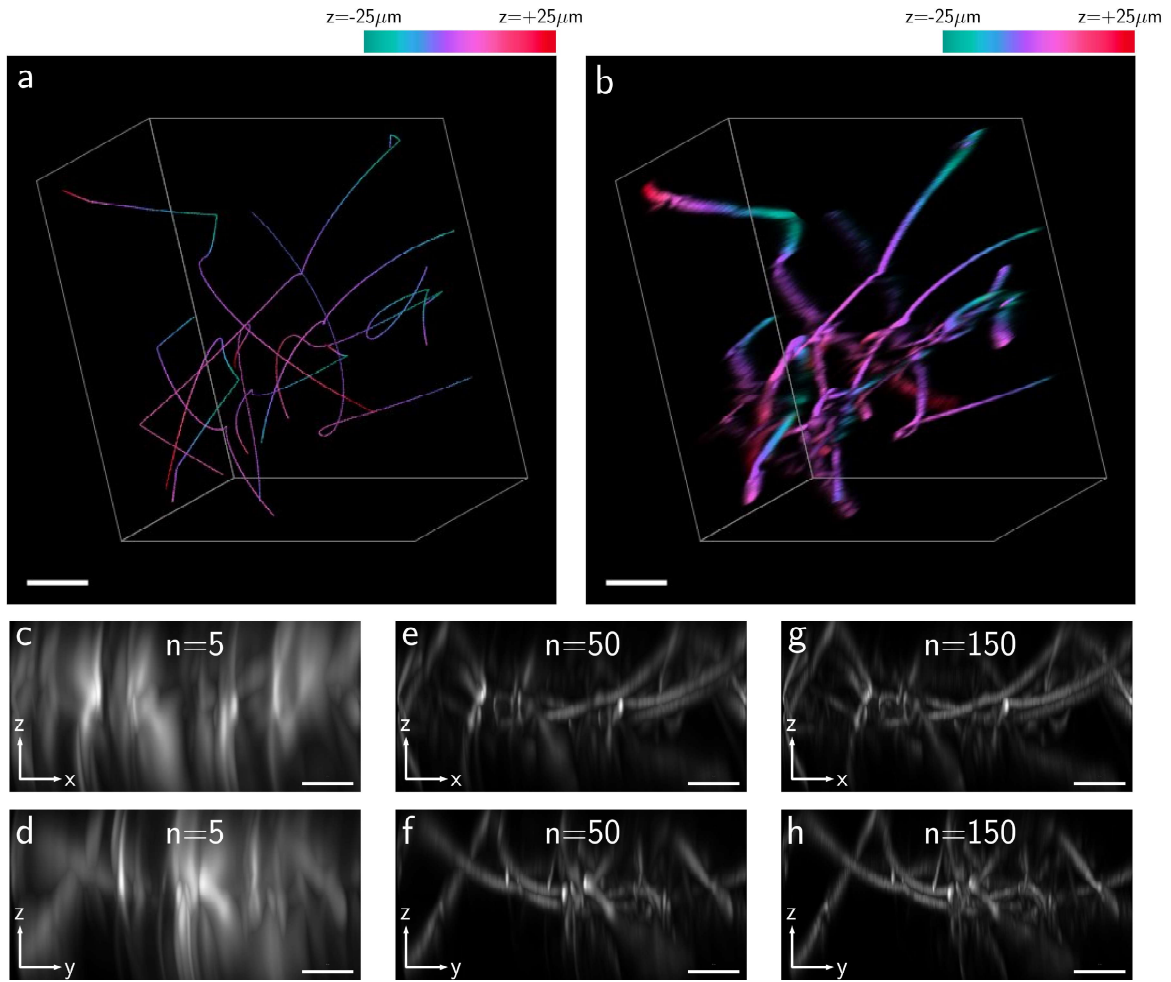


Figure 3.11: Simulated 0.5NA 3D-EPM imaging of a network of filaments inside a $102 \times 102 \times 50 \mu\text{m}$ volume. Panels (a) and (b) display 3D renderings of the ground-truth and reconstructed volumes respectively, colour-coded for depth. Panels (c) and (d) display zx and zy maximum-intensity projections through the reconstruction after 5 iterations respectively, where the solution has clearly not yet converged. The same projections through the solution after 50 iterations are displayed in (e) and (f), and after 150 iterations in (g) and (h). Scale-bar is $15 \mu\text{m}$.

depth could be imaged within the EDoF of the imaging system. The raw images therefore contain little contributions from out-of-focus light, which is beneficial for deconvolution performance. Figures 3.12c-k display the reconstructed volume after 150 deconvolution iterations. To effectively visualise the full dynamic range of the solution, the reconstructed volume was subject to a gamma intensity correction described by

$$I_{out} = I_{in}^{\gamma} \quad (3.18)$$

where I_{out} and I_{in} are the input and output intensity values respectively, and γ is the gamma-correction parameter. In this case the value of γ was selected empirically to be 0.8.

Figure 3.12c displays an xy summed-intensity projection through the reconstructed volume, which has been color-coded to display depth. Here, the 3D nature of the

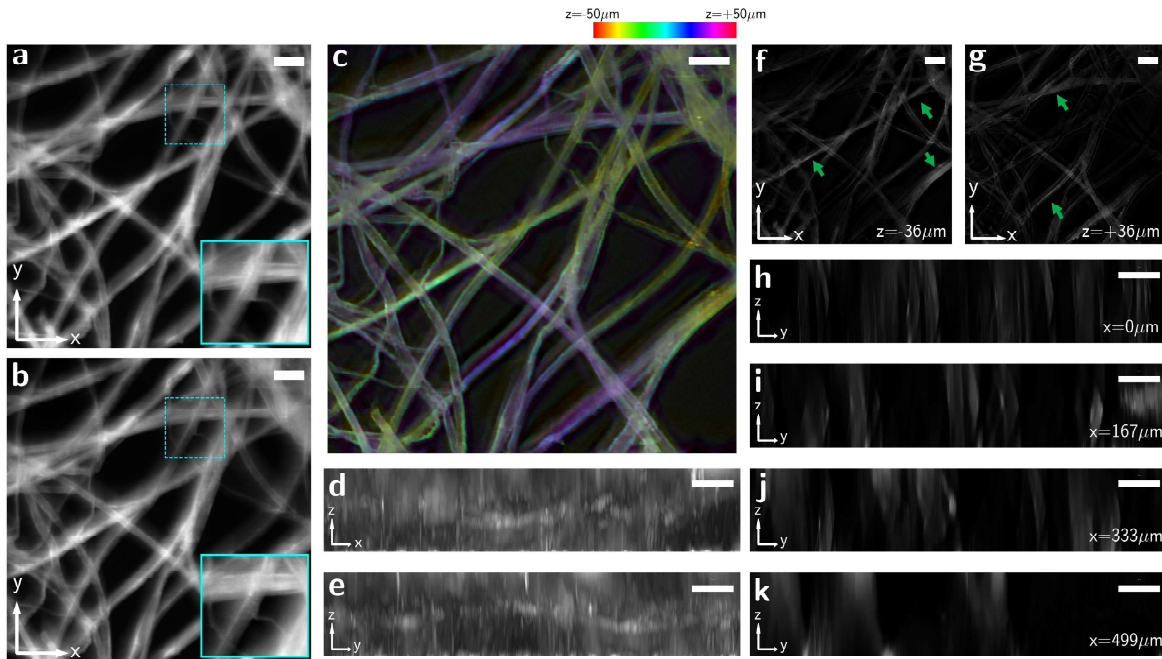


Figure 3.12: Experimental 3D-EPM imaging of a sample of lens-tissue, demonstrating how volumes of extended fluorescent structure are imaged in all three spatial dimensions in a single snapshot. (a) and (b) display the raw cubic-encoded images, acquired simultaneously in a single snapshot. Cyan insets show zoomed in versions of the region bound by the dashed cyan boxes, highlighting the disparity between images. Panel (c) shows a summed-intensity projection through the reconstructed volume, colour-coded for depth. Panels (d) and (e) display zx and zy maximum-intensity projections through the reconstruction respectively. (f) and (g) xy planes in the reconstructed volume at depths of $z = -36 \mu\text{m}$ and $z = 36 \mu\text{m}$ respectively. (h-k) zy planes in the reconstructed volumes at the indicated positions in x . Scale-bar is $50 \mu\text{m}$.

sample volume is evident, with filaments overlapping and running from one side of the volume depth to the other. Max-intensity projections along the other axes through the reconstruction are displayed in figure 3.12d and 3.12e. To further demonstrate the 3D-resolved structure in the reconstructed volume, figures 3.12f and 3.12g display xy planes at depths of $-36 \mu\text{m}$ and $36 \mu\text{m}$ respectively, where different sample features are visible, with prominent examples indicated by the green arrows. Figures 3.12h-k display zy slices through the reconstruction at different locations in x , where the different depths of the various fibres is evident. To regularise the deconvolution, a regularisation parameter of $\lambda_{TV} = 0.001$ was used according to equation 2.19. Full details of the microscope and deconvolution parameters are provided in table 3.2.

3.6 Dynamic Samples

The results presented so far in this chapter demonstrate that the 3D-EPM imaging modality is capable of extracting the depth-information encoded by a differential-Airy PSF by reconstructing a volume of fluorescent-beads, and comparing their estimated positions in the reconstructed volume to those in a light-sheet ground-truth image of

the same volume. We've also demonstrated the capability of the method to reconstruct extended fluorescent structure, as well as sparse point-sources. However, thus far, we haven't demonstrated the benefit of performing volumetric imaging in a snapshot. Indeed, the reconstructed volumes displayed in both the beads and lens-tissue imaging experiments could also have been acquired via the conventional methods of light-sheet or confocal imaging, which naturally would have provided better optical sectioning. To demonstrate the full benefit of snapshot volume imaging, we now turn our attention to samples that undergo motion during acquisition, where scan-based volume-imaging methods could not be applied.

3.6.1 Simulated filament growth

Again, we begin with a simulated example to demonstrate the application of 3D-EPM to this type of sample. The filaments displayed in figure 3.13 were simulated to grow in time throughout all three spatial dimensions. Differential-Airy images were simulated at 60 discrete time-points. The acquired data-set, therefore was three-dimensional (txy). From each pair of differential-Airy images the sample volume was reconstructed with 150 3D-EPM deconvolution iterations. Figure 3.13a-f displays the raw snapshot differential-Airy images at three distinct time-points ($t = 20, t = 35$ and $t = 50$ AU). Figures 3.13g-i displays 3D renderings of the reconstructed volume at the same time-points, which are colour-coded to display depth. The reconstructed data-set is four-dimensional ($tzxy$). The simulated example displayed in figure 3.13 demonstrates the benefit provided by the snapshot imaging capability offered by 3D-EPM: if this simulated motion occurred over sub-second timescales, this sample would not be possible to probe with methods such as confocal or light-sheet microscopy. It is here that the benefits of the imaging method developed in this chapter are most clear: by removing the need to scan during acquisition of the 3D image, the temporal resolution of the system is greatly increased, permitting the study of dynamic samples where the observed motion occurs too quickly to probe with conventional scan-based methods.

3.6.2 Chloroplast motion in *Egeria densa*.

The final result presented in this chapter is an experimental demonstration of the application of 3D-EPM to a biological sample that undergoes motion during acquisition. The 3D the rapid movements of chloroplasts in a sample of *Egeria densa*, a type of pond-weed whose chloroplasts undergo fast motion resulting from cytoplasmic streaming, were imaged via 3D-EPM. In a similar fashion to the other experimental demonstrations presented in this chapter, the sample was imaged via a custom-built upright microscope based around a 0.8NA water-immersion objective lens, with a CPM placed at the re-imaged pupil plane. The sample was prepared by resting a leaf-cutting on a bed of low gelling-point, 1% agarose in a petri-dish, which, when solidified, ensured the sample did not move during imaging. The immobilised sample was then

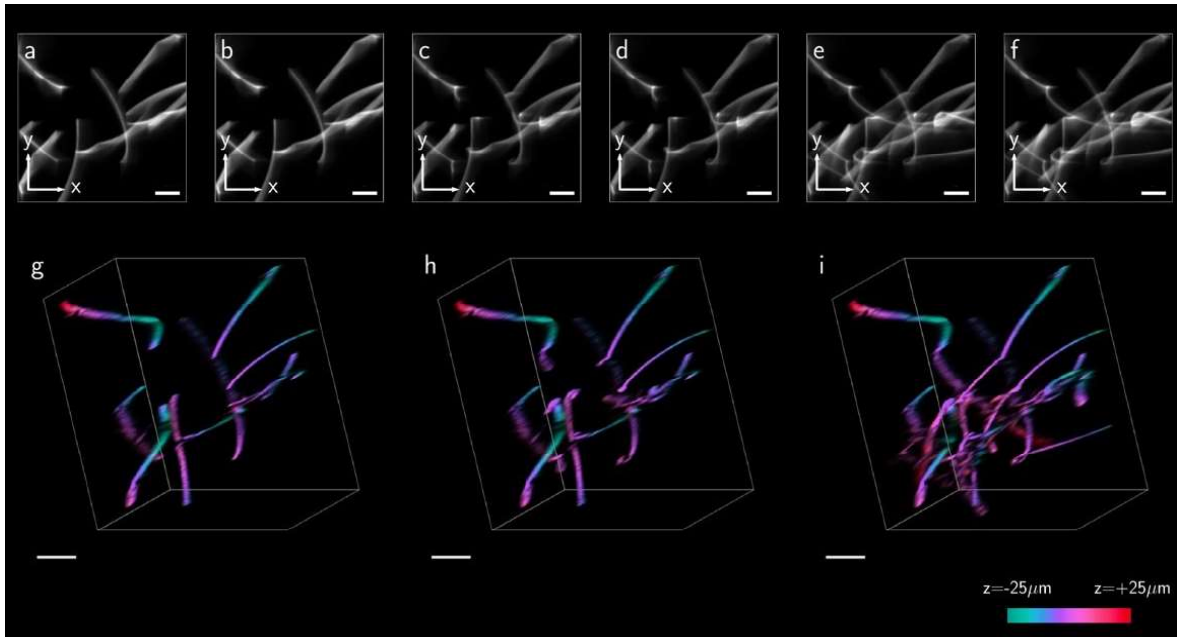


Figure 3.13: Simulated 3D-EPM imaging of filament growth. Panels (a) and (b) display the simulated snapshot images at $t=20$ AU. The raw data at $t=35$ AU and $t=50$ AU is displayed in (c,d) and (e,f) respectively. Panels (g-i) display 3D renderings of the reconstructed volumes at the same time-points, colour-coded for depth. Scale-bar is $15 \mu\text{m}$.

covered in water. Auto-fluorescence emission from the chlorophyll in the chloroplasts was excited via illumination with a 488 nm laser illumination and detected via a 630 nm narrow-band fluorescence emission filter in the imaging path. The sample was imaged at a frequency of 5 Hz over a period of 80 s. Figure 3.14 displays a 3D visualisation of the reconstruction of a single time-point of the 4D-timelapse, as well as motion tracks of the chloroplasts throughout the 80 s acquisition, estimated using the TrackMate ImageJ plugin [94].

3.7 Performance and limitations

The previous sections demonstrated the application of 3D-EPM to a variety of different samples. In this section, we further explore the performance of the imaging method, focusing on the achievable resolution, factors affecting the convergence of the deconvolution and the effect of signal-to-noise ratio and regularisation. We also address the limitations of the method and explore where it might fall down.

3.7.1 Resolution

The goals set out at the start of this chapter were to develop an imaging method that could image a volume in a single snapshot, while maintaining high spatial resolution. The first of these goals has been addressed with the development of 3D-EPM, which we have seen to be capable of imaging both point-like and extended structure in a single

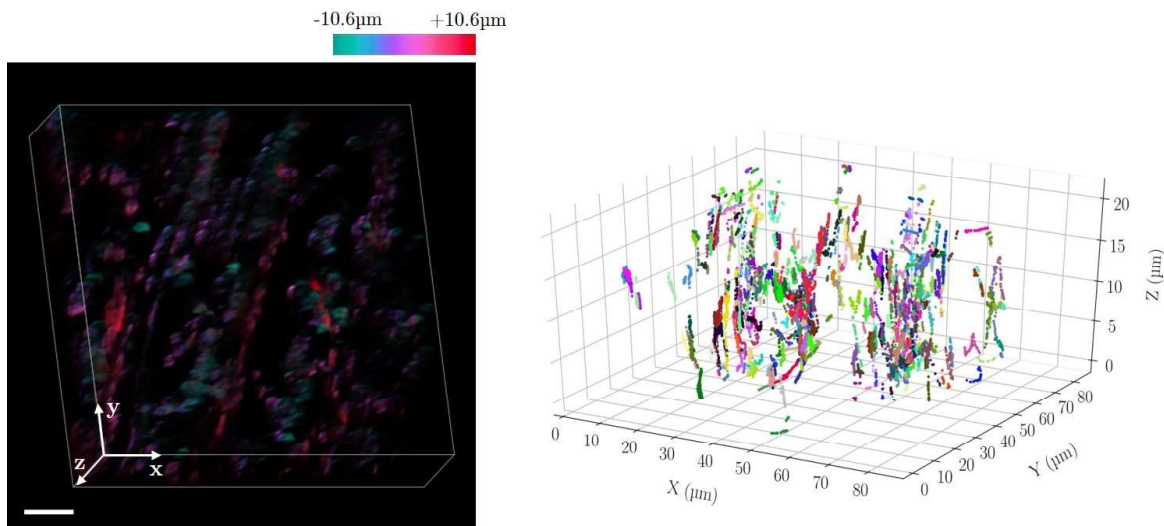


Figure 3.14: 3D-EPM imaging of chloroplasts undergoing cytoplasmic streaming in a sample of *Egeria Densa*. Snapshot imaging was performed at a frequency of 5 Hz over a duration of 80 s. Left: a 3D rendering of the $t = 0$ s reconstructed volume, colour-coded for depth. Right: estimated motion tracks of chloroplasts over the 80 s.

snapshot. However, the second of these aims, maintaining high spatial resolution, has not yet been addressed. We saw in figure 3.9 that traces through the reconstructed points had full-widths at half-maximums of $0.3 \mu\text{m}$ laterally and $4.3 \mu\text{m}$ axially. However, it is important to note that these quantities do not directly correspond to the resolution of the reconstructed image, or indeed the resolution provided by the imaging system, in terms of its ability to distinguish closely-spaced points. This is because RL deconvolution will, if left to iterate, continue to shrink the appearance of a single point; increasing the intensity of the central pixel(s) while decreasing the intensity of the neighbouring pixels. This behaviour is exemplified in figure 3.15, where the size of a single reconstructed point is shown to decrease as the solution is left to iterate, and explained by the noise amplification phenomena discussed in chapter 2. Therefore, the size of a single point in a deconvolved image tells us almost nothing about the resolution of the imaging system, which is the closest distance two points can be spaced while still being identifiable as two distinct points. The Rayleigh criterion requires two such points to be separated such that the first diffraction minimum of one point coincides with the first diffraction maximum of the other: separations beneath this distance are not resolved¹. Therefore, to determine the resolution of a snapshot volumetric imaging system, the separation between two points must be increased until they are resolved in the reconstructed volume. In 2D, resolution test-targets make this process trivial, however in 3D fluorescence imaging it is more challenging, especially when the method benefits from sample sparsity: to my knowledge no suitable 3D fluorescent resolution test target exists for such a method. Therefore, to assess the resolution offered by 3D-EPM, this process of separating two points was performed in simulation.

¹according to the Rayleigh Criterion. Other definitions of resolution exist, however Rayleigh is the most common.

Figure 3.16 displays the results of a simulation where the resolution of a reconstructed volume was explored by increasing the distance between two points in the ground-truth image until they were clearly resolved in the reconstruction. The simulation was performed at a NA of 0.5 and with an emission wavelength of $\lambda=525$ nm. Figure 3.16a shows the results as the lateral separation is increased, where it is clear that at a separation distance of $0.5\ \mu\text{m}$, the points are not resolved in the reconstructed volume, however at a separation of $0.6\ \mu\text{m}$ the points are clearly resolved as two distinct points. Therefore, the lateral resolution of the reconstructed volume falls between $0.5\ \mu\text{m}$ and $0.6\ \mu\text{m}$. The theoretical lateral resolution limit R_{xy} of a 0.5NA system at this wavelength, calculated by (1.6) is R_{xy} is $0.525\ \mu\text{m}$. It is apparent, therefore, that when resolving two points, the lateral resolution of the reconstructed volume approaches that of the parent imaging system. Figure 3.16b shows the results of the same simulation, but where the points were separated axially instead of laterally. Here, we require a separation of approximately $4.5\ \mu\text{m}$ to resolve the points. The Abbe resolution limit for the axial direction, R_z is given by

$$R_z = \frac{2\lambda}{NA^2} \quad (3.19)$$

which, in this simulation, gives a value of $R_z = 4.2\ \mu\text{m}$.

The fact that 3D-EPM can approach the resolution of the parent imaging system represents a major advantage over LFM. LFM also enables snapshot volumetric imaging, however the volumetric imaging capabilities are provided at the expense of a significant loss of lateral resolution. To date, the highest-resolution implementation of conventional LFM reported the ability to resolve structure separated by 757 nm, 711 nm and 790 nm in xyz respectively [46]. However, these results were obtained with an emission wavelength of 680 nm on a system with a primary NA of 1.45 , which according to the Abbe resolution limit is capable of achieving a lateral resolution of 230 nm and an axial resolution of 650 nm, representing a reduction in lateral resolution by a factor of approximately three compared with the parent microscope. With 3D-EPM, because the achievable resolution is similar to that of the parent microscope, sub-micron resolution is available at far lower NA, representing a far more versatile high-resolution imaging technique.

3.7.2 Sample sparsity

It was mentioned in section 3.5.3 that while 3D-EPM is able to reconstruct samples exhibiting extended fluorescent structure, there remains a modest requirement for sparsity in the sample for the reconstruction to be effective. This raises the question: what is the upper limit on how dense a sample can be before the 3D-EPM reconstruction no longer works? This question is somewhat answered through consideration of how depth information becomes encoded into the snapshot differential-Airy images. The differential defocus between imaging channels manifests as disparity between images.

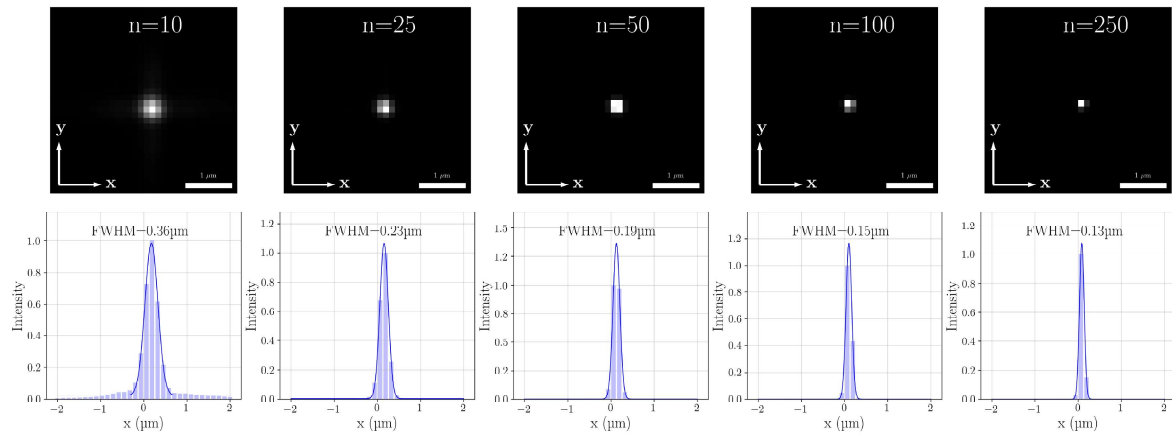


Figure 3.15: Convergence of a single point-source with increasing iterations. Images show normalised xy planes through a simulated reconstruction of a single point-source with increasing iterations. Plots show $y=0 \mu\text{m}$ traces through the images, which show the full-width at half-maximum continues to decrease while the intensity of the central pixel increases as the solution iterates. This does not necessarily correspond to an improvement in resolution of the reconstructed image: it is instead related to the overfitting of high-frequency components in the image data.

It follows, therefore, that if a certain sample appears the same in both imaging channels, the reconstruction will be unable to extract depth information and is unlikely to work effectively. An example of such a sample includes one that features large areas of smooth, uniform fluorescence emission that are larger than the cubic PSF. Imaging such a sample with a differential-Airy PSF will yield images with large regions that appear broadly similar. This lack of disparity causes the 3D-EPM reconstruction to break down.

The challenges associated with a lack of disparity in projection images are not unique to 3D-EPM: instead it is a general problem associated with the calculation of 3D information from 2D images. In the differential-Airy implementation of 3D-EPM, as with CKM, depth information is encoded through image disparity, related to the lateral translation of the cubic PSF. In LFM, on the other hand, disparity is generated from the varying angular perspectives provided by the MLA. Similarly, we shall see in chapter 5 that 3D reconstruction from multi-angle projection data is made possible by exploiting the information from projecting the sample at different angles. Because each of these methods exploit image disparity to calculate depth, they all share the same limitations with respect to sample sparsity, with none being able to accurately reconstruct extremely dense fluorescent structures, or samples with large regions of uniform fluorescence emission. The acceptable density of the sample could, in principle, be increased by either acquiring a greater number of projection images, or by artificially introducing structure onto the sample with e.g. structured laser illumination. The effect of reconstructing from a greater number of projection images is explored quantitatively in chapter 5 in the context of multi-angle projection imaging [5]. The concept of using structured illumination to improve the versatility of 3D-EPM is discussed at the end of this thesis in the context of possible future work.

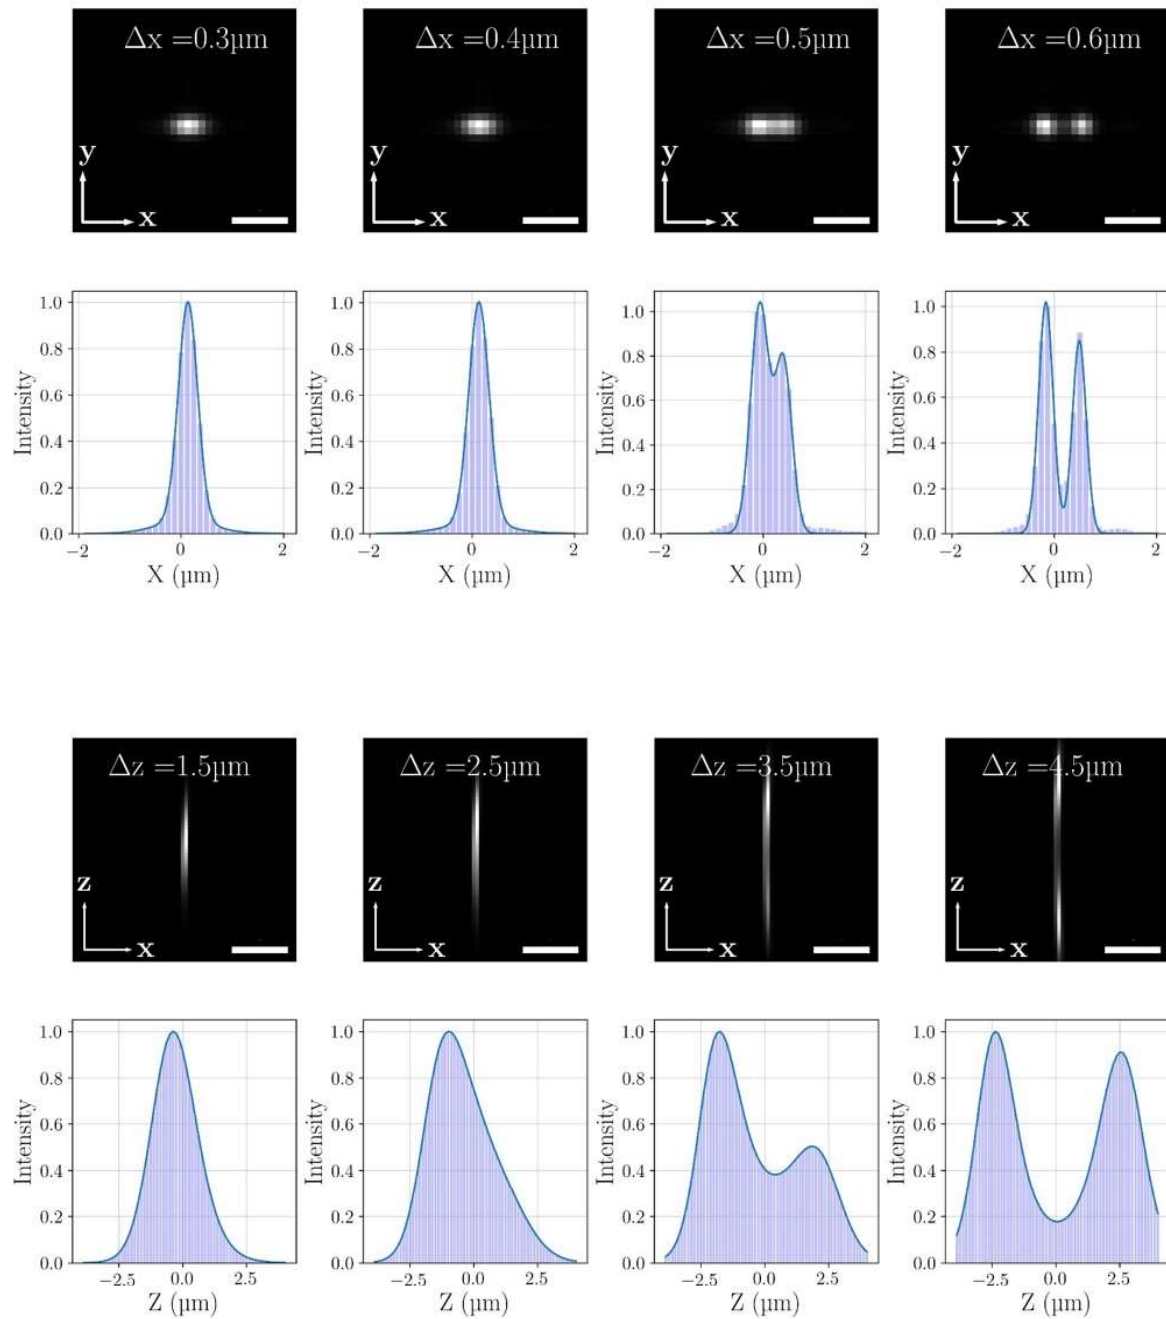


Figure 3.16: Simulated resolution test of 3D-EPM reconstruction. Images show $z=0$ (top) and $y=0$ (bottom) slices through a reconstructed volume containing two point sources separated laterally (top) and axially (bottom).




Figure 3.17: Plot of convergence with varying differential defocus.

3.7.3 The influence of differential defocus

We have seen that the differential defocus between the two cubic PSFs is central to the imaging method presented in this chapter, through its introduction of disparity into the differential-Airy images. However, thus far, there has been no theoretical guidance provided as to what the magnitude of the difference should be for optimal 3D reconstruction. Figure 3.17 shows the results of a simulation where the mean-square error between the solution and ground-truth was measured with increasing iterations, at different defocus separations between the two 0.5NA cubic PSFs. The ground-truth was a volume densely seeded with point-sources, and the cubic PSFs were simulated such that the overall DoF was approximately $40\mu\text{m}$. The plot demonstrates that in this particular imaging scenario, which was simulated to be free of noise, a differential defocus $\delta z = \pm 8\mu\text{m}$ yielded the best solution after 50 deconvolution iterations. At much larger values of δz , the solution stopped converging after very few iterations and quickly began to diverge. This behaviour is explained by considering the scenario where a differential defocus is used that is too large, leading to sample features that are imaged within the DoF of one of the imaging channels being positioned outwith the DoF of the other channel, and therefore different features being imaged by each camera. The deconvolution, therefore, is unable to effectively reconstruct the sample volume because depth information has not been effectively encoded into images. Conversely, at defocus separations that are too small, minimal disparity between images is introduced, again not effectively encoding depth. In practice we find a differential defocus of approximately 25% of the EDoF to be optimal for encoding depth into differential-Airy images: this value is sufficiently large such that enough substantial disparity between images is introduced, but not so large that the two cameras image entirely different volumes.

3.7.4 Convergence and regularisation

The number of iterations required to reconstruct the sample volume depends on a number of factors including the signal-to-noise ratio of the raw data and the type of sample

being imaged. Denser samples generally require a greater number of iterations before convergence is reached. However, the deconvolution should be stopped before the noise amplification discussed in chapter 2, caused by overfitting to high frequencies dominated by noise, starts to degrade the solution. Since the solution generally converges faster laterally than it does axially, the total variation regularisation term described in chapter 2 is used to suppress the noise amplification, allowing the deconvolution to iterate further. Generally, for experimental data we find the optimal number of iterations to lie in the range 50-500. The number of iterations used for all of the experimental results shown in this chapter, as well as details of all other experimental parameters including the strength of the total-variation regularisation, are listed in table 3.2

3.8 Chapter 3 summary

In this chapter we developed 3D-EPM, a microscopy method that is able to image entire volumes in a single snapshot at high resolution. We saw that when the imaging method of CKM is used to simultaneously extend the DoF and encode depth information into images, the volume-reconstruction pipeline developed in chapter 2 may be applied to the snapshot-acquired data to reconstruct a full 3D volume. We validated the method by comparing a reconstructed volume to a light-sheet ground-truth image of the same sample, where we saw good agreement. We then demonstrated the application of 3D-EPM to samples that were extended in nature and therefore could not be imaged with SMLM methods, as well as dynamic samples that could not be imaged with conventional scan-based volumetric imaging techniques. Finally, we explored the performance and limitations of 3D-EPM, where we demonstrated that the native resolution of the microscope in both time and space is preserved under ideal imaging conditions, offering significant advantages over LFM, which requires a significant compromise in lateral resolution. In summary, 3D-EPM is a general microscopy technique that can be performed on any widefield microscope through simple modification of the imaging path, which enables volumetric imaging at a temporal resolution far greater than scan-based methods, while maintaining high spatial resolution.

	Parameter	Beads	Experiment Lens Tissue	<i>E. densa</i>
Microscope	Objective	Nikon CFI Apo NIR 40X W	Nikon CFI Plan Fluor 10X	Nikon CFI Apo NIR 40X W
	Tube Lens focal length (mm)	150	100	200
	L1 focal length (mm)	100	75	100
	L2 focal length (mm)	100	150	100
	Emission Filter peak (nm)	525	525	630
	Excitation wavelength	488	488	488
	Excitation Power (mW)	50	20	50
	Camera	Ximea USB 3.0 MD028MU-SY CCD	Andor Zyla sCMOS	Andor Zyla sCMOS
PSF Measurement	Bead	Bangs- Laboratories, 0.19 μm Dragon- Green, Volume- Labelled	Bangs- Laboratories, 0.4 μm Dragon- Green, Volume- Labelled	Mesolight Quantum Dots FWHM 26 nm Peak emission 623 nm
	Exposure Time (ms)	200	400	400
	Z-interval (nm)	150	1000	200
	Num Zernike Modes	20	20	20
Snapshot Acquisition	Exposure Time (ms)	100	50	50
	Frame Rate (Hz)	-	-	5
Deconvolution	Iterations	200	500	300
	λ_{TV}	0	0.001	0.0005

Table 3.2: Acquisition parameters for each experiment detailed in this chapter.

Chapter 4

3D-EPM localisation

In the previous chapter 3D-EPM was introduced, where the volume reconstruction framework developed in chapter 2 was applied to PSF-engineered images. This enabled the snapshot volumetric imaging of both point-like and extended fluorescent structures, which we saw to be beneficial in the 3D imaging of dynamic samples. In this chapter, we turn our attention away from extended fluorescent structures, instead focusing on the benefits that the volume-reconstruction pipeline can introduce to point-localisation applications. Additionally, where up to now our discussion has focused on implementing the volume-reconstruction from cubic-encoded images, in this chapter we demonstrate its implementation with other PSFs. Indeed, we shall see that the method is applicable to any engineered PSF, providing certain key criteria are met. We explore through simulation the application of 3D-EPM to point-localisation and demonstrate significant advantages over similar localisation frameworks. We begin by discussing the concepts of point-localisation in microscopy, before setting out how 3D-EPM localisation works. We then assess its performance and compare it to a previously-developed point-localisation routine designed for the twin-Airy PSF [68].

4.1 Point localisation in microscopy

Accurately and precisely determining the spatial coordinates of point-like emitters underpins several different microscopy methods across different spatial scales. At the organ to organism scale, example applications include mapping cardiovascular blood-flow in the embryonic zebrafish [54, 95], mapping zebrafish heart-wall trajectories [96] as well as blood-plasma velocity measurement in the embryonic avian heart [97]. Towards the cellular level, traction-force microscopy exploits point localisation to enable the study of the traction, or force, at the interface between a cell and the surrounding medium [98, 99], providing insight into cellular mechanics in both two and three dimensions [100]. Perhaps the most powerful application of point localisation in microscopy, however, is seen at the sub-cellular scale, where SMLM methods permit the nano-scale imaging of fluorescent structure through the localisation of individual fluorescent

molecules. In SMLM, the spatial coordinates of individual fluorescent molecules are determined from sparse diffraction-limited images. The localisations of many individual molecules are then combined into a single reconstructed image, with a resolution surpassing that dictated by the optical diffraction limit [60, 61, 62]. To enable effective localisation from PSFs, SMLM requires either temporal or spatial separation between neighbouring points in images to ensure PSFs do not significantly overlap. In fluorescence microscopy, this is commonly achieved by exploiting the concepts of photo-switching or photo-activation. Photoswitchable fluorophores stochastically switch between emitting (on) and non-emitting (off) states under irradiation at an appropriate excitation wavelength [62]. Photoactivatable fluorophores, on the other hand, irreversibly change from a non-emitting state to an emitting state, and therefore must be bleached prior to new fluorophores changing to an “on” state to ensure temporal separation between images of neighbouring fluorophores being acquired [60, 61]. In both cases, the lifetime of emitting fluorophores and number of fluorophores active simultaneously may be controlled through choice of appropriate dyes and/or modulation of the chemical/irradiation conditions [59].

Individual fluorophores are imaged as individual PSFs, with raw SMLM images featuring a resolution determined by the diffraction limit, related to the size of the PSF. By estimating the positions of the underlying molecules, however, single molecules may be localised with a precision far better than this limit, thus SMLM methods are said to achieve super-resolution. Position estimation may be performed in a number of ways: for instance in 2D-SMLM, the position of fluorophores may be estimated through least-squares fitting of 2D Gaussian profiles to PSFs identified in images [60, 61, 62]. Alternatively, maximum-likelihood estimation-based approaches may be used to determine the position of a point-emitter [101]. The resolution of SMLM then becomes associated with the uncertainty of the measurement, or the localisation precision, rather than being solely determined by the optical diffraction limit. Quantifying the resolution achieved by SMLM is not a trivial task: there exists a theoretical upper limit to the achievable localisation precision, determined by the Cramer-Rao lower bound (CRLB), discussed further in section 4.5. However, in practice, experiments may not achieve this upper limit for several reasons including optical aberrations or imperfect localisation routines. Estimating the resolution from super-resolved images directly, therefore, is a more informative approach: example methods include measuring the uncertainty associated with repeated position measurements of a single fluorophore, or through calculation of e.g. the Fourier ring correlation [102]. Alternatively, imaging nano-ruler samples, which feature fluorophores spaced at regular intervals, offer direct measurement of image resolution, albeit from a different sample [103]. Despite these complexities, lateral resolutions of a few tens of nanometers are possible with SMLM. For this reason, SMLM has become a pivotal tool in the life sciences through its ability to resolve structure that otherwise is too small to resolve with conventional microscopy methods, with notable applications including the mapping of nanoscale protein archi-

texture in focal adhesions [104], imaging chromatin organisation and conformity [105], and discovery of periodic ringlike actin structures around the circumference of axons [106].

4.2 3D-SMLM

While super-resolution is easily achieved in two dimensions through SMLM, obtaining images that are super-resolved in all three spatial dimensions is more challenging because of the requirement for additional optical components and more sophisticated image processing. This is problematic, since most biological structures and processes of interest are inherently three dimensional. SMLM methods generally require imaging with extremely high NA objective lenses for optimal light collection and image resolution. However, a compromise that must be made when imaging with high-NA lenses is that the associated DoF is constrained to just a few hundred nanometers. Additionally, as we saw in the previous chapter, the shape of the diffraction limited PSF is approximately symmetric about the focal plane, meaning that without combining SMLM with other optical techniques, the axial coordinate of fluorophores is not possible to determine from SMLM image data.

Several methods exist that extend the super-resolution capability of SMLM into all three spatial dimensions. For instance, bi-plane fPALM combines bi-plane imaging with photo-activated fluorophore imaging to achieve sub-100 nm resolution in each spatial dimension [107]. Here, the 3D spatial coordinates of fluorophores are estimated by fitting an experimentally-measured 3D PSF to PSFs in the 3D image data (formed by axially stacking the two imaged planes). Similar methods include biplane spectroscopic SMLM [108] and multi-plane spectrally-resolved SMLM [109]. Alternatively, astigmatic imaging may be combined with point-localisation through the introduction of a weak cylindrical lens into the imaging path. This results in a depth-dependent change in the ellipticity of the PSF owing to a different axial focus in x and y , allowing 3D particle tracking [110, 111] and, when combined with STORM imaging, three dimensional super-resolved imaging [63]. Here, the position of fluorophores may be estimated again through fitting a 2D Gaussian profile to PSFs in images, however unlike in conventional SMLM, the axial coordinate may be estimated through the difference between the Gaussian widths in x and y [63].

A limitation of these methods, however, is the limited depth-range over which fluorophores can be localised. Without scanning the sample axially, both methods provide 3D super-resolution over a depth-range of less than approximately 1 μm , which is insufficient when imaging typical 3D samples such as mammalian cells. To image beyond this depth range, the DoF of the imaging system must be extended. While this has been demonstrated with, for instance LFM [112], PSF engineering is the most common strategy for 3D SMLM across an EDoF. Several pupil functions have been developed for these purposes including the double-Helix [64], Corkscrew [113], Tetrapod

[65], and twin-Airy [68]. PSF engineering methods extend the depth over which 3D-SMLM may be performed: for instance, a precision of 30 nm over a depth range of 7 μm was reported in [68].

3D position estimation from PSF engineered images may be performed in a number of ways, e.g. maximum-likelihood estimation [65], custom algorithms exploiting particular features of given PSFs, for instance the rotation or separation of PSF lobes [64, 68], and neural network or deep-learning based methods [114, 115]. However, similar to 2D-SMLM, most methods require either temporal or spatial separation between images of neighbouring points for effective localisation. Since the lateral footprint of engineered PSFs is typically far larger than their diffraction-limited equivalents, the constraints on emitter density are often far stricter in 3D-SMLM. While various methods have been reported to enable the localisation from overlapping PSFs, they are each restricted for use with only a single type of engineered PSF [114, 68]. There is, therefore, a requirement for a general localisation method that is not limited to a single type of engineered PSF, that can cope with overlapping emitters. Such a method would offer both versatility in terms of PSF choice and an improvement in achievable temporal resolution for 3D-SMLM methods by allowing samples to be imaged with higher labelling densities than is currently possible.

4.3 Localising with volume reconstruction

In this section, the process of point-localisation performed via the volume reconstruction method embodied by 3D-EPM is outlined. For simplicity, we initially restrict our discussion to the high signal-to-noise ratio regime, although the influence of signal-to-noise ratio is investigated later in this chapter. We also restrict our initial discussion to the case of imaging with the twin-Airy PSF, although it will be shown later that we can in fact use any PSF that uniquely encodes depth information into images.

The localisation procedure begins with acquisition of the raw data. Data acquisition is performed in the same fashion as other 3D point-localisation techniques that utilise PSF engineering: typically using epi-illumination on a commercial microscope, although we are not limited to this case. Usually, SMLM experiments are performed using extremely high NA objective lenses that meet the resolution requirements, although some other point-localisation experiments, e.g. particle-tracking applications may not share this requirement. The objective lens and phase-mask should be chosen such that the EDoF is large enough to cover the depth-range that is to be imaged. The choice of camera used for imaging should be appropriate to the experiment: typically for single-molecule imaging an electron-multiplying charge-coupled device (EMCCD) camera, capable of single-photon detection, is used; for particle-tracking experiments where the photon-budget is higher a charged-coupled device (CCD) camera may be used instead. In addition to the raw data, a calibration PSF z -stack should be acquired, which may be refined via the same phase-retrieval and Zernike filtering procedures outlined in the

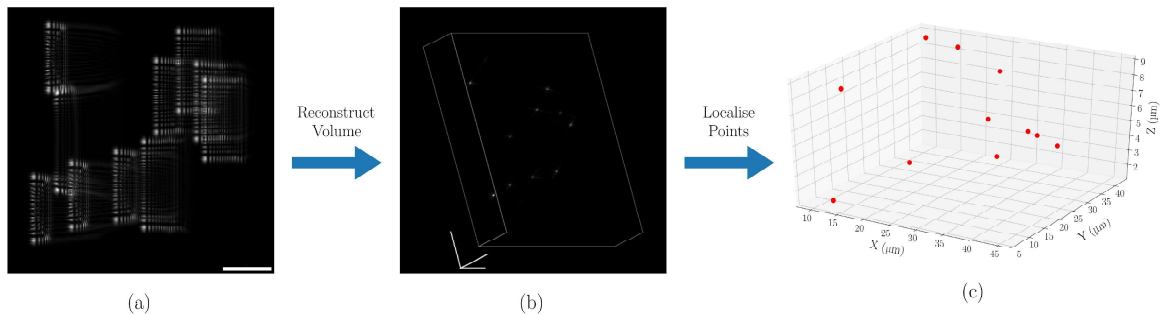


Figure 4.1: Schematic of localisation performed with 3D-EPM localisation. (a) PSF-engineered images are acquired of the sample, which comprises point-emitters that appear as PSFs in the raw images. Scale-bar is $10\ \mu\text{m}$. (b) The volume reconstruction employed in 3D-EPM is applied to reconstruct the 3D sample volume. (c) Points are identified and localised within the reconstructed 3D volume.

previous chapter.

Once the raw data and calibration PSF have been acquired, the 3D-EPM localisation differs from the norm of localising from the raw 2D data. Instead, the 3D sample volume is reconstructed according to equation 2.14, where the encoded depth information is extracted during deconvolution and the PSFs in the raw image data converge back to points in 3D space. Localisation is then performed in the reconstructed 3D volume, by fitting 3D Gaussian profiles to points identified within the reconstructed volume. For the simulations and analysis performed in this chapter, the 3D localisation was performed using the *trackpy* python library [116], however use of more sophisticated fitting regimes may further improve localisation performance by, for instance, more effectively identifying closely-spaced points as separate. Figure 4.1 displays a schematic outlining the reconstruction and localisation procedure.

In statistical terms, a procedure for estimating underlying parameters (for instance the spatial coordinates of point-emitters) from raw data is called an estimator. In our case, the estimator is the procedure of 3D-reconstruction and localisation. A good estimator should exhibit a number of features. Firstly, it should be unbiased: i.e, on average, it should recover the true underlying parameters. Additionally, it should be precise: there should be a small variance in repeated measurements of the same parameter. There is a theoretical limit to the precision of an estimator, given by the CRLB, which depends on the information content of a given system about the parameter(s) of interest; this is discussed further in section 4.5. A good estimator of a parameter of interest should, therefore, estimate the parameter with a precision approaching the theoretical limit dictated by the CRLB. In the following sections these characteristics are explored further.




Figure 4.2: Reconstruction-localisation, on average, recovers close to the true value of position for a single point-source.

4.4 3D-EPM localisation as an unbiased estimator

Figure 4.2 displays the results of 1000 repeats of a noise-free simulation where the 3D spatial coordinates of a single point source, imaged with a twin-Airy PSF, were estimated with the method outlined in figure 4.1. The simulation was performed at 0.5 NA with an emission wavelength of $\lambda=660$ nm. In each repeat, a $65 \times 65 \times 15 \mu\text{m}$ volume was seeded with a single point at randomly determined coordinates. To allow sub-pixel coordinates, the point was simulated as a 3D Gaussian distribution with a standard deviation $\sigma = 32.5$ nm. The point was then simulated to be imaged with a twin-Airy PSF according to equation 2.10. For this initial discussion, no noise was added to the images. The 3D sample volume was then reconstructed with 10 iterations of the 3D-EPM reconstruction algorithm, before the 3D position of the point was estimated using *trackpy* software. Only 10 iterations were required in this simulation since the images were free of noise and the sample comprised only a single point-source. Figure 4.2 shows histograms of the error in recovered position, defined as *true position - estimated position*, along each axis. We see that, on average, the localisation procedure estimates the position of the point with close to zero error: in this simulation we saw mean errors of -9.2 nm, 4.8 nm and -12.6 nm along *xyz* respectively, which are far beneath the diffraction-limited PSF width of such a system (as determined by equations 1.2 and 1.3). Since these error values are far beneath the theoretical resolution limit of such a microscope, we can conclude from the results shown in this section that estimating the 3D position of a point-emitter from within the 3D reconstructed volume does not introduce significant bias into the estimated spatial coordinates.

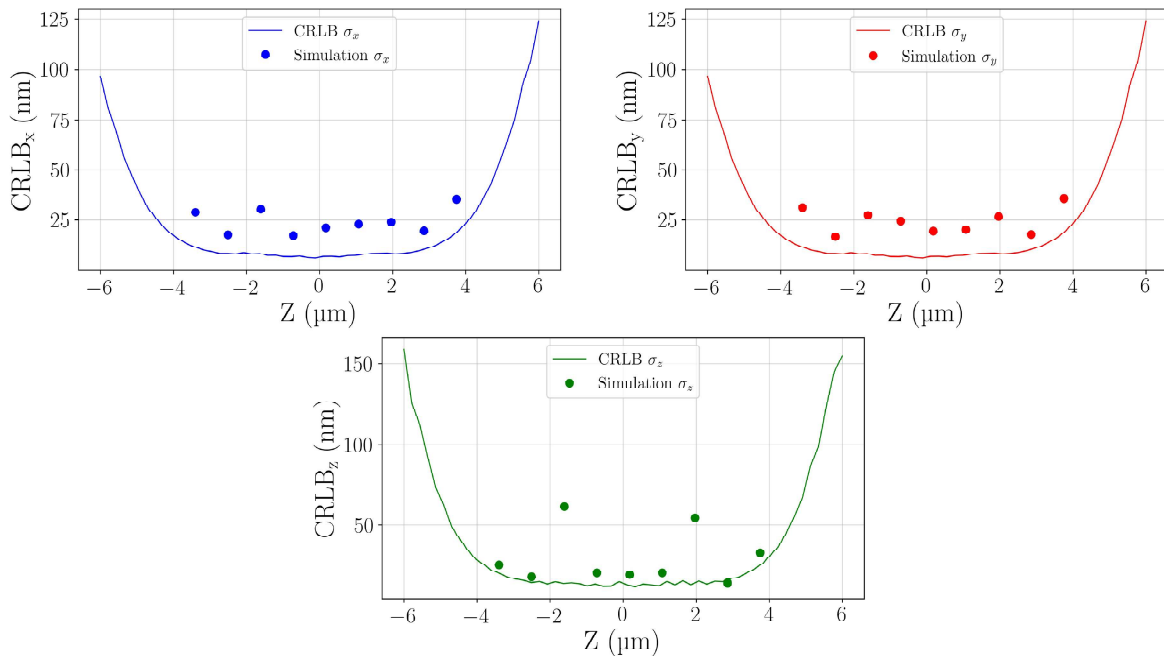


Figure 4.3: E3D-EPM localisation approaches the theoretical limit on precision set out by the CRLB. Plots show the theoretical limit (solid line) on precision along x (blue), y (red) and z (green), as determined by the CRLB for a simulated oil-immersion 1.4NA $\alpha = 5$ COSA PSF, assuming the Poisson data model [117]; a detector with $13\ \mu\text{m}$ pixels; system magnification of $60\times$; 3500 signal photons per localisation and mean 50 background photons per pixel. Each point shows the standard deviation on the error of 50 measurements of the position of a single point via 3D-EPM localisation over a depth range of $7\ \mu\text{m}$, under the same imaging conditions.

4.5 Fisher Information and the Cramer-Rao Lower Bound

As well as providing accurate measurements to the position of point-sources, an estimator should also be precise. The CRLB is a statistical framework that provides a theoretical maximum to the precision that may be achieved by a given estimator when measuring a parameter of interest. The CRLB is related to the Fisher information, a metric that quantifies the sensitivity of a given system to changes in its underlying parameters. In point-localisation applications, Fisher information can be used to assess the information content of a given PSF about parameters of interest, for instance spatial coordinates [117]. Because the Fisher information and CRLB of a given PSF are independent of the estimator used in data analysis, they can be used as comparison metrics to compare the theoretical performance of different PSFs. For this reason, the CRLB is often used in the field of PSF engineering, both in the design of PSFs [65, 118, 119] and in assessing the performance of existing PSFs [66, 68, 113].

While PSF design is not the topic of this thesis, the CRLB is still important to our discussion. Regardless of the PSF chosen for imaging, a good estimator should estimate the coordinates of a point-emitter with a precision approaching the theoretical maximum imposed by the CRLB. Strictly speaking, the CRLB is the lower bound on

the variance that an unbiased estimator may achieve when estimating an underlying parameter θ from the data. In the context of our discussion, the parameters of interest are the spatial coordinates of a single molecule. The square root of the CRLB is therefore the lowest achievable standard deviation of repeated measurements of the spatial coordinates of a point.

The CRLB is calculated from the Fisher information matrix for a given PSF, which requires the probability distribution for the pixel count on the k^{th} pixel to be calculated. The exact form of the probability distribution depends on the mathematical model used to characterise the imaging system: Chao *et al.* described a number of such models including the idealised Poisson data model as well as models for image detection with charged coupled devices and EMCCD cameras [117]. Nevertheless, regardless of the data model chosen, the Fisher information matrix \mathbf{I} for a parameter $\theta \in \Theta$ from data τ is calculated as [117]:

$$\mathbf{I}(\theta) = E \left[\left(\frac{\partial}{\partial \theta} \ln f_{\theta,k}(\tau) \right)^T \left(\frac{\partial}{\partial \theta} \ln f_{\theta,k}(\tau) \right) \right] \quad (4.1)$$

where E denotes the expected value, Θ is the set of underlying parameters (in our case $\Theta = [x, y, z]$) and T denotes transpose. In this discussion, in line with other works in the field, the Poisson data model is considered when calculating the CRLB. The Poisson data model accounts for the shot noise of fluorescence emission and a pixelated detector [117]. However, it does not account for noise introduced during the readout process, and as such is an idealised data model. In the following sections, a data model that does account for readout noise is used for all simulations. However, for this general discussion, the Poisson data model is considered adequate. For this description of image formation, the probability distribution for the k^{th} pixel is given by

$$p_{\theta,k}(N) = \frac{\nu_{\theta,k}^N e^{-\nu_{\theta,k}}}{N!} \quad (4.2)$$

where N denotes the photon count and $\nu_{\theta,k}^N$ is the expected photon count at the k^{th} pixel. The expected photon count is the sum of the expected signal photon count $\mu_{\theta,k}$ and expected background photon counts $B_{\theta,k}$:

$$\nu_{\theta,k} = B_{\theta,k} + \mu_{\theta,k} \quad (4.3)$$

With this data model, and assuming a constant value for the background photon count, the Fisher information matrix for a K pixel image is then given by [117]:

$$\mathbf{I}(\theta) = \sum_{k=1}^K \frac{1}{B_{\theta,k} + \mu_{\theta,k}} \left(\frac{\partial \mu_{\theta,k}}{\partial \theta} \right)^T \left(\frac{\partial \mu_{\theta,k}}{\partial \theta} \right) \quad (4.4)$$

The best-achievable precision, i.e. standard deviation of a set of repeated measure-

ments of the parameter θ_i , that an unbiased estimator can make is then bounded by the CRLB, which is calculated as the inverse of the diagonal element of the Fisher information matrix:

$$CRLB_{\theta_i} = \sqrt{\frac{1}{I_{ii}(\theta_i)}} \quad (4.5)$$

To test the precision of 3D-EPM localisation, the CRLB for a simulated oil immersion 1.4NA $\alpha = 5$ COSA PSF was calculated, assuming the Poisson data model, imaged onto a detector with a pixel size of $13 \mu\text{m}$ at $60\times$ magnification and an emission wavelength of $\lambda = 660 \text{ nm}$. The CRLB for $\Theta = [x, y, z]$ was calculated across an axial range of $12 \mu\text{m}$. The calculated CRLB for each coordinate is plotted in figure 4.3 (solid lines). A point-emitter was then simulated to be imaged throughout the central $6.5 \mu\text{m}$, before the position of the point was estimated with 3D-EPM localisation. The number of signal photons per localisation event was 3500, and the background photon count was 50. This process was repeated 50 times at each z position. The standard deviation in the error between the true and estimated position for each coordinate is plotted on the same graphs as the theoretical limit as determined by the CRLB for each axis. We see that throughout the tested depth-range, when measuring the 3D position of a single point-emitter with this PSF and under the tested noise conditions, the standard deviation along each axis approaches the limit imposed by the CRLB.

Various studies have shown that maximum-likelihood estimation techniques can achieve the CRLB when estimating the position of a single point [101]. However, maximum-likelihood approaches rely on identifying individual PSFs from the raw 2D image data and as such are fundamentally unsuited for dense images that feature significantly overlapping PSFs. While some other approaches enable localisation of overlapping PSFs in 2D [120, 121], the progress in 3D has been more limited, with existing approaches tailored to use with specific types of PSF [114, 68].

In the following sections we will see that 3D-EPM localisation introduces significant advantages when images feature overlapping PSFs. We also explore in simulation its performance as other imaging parameters are varied, including signal-to-noise ratio and the proximity of point-emitters to the edges of the imaging domain. We start, however, by setting out the parameters used for the simulations performed in the results section of this chapter.

4.6 Simulation parameters

Now that we've seen 3D-EPM localisation is a viable position estimation routine, in that (i) it is, to a good approximation, an unbiased estimator, and (ii) the precision approaches the upper limit imposed by the CRLB, we can begin to assess its performance in different imaging conditions. In the following section, we evaluate the

performance of 3D-EPM localisation in simulation and, for context, compare it against a previously-developed point-localisation routine. In the current section, the details of the simulation are set out and the localisation routine used for comparison is described.

4.6.1 Performance metrics

The performance of a given localisation routine that estimates the 3D position of points from raw point-localisation data may be assessed in a number of ways. This is because a good localisation routine (estimator) should exhibit a number of key characteristics. Firstly, the estimator should correctly identify a high proportion of points present in the data. Points that are correctly identified are referred to as true positives, whereas points that are present in the data that the estimator fails to identify are referred to as false negatives. Secondly, the estimator should not incorrectly identify points that are not present in the sample. Points that are incorrectly identified are false positives. A metric that is often used in SMLM (e.g. in [122, 123]) that combines these quantities is the Jaccard index J , defined as:

$$J = 100 \frac{TP}{TP + FP + FN} \quad (4.6)$$

where TP , FP and FN are the number of true positives, false positives and false negatives respectively. By this definition, a Jaccard index of 100 indicates perfect localisation, a lower value reflects worse localisation performance.

In addition to these performance metrics that relate to the proportion of points identified correctly, the accuracy of the coordinates assigned to localised points should also be assessed. In the following analysis, the accuracy is quantified by calculating the root mean-square error (RMSE) between the estimated location of points that have been localised correctly, and the true location. Because the performance of point-localisation routines is often different in z than it is in xy , the RMSE is separated into (i) the lateral RMSE, and (ii) the axial RMSE. The lateral RMSE is defined as:

$$RMSE_{xy} = \sqrt{\frac{1}{TP} \sum_{i=1}^{TP} (\hat{x}_i - x_i^{GT})^2 + (\hat{y}_i - y_i^{GT})^2} \quad (4.7)$$

where \hat{x}_i and \hat{y}_i are the coordinates of the i^{th} point in x and y as estimated by a given localisation routine, and x_i^{GT} and y_i^{GT} are the ground-truth (true) positions of the same point. Similarly, the axial RMSE is defined as:

$$RMSE_z = \sqrt{\frac{1}{TP} \sum_{i=1}^{TP} (\hat{z}_i - z_i^{GT})^2} \quad (4.8)$$

for the estimated and true positions for the i^{th} point in z , \hat{z}_i and z_i^{GT} , respectively. Together, the performance metrics outlined in this section quantify both the effectiveness

and accuracy of a localisation routine in a given simulation.

The question of what qualifies as a “correct” localisation, or true positive, is an important one, as each performance metric defined above depends on its the definition of TP . For each simulation detailed in this chapter, the set of estimated coordinates of point-emitters was compared against the set of true coordinates by means of a custom nearest-neighbour searching algorithm, which was implemented as follows:

1. For each point in the set of true points, the separation distance between it and each point in the set of estimated points was calculated.
2. The closest point in the set of estimated points was selected. If the separation distance between this point and the point being considered in the set of true points was beneath a given threshold value T , the points were deemed to have been matched, i.e a true positive has been found.
3. Matched points were then removed from consideration for all future points to avoid points being matched twice.
4. This process was then repeated for every point in the set of true points.

At the end of the point-matching process, points that were left unmatched in the set of true points are false negatives (i.e. points that are present in the ground-truth data-set but have been missed by the localisation routine), and points that were left unmatched in the set of estimated points were false positives (points that *were* estimated by the localisation routine but were not present in the ground truth data). The threshold separation value T that determines whether points are close enough to be matched therefore affects the results and must be kept constant for a given test. Unless otherwise specified in this section, the value used for T was twice the width of a single pixel in the image.

4.6.2 Twin-Airy comparison

In the following section, the performance of 3D-EPM localisation is evaluated in simulation in varying imaging conditions. Imaging was simulated to be implemented using either a twin-Airy PSF or a COSA PSF, since these PSFs both provide large extensions to the native DoF [68, 66]. For context, the localisation results were compared against those provided by the localisation routine designed for twin-Airy data described in [68], using a python implementation of the routine with minor refinements made, written by Michael Handley, a PhD student in the Imaging Concepts Group at the University of Glasgow. This algorithm exploits the fact that the twin-Airy PSF is conceptually similar to a pair of cubic PSFs that share a common axis and are oriented opposite to one another, see figure 4.1 for an example. The lateral separation between the two

sides of the twin-Airy PSF in images depends on the depth of the point-emitter. Localisation works by matching points between two Wiener-restored images that have been deconvolved with either side of the $z = 0$ in-focus PSF. Full description of the algorithm can be found in ref [68], however a summary of the algorithm is as follows:

1. The raw, PSF-engineered images are Wiener-deconvolved twice: once with each half of the in-focus twin-Airy PSF (each deconvolution kernel corresponds to approximately a cubic PSF). This restores PSFs to approximately point-sources: those that were deconvolved with the correct side of the twin-Airy PSF are bright spots, the shape of which depends on the mismatch in defocus between the PSF and deconvolution kernel; whereas those that were deconvolved with the incorrect side of the twin-Airy PSF are far less intense.
2. The restored images are then subject to an intensity threshold, where the lower-intensity spots corresponding to points restored with the incorrect side of the deconvolution kernel are disregarded. Gaussian fitting is then performed on the points that remain.
3. Points are then matched between the two restored and thresholded images. For two lobes to be matched, they must satisfy certain criteria including being located on the same axis, sharing similar intensities and having a separation distance that corresponds to a physically-possible value of defocus.
4. Matched points are then compared with calibration curves to calculate their 3D position.

The calibration curves link the separation distance of deconvolved lobes with the corresponding defocus, and the defocus with the translation of the PSF along the other lateral axis. They are calculated in advance by deconvolving each xy slice in a calibration PSF z -stack via the same method described above, where the defocus is a known prior. The resulting change in separation and translation of the Wiener-restored lobes as defocus is varied serves as the necessary calibration data. A schematic demonstrating the workings of this algorithm is shown in figure 4.4. It should be emphasized, however, that this twin-Airy lobe-matching algorithm is not the only choice of localisation routine: several different frameworks to estimate 3D position from PSF-engineered data exist, each with their own merits and drawbacks. However, to my knowledge it is the only estimator that has been demonstrated to be capable of localising from twin-Airy data where the PSFs overlap.

4.6.3 Image formation model

The final part of this section is to introduce the model of image formation used in the simulations detailed in the remainder of the chapter. While point-localisation

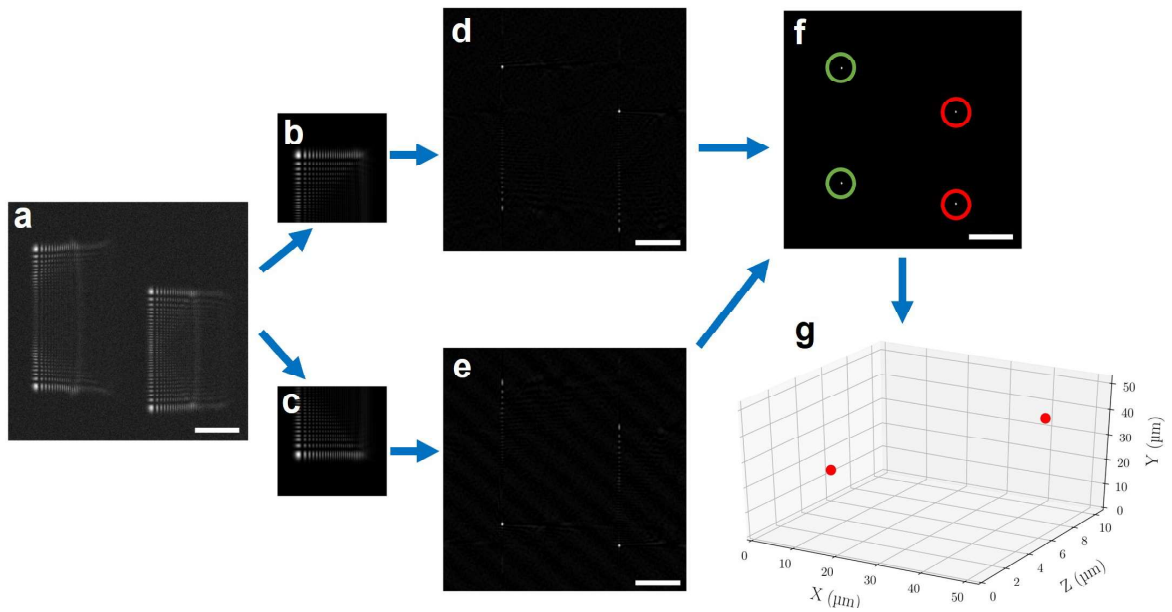


Figure 4.4: Schematic of the point-matching approach for twin-Airy point localisation. Full details can be found in [68]. The raw, PSF-engineered image (a) is Wiener-deconvolved separately with the upper (b) and lower (c) halves of the $z = 0 \mu\text{m}$ PSF. The corresponding Wiener-restored images (d-e) are subject to an intensity threshold to remove dimmer lobes caused by deconvolution with the wrong side of the PSF, before being summed to recombine into a single restored image (f). Points are then identified in this image and matched according to several criteria, including their intensity, x position and ellipticity. Finally, matched points are compared with pre-determined calibration curves that measure both the translation and separation of lobes with defocus, to calculate the 3D position of the corresponding point. Scale-bar $10 \mu\text{m}$.

has applications across many spatial scales in microscopy, it is perhaps at the single-molecule level where it has had the most impact in the life sciences through its ability to image structure with a resolution surpassing the diffraction limit. For this reason, the simulations listed in this chapter are performed at the spatial scales associated with single-molecule imaging. However, single-molecule imaging is often subject to relatively low numbers of signal photons, which, when combined with PSF engineering techniques, may be spread over a much larger area than the diffraction-limited PSF. This means that ordinary charged-coupled device cameras, for instance, are not well-suited to single molecule imaging since the noise introduced during readout can overwhelm the weak signal. Advances in single-molecule imaging have been aided by the improvement of detector technology; for instance the EMCCD camera, which is able to detect single photon events and offer improved imaging performance in extremely low-light imaging conditions. EMCCD cameras feature on-chip multiplication, where the number of electrons generated from a single photon is amplified prior to being read out, thereby elevating the signal level above the noise floor. EMCCD cameras are widely used in SMLM and therefore the simulations discussed in this chapter also consider image detection with an EMCCD camera. However, because single-molecule imaging is typically performed in such a low-light regime, the influence of noise plays a crucial role in the performance of SMLM. As such, it is necessary to amend our model of image formation to reflect the noise characteristics of EMCCD image detection.

EMCCD data model

For consistency with the literature, the model used for EMCCD image formation in the remainder of this section is the same as that used in [122] and derived in [124]. The model describes how EMCCD images are subject to numerous noise contributions including the shot noise of fluorescence emission from both the sample and any fluorescent background, described by a Poisson noise distribution; the read noise of the camera, described by a Gaussian noise distribution; and the noise introduced during the electron multiplication, described by a Gamma distribution. The value recorded by a given pixel during image formation is described in terms of the number of input photons, n_{PI} , the resulting number of input electrons, n_{EI} , and the number of output electrons after on-chip multiplication and readout, n_{EO} . The number of input electrons is calculated as:

$$n_{EI} = \text{Pois}\{QE \cdot n_{PI} + c\} \quad (4.9)$$

where QE is the quantum efficiency of the camera, c is the number of spurious charge electrons introduced per pixel (that are multiplied during the on-chip amplification) and Pois denotes the Poisson distribution. The number of output electrons, after electron-multiplication, is then:

$$n_{EO} = \Gamma(n_{EI}, EM_{gain}) + \text{Gaus}(0, \sigma_R) \quad (4.10)$$

where EM_{gain} is the electron-multiplying gain, Γ denotes the Gamma distribution, σ_R is the root mean-square noise and Gaus denotes the Gaussian distribution. The final pixel count P is then determined by the camera baseline BL and the analog-to-digital conversion factor E_{ADU} :

$$P = \text{floor} \left(BL + \frac{n_{EO}}{E_{ADU}} \right) \quad (4.11)$$

where the ‘‘floor’’ operation replaces the input value with the greatest integer beneath the value in the brackets. The model is completed by clipping the maximum pixel value according to the bit-depth of the camera: e.g. for a 16-bit image the maximum value is capped at 65,535.

For consistency with the literature, the values used for the parameters in this model in the simulations described in this chapter are the same as those used in [122], which were taken as the manufacturer-quoted values for the Photometrics Evolve Delta 512 EMCCD camera. The exact values used are listed in table 4.1.

Parameter	Value
QE	0.9
σ_R	74.4 e^-
EM_{gain}	300
BL	100
E_{ADU}	45 e^-
c	0.002 e^-

Table 4.1: EMCCD camera parameters used in simulation. Values are the same as those used in [122] and are the manufacturer-quoted values for the Photometrics Evolve Delta 512 EMCCD camera.

4.6.4 Iterations used for 3D-EPM localisation

Determining the optimal number of iterations used for RL deconvolution is not trivial when the ground-truth solution is not known. This is because the convergence of the solution depends on several factors including the signal-to-noise ratio of both the image and the deconvolution kernel. When deconvolving noisy images of extended fluorescent structure, the noise amplification phenomenon discussed in section 1.2.5 limits how far the deconvolution can be left to iterate, unless controlled with regularisation. However, when reconstructing volumes comprising point-emitters, as implemented in the current chapter, the phenomenon of noise amplification becomes less problematic, since the reconstructed volume is an intermediate result and further processing is required to localise each point. For optimal localisation performance, a sufficient number of iterations is required to ensure that each PSF has converged back to a point in the

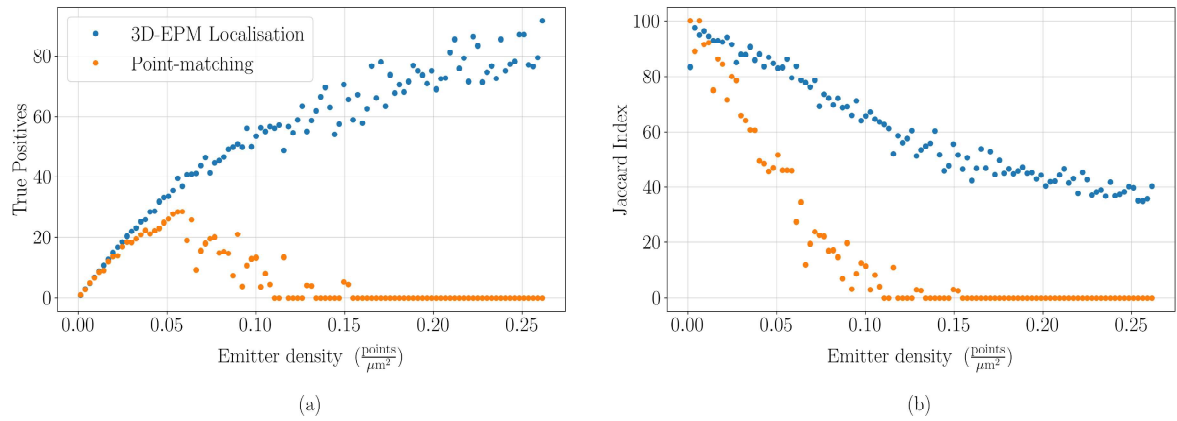


Figure 4.5: 3D-EPM localisation performs well with increasing emitter density. (a) The number of true positives identified with increasing emitter density by 3D-EPM localisation (blue) and the point-matching approach from [68] (orange). (b) The Jaccard index from the same set of results.

reconstructed 3D volume. The localisation can then be instructed to ignore points identified within the reconstructed volume that are smaller than the diameter of a deconvolved PSF, which are likely to be artefacts of noise amplification. Unless otherwise stated in this chapter, 250 deconvolution iterations were employed in the intermediate volume-reconstruction step.

4.7 Results

In this section, the performance of 3D-EPM localisation is assessed. In each case, twin-Airy EMCCD images are simulated using the model described in the previous section. The performance of 3D-EPM localisation is compared against that achieved with the twin-Airy point-matching localisation procedure [68]. Imaging is simulated with an oil-immersion 1.4NA objective lens, a fluorescence emission wavelength of 660 nm and an $\alpha = 5$ twin-Airy phase-mask at the exit pupil.

4.7.1 Emitter density

The first results we discuss are the performance of 3D-EPM localisation with increasing emitter density. Figure 4.5 shows the results of a simulation that was performed where the number of points present within a $52 \times 52 \times 8 \mu\text{m}$ ground-truth volume was increased. The ground-truth sample volume was seeded with point-emitters by adding a 3D Gaussian spot with a full-width at half-maximum of 150 nm at randomly generated sub-pixel coordinates. EMCCD images of the ground-truth volume were then simulated according to the EMCCD image formation model in equation 4.11. The total system magnification was $100\times$ and the EMCCD pixel size was simulated to be $13 \mu\text{m}$, with a 100% fill factor. To ensure the entire lateral extent of PSFs was captured within the FoV, the points were confined to a smaller central lateral region of the ground-truth

volume, half the dimensions of the larger volume. Points were localised from the simulated EMCCD images using the 3D-EPM localisation method described in section 4.3, and then for comparison the twin-Airy point-matching localisation approach summarised in section 4.6.2 was used to localise the same data. The set of coordinates of the localised points from both localisation approaches was then compared against the set of true coordinates of point-emitters in the ground-truth images. This process was then repeated as the number of points in the ground-truth volume was increased.

Figure 4.5a shows how the number of true positives identified via 3D-EPM localisation (blue) and the point-matching approach (orange) as the density of emitters in the ground-truth volume is increased. Figure 4.5b shows the Jaccard index for the same data. Each point is the mean of ten repeats at the corresponding emitter density. We see that the number of true positives identified by the point-matching approach increases with the number of points at the start of the tested range, but then begins to fall again once the emitter density increases above a density of approximately 0.6 emitters per square micron. With 3D-EPM localisation, however, the number of true positives continues to rise with emitter density and does not see the same drop in performance above a given threshold. This same trend is observed in the Jaccard index J displayed in figure 4.5b: J falls much quicker with the point-matching approach than with 3D-EPM, where J remained significantly higher as the emitter density increased, reflecting better localisation performance. The value of J fell to and remained at approximately 0 with the point-matching approach above emitter densities of 0.12 points per square micron, whereas with 3D-EPM it remained above 35 throughout the entire test range. This means that 3D-EPM localisation performs significantly better than the point-matching approach as the fluorophore density is increased.

The improved performance at higher emitter density of 3D-EPM localisation when compared with the point-matching approach can be attributed to several factors. Firstly, 3D-EPM localisation performs the localisation of points in the reconstructed 3D volume, whereas the point-matching approach localises within the Wiener-restored 2D images. This means that in 3D-EPM localisation, the effective sparsity of the fitting domain is increased and therefore points are less likely to coincide. Secondly, while Wiener restoration of the twin-Airy PSFs in the point-matching approach yields a bright spot when the correct side of the in-focus PSF is used as the deconvolution kernel, restoration with the incorrect side of the PSF introduces deconvolution artefacts in the form of several dimmer spots in the restored image, see figure 4.4d and 4.4e. Despite steps being taken to remove these artefacts from consideration when matching points, such as the application of an intensity threshold and checks on e.g. the ellipticity of any matched points [68], such artefacts can inhibit the performance of the localisation routine through the identification of false positives, especially when significant variation in the intensity of PSFs is observed in the raw images.

Emitter density definition

Defining the metric used to quantify emitter density in 3D SMLM is not trivial. This is because an acquired 2D image in 3D-SMLM implemented with engineered PSFs is a projection of the 3D sample volume. However, the depth of the volume being projected will vary between different experiments and different PSFs. An obvious choice of metric, for instance, might be the number of emitters per cubic micron in the ground-truth volume. Yet, an important feature of PSFs designed for 3D SMLM is the depth-range over which they can be used, with PSFs such as the tetrapod and twin-Airy being designed to maximise this range [65, 68]. However, a deeper volume with the same emitter density as quantified by this metric will lead to a greater number of emitters being imaged. In other words, a PSF that provides a greater extension to the DoF will yield denser projection images than a PSF that provides a smaller extension for the same number of emitters per cubic micron. This choice of emitter density metric, therefore, does not adequately reflect the challenge of localising points from images acquired with engineered PSFs. For this reason, a custom metric that calibrated the emitter density based on the lateral extent of a given PSF with respect to that of the diffraction-limited PSF was used in [114]. However, since the focus of this chapter is on the localisation routine rather than PSF design a metric that quantified 2D density of emitters in projection images was deemed sufficient. The definition of emitter density used in this chapter, therefore, is the number of emitters present within a volume divided by the lateral area of the projected volume, giving units of emitters per square micron. This is deemed an appropriate unit for the comparisons shown in this chapter since the imaging depth range was kept constant in each simulation, however care should be taken before directly comparing these results with those presented in other works.

4.7.2 Signal-to-noise ratio

The next factor we consider is the influence of signal-to-noise ratio on localisation performance. A versatile localisation routine should perform well at varying noise levels. This is because while some point-localisation applications, e.g. particle-tracking experiments, are not photon limited, the photon budgets associated with SMLM are typically much lower. To verify 3D-EPM localisation performs well at varying signal-to-noise ratio, a simulation was performed where a volume seeded with a constant number (15) of point-sources was imaged while varying the number of signal photons associated with each localisation event. Imaging was simulated again with a twin-Airy PSF onto an EMCCD camera, with a relatively high background of 50 photons per pixel also incident. Localisation from the twin-Airy images was then performed with both the point-matching approach and 3D-EPM localisation. Figure 4.6a shows the Jaccard index of the resulting localisations, and figure 4.6b shows the number of true positives achieved as the number of signal photons per localisation event was increased.

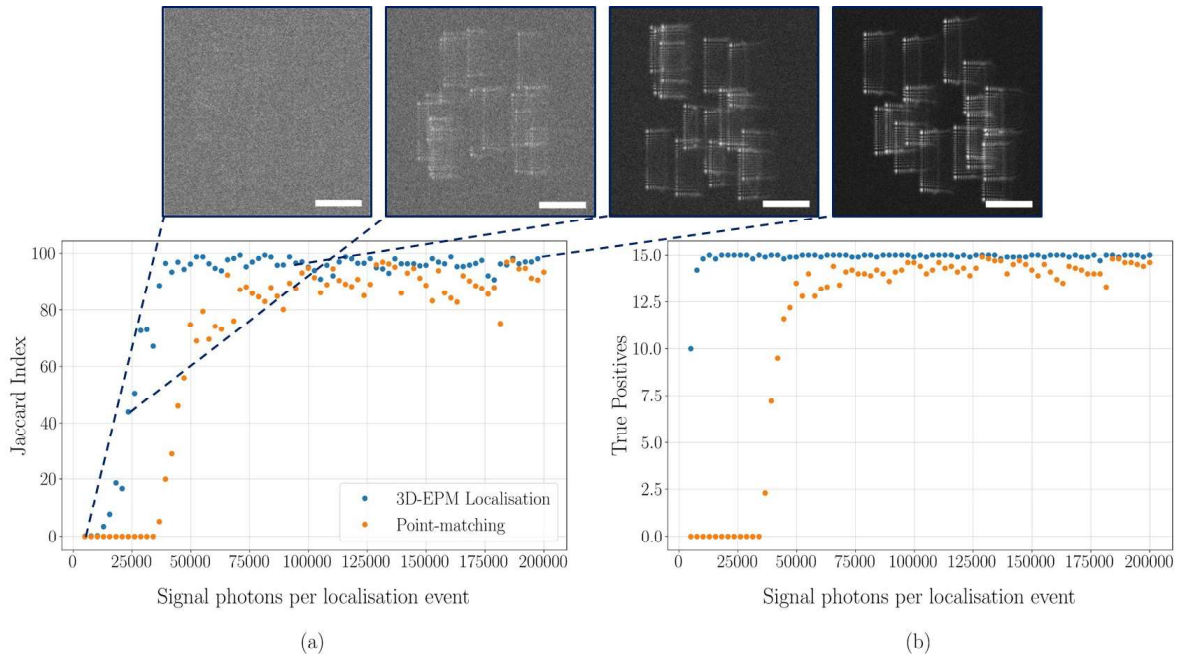


Figure 4.6: 3D-EPM localisation performs well with varying signal-to-noise ratio. (a) the Jaccard index achieved by 3D-EPM localisation (blue) and the point-matching approach (orange) as the number of signal photons per localisation event was increased. (b) The number of true positives identified by each estimator. Images were formed using the EMCCD data model with a background photon count of 50 photons per pixel. Inset panels show example images at various signal-to-noise ratios in the tested range. At the lowest signal levels no PSFs are evident in the simulated images.

Each point in figures 4.6a and b is the mean of ten repeats. The inset panels at the top of figure 4.6 display example frames at various numbers of signal photons.

At the lowest-tested signal photon levels, neither localisation routine achieves a high Jaccard index (both score close to zero). This is not unexpected, since the signal is lost in the background noise (see left inset panel, for instance). Yet, the number of true positives with 3D-EPM is not zero at these low signal levels. Indeed, figure 4.6b seems to suggest that 3D-EPM is performing better in this region. How can it be, then, that 3D-EPM localisation still identifies some correct coordinates of points at such low signal levels? This is explained by inspection of a reconstructed volume from a low signal-to-noise ratio image. Figure 4.7a displays an xy maximum intensity projection through a reconstructed volume at the lowest-tested number of signal photons per localisation event. We see that the volume appears to consist of only noise artefacts. The point-detection algorithm, therefore, identifies many thousands of points when applied to such a reconstructed volume. Figure 4.7b displays a plot of the number of false positives localised with varying signal levels. At low signal-to-noise ratio, we see many thousands of false positives identified with 3D-EPM localisation. The apparent success in identifying true positives at these signal levels, therefore, can be attributed to a proportion of these many thousands of noise artefacts in the reconstructed volume being within the distance threshold to the true location of a point in the sample volume and therefore being identified as true positives.

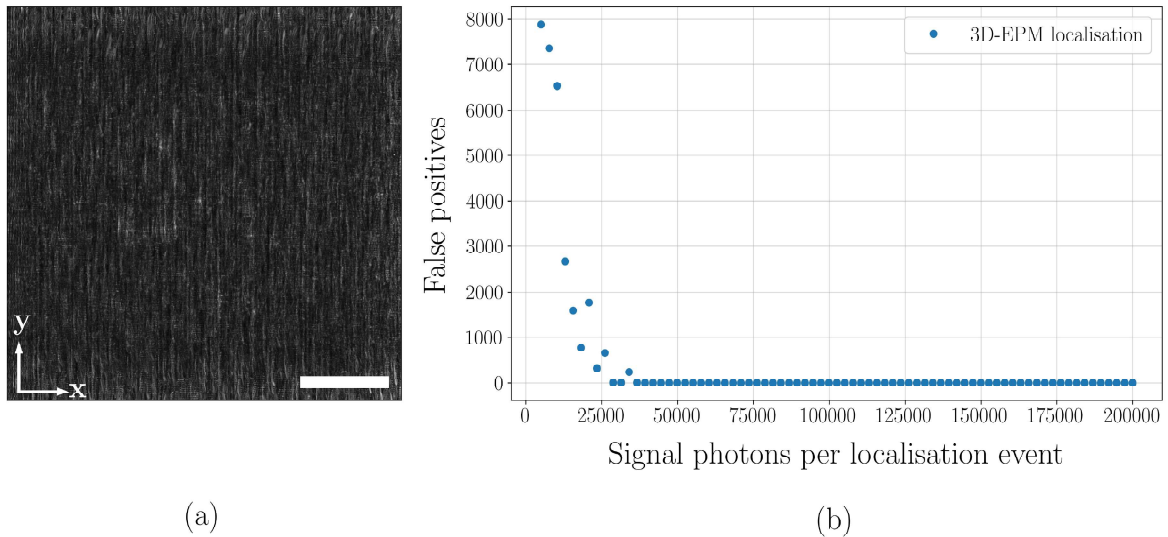


Figure 4.7: 3D-EPM localisation identifies lots of false positives at extremely low signal-to-noise ratio. (a) xy maximum-intensity projection through a reconstructed volume at the lowest-tested signal levels. Noise artefacts are evident, which the localisation routine treats as points. (b) Plot of the number of false positives identified with increasing signal photons. At the lowest signal levels, thousands of false positives are identified in the reconstructed volume.

As the number of signal photons increases, the Jaccard index of both localisation routines rises, as does the number of true positives identified. Above the lowest signal levels, 3D-EPM localisation on average identifies almost all 15 points correctly, while the point-matching approach identifies marginally fewer. The Jaccard index for 3D-EPM localisation rises faster than the point-matching localisation routine, suggesting better performance at intermediate signal levels. From the results presented in this section, therefore, we see that 3D-EPM localisation performs well at all but the lowest signal levels, where the signal is lost in the background noise.

4.7.3 3D-EPM localisation is not limited to a single PSF

The next results we show in this section concern the choice of PSF used for imaging. So far in this chapter, we have demonstrated 3D-EPM localisation using the twin-Airy PSF. However, when implemented with the twin-Airy PSF, our 3D-EPM deconvolution algorithm requires only a calibration PSF stack and the raw PSF-engineered image data to perform volumetric reconstruction. Does this mean, therefore, that *any* engineered PSF is suitable for 3D-EPM localisation? In chapter 3, we demonstrated 3D-EPM using the differential-Airy PSF because it effectively encoded depth information into projection images through disparity. This imaging modality required us to image simultaneously with two cubic PSFs that differed by a small defocus offset. The differential defocus was necessary because although the cubic PSF exhibits a lateral translation with defocus, its appearance remains largely unchanged over the EDoF, and the depth-dependent translation is symmetric about the focal plane. To determine the 3D position of a point, therefore, two projections were needed. The

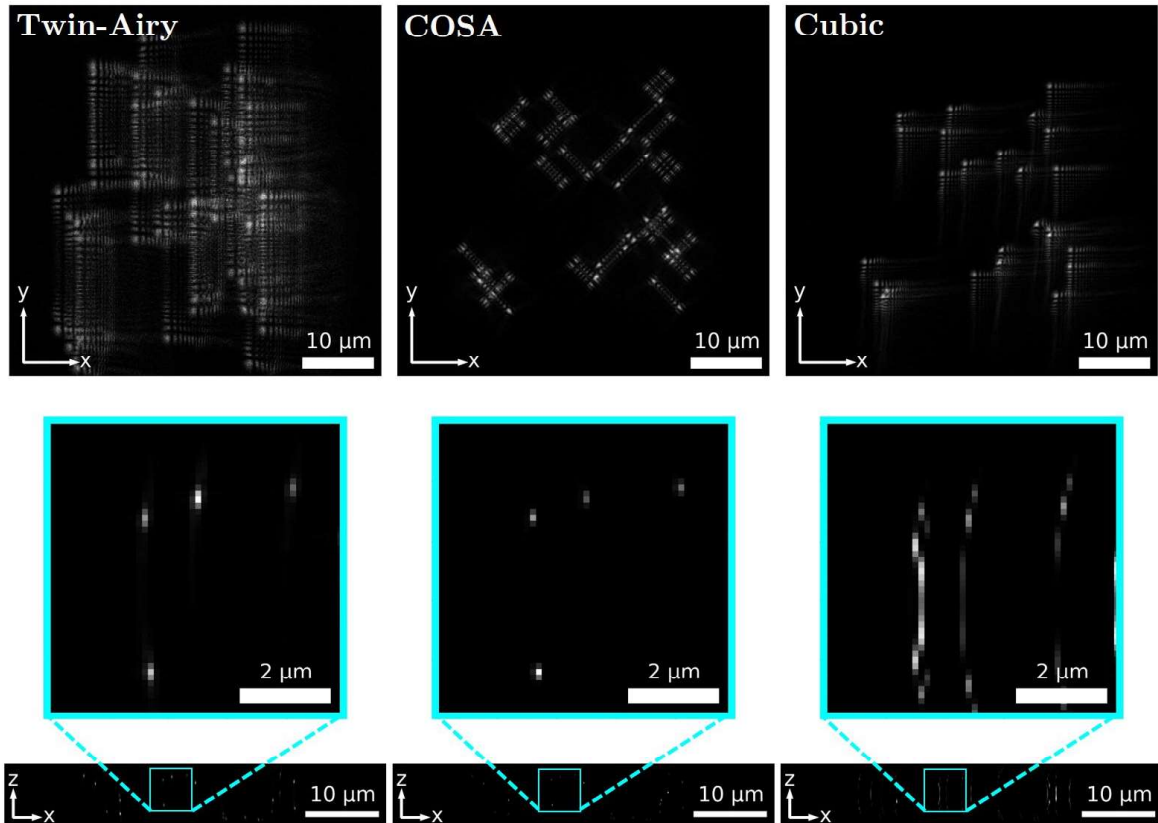


Figure 4.8: 3D-EPM localisation requires images to be acquired with a PSF that effectively encodes depth information across the EDoF. Top row: simulated images of a volume seeded with point-sources using a twin-Airy PSF (left), a COSA PSF (middle) and cubic PSF (right). Bottom row: ZX maximum-intensity projections through the respective reconstructed volumes after 50 3D-EPM deconvolution iterations. Insets show the same region in each. From the twin-Airy and COSA images, the PSFs converge back to points in 3D space. This is not possible from the cubic-encoded image, however, owing to the lack of depth information in a single cubic PSF.

twin-Airy PSF, on the other hand, effectively encodes depth information without requiring multiple perspectives. This is because, unlike the cubic PSF, the appearance of the twin-Airy PSF does change throughout the EDoF and is asymmetric about the focal plane. When compared with a reference PSF z-stack, therefore, the image of a single point does contain enough information to determine its 3D position. This concept is illustrated in figure 4.8, where a simulation was performed in which a volume seeded with point-sources was imaged with a twin-Airy PSF (left), a COSA PSF (centre), and a cubic PSF (right). In each case, the projection image and corresponding PSF stack were used to attempt to reconstruct the sample volume. The bottom row of figure 4.8 displays xz maximum-intensity projections through the reconstructed volumes after 50 deconvolution iterations, with the insets showing zoomed-in versions of the same region. We see that with the twin-Airy and COSA PSFs, the deconvolution is able to successfully extract the encoded depth information and the PSFs in the raw projection data converge back to points in 3D space. With the cubic PSF, however, the PSFs do not converge back to a single point, instead reconstructing to curved streaks, owing to the lack of depth information contained in a single cubic PSF. Therefore, just as with




Figure 4.9: 3D-EPM localisation performs well when point-emitters are allowed close to the edges of the image domain. Plot shows the Jaccard index achieved by 3D-EPM localisation (blue) and point-matching (orange), as the minimum distance from the edges of the image domain that points are allowed to be randomly seeded. 3D-EPM localisation achieves a higher Jaccard index when points are allowed closer to the edge of the imaging domain. Each point is the mean of ten repeats.

3D-EPM in the previous chapter, for a PSF to be effective in 3D-EPM localisation it must effectively encode depth information across the required EDoF. If imaging is performed with only a single PSF, rather than the dual-PSF differential-Airy approach, the appearance of the PSF must change over the required depth range and be asymmetric about the focus to avoid a sign ambiguity. As well as the twin-Airy PSF, other suitable PSFs that satisfy these requirements include the Tetrapod [65], double-Helix [64] and astigmatic PSFs.

4.7.4 3D-EPM localisation performs well close to the edges of the imaging domain

The final results we consider in this chapter concern the performance of 3D-EPM localisation with point-sources that are located close to the edges of the imaging domain. So far in this chapter, the points have been confined to a central region of the sample volume, to ensure the lateral extent of the PSFs does not exceed the edges of the image. However, in practice this is difficult to achieve: the sample is likely to fill the entire FoV of the microscope, and therefore it is likely that some PSFs won't be imaged entirely within the lateral confines of the imaging domain. How then, does this affect the performance of point localisation?

Let us first consider the point-matching localisation approach for the twin-Airy PSF [68]. Here, a point is localised by deconvolving both sides of the PSF with the respective side of the in-focus twin-Airy PSF. This leaves a pair of lobes, the separation between which allows us to determine the points 3D coordinates. However, if a point is positioned such that its image extends beyond the lateral confines of the imaging

domain, this becomes problematic. Consider, for instance, a PSF that is positioned such that only one half of the twin-Airy PSF is present in the image, with the other half falling beyond the image boundary. When deconvolved with the respective sides of the deconvolution kernel, only one bright lobe will be present in the deconvolved images. This means the localisation routine will be unable to match the two lobes necessary for localisation, and the point will therefore not be localised. This is also true for PSFs that only partially spill over the edges of the imaging domain: if the deconvolution doesn't yield two bright deconvolved lobes, the point-matching approach will be unable to successfully localise the point. With 3D-EPM localisation, however, we made explicit allowances in the deconvolution algorithm to account for the fact that some light emitted by points located within the imaging volume may not be recorded within the lateral confines of the FoV. It follows, therefore, that a point should still be reconstructed even if some of its image extends beyond the image boundaries. Therefore, we expect 3D-EPM to perform better when points are located close to the edges of the imaging domain.

To test this hypothesis, a further simulation was performed where the distance from the edge of the image domain that points were allowed to be placed was varied. For each test, a volume was seeded with 15 points placed at randomly generated coordinates, above a certain minimum distance from the edges of the imaging domain. This minimum distance was increased and the performance of 3D-EPM localisation and the point-matching approach was assessed. Figure 4.9 displays the Jaccard index of both methods as the minimum distance was increased. At smaller distances, i.e. when points were allowed closer to the edges of the image domain, 3D-EPM achieves a significantly higher Jaccard index than the point-matching approach, confirming our hypothesis that 3D-EPM better handles points that are close to the edge of the image, where some of the PSF may extend beyond the FoV.

4.8 Discussion

We have seen in this chapter how the theme of this thesis, 3D reconstruction from 2D projection data, introduces significant advantages when localising points in all three spatial dimensions from PSF-engineered images. Because the localisation is performed in the reconstructed 3D volume, as opposed to the raw 2D projection images, the effective sparsity of the fitting domain is increased. We saw in section 4.7.1 how this enables localisation to be performed in dense images featuring significantly overlapping PSFs, as well as sparser images of isolated points. Additionally, we saw how 3D-EPM performs well with varying signal-to-noise ratio and performs better than the point-matching algorithm when emitters are located close to the edge of the fitting domain.

These findings demonstrate that 3D-EPM localisation is an effective, versatile localisation method for 3D point localisation applications in microscopy where PSF engineering is employed to extend the DoF and encode the axial coordinate of emitters




Figure 4.10: 3D-EPM takes longer to implement than the twin-Airy point-matching approach.

into images. By offering improved localisation performance at higher emitter densities, the localisation method developed in this chapter offers the ability to image samples that are more densely labelled than permitted by other analysis frameworks. This has major implications for 3D-SMLM, where experiments typically suffer from poor time resolution, requiring long exposures and high laser intensities, raising the likelihood of photodamage in live biological samples. By localising more emitters per frame with 3D-EPM localisation, the temporal resolution of 3D-SMLM may be improved, requiring fewer frames, shorter acquisitions, and less exposure to harmful laser excitation illumination for a reconstructed SR image. Additionally, the generality of the volume-reconstruction method offers flexibility in the choice of PSF used for imaging, allowing the user to choose a phase function that best suits their experiment.

A potential drawback of 3D-EPM localisation, however, is the computational time needed to perform volume reconstruction and localisation within the 3D volume, compared with the time needed to perform localisation by other methods. Because the volume reconstruction is implemented as a series of 2D convolutions, the time needed to reconstruct scales with image size, the depth of the volume being reconstructed and the number of RL iterations required. The twin-Airy localisation routine, on the other hand, is implemented through non-iterative Wiener deconvolution followed by least-squares fitting of points in the deconvolved images, which is far quicker to implement. Figure 4.10 displays a plot of the time needed to localise 15 points within a volume as the number of pixels increases, using a maximum 100 RL iterations, on a machine equipped with a single Nvidia GeForce GTX 1080 GPU. In this test, all the convolutions required for 3D-EPM localisation were accelerated using GPU, whereas the twin-Airy localisation code is implemented entirely on the CPU. Despite this, the 3D-EPM localisation routine required approximately an order of magnitude greater run-time to implement when localising the largest-tested frames. While this does not affect the time needed for data acquisition (3D-EPM localisation permits faster data

acquisition through its ability to localise dense images), this may be a potential limitation of the technique if real-time data processing is required.

For the results presented in this chapter, the open-source *trackpy* python library was used to detect and localise points within the reconstructed 3D volumes. This software is quick to implement and user-friendly, performing initial feature identification and then least-squares fitting to further refine the sub-pixel coordinates of points in images [116]. However, it is thought that further performance gains may be achieved if fitting routines that use some prior knowledge of the shape of the reconstructed point is used instead.

4.9 Simulated microtubule imaging

We conclude this chapter with a demonstration of the application of 3D-EPM localisation to simulated SMLM images of a sample of synthetic microtubules. Each microtubule in the ground-truth volume was simulated with three 5th order polynomials with randomly generated coefficients for the true xyz coordinates. A total of 1×10^6 fluorophore coordinates were then generated along the trajectory of the microtubules. A SMLM dataset comprising 2500 individual frames was then generated, with a randomly selected subset of the fluorophores “active” each frame. Each frame, a new sample volume was generated comprising the active fluorophores, each a 3D Gaussian spot with $\sigma_{xyz} = 100$ nm placed at its corresponding coordinates. From each sample volume, an EMCCD image was simulated according to the EMCCD data model described earlier in this chapter, with 7000 signal photons per localisation event. Imaging was simulated with a 1.4NA $\alpha = 6$ twin-Airy PSF, a magnification of $100\times$, an emission wavelength of 660 nm over a depth range of 8 μm . An example frame is displayed in figure 4.11a. From each EMCCD image, 3D-EPM localisation was then implemented to localise the active fluorophores. In total, 170196 emitters were localised, an average of 68 per frame. The localised fluorophore coordinates were then used to render a final 3D image, see 4.11b-d.

4.10 Chapter 4 summary

In this chapter, we applied the 2D-3D volume reconstruction pipeline that was developed in chapter 2 and embodied by the 3D-EPM snapshot volumetric imaging method introduced in chapter 3 to the different application of point localisation, where we saw the combination of 3D-EPM reconstruction with 3D fitting (which we named 3D-EPM localisation) to be beneficial. We saw how 3D-EPM localisation performed better than a previously-developed localisation routine when localising from dense images with significantly overlapping PSFs. We also saw how 3D-EPM performs with varying signal-to-noise ratio, and saw further benefits when emitters were in close proximity

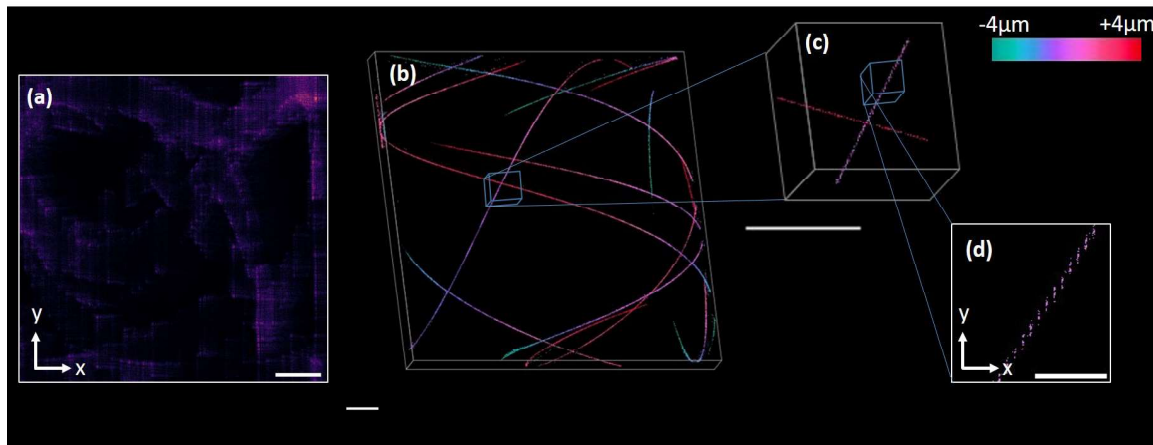


Figure 4.11: Simulated microtubule imaging with 3D-EPM localisation. (a) an example raw frame, featuring a subset of active fluorophores imaged with a twin-Airy PSF. (b) 3D rendering of the reconstructed 3D volume, colour-coded for depth. (c) Zoomed-in version of the region bound by the blue box in (b). (d) xy maximum-intensity projection of the region bound by the blue box in (c). Scale-bars for (a-c) are $10\ \mu\text{m}$, scale-bar for (d) is $2\ \mu\text{m}$.

to the edges of the image domain. Finally, we demonstrated the method by simulating SMLM imaging of a sample of synthetic microtubules. In the next chapter, we turn our attention away from the field of PSF-engineering, and instead apply the volume-reconstruction method to projection images acquired with a different imaging modality.

Chapter 5

Other applications of volumetric reconstruction

In the previous chapters we've seen how the volume reconstruction pipeline developed in chapter 2 may be combined with PSF-engineering methods to perform snapshot volumetric imaging, enabling the 3D study of samples exhibiting extended fluorescent structure. We saw how this is particularly beneficial for samples that undergo motion during acquisition, such as the *Egeria densa* imaged in section 3.6.2. Additionally, we saw how the process of volume reconstruction embodied by 3D-EPM is beneficial in the field of 3D point-localisation, where the acceptable density of emitters can be increased when localisation is performed in the reconstructed volume, compared with the convention of localising directly from the acquired image data. However, up to this point, there remain some limitations of 3D-EPM that we have not yet addressed. These include:

- The depth of sample that can be imaged via 3D-EPM is limited by the DoF of the imaging system. While this may be extended by using an engineered PSF, there still remains a finite limit to the depth range over which the MTF remains useful.
- The resolution of the reconstructed image is fundamentally limited by the numerical aperture of the objective lens. This leads to an anisotropic resolution: by equations 1.2 and 1.3, the axial resolution is always worse than the lateral resolution.
- The density of the sample remains a limiting factor to the effectiveness of the reconstruction. For instance, when imaging with the differential-Airy PSF as implemented in chapter 3, dense samples lead to minimal depth-dependent disparity between the two snapshot images.

What happens, therefore, if we wanted to image a sample that is deeper than the EDoF provided by a cubic PSF? In a light-sheet system this is achieved by simply

extending the depth range over which the sample is scanned. In our snapshot imaging modality, however, this problem is not solved as easily. Of course, the DoF may be extended by decreasing the numerical aperture, but the associated loss of resolution means this solution is far from ideal. Alternatively, increasing the strength of the CPM also extends the DoF, however we saw in figure 3.1d that this extension is provided at the cost of a decreased peak signal-to-noise ratio, as well as a larger PSF footprint, which acts to decrease sparsity in the raw snapshot data. Is this it then? Have we reached the limitations of the method? Fortunately this is not the case. In chapter 2, we developed a generalised mathematical framework for volumetric reconstruction from 2D data. We noted that for the reconstruction to work effectively, depth information must be effectively encoded into the image data across the required DoF. However, up to this point, we have only applied the reconstruction pipeline to data acquired using PSF-engineering methods. In the current chapter, we explore other imaging modalities where the reconstruction pipeline developed in chapter 2 can be applied, where we will see that the above limitations can be addressed.

5.1 Projection imaging as a general description

To begin this discussion, we initially turn our attention back to the model introduced in chapter 2 that describes EDoF image formation. For ease, the equation describing the formation of such an image I is repeated here:

$$I(x, y) = \text{Pois} \left\{ \sum_k O(x, y, z = k) \otimes \text{PSF}(x, y, z = k) \right\} + N(x, y) \quad (5.1)$$

where x, y and z are 3D coordinates, O is the 3D sample object, N is a noise term, \otimes denotes convolution, and Pois denotes the Poisson noise statistics that describe the detection of fluorescence emission, as described in section 1.2. In chapter 3, this model was used to describe the image formed on an epi-fluorescence system with an engineered pupil function. However, equation 5.1 is not unique to this particular imaging modality. In fact, it describes *any* imaging format where an extended depth is projected into a single image, providing the PSF can be approximated as spatially invariant across the FoV. That is to say, providing the 2D image formation of an infinitely thin xy plane can be approximated as a 2D convolution, the above relationship holds for an EDoF projection image.

Many different techniques exploit projection imaging as a concept to improve temporal resolution [52, 71]. By projecting the entire sample depth onto a single 2D image, the entirety of the sample can be imaged within the course of a single camera exposure. Notable applications include calcium imaging in the zebrafish brain using a scanned light-sheet and cubic-encoded detection [125], and functional brain-imaging using elongated Bessel focus [126]. The disadvantage of projecting the entire sample depth into a 2D image, however, is that no information about the position of individ-




Figure 5.1: Schematic of a multi-angle projection microscope, introduced by Chang et al [5], which enables real-time viewing of samples from different projection angles. A scan unit, comprising a pair of galvanometric mirrors, is introduced into the imaging path of a microscope capable of optical sectioning. As the focus is swept along the optical axis during acquisition of a z -projection, the angle of the mirrors θ is changed, introducing a shear at the image. The resulting projection is a stretched version of the projection that would be acquired if one was to physically rotate the sample.

ual features along the optical axis is obtained, therefore any analysis is constrained to two dimensions. While this may not be too problematic for e.g. functional brain imaging, projection imaging methods are not normally a viable solution wherever 3D information about the sample is needed, e.g. in the study of 3D cell morphology [127] or mapping blood-flow in 3D [128, 129].

However, using the deconvolution framework developed in chapter 2, we demonstrated in chapter 3 how a 3D volume may be reconstructed from projection images, providing depth information is present in the acquired projection data. It follows, therefore, that if another projection image modality existed where this requirement is satisfied, the same mathematical framework could be applied to that imaging method also. In the following section we examine one such imaging modality, multi-angle projection imaging [5], and demonstrate how the volumetric reconstruction may also be applied to data acquired on a multi-angle projection imaging system, permitting the reconstruction of larger sample depths across a broad range of microscope systems.

5.2 Multi-angle projection imaging

Multi-angle projection imaging is a recently-developed imaging modality that enables 3D samples to be rapidly interrogated from varying angular perspectives [5]. By incorporating a scan-unit, comprising a pair of galvanometric mirrors, into the imaging path of any microscope that inherently provides optical sectioning, the system is transformed into a multi-angle projection microscope capable of interrogating the sample from varying projection angles without physically rotating the sample. A schematic of the imaging path of a multi-angle projection microscope is shown in figure 5.1.

In conventional projection imaging modalities, images are acquired by rapidly sweeping the focus through the sample depth during the course of a single camera exposure, resulting in a 2D projection of the 3D sample along the optical axis. This focal-sweep may be achieved with, for instance, a tunable lens in the optical path that rapidly adjusts the axial focus of the microscope [69, 70]. Alternatively, phase-modulating devices such as deformable mirrors that are placed at a conjugate pupil plane may be used to rapidly sweep the focal plane through the sample [71]. In the absence of noise, the resulting projection image is numerically equivalent to the image that would result from summing the intensity along z in a 3D image of the sample acquired on the same microscope. Such imaging modalities benefit from fast acquisition speeds and reduced data overheads compared with 3D z -stacks acquired via confocal or light-sheet microscopy. However, as mentioned in the previous section, projection images acquired in this fashion suffer from a lack of any information about the position of sample features along the optical axis, constraining any further image processing or analysis to two dimensions.

Multi-angle projection imaging, on the other hand, overcomes this limitation by enabling the sample to be optically projected from varying angular perspectives [5]. By controlling the angles of the galvanometric mirrors, which are synchronised to step with the axial sweeping of the focal plane, the angle at which the sample is projected may be adjusted. This enables information about the 3D structure of the the sample to be captured by means of a projection series, where the sample is optically rotated between frames, representing a conceptually different way of interrogating a 3D sample compared with conventional z -stack imaging. Probing a sample by means of a projection series is likely to benefit from reduced acquisition times compared with z -stack imaging, since fewer projection images are likely to be required than the number of xy slices necessary for a z -stack. Additionally, while it remains true that a single projection image contains no information about position along the projection axis, acquiring two or more projection images from different angular perspectives *does* provide 3D information, in much the same way that stereoscopic imaging enables calculation of a depth map of the observed scene [130]. Indeed, the authors of ref. [5] demonstrated 3D particle tracking performed on a multi-angle projection microscope by imaging the sample from just two angular perspectives. However, this analysis was constrained to the localisation of sparse point-sources rather than a full 3D volume reconstruction. Nonetheless, this raises the question: since depth information is present in multi-angle projection data, can the volume-reconstruction framework developed in chapter 2 also be applied here? We shall see in the following sections that the answer to this question is yes: volumetric reconstruction is possible from a small number of angular projections acquired via multi-angle projection imaging. We demonstrate this across a number of samples, microscopes and spatial scales in the following sections. However, just as was the case with 3D-EPM, before we can begin reconstructing a volume from multi-angle projection data, we require a model for image formation that we can then attempt to

invert.

5.2.1 Multi-angle projection image formation

Image formation on a multi-angle projection microscope can be treated as two distinct steps: shearing and projection. In practice, the shearing and the projection happen in synchrony, however in this section we treat them independently. The shearing is performed by introducing a depth-dependent lateral translation of the image, i.e. as the focus is swept through the sample, the image at the camera is translated laterally along one axis. For our analysis, we constrain this translation to the y axis. This translation is introduced by rotating the galvanometric mirrors in synchrony with the sweeping of the focus. The translation must be performed in a way that does not impart an angular change on the rays incident on the camera: the authors of ref. [5] implemented this using a dual-mirror approach where both mirrors move with equal velocity, a schematic of which is shown in figure 5.1, however they noted is is also possible to perform using only a single galvanometric mirror at the Fourier plane of a $4f$ relay lens system. The result of the depth-dependent image translation is to optically shear the volume: for instance, if a z -stack was acquired in this fashion, with a z -step-size ΔZ and an angular increment imposed by the scan unit of θ_{proj} , through simple geometric arguments each xy plane in the resulting volume would be translated in the y direction by an amount ΔY given by:

$$\Delta Y = \Delta Z \tan(\theta_{proj}) \quad (5.2)$$

compared with the previous xy plane, see figure 5.2 for a schematic of the associated geometry. The number of pixels this translation corresponds to is calculated as $\frac{\Delta Y}{P_{xy}}$, where P_{xy} is the dimension of the camera pixels in object space.




Figure 5.2: Geometry of optically-induced shear in multi-angle projection imaging.

The second step in multi-angle projection image formation is then the projection of the sheared volume into a single 2D image. Again, this projection is performed optically, by stepping the focus through the sample as the galvanometric mirrors are rotated. Multi-angle projection angle image formation can thus be described by the amended model for image formation given by:

$$I(x, y) = \text{Pois} \left\{ \sum_{k=-\frac{N_z}{2}}^{\frac{N_z}{2}} O\left(x, y - k \frac{\Delta Y}{P_{xy}}, z = k\right) \otimes \text{PSF}(x, y, z = k) \right\} + N(x, y) \quad (5.3)$$

where ΔY is the shear increment given by equation 5.2 and N_z is the total number of xy planes in the volume. $N(x, y)$ corresponds to a noise term reflecting background and readout sources, as described in section 1.2, approximated by a Gaussian distribution. Equation 5.3 has been constructed such that regardless of the projection angle θ_{proj} , no shear is introduced to the central xy plane within the sample volume. We shall see that this corresponds to a sample rotation where the axis of rotation is centered in y .

In a multi-angle projection microscope, the shearing and projection are synchronised to occur simultaneously and rapidly, sweeping the entire sample depth within the time-window of a single camera exposure. Therefore, in the absence of noise, the acquired image is equivalent to the summed z -projection of the sheared volume. The effect of projecting a sheared volume in this fashion is understood by considering the shear-warp transform, derived and described in full in [131]. The shear-warp transform demonstrates that the projection acquired from a volume that has been sheared by a plane-wise shear increment ΔY , as defined by equation 5.2, is identical except for a scaling factor to a summed projection of the volume that has been rotated by an angle θ_R . The scaling factor S is given by:

$$S = \frac{1}{\cos \theta_R} \quad (5.4)$$

Therefore, under projection and after correcting for the scale factor, the operations of rotating a volume and shearing a volume yield the same result. The scale factor S imposes physical limits on the angular range that can be achieved by projection imaging, such that the image of the sheared volume does not exceed the lateral extent of the camera sensor. For instance, a projection angle of 70° introduces a stretch factor of almost three along the shear axis; as the projection angle approaches 90° the stretch factor approaches infinity. The exact limitations of the angular range depend on the camera chip dimensions and the size of the sample being imaged: for the samples considered in this chapter an angular range limit of $\pm 60^\circ$ was deemed appropriate, which introduces a maximum stretch factor of two. The volumetric reconstruction from multi-angle projection data discussed in the following sections requires this stretching to be present in the raw data, as the deconvolution attempts to invert the sheared volume image formation rather than the rotated volume image formation. However, in conventional multi-angle projection imaging, the stretching may be computationally removed post-acquisition to return the rotated projection image by means of an scaling operation, implemented via an affine transform defined by the matrix A :




Figure 5.3: Equivalence of shearing and rotation after projection and scaling. (a) shows an example sample volume with isotropic voxel dimensions. (b) shows an xy summed-intensity projection of the volume in (a). (c) shows a summed-intensity projection after rotating the volume by 45° about the x axis. (d) shows an xy summed projection after shearing the volume by a shear-increment of 1 pixel per xy plane. (e) shows the projection in (d) after stretching in the y direction by a scale factor $\cos(45^\circ)^{-1}$. (e) and (c) are equivalent images.

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_R & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (5.5)$$

Figure 5.3 depicts a schematic demonstrating the equivalence between the projections acquired from a rotated volume and a scaled projection of the same volume that has been sheared.

5.2.2 Multi-angle projection reconstruction

In the previous sections we modelled image formation on a multi-angle projection imaging system and discussed the 3D information content contained in two or more angular projections of the same sample. In this section we modify the volumetric reconstruction pipeline developed and applied in the previous chapters for application to multi-angle projection data. To reconstruct a volume from multi-angle projection data, we require at least two projection images to encode the necessary 3D information.

We shall see in later sections, however, that for most samples other than those that are extremely sparse, we generally require a higher number. When reconstructing a 3D volume from the 2D projection data, we are free to decide the number of xy planes we wish to reconstruct, as well as the voxel depth in the reconstructed volume. A convenient choice is often to match the voxel depth with the pixel dimensions in the image, however we are not confined to this choice. Indeed, it often makes sense to choose a larger voxel depth, since the resolution of any single projection image is limited by the system NA and therefore the axial component of the resolution is worse than the lateral, meaning the sampling requirements are different in a 3D image. The total reconstructed depth range should be at least as large as the focus sweep-range in the projection data to ensure the entire sample depth is accounted for in the reconstruction.

To begin considering the volumetric reconstruction, it is convenient to rewrite equation 5.3 using the sifting property of the unit impulse response function to represent the shear introduced [132]. The image of an object xy plane $O(x, y)$ that is translated in the y direction by some amount $\frac{k\Delta Y}{P_{xy}}$, may be written as the discrete convolution given by:

$$O(x, y - \frac{k\Delta Y}{P_{xy}}) = O(x, y) \otimes \delta(x = 0, y = -\frac{k\Delta Y}{P_{xy}}) \quad (5.6)$$

where $\delta(x = 0, y = -\frac{k\Delta Y}{P_{xy}})$ is a 2D function defined by

$$\delta(x = 0, y = -\frac{k\Delta Y}{P_{xy}}) = \begin{cases} 1, & x = 0, y = -\frac{k\Delta Y}{P_{xy}} \\ 0, & \text{otherwise} \end{cases} \quad (5.7)$$

Using the associativity of the convolution operation, the model of image formation can then be written as:

$$I(x, y) = \text{Pois} \left(\sum_{k=-\frac{N_z}{2}}^{\frac{N_z}{2}} \left[(O(x, y, z = k) \otimes PSF(x, y, z = k)) \otimes \delta(x = 0, y = -\frac{k\Delta Y}{P_{xy}}) \right] \right) + N(x, y) \quad (5.8)$$

This amended definition of image formation is beneficial in multi-view reconstruction, since it allows volume reconstruction to be performed while attempting to invert only the second convolution operation rather than both convolutions in equation 5.8. This is advantageous, since by equation 5.6, we see that the reconstruction process may be implemented through a series of 2D image translation operations instead of convolution operations. Image translation operations are computationally less burdensome to implement than convolutions, speeding the reconstruction process. However, it is worth noting that if the reconstruction is implemented in this fashion, the optical blurring

of the PSF is not inverted, since our processing does not consider the first convolution operation in equation 5.8 and therefore the algorithm should no longer be described as a deconvolution in the conventional sense¹. However, the depth information encapsulated in the multiple angular perspectives may still be extracted to reconstruct a full 3D volume.

To reconstruct a 3D volume from multi-angle projection data, the same RL deconvolution framework is used that is effective in reconstructing volumes in 3D-EPM. As with 3D-EPM, this requires a forward operation, that maps the 3D volume to the 2D projection image, and a backward operation, that maps the 2D projection image back to the 3D volume.

The forward operation

The forward operation for the adapted RL volume reconstruction, just like in 3D-EPM, consists of our model for image formation. Since we are, for now, only concerned with inverting the second convolution in equation 5.8, we can define the quantity $O'(x, y, z)$, the object we wish to reconstruct, as:

$$O'(x, y, z) = O(x, y, z) \otimes PSF(x, y, z). \quad (5.9)$$

We can rewrite the forward projection for multi-angle projection imaging as

$$I(x, y) = \sum_{k=-\frac{N_z}{2}}^{\frac{N_z}{2}} O'(x, y, z = k) \otimes \delta(x = 0, y = -\frac{k\Delta Y}{P_{xy}}) \quad (5.10)$$

where the Poisson and Gaussian noise terms have been removed². Since this operation represents a translation in one direction, it is computationally more efficient to implement this as a simple image translation operation instead of a convolution. In this thesis the image translation is implemented using a third-order flux-conserving spline interpolation to calculate the pixel values for an image shifted by a non-integer number of pixels, using functions from the `scipy` and `cupy` python libraries [133]. This shifting of a single 2D frame can be represented by the affine transform matrix:

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -\frac{k\Delta Y}{P_{xy}} \\ 0 & 0 & 1 \end{bmatrix} \quad (5.11)$$

¹since deconvolution typically refers to inverting the blur induced by the system PSF.

²It is not necessary to include the Poisson term in the forward projection of the RL algorithm, since the forward operation should yield the expected (mean) pixel values in the image for a given image formation model and the current estimate to the sample object. See chapter 2 for a full description of the RL algorithm.

The backward operation

Similarly, the backward operation is also implemented by translating a 2D image to map the 2D projection back to the 3D sample volume. In 3D-EPM, the backward projection consisted of convolving the discrepancy ratio between the recorded data and the forward projection with the flipped PSF at every z position across the sample depth, see equation 2.15. The same is true in this case, except instead of convolving with the flipped PSF, we instead use the shifted impulse function $\delta(x = 0, y = k\Delta Y)$, such that the back-projected volume at iteration n and corresponding to the m^{th} projection view B_m^n is formed by the series of 2D operations given by:

$$B_m^n(x, y, z = k) = \frac{I_m(x, y)}{\sum_k \delta(x = 0, y = -\frac{k\Delta Y}{P_{xy}}) \otimes \hat{O}^n(x, y, z = k)} \otimes \delta(x = 0, y = \frac{k\Delta Y}{P_{xy}}) \quad (5.12)$$

where \hat{O}^n is the estimate to O' at iteration n . We see that the denominator of the fraction in equation 5.12 is the forward projection defined above. Just like the forward projection, the second convolution operation in equation 5.12 performs a 1D shift along the y axis, and therefore may also be implemented as a 2D image translation operation which can be represented by the affine transform matrix:

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & \frac{k\Delta Y}{P_{xy}} \\ 0 & 0 & 1 \end{bmatrix} \quad (5.13)$$

The reconstruction algorithm, adapted for multi-angle projection data, from M 2D projections and expressed in terms of convolution operations, is then:

$$\hat{O}(x, y, z = k)^{n+1} = \hat{O}(x, y, z = k)^n \cdot \frac{1}{M} B(x, y, z = k)_m^n \quad (5.14)$$

The deconvolution may also be subject to the total variation regularization term defined in equation 2.16, however the multi-angle reconstruction typically converges faster than 3D-EPM deconvolution and therefore, at least in the high signal-to-noise ratio regime, noise-amplification is not as problematic.

5.2.3 Simulated HeLa cell organelles

We first demonstrate volumetric reconstruction applied to multi-angle projection imaging with synthetic projection data, generated from a 3D confocal image of labelled mitochondria in a HeLa cell. The confocal image serves as a ground-truth volume against which the reconstructed volume can be compared. The ground-truth volume has isotropic voxel dimensions of 64 nm. A projection series was simulated by shearing the volume in y according to equation 5.2 before summing the sheared volume into a

single xy plane. For the initial case, no noise was added to the simulated projection images. The data shown here was simulated by James D. Manton and is used in this thesis with permission. Initially, five projection views were selected for reconstruction, at angles of -60° , -30° , 0° , $+30^\circ$, and $+60^\circ$ with respect to the ground-truth volume. Figures 5.4a-e display the raw 2D projections used for reconstruction, where the apparent rotation of the sample around the x axis is evident, as well as the shear-induced stretching in y . The sample volume was then reconstructed according to equation 5.14 using 25 reconstruction iterations. Figures 5.4f-h show maximum-intensity projections through the reconstructed volume along each axis, where the 3D structure of the mitochondrial networks is apparent. Figures 5.4i-k show the same projections through the ground-truth volume, where we see a high degree of similarity between the five-view reconstruction and the confocal stack. This simulation demonstrates how it is possible to perform 3D reconstruction using the same mathematical framework developed for 3D-EPM, using multi-angle projection data. In this example, only five 2D projection images were used to perform the 3D reconstruction, compared to 100 xy planes in the ground-truth data stack, representing a possible $20\times$ speedup in acquisition time for the full 3D volume.

5.2.4 Varying the number of angular projections used for reconstruction

The reconstruction displayed in figure 5.4 was performed using five raw projection images. However, it has not yet been discussed how this number is determined, or the effect on the reconstruction if a different number are used. Obviously, the number used should be far less than the number of xy slices in a conventionally-acquired z -stack for there to be some benefit in temporal resolution. The number should be as low as possible, since for sequentially acquired projections, the temporal resolution drops linearly with the number needed for reconstruction. While two projections is the theoretical minimum needed to encode the 3D position of a single point, we find that for most samples this is not enough to perform a good reconstruction, with ghost artefacts and elongation of image features being introduced into the reconstructed volume. Ghost artefacts are introduced into the solution when there is ambiguity about the position of a particular feature and the RL deconvolution algorithm is unable to distinguish between possible solutions. An example of such ambiguity is illustrated in figure 5.5, where a two-view reconstruction of a simulated sample volume containing two point-sources introduces ghost artefacts into the reconstructed volume. Figures 5.5a and 5.5b display xy and yz maximum-intensity projection of the ground-truth sample volume, which is then projected at angles of $\pm 45^\circ$. The projection angles are indicated with the green and magenta dashed lines in figure 5.5b, and the resulting projections are displayed in 5.5c and d. Figures 5.5e and 5.5f display xy and yz maximum-intensity projections through the reconstructed sample volume using only these two angular pro-




Figure 5.4: Simulated multi-angle projection imaging and volume reconstruction of mitochondria in a HeLa cell. (a-e) Simulated multi-angle projection images from real confocal data used for reconstruction, at angles ranging from -60° to $+60^\circ$. (f-h) Maximum intensity projections along each axis in the reconstructed volume after 25 volume-reconstruction iterations. (i-k) same projections as (f-h) but through the ground-truth data. Scale bar is $10\ \mu\text{m}$. Input projection data was simulated by James D. Manton from a ground-truth confocal z -stack and used here with permission. The reconstruction was performed by myself according to equation 5.14.

jections. We see that ghost artefacts, or repeated sample features, are introduced into the reconstruction. These are the result of ambiguity in the problem: either pair of beads indicated by the cyan and light-green circles could form the projections displayed in figures 5.5c-d, and the deconvolution is therefore unable to correctly reconstruct the sample. However, we see in figures 5.5g-h that by including a third projection angle (indicated by the white line in 5.5b), this ambiguity is completely removed, and the reconstruction is able to reconstruct the correct sample volume.

To further explore the influence of the number of projections used to reconstruct, the mitochondria sample reconstructed in section 5.2.3 was again reconstructed while the number of projections used was varied. Figure 5.6 displays xy maximum-intensity projections of the reconstructed volume after 25 iterations using two, three and five projection views (top, l-r), as well as a plot of the convergence of the solutions with increasing iterations. We see in the two-view solution that the reconstructed volume features lots of artefacts; features appear elongated and the structure of the sample is not well-resolved. Here, two angular-projections (-60° and $+60^\circ$) did not provide enough 3D information for the reconstruction to perform well. In the three-view so-




Figure 5.5: Ghost artefacts may be introduced into the reconstructed volume when reconstructing with only two projections because of ambiguity in the inverse problem. (a) and (b) show xy and yz maximum-intensity projection through simulated volume containing two point-sources. (c) and (d) show the projections that would be recorded if the sample was projected at angles of 45° and -45° . The projection angles are indicated by the magenta and green dashed lines in (b). (e) and (f) show maximum-intensity projection through the reconstructed sample volume if only these two projections are used to reconstruct. Ghost features are introduced because of ambiguity in the problem: either pair of points indicated by the green and cyan circles in (f) would produce the projections in (c) and (d), and the deconvolution is therefore unable to reconstruct the correct sample. (g) and (h) show the reconstructed volume when an additional projection is used (0° , indicated by the white dashed line in (b)). Including a third projection completely removes this ambiguity. This figure was inspired by helpful discussion with James D. Manton who pointed out this issue.

lution, the sample structure is much-better resolved and the reconstruction features far fewer artefacts. The solution converges faster than the two-view solution and the mean-square error after 25 iterations is lower, representing a solution that is closer to the ground-truth. The extra 3D information provided by introducing the third angular projection (0°), therefore, significantly improves the performance of the deconvolution. With five views (-60° , -30° , 0° , $+30^\circ$ and $+60^\circ$), the solution improves further still, though the improvement is not as significant as that seen between two and three projection views. For most samples that are well-suited for volumetric reconstruction from multi-angle projection imaging, we find five projections across a similar angular range to that used here to be a good compromise between imaging speed and reconstruction quality: generally, not much benefit is seen by using a greater number of projections. However, as we will see, some samples are fundamentally not well suited for reconstruction with this method. Similar to the limitations of 3D-EPM, projection images of samples that are particularly dense, or those that feature large regions of texture-free uniform fluorescence emission, do not reconstruct well. Samples that are





Figure 5.6: Reconstruction performance with increasing numbers of projection angles. Scale bar is $10\ \mu\text{m}$.

best suited to this method, therefore, are those that are reasonably sparse (such as the mitochondria shown in figure 5.4). An example of a sample that is too dense for reconstructing with this method is provided in figure 5.7. The sample shown here is chromatin, imaged in the same cell as the mitochondria shown in figure 5.4. Figure 5.7 displays maximum-intensity projections through the 3D ground-truth volume, as well as the same projections through the reconstructed sample volume using both five and ten projections to reconstruct. Even when using ten projections, this sample did not reconstruct well: while the general shape of the sample is apparent, the finer sample structure is not well resolved in the reconstruction.

5.3 OPM Multi-angle projection imaging

Next, we demonstrate the reconstruction of multi-angle projection images with experimental data, acquired on an oblique-plane microscopy (OPM) system. The imaging modality used to acquire this data introduces additional complexity into the volume reconstruction pipeline, owing to the geometry of an OPM. Processing multi-angle projection data that has been acquired in this fashion, therefore, requires us to take a closer look at the geometry of an OPM system.

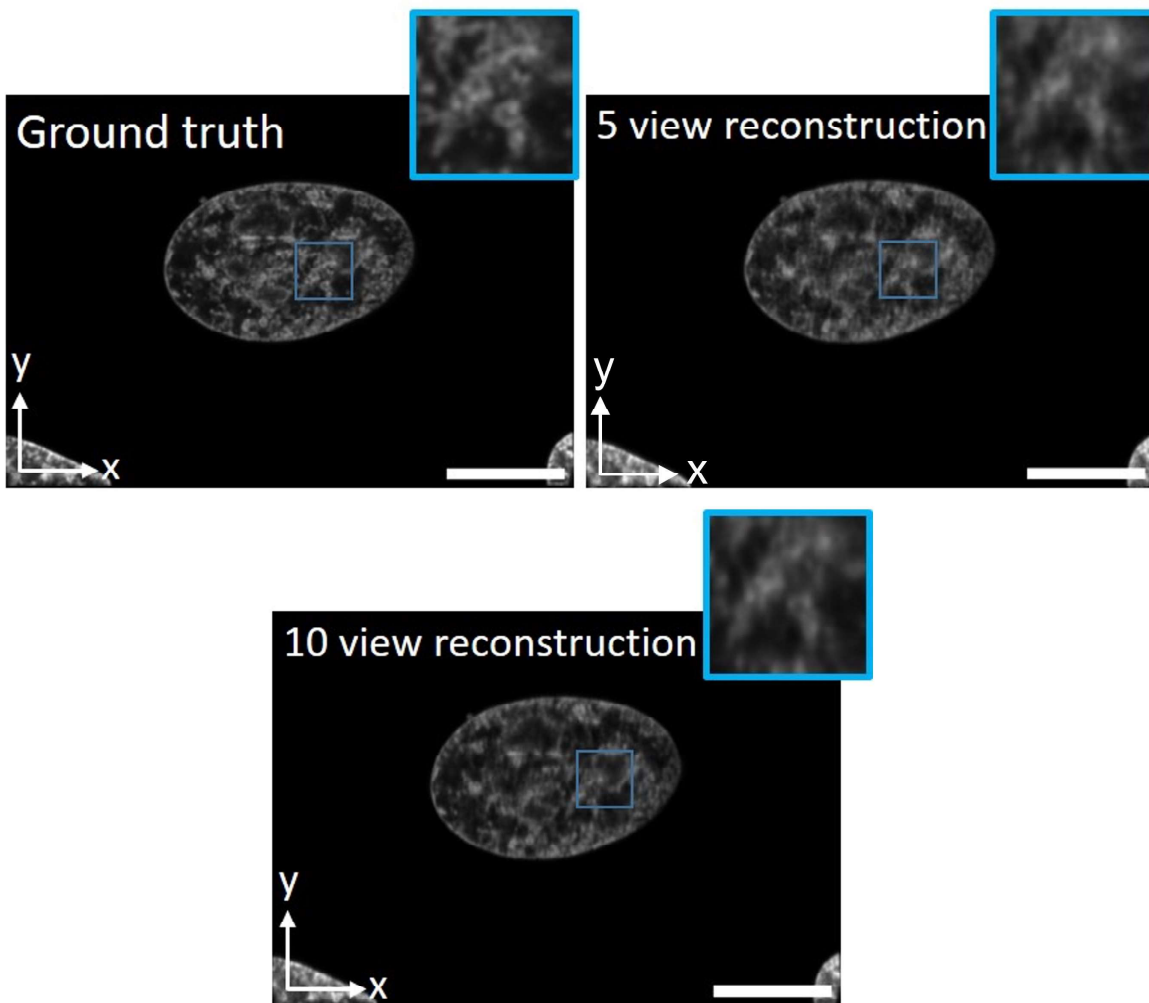


Figure 5.7: xy planes from a 3D reconstruction of simulated multi-angle projection images where the sample is too dense to reconstruct effectively. The sample is chromatin imaged in the same cell as shown in figure 5.4 and the raw data was also provided by James D. Manton. Top-left panel shows a single xy plane in the ground-truth image. Top-right panel shows the same plane in a volume reconstructed from five angular projections, and the bottom panel shows the same plane in a ten-view reconstruction. The finer sample features are not well-resolved in either the five or ten view reconstruction. Scale-bar is $10\ \mu\text{m}$.

5.3.1 OPM Geometry

OPM is a form of light-sheet imaging that overcomes some of the limitations of conventional light-sheet microscopy [33]. In OPM, the light-sheet is delivered to the sample at an oblique angle, launched from the same primary objective lens that is used to image the sample, thereby removing the requirement to have two objective lenses arranged perpendicular to one another. This allows OPM to image samples, such as those mounted on glass slides, that could not be imaged with conventional light-sheet microscopy due to constraints on how samples must be mounted.

OPM relies on the principle of remote refocusing to form an intermediate 3D replica of the fluorescence collected by the primary imaging objective that is free of spherical aberrations [134]. The remote-refocused image is formed by de-magnifying the image




Figure 5.8: Schematic of an oblique plane microscope (not to scale). The light sheet is delivered to the sample through the primary imaging objective at an angle θ from the y axis in the normal coordinate system, Y_N . The first microscope (M1) forms an image of the illuminated oblique plane, which is then demagnified by the second microscope (M2) to form a remote image at the focus of objective 2. A third microscope (M3), aligned along Z_{opm} , then re-magnifies the remote image onto a detector placed at the final image plane. To acquire a 3D image the sample is scanned along Y_N during acquisition.

formed at the native image plane with a second microscope, see figure 5.8. Finally, a third microscope, tilted from the optical axis shared by the primary and secondary objectives by an angle θ , is then used to re-magnify and image the remote volume onto a detector. Unlike conventional light-sheet imaging, in OPM the light-sheet is delivered at an angle that is not perpendicular to the optical axis of the primary imaging objectives. Instead, the light-sheet is launched along the Y_{opm} direction, tilted by the angle θ from the conventional launch direction Y_N . 3D images are then acquired by stacking sequential 2D acquisitions while scanning along Y_N . The scanning is implemented by translating the sample through the light-sheet, or by moving the light-sheet across the FoV, or via a hybrid approach [39]. A schematic of the geometry of an OPM system at the sample is shown in figure 5.9.

The optical axis of the OPM, Z_{opm} is rotated from optical axis of the primary objective lens Z_N by the angle θ_{opm} . However, volumetric imaging is performed by sequentially acquiring images while translating the sample or light-sheet laterally in the normal coordinate system along Y_N . Because the normal coordinate system is rotated from the OPM coordinate system, this means a component of the scan direction lies along Z_{opm} . The resulting OPM volume, therefore, is sheared laterally in one direction [33]. In a similar fashion to the shearing imposed by the multi-angle projection scan unit, this means that each 2D xy slice in the acquired volume is laterally translated along the shear axis with respect to the previous slice. For an OPM system with the angle θ_{opm} defined as shown in figure 5.9 and a volume acquired by scanning along the Y_N direction with a step-size ΔY_N , the resulting step-size along the OPM optical axis, ΔZ_{opm} is given by:

$$\Delta Z_{opm} = \Delta Y_N \sin \theta_{opm} \quad (5.15)$$




Figure 5.9: Schematic illustrating the geometry of an opm system. Not to scale. The light sheet (blue) is delivered through the primary objective at an oblique angle, perpendicular to the optical axis of the OPM (Z_{opm}), which is rotated from the “normal” optical axis of the primary objective Z_N by an angle θ . Volumes are acquired by scanning the sample or light-sheet across Y_N , which imposes an optical shear in the OPM coordinate system.

and the shear introduced along Y_{OPM} is given by:

$$\Delta Y_{opm} = \Delta Y_N \cos \theta_{opm} \quad (5.16)$$

Therefore, if we wish to return a volume acquired using an OPM system to the (normal) coordinate system of the primary objective, the OPM-induced shear must be removed computationally. This is normally performed with a simple linear translation operation, shifting each plane in the acquired volume by increasing increments of $-\frac{\Delta Y_{opm}}{P_{xy}}$ [33]. What happens, however, if the OPM is converted into a multi-angle projection microscope?

OPM multi-angle projection imaging

We have seen how both OPM and the multi-angle projection imaging modality impose lateral depth-dependent shears along one axis during volumetric acquisition. It follows, then, that if the scan-unit is aligned such that the projection shear is imposed along the same axis as the OPM shear, the two shears are combined. Indeed, the authors of Ref [5] noted that the scan unit could be used to optically remove the OPM-imposed shear. However, for volumetric reconstruction of multi-angle OPM projections we require more than one angular perspective of the sample and therefore we need to calculate the total shear imposed during imaging. The total lateral shear, ΔY_{tot} , imposed during volumetric acquisition is calculated as the sum of the shear imposed by the OPM, ΔY_{opm} and that imposed by the projection scan-unit, ΔY_{proj} :

$$\Delta Y_{tot} = \Delta Y_{opm} + \Delta Y_{proj} \quad (5.17)$$

The volumetric reconstruction from M multi-angle projections acquired on an OPM system can therefore be implemented, as before, according to equation 5.14, replacing

ΔY with ΔY_{tot} to reflect the additional shear introduced by the OPM.

5.3.2 Reconstruction of experimental multi-angle projection data

In this section we apply the theory developed in the previous section to reconstruct experimentally-acquired multi-angle OPM projections. Multi-angle projection images of the fluorescently-labelled vasculature and blood-cells in the tail of a fixed zebrafish were acquired and used for reconstruction. The raw projection images were acquired by Bingying Chen and used here with permission. The OPM used for data acquisition featured a tertiary arm tilted by $\theta_{opm} = 45^\circ$. Alongside the raw projection data, an OPM z -stack was acquired in the conventional fashion, which was computationally de-sheared back to the normal coordinate system and used as a ground-truth image, against which the reconstructed volume could be compared. The ground-truth z -stack image was acquired with a step size of 500 nm, which, when de-sheared, gives a Z_N step size of 354 nm. The OPM used for acquisition had a primary NA of 1.1. Figure 5.10a shows maximum-intensity projections through the ground-truth image of the vasculature, where the 3D structure of a section of the dorsal aorta and intersegmental vessels is clear. The zebrafish was oriented approximately vertically, i.e. the dorsal aorta is roughly parallel to the X_N axis. Figure 5.10b shows the five angular projections of the vasculature used to reconstruct the sample. The angular projections were acquired as described in section 5.2 and 5.4: for each projection, the focus was rapidly swept through the sample during the course of a single camera exposure while the scan-unit introduced the depth-dependent shearing in synchrony with the focus sweep. For the reconstruction shown here, five projections were used at angles of -60° , -30° , 0° , 30° , and 60° . Figure 5.10c shows the same projections as 5.10a but through the reconstructed volume, after 25 iterations of equation 5.14. In the reconstruction, the 3D nature of the sample is clear and the reconstructed volume appears of similar quality to the ground-truth volume. Figure 5.10d shows maximum-intensity projections of the ground-truth image of the blood cells, and figure 5.10e shows the same projections through the reconstructed blood cell volume. The blood cell volume was also reconstructed from five angular projections acquired at the same angles as the vasculature. Figure 5.10f shows a composite xy maximum-intensity projection of both reconstructed volumes, where the position of the blood cells is clearly in agreement with that of the blood vessels, showing good agreement between the two reconstructed volumes. A composite 3D rendering of the reconstructed volumes is shown in figure 5.10g, emphasizing the 3D dimensionality of the reconstructed volumes, compared with the raw 2D projection data.




Figure 5.10: Demonstration of multi-angle projection volume-reconstruction on experimental images of a biological sample, exhibiting extended regions of fluorescence emission. Figure shows the reconstruction of fluorescence projection images of vasculature (yellow-green) and blood cells (red) in the same region of the tail of a fixed zebrafish. (a) and (d): maximum-intensity projections along each axis in a light-sheet ground truth image of the vasculature and blood-cells respectively. (b) shows the five raw projection images used for reconstruction of the vasculature. (c) and (e) show the same projections as (a) and (d) but through the reconstructed 3D volumes after 25 reconstruction iterations. (f) a composite xy projection through the reconstructed volumes confirming agreement between the two channels. (g) composite 3D rendering of the reconstructed volumes. Raw data acquired by Bingying Chen and used here with permission. Scale-bar is $10\ \mu\text{m}$.

5.3.3 Live-cell multi-angle projection reconstruction

A further experimental example of volumetric reconstruction from multi-angle projection data is provided in figure 5.11. OPM multi-angle projection images, along with ground-truth z -stacks, were acquired of single-cells from the COS7 line (labelled mitochondria, blue), and A375 line (labelled tractin, orange). Data was acquired by Bo-jui Chang and used here with permission. Figure 5.11a-e shows five angular projections of the mitochondria used for reconstruction, at angles ranging from -60° to 40° . Figures 5.11f-h displays maximum-intensity projections along each axis through the reconstructed volume, and figures 5.11i-k show the same projections through the computationally de-sheared ground-truth z -stack. The high degree of similarity between the reconstruction and ground-truth confirms that mitochondria projection images are well-suited to volumetric reconstruction. Conversely, the tractin label was too dense to properly reconstruct: figure 5.11l-n shows maximum-intensity projections through the reconstructed volume. While the general shape of the cell is apparent, some of the finer




Figure 5.11: Multi-angle volume reconstruction of OPM images of organelles in live cells, confirming the high-resolution capability of the method and its feasibility for live-cell imaging. Raw data acquired by Bo-Jui Chang and used here with permission. (a-e): raw multi-angle projection data of labelled mitochondria used for reconstruction. (f-h) maximum-intensity projections through the reconstructed volume. (i-k): same projections as (f-h) but through a ground-truth z -stack image that has been computationally de-sheared. (l-m) maximum-intensity projection through a reconstructed volume from multi-angle projection images of a A375 cell with labelled tractin. (o-q) maximum-intensity projections through the de-sheared ground-truth image of the same sample. Scale-bar is $20\ \mu\text{m}$.

structure visible in the ground-truth volume (figure 5.11o-q) is not as apparent in the reconstruction. This is particularly apparent in the zy projection and is a result of the fluorescence emission being largely uniform across large areas of the sample, leading to little disparity in the raw multi-angle projection data.

5.3.4 Discussion

So far in this chapter we have demonstrated the ability to perform volumetric reconstruction from both simulated and experimental multi-angle projection data, acquired via a number of different imaging modalities. The fact that the multi-angle scan unit can be easily installed on *any* optical-sectioning-capable microscope means that the method of image acquisition can be tailored to suit the needs of a particular experiment. For instance, for live-cell imaging, where extremely high resolution is required, a confocal imaging system may be the optimal choice, whereas for imaging at the multi-cell to whole organism scale, a light-sheet imaging modality may be better suited. Volume reconstruction from multi-angle projection data represents a general method

where, providing the sample features a reasonable degree of sparsity, the exact choice of image acquisition method is unimportant. Multi-angle projection data shares the spatial resolution of the parent imaging system [5]: just as in 3D-EPM, no sacrifice in resolution is required for encoding 3D information and performing volumetric reconstruction, again offering significant advantages over LFM.

Performing 3D reconstruction from sequences of angularly-separated 2D projection images underpins several imaging techniques, perhaps most notably in computed tomography (CT) imaging [135]. In CT imaging, a sequence of x-ray images is acquired at varying projection angles, commonly implemented by rotating an x-ray tube and detector array around a stationary sample. Tomographic projection image formation more generally is described mathematically by the Radon transform [136], with the resulting projection data referred to as a sinogram. Reconstructing the 3D sample from an acquired sinogram, therefore, requires inversion of the Radon transform. This is commonly approached via a filtered back-projection operation [136], or through application of more complex iterative reconstruction algorithms [137, 138]. Such methods could, in principle, be applied to the fluorescence projection images acquired via multi-angle projection imaging, with appropriate consideration made for the differences in which contrast is achieved in fluorescence imaging compared with x-ray imaging. However, these methods generally require many projection images to accurately reconstruct the sample volume, with CT datasets often comprising several tens, hundreds or even thousands of images [139]. We have seen in this chapter how, in principle, depth information may be encoded into far fewer projections of sparse fluorescent samples. Indeed, the authors of [5] noted that just two projection images of the same sample acquired from different projections does, in principle, encode 3D information and permit 3D reconstruction. However, this concept was initially demonstrated with the 3D localisation of sparse, point-like nano-particles in a fixed cell via a simple triangulation approach [5], a method that is not applicable to the more general 3D reconstruction of extended fluorescent structure. To the best of the authors knowledge, the results presented in this chapter represent the first full 3D reconstruction of fluorescence images acquired via real-time multi-angle projection imaging. Additionally, the reconstruction methods presented here are applicable to both extended fluorescent structure as well as more sparse, point-like samples, while requiring fewer projection images than other approaches used to reconstruct tomographic image data.

The temporal resolution of volumetric imaging performed by reconstructing multi-angle projection data is limited by how quickly the necessary projections can be acquired. Acquiring a single projection image requires the focus to be swept through the entire sample depth, or vice versa. If this is implemented mechanically, e.g. with a motorised stage scan, this process is likely to be slow. However, if the focus scanning is implemented by other rapid means, the temporal resolution is significantly improved. For instance, in OPM, the focus-sweeping may be performed by scanning the light-sheet across the sample with a galvanometric mirror instead of translating the sample

with a motorised stage. In this case, the entire sample depth can be imaged extremely quickly (scan-rates of several kHz are possible), and the time required to acquire a single projection becomes limited in principle only by the frame-rate of the camera. In practice, the achievable signal-to-noise ratio may further limit the possible acquisition speeds: in fluorescence imaging of live biological samples this is often determined by the sensitivity of the sample to the laser excitation. For volumetric reconstruction, the signal-to-noise ratio should be kept as high as possible to avoid the noise amplification phenomena discussed in section 1.2. Despite these considerations, the results presented in this chapter still represent a major improvement in volumetric acquisition speeds. For instance, in the OPM imaging of the zebrafish vasculature and blood-cells displayed in figure 5.10, the ground-truth z -stack images comprised 300 stacked 2D acquisitions. The volume reconstruction, on the other hand, required only five 2D acquisitions, acquired sequentially at a speed limited by the camera frame-rate. For this particular example, therefore, the process of performing volumetric imaging by means of reconstructing multi-angle projection data represents an improvement in acquisition speed of approximately $60\times$. However, if enough beam-splitters and galvanometric mirrors are available, and for high-enough photon budgets, the five projection images needed for volumetric reconstruction could, in principle, be acquired simultaneously. This would mean that similar to the 3D-EPM method developed in chapter 3, the volume-reconstruction from multi-angle projection imaging could also be implemented as a snapshot volume-imaging technique.

The improvement in temporal resolution offered by performing volumetric reconstruction from multi-angle projection data has significant implications in the field of bio-imaging. The temporal resolution of volumetric imaging methods such as light-sheet and confocal microscopy has thus far prevented their application to the 3D study of dynamic samples where the dynamics occur over millisecond timescales, as the time needed to acquire a volume is too great. However, by decreasing the time needed to acquire a volume through combining projection-imaging with the volume reconstruction pipeline developed in this thesis, the possibility of applying these methods to highly dynamic 3D samples is opened, potentially offering 3D insight into biological events that have previously only been possible to directly image in 2D. We therefore expect the volumetric imaging methods developed here to find broad applicability across multiple scales, including the 3D study of cellular morphology, organelle dynamics and interactions, and calcium signalling.

The ability to perform volumetric reconstruction from multi-angle projection images as well as PSF-engineered images demonstrates the generality of the mathematical framework developed in chapter 2 of this thesis. The optimal choice of projection imaging modality depends on the requirements of a particular experiment or sample. For instance, if the sample moves extremely rapidly, where snapshot imaging is necessary to properly capture the observed dynamics, 3D-EPM is better suited, as the temporal resolution is in principle limited only by the frame-rate of the camera. On the other

hand, for samples that are deeper than the EDoF provided by an engineered PSF, multi-angle projection imaging modalities are better suited, where the depth-range projected over can be chosen to meet the thickness of the sample.

At this point, we turn our attention back to the challenges set out at the start of this chapter. For volumetric reconstruction from projection data, samples that are very dense will, unfortunately, always present a challenge for the type of volume reconstruction developed in this thesis, since the reconstruction depends on extracting 3D information from disparity between projection images. Samples that feature large areas of texture-free, uniform fluorescence emission will not work with the methods presented here (although there is scope to artificially introduce texture onto such samples with structured illumination). But what about the other two challenges? The first challenge, imaging samples that are deeper than the EDoF provided by the cubic PSF is solved by using multi-angle projection imaging, where the scan-range can be adjusted to capture the entire sample depth. However, we have not yet addressed the other challenge of reconstructed volumes suffering from anisotropic spatial resolution.

In both the 3D-EPM and multi-angle projection imaging scenarios, the resolution of the reconstructed volumes will, at best, be determined by the resolution of the parent imaging system. This means that the lateral resolution will always exceed the axial resolution in any reconstruction. For extremely high NA objective lenses, the axial resolution approaches (but does not equal) the lateral resolution, however at lower NA the anisotropy gets worse. This is because the lateral resolution is proportional to $\frac{1}{\text{NA}}$, whereas the axial resolution is proportional to $\frac{1}{\text{NA}^2}$. This resolution anisotropy may be undesirable for certain samples. This raises the question: is there anything we can do to improve the axial resolution, such that the resolution in the reconstructed volume becomes isotropic? In the final section of this chapter, an imaging modality is proposed that solves this problem, offering isotropic spatial resolution across all three spatial dimensions.

5.4 Tomographic 3D-EPM

When imaging under normal diffraction-limited conditions (i.e. without the aid of any “super-resolution” imaging techniques), the resolution of acquired 3D images is always worse axially than it is laterally, because of the limited collection angle possible through a single objective lens. However, the effective collection angle may be increased if information is fused together from two or more objective lenses imaging the sample at different angles. This concept is exploited in imaging modalities such as 4-pi microscopy [140], which features two opposing objective lenses along the same optical axis, or dual/multi-view light-sheet microscopy [141, 78, 142]. If instead imaging is performed simultaneously along orthogonal axes, it is possible to fuse images during image processing to yield a deconvolved image featuring isotropic 3D spatial resolution. This has been demonstrated with several different imaging modalities including




Figure 5.12: Schematic of tomographic EPM setup. The sample is probed by two objectives simultaneously, arranged orthogonally to one another such that their focal planes intersect. The sample is mounted such that it can be probed from multiple directions simultaneously.

light-sheet imaging and LFM [50]. Isotropic 3D spatial resolution is possible in these imaging modalities because of the orthogonality between acquired images: the z -axis of one image becomes the x -axis in the other. However, while such techniques do provide isotropic spatial resolution, they still suffer some limitations of their respective related techniques. Multi-view light-sheet imaging, for instance, still requires the sample to be scanned relative to the illumination, meaning the method is unsuited for highly dynamic scenes. Dual-view LFM, on the other hand, requires a compromise in spatial resolution compared to that of the parent imaging system. In chapter 3 we saw how these drawbacks may be avoided with 3D-EPM, which enables snapshot volumetric imaging without requiring a compromise in spatial resolution. Is it possible, therefore, to combine the benefits of 3D-EPM with those of orthogonal dual-view image acquisition, to provide snapshot volume imaging with an isotropic 3D resolution?

Proposed design

Figure 5.12 displays a schematic of part of a proposed design for an experimental setup that enables orthogonal PSF-engineered images to be acquired simultaneously. The design images the sample through two objective lenses that are mounted orthogonally to one another. To achieve isotropic resolution in a reconstructed 3D image, the objective lenses should feature the same NA. The sample is centered on a point that coincides with the front focal plane of both objectives. Behind each objective is an optical train similar to that shown in figure 1.10: a tube lens forms an image at an intermediate image plane, which is then re-imaged with a $4f$ relay, allowing a phase-mask

to be introduced at a plane conjugate to the exit pupil. The objective arrangement is analogous to that seen in light-sheet microscopy, however in this case both objectives are used for imaging. The dual-objective imaging modality described here, where the objectives are arranged orthogonally, is referred to as tomographic 3D-EPM. Figure 5.12 displays the illumination being delivered through only one objective in an epi-illumination fashion, however the design is not limited to this choice: delivering the illumination through both objectives may be advantageous in countering shadow artefacts. Alternatively, selective-volume illumination could be implemented to reduce the contributions from out-of-focus fluorescence [50]. While there remains flexibility in the design of such a tomographic experiment, the optical setup is subject to the following additional constraints compared to 3D-EPM:

- The magnification must be chosen such that, in addition to properly sampling the image, the FoV matches the EDoF provided by engineered PSF. This is because the volume imaged by both cameras must be the same.
- The sample must be mounted such that it can be probed from multiple angles simultaneously. This means that conventional sample-mounting methods, such as on glass-slides, are not suitable for tomographic 3D-EPM. Instead, mounting protocols common in light-sheet imaging, such as suspending the sample in an agarose gel inside optically transparent fluorinated ethylene-propylene tubing, should be used.
- The sample must be placed in the focal plane of both objective lenses, which restricts the type of lens that can be used with 3D-EPM. For instance, high NA oil-immersion lenses, which require the sample to be mounted within a few-hundred microns of the front face of the objective lens, are not suited for tomographic 3D-EPM. Instead, lenses with a longer working distance must be used.

Deconvolution algorithm

Volumes may be reconstructed from data acquired on an experimental setup such as that shown in figure 5.12 in the same way as in 3D-EPM, with the only required modification being to reflect the orthogonal perspectives of the two images. This modification is implemented by rotating the volume through 90° prior to calculating the forward projection for the second snapshot image each iteration. All other aspects of the deconvolution algorithm, including refining the PSFs and calculation of the edge corrections, are identical to those implemented in 3D-EPM.

Proof of concept results

To demonstrate the potential of tomographic EPM to provide isotropic 3D resolution, a simulation was performed where a volume containing two point-sources was reconstructed from snapshot EPM images. The simulation was performed at 0.5NA with

$\alpha = 3$ cubic PSFs. The sample volume comprised two point-sources, separated axially by $1.25\ \mu\text{m}$; a separation distance far beneath the axial resolution limit of a 0.5NA objective (which from equation 3.19, we see to be $4.2\ \mu\text{m}$ for an emission wavelength of $0.525\ \mu\text{m}$.) Therefore, if the sample was imaged via a widefield technique through a single 0.5NA objective, we expect the two points to not be axially resolved. The simulation was performed in two distinct imaging scenarios: firstly, the synthetic sample was imaged via a differential-airy system PSF, implementing 3D-EPM as described in chapter 3, using only a single imaging objective. The simulation was then repeated, where the sample was imaged via two objectives arranged orthogonally as shown in figure 5.12.

Figures 5.13a and 5.13b show the raw images that would be acquired via a single-objective differential-Airy PSF as described in chapter 3, with a differential defocus of $\pm 7\ \mu\text{m}$. From the raw xy images it is not apparent that there are two distinct points being imaged. Figure 5.13c displays a zx maximum-intensity projection through the reconstructed volume after 25 deconvolution iterations. The inset in figure 5.12c shows a zoomed-in version of the region bound by the cyan box, and the plot shows a z profile along the dashed line. As expected, the two points are not resolved in the reconstructed volume: the profile confirms only one point is apparent in the reconstruction. Figure 5.13 displays the snapshot images that would be acquired in tomographic fashion. While the raw images are still 2D, in this imaging configuration one camera records an xy image, while the other camera records an zy image. That is to say, the z axis of the first camera becomes the x axis of the second. Figure 5.13f displays the reconstructed volume from the tomographic imaging configuration, with the inset and plot showing the same regions as in figure 5.13c. Here, the two points are clearly resolved in the reconstruction, which is confirmed by the plot along the blue dashed lines. We see, therefore, that by acquiring images from orthogonal perspectives, the resolution along the axial dimension is increased.

5.5 Chapter 5 summary

In this chapter we have seen how the deconvolution methodology developed in chapter 2 to that enable the full 3D reconstruction of a volume from multiple projection images represents a general framework, applicable to a multitude of projection imaging modalities. In chapter 3 we applied the reconstruction framework to simultaneously-acquired PSF-engineered images, enabling snapshot volumetric imaging of samples exhibiting extended regions of fluorescence emission. However, there were a number of limitations of the method, including the depth over which we were able to image, and the difference between the axial and lateral resolution in the reconstructed volumes. In the current chapter we attempted to overcome these limitations by exploring the application of the volume-reconstruction framework to other projection-imaging modalities. In particular, we have applied the reconstruction to two distinct imaging methods, each of which

offers different advantages. Firstly, we applied the volume-reconstruction pipeline to data acquired via multi-angle projection imaging, where the sample is probed rapidly from multiple different projection angles. By employing this imaging modality, the depth-range over which we can image is increased beyond the EDoF provided by an engineered PSF, while still offering volumetric imaging at a significantly improved temporal resolution compared with conventional imaging modalities such as light-sheet or confocal microscopy. We explored the performance of the method with varying numbers of projection angles, and furthered the mathematical framework to allow its application to multi-angle projection images acquired on an oblique-plane light-sheet microscope. We also explored its limitations in terms of the density of the sample that the method could be applied to. We then saw how the same volume-reconstruction framework could be applied to a tomographic implementation of 3D-EPM, where the sample is imaged simultaneously through two orthogonal objectives. In this implementation an isotropic 3D resolution is possible, overcoming the resolution anisotropy seen in 3D-EPM reconstructions and increasing the axial resolution beyond the native resolution provided by a single-objective widefield system.

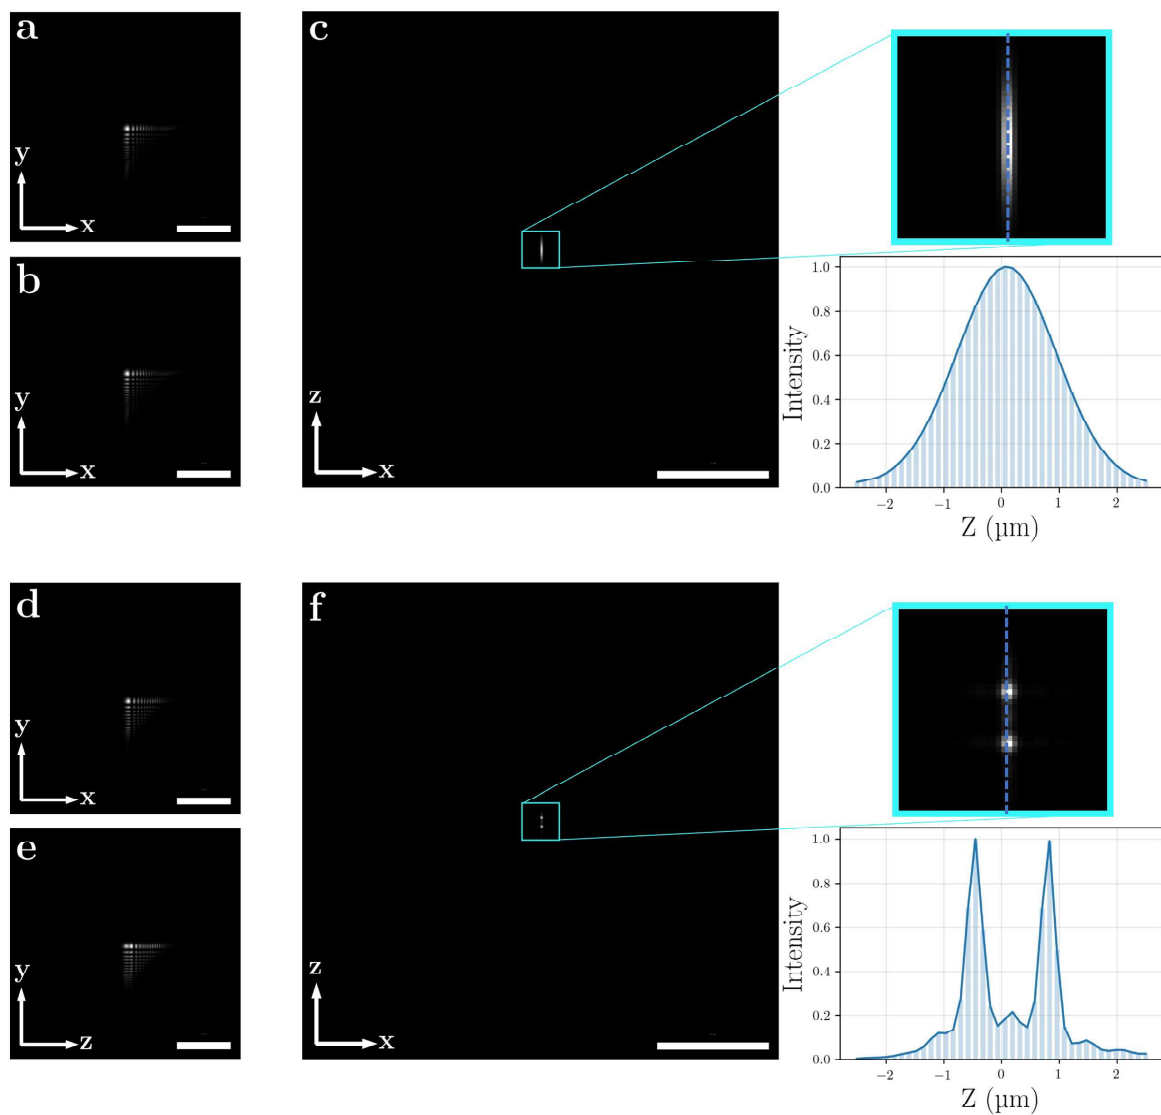


Figure 5.13: Proof of concept results for tomographic EPM. (a) and (b) show simulated differential-airy images of a volume containing a pair of point-sources, separated axially by a distance of $1.25 \mu\text{m}$. (c) the reconstructed volume after 25 deconvolution iterations. The points are not resolved in the reconstructed volume. (d-f) show the same as (a-c), but where the experiment was implemented via tomographic EPM. By imaging from two orthogonal angles simultaneously we are able to resolve the points in the reconstructed volume. Scale-bar is $15 \mu\text{m}$.

Chapter 6

Conclusion and future Work

Computational post-processing has become an integral part of modern fluorescence microscopy. In particular, deconvolution has become a staple tool, not only in its conventional application of improving image contrast, but also in the restoration of fluorescence images acquired via unconventional methods. Broadly termed hybrid imaging, this approach of acquiring intermediate images using non-conventional optics that appear inferior when compared with conventional images, but that contain additional information that may be extracted through post-processing to restore an image of superior quality, is now commonplace in fluorescence imaging. Its use spans applications ranging from light-sheet microscopy, where non-conventional beam profiles are used to improve various aspects of imaging performance, to single-molecule imaging, where post-processing enables precise localisation of fluorophores from PSF-engineered images, yielding reconstructed images with a resolution surpassing the optical diffraction limit.

We saw in chapter 1 how a particular limitation of fluorescence microscopy is its inability to perform 3D imaging while maintaining high resolution in both the temporal and spatial domains. This is significant, since the structure and dynamics of biological samples is often 3-dimensional. The focus of this thesis has been on addressing this limitation, through the application of novel computational processing pipelines that are based on conventional deconvolution frameworks. In chapter 3, we sought to address the question: is it possible to image an entire volume in a single snapshot, while maintaining high spatial resolution? In answering this question, we saw how PSF engineering techniques could be used to (i) extend the DoF of an imaging system, such that in-focus light can be collected from a larger depth-range, and (ii) encode the depth of sample features into images. By modelling image formation as a projection across the EDoF, we demonstrated how the RL deconvolution algorithm could be amended to map between a 3D sample and 2D projection images, allowing full 3D reconstruction of the sample volume from snapshot 2D data. This microscopy technique, which we named 3D-EPM, introduces the following key advantages:

1. Unlike other popular snapshot volume imaging techniques, 3D-EPM does not

inherently require a compromise in the spatial resolution of the parent microscope. Additionally, 3D-EPM may be implemented by the simple addition of a $4f$ relay behind the image plane of any standard fluorescence imaging system. These two factors make 3D-EPM a versatile technique that can be applied across many different spatial scales.

2. Where previously the application of PSF engineering to snapshot 3D imaging has been largely focused on SMLM, 3D-EPM permits snapshot volumetric imaging of samples with extended regions of fluorescent structure. We demonstrated this in both simulation and experiment by imaging filament-like structures via 3D-EPM.
3. Because 3D-EPM provides true volumetric imaging in a single snapshot, it enables the study of rapidly-moving scenes in all 3 spatial dimensions. We demonstrated its application to the study of 3D dynamics by imaging the motion of chloroplasts in a sample of *Egeria densa*.

In chapter 4, we turned our attention to the field of point-localisation, where the spatial coordinates of individual point emitters are estimated from 2D image data. PSF-engineering has found popularity in 3D-SMLM, however a fundamental limitation is the inability of many localisation algorithms that estimate the 3D position of emitters directly from the raw 2D images to cope with PSFs that significantly overlap. The limits on acceptable density of emitters, therefore, places constraints on the number of localisations that can be calculated per frame, which in turn restricts the temporal resolution of 3D-SMLM where photo-switchable fluorophores are used. We saw that when the volume-reconstruction process of 3D-EPM was applied to PSF-engineered point-localisation data, performing the localisation within the reconstructed 3D volume as opposed to the raw 2D image data meant significantly higher densities of emitters could be effectively localised per frame than with a previously-developed point-matching approach. We also saw how 3D-EPM localisation performed better at lower signal-to-noise ratios, and when emitters were positioned close to the edge of the imaging domain. Additionally, in contrast with many localisation algorithms that are designed to work with specific PSFs, 3D-EPM localisation is a versatile approach that makes no prior assumptions about the nature of the PSF used for imaging. 3D-EPM localisation will, therefore, work with *any* engineered PSF providing it effectively encodes depth information over the required axial range.

Finally, in chapter 5, we saw how the mathematical framework used for volume-reconstruction from 2D projection images in 3D-EPM could be applied to other projection imaging modalities, wherever some form of latent 3D information is present within the raw projection data. We adjusted our model of image formation to reflect image formation on a multi-angle projection imaging system, and saw how the model could be inverted via our amended RL deconvolution framework to provide full 3D reconstruction from few projection views. We demonstrated its application to a number of samples, both in simulation and with experimentally acquired data, including that




Figure 6.1: Top-down view of an experimental setup designed to implement the tomographic-EPM setup introduced in chapter 5. The experimental setup was designed by myself and built by Katarzyna Glinka during a summer research project within the Imaging Concepts Group.

acquired on an OPM system which introduced additional complexities to the reconstruction process. Additionally, a design for a tomographic EPM system was proposed that probes the sample via two objectives at orthogonal angles. Proof-of-concept results were presented that demonstrated how isotropic 3D resolution can be achieved with such a system, overcoming the anisotropy between lateral and axial resolution that exists when imaging solely through a single objective lens.

In summary, this thesis has presented a mathematical framework that enables the 3D reconstruction of a sample volume from 2D projection data, thereby removing the need for many sequential 2D acquisitions in order to obtain high-resolution 3D images. We have modified a model of image formation to describe image formation across a number of different projection imaging modalities, and demonstrated the application of the volume reconstruction framework to those imaging methods in both simulation and experiment. By offering volumetric imaging with high resolution in time and in space, it is hoped that the imaging methods developed in this thesis will find broad applicability in the 3D study of fast biological dynamics.

Ongoing work

While proof-of-concept results for the tomographic 3D-EPM implementation were presented in chapter 5 of this thesis, there was sadly not enough time available during my PhD to implement this work experimentally. This work is in progress, and I hope to obtain experimental results from this implementation in the near future. Figure 6.1 displays an experimental setup designed by myself to implement this experiment, built by an undergraduate student, Katarzyna Glinka, during a summer research project within the Imaging Concepts Group. With this setup we hope to image the beating zebrafish heart with an isotropic 3D resolution.

Additionally, we have ongoing collaborations with the authors of [5], with whom we hope to demonstrate the 3D reconstruction from multi-angle projection image sequences of dynamic biological processes that occur too quickly to probe with conventional z -stack imaging.

Contributions

Parts of the work presented in this thesis involve collaboration with others and their contributions are detailed below.

Michael Handley wrote the python code for the twin-airy point-matching localisation routine which was originally developed by Dr Yongzhuang Zhou, against which the 3D-EPM localisation routine was compared in chapter 4.

The results presented in chapter 5 form part of an ongoing collaboration with Professor Reto Fiolka, Dr Bingying Chen and Dr James D. Manton. The raw projection images shown in figures 5.4, 5.6 and 5.7 were simulated from ground-truth images by Dr James D. Manton. The raw projection data shown in figure 5.10 was acquired by Dr Bingying Chen. The raw projection data shown in figure 5.11 was acquired by Professor Bo-Jui Chang. All data was used with permission, and all 3D reconstructions shown were performed by myself. The phenomena of ghost artefacts being introduced into reconstructed volumes when reconstructing multi-angle projection data using only two views, discussed in section 5.5, was helpfully pointed out by Dr James D. Manton.

The tomographic 3D-EPM setup shown in figure 6.1 was designed by myself and built by an undergraduate student, Katarzyna Glinka, as part of a summer research project.

The *Egeria densa* sample imaged in chapter 3 was suggested by Professor Michael Blatt.

All of the work presented in this thesis was carried out under the invaluable supervision and guidance provided by Dr Jonathan M. Taylor and Professor Andrew R. Harvey.

Bibliography

- [1] Hooke, R. *Micrographia, or, Some physiological descriptions of minute bodies made by magnifying glasses : with observations and inquiries thereupon* (J. Martyn and J. Allestry, 1665).
- [2] Levoy, M., Ng, R., Adams, A., Footer, M. & Horowitz, M. Light field microscopy. In *ACM SIGGRAPH 2006 Papers*, 924–934 (2006).
- [3] Richardson, W. H. Bayesian-based iterative method of image restoration. *JoSA* **62**, 55–59 (1972).
- [4] Lucy, L. B. An iterative technique for the rectification of observed distributions. *The astronomical journal* **79**, 745 (1974).
- [5] Chang, B.-J. *et al.* Real-time multi-angle projection imaging of biological dynamics. *Nature methods* **18**, 829–834 (2021).
- [6] Chalfie, M., Tu, Y., Euskirchen, G., Ward, W. W. & Prasher, D. C. Green fluorescent protein as a marker for gene expression. *Science* **263**, 802–805 (1994).
- [7] Lipson, A., Lipson, S. G. & Lipson, H. *Optical physics* (Cambridge University Press, 2010).
- [8] Shannon, C. E. Communication in the presence of noise. *Proceedings of the IRE* **37**, 10–21 (1949).
- [9] Snyder, D. L., Helstrom, C. W., Lanterman, A. D., Faisal, M. & White, R. L. Compensation for readout noise in ccd images. *JOSA A* **12**, 272–283 (1995).
- [10] Pawley, J. *Handbook of biological confocal microscopy*, vol. 236 (Springer Science & Business Media, 2006).
- [11] Goodman, J. W. *Introduction to Fourier Optics. Goodman* (McGraw-Hill, 1968).
- [12] Hanser, B. M., Gustafsson, M. G., Agard, D. & Sedat, J. W. Phase-retrieved pupil functions in wide-field fluorescence microscopy. *Journal of microscopy* **216**, 32–48 (2004).
- [13] Hanser, B. M., Gustafsson, M. G., Agard, D. A. & Sedat, J. W. Phase retrieval for high-numerical-aperture optical systems. *Optics letters* **28**, 801–803 (2003).

- [14] Lee, J., Kim, S., Yoon, C. H., Kim, M. J. & Kim, K. H. Moxifloxacin based axially swept wide-field fluorescence microscopy for high-speed imaging of conjunctival goblet cells. *Biomedical optics express* **11**, 4890–4900 (2020).
- [15] Forsgren, E. *et al.* High-throughput widefield fluorescence imaging of 3d samples using deep learning for 2d projection image restoration. *Plos one* **17**, e0264241 (2022).
- [16] Minsky, M. Memoir on inventing the confocal scanning microscope. *Scanning* **10**, 128–138 (1988).
- [17] Shaw, P. J. Comparison of widefield/deconvolution and confocal microscopy for three-dimensional imaging. In *Handbook of biological confocal microscopy*, 453–467 (Springer, 2006).
- [18] Petráň, M., Hadravský, M., Egger, M. D. & Galambos, R. Tandem-scanning reflected-light microscope. *JOSA* **58**, 661–664 (1968).
- [19] Koester, C. J. Scanning mirror microscope with optical sectioning characteristics: applications in ophthalmology. *Applied Optics* **19**, 1749–1757 (1980).
- [20] Denk, W., Strickler, J. H. & Webb, W. W. Two-photon laser scanning fluorescence microscopy. *Science* **248**, 73–76 (1990).
- [21] Goppert-Mayer, M. Two-quantum processes. *Ann Phys (Leipzig)* **9**, 273–294 (1931).
- [22] Doi, A., Oketani, R., Nawa, Y. & Fujita, K. High-resolution imaging in two-photon excitation microscopy using in situ estimations of the point spread function. *Biomedical Optics Express* **9**, 202–213 (2018).
- [23] Hwang, J. Y. *et al.* Multimodal wide-field two-photon excitation imaging: characterization of the technique for in vivo applications. *Biomedical optics express* **2**, 356–364 (2011).
- [24] Amor, R. *et al.* Widefield two-photon excitation without scanning: Live cell microscopy with high time resolution and low photo-bleaching. *PLoS One* **11**, e0147115 (2016).
- [25] Hell, S. W. *et al.* Three-photon excitation in fluorescence microscopy. *Journal of Biomedical Optics* **1**, 71–74 (1996).
- [26] Wokosin, D. L., Centonze, V. E., Crittenden, S. & White, J. Three-photon excitation fluorescence imaging of biological specimens using an all-solid-state laser. *Bioimaging* **4**, 208–214 (1996).
- [27] Bakker, G.-J. *et al.* Intravital deep-tumor single-beam 3-photon, 4-photon, and harmonic microscopy. *Elife* **11**, e63776 (2022).

- [28] Horton, N. G. *et al.* In vivo three-photon microscopy of subcortical structures within an intact mouse brain. *Nature photonics* **7**, 205–209 (2013).
- [29] Wang, T. *et al.* Three-photon imaging of mouse brain structure and function through the intact skull. *Nature methods* **15**, 789–792 (2018).
- [30] Voie, A. H., Burns, D. & Spelman, F. Orthogonal-plane fluorescence optical sectioning: Three-dimensional imaging of macroscopic biological specimens. *Journal of microscopy* **170**, 229–236 (1993).
- [31] Keller, P. J. & Stelzer, E. H. Quantitative in vivo imaging of entire embryos with digital scanned laser light sheet fluorescence microscopy. *Current opinion in neurobiology* **18**, 624–632 (2008).
- [32] Keller, P. J., Schmidt, A. D., Wittbrodt, J. & Stelzer, E. H. Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy. *science* **322**, 1065–1069 (2008).
- [33] Dunsby, C. Optically sectioned imaging by oblique plane microscopy. *Optics express* **16**, 20306–20316 (2008).
- [34] Vettenburg, T. *et al.* Light-sheet microscopy using an airy beam. *Nature methods* **11**, 541–544 (2014).
- [35] Fahrbach, F. O., Simon, P. & Rohrbach, A. Microscopy with self-reconstructing beams. *Nature photonics* **4**, 780–785 (2010).
- [36] Planchon, T. A. *et al.* Rapid three-dimensional isotropic imaging of living cells using bessel beam plane illumination. *Nature methods* **8**, 417–423 (2011).
- [37] Yang, B. *et al.* Epi-illumination spim for volumetric imaging with high spatial-temporal resolution. *Nature methods* **16**, 501–504 (2019).
- [38] Sapoznik, E. *et al.* A versatile oblique plane microscope for large-scale and high-resolution imaging of subcellular dynamics. *Elife* **9**, e57681 (2020).
- [39] Yang, B. *et al.* Daxi—high-resolution, large imaging volume and multi-view single-objective light-sheet microscopy. *Nature methods* **19**, 461–469 (2022).
- [40] Prabhat, P., Ram, S., Ward, E. S. & Ober, R. J. Simultaneous imaging of different focal planes in fluorescence microscopy for the study of cellular dynamics in three dimensions. *IEEE transactions on nanobioscience* **3**, 237–242 (2004).
- [41] Geissbuehler, S. *et al.* Live-cell multiplane three-dimensional super-resolution optical fluctuation imaging. *Nature communications* **5**, 1–7 (2014).
- [42] Xiao, S. *et al.* High-contrast multifocus microscopy with a single camera and z-splitter prism. *Optica* **7**, 1477–1486 (2020).

- [43] Blanchard, P. M. & Greenaway, A. H. Simultaneous multiplane imaging with a distorted diffraction grating. *Applied optics* **38**, 6692–6699 (1999).
- [44] Levoy, M., Zhang, Z. & McDowall, I. Recording and controlling the 4d light field in a microscope using microlens arrays. *Journal of microscopy* **235**, 144–162 (2009).
- [45] Broxton, M. *et al.* Wave optics theory and 3-d deconvolution for the light field microscope. *Optics express* **21**, 25418–25439 (2013).
- [46] Li, H. *et al.* Fast, volumetric live-cell imaging using high-resolution light-field microscopy. *Biomedical optics express* **10**, 29–49 (2019).
- [47] Cohen, N. *et al.* Enhancing the performance of the light field microscope using wavefront coding. *Optics express* **22**, 24817–24839 (2014).
- [48] Prevedel, R. *et al.* Simultaneous whole-animal 3d imaging of neuronal activity using light-field microscopy. *Nature methods* **11**, 727–730 (2014).
- [49] Nöbauer, T. *et al.* Video rate volumetric ca²⁺ imaging across cortex using seeded iterative demixing (sid) microscopy. *Nature methods* **14**, 811–818 (2017).
- [50] Wagner, N. *et al.* Instantaneous isotropic volumetric imaging of fast biological processes. *Nature methods* **16**, 497–500 (2019).
- [51] Guo, C., Liu, W., Hua, X., Li, H. & Jia, S. Fourier light-field microscopy. *Optics express* **27**, 25573–25594 (2019).
- [52] Dowski, E. R. & Cathey, W. T. Extended depth of field through wave-front coding. *Applied optics* **34**, 1859–1866 (1995).
- [53] Bradburn, S., Cathey, W. T. & Dowski, E. R. Realizations of focus invariance in optical–digital systems with wave-front coding. *Applied optics* **36**, 9157–9166 (1997).
- [54] Zhou, Y., Zickus, V., Zammit, P., Taylor, J. M. & Harvey, A. R. High-speed extended-volume blood flow measurement using engineered point-spread function. *Biomedical Optics Express* **9**, 6444–6454 (2018).
- [55] Demenikov, M. & Harvey, A. R. Image artifacts in hybrid imaging systems with a cubic phase mask. *Optics express* **18**, 8207–8212 (2010).
- [56] Demenikov, M. & Harvey, A. R. Parametric blind-deconvolution algorithm to remove image artifacts in hybrid imaging systems. *Optics Express* **18**, 18035–18040 (2010).
- [57] Olarte, O. E., Andilla, J., Artigas, D. & Loza-Alvarez, P. Decoupled illumination detection in light sheet microscopy for fast volumetric imaging. *Optica* **2**, 702–705 (2015).

- [58] Zammit, P., Harvey, A. R. & Carles, G. Extended depth-of-field imaging and ranging in a snapshot. *Optica* **1**, 209–216 (2014).
- [59] Lelek, M. *et al.* Single-molecule localization microscopy. *Nature Reviews Methods Primers* **1**, 1–27 (2021).
- [60] Betzig, E. *et al.* Imaging intracellular fluorescent proteins at nanometer resolution. *science* **313**, 1642–1645 (2006).
- [61] Hess, S. T., Girirajan, T. P. & Mason, M. D. Ultra-high resolution imaging by fluorescence photoactivation localization microscopy. *Biophysical journal* **91**, 4258–4272 (2006).
- [62] Rust, M. J., Bates, M. & Zhuang, X. Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (storm). *Nature methods* **3**, 793–796 (2006).
- [63] Huang, B., Wang, W., Bates, M. & Zhuang, X. Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy. *Science* **319**, 810–813 (2008).
- [64] Pavani, S. R. P. *et al.* Three-dimensional, single-molecule fluorescence imaging beyond the diffraction limit by using a double-helix point spread function. *Proceedings of the National Academy of Sciences* **106**, 2995–2999 (2009).
- [65] Shechtman, Y., Weiss, L. E., Backer, A. S., Sahl, S. J. & Moerner, W. Precise three-dimensional scan-free multiple-particle tracking over large axial ranges with tetrapod point spread functions. *Nano letters* **15**, 4194–4199 (2015).
- [66] Zhou, Y. & Carles, G. Precise 3d particle localization over large axial ranges using secondary astigmatism. *Optics Letters* **45**, 2466–2469 (2020).
- [67] Jia, S., Vaughan, J. C. & Zhuang, X. Isotropic three-dimensional super-resolution imaging with a self-bending point spread function. *Nature photonics* **8**, 302–306 (2014).
- [68] Zhou, Y., Zammit, P., Zickus, V., Taylor, J. M. & Harvey, A. R. Twin-airy point-spread function for extended-volume particle localization. *Physical Review Letters* **124**, 198104 (2020).
- [69] Liu, S. & Hua, H. Extended depth-of-field microscopic imaging with a variable focus microscope objective. *Optics express* **19**, 353–362 (2011).
- [70] Grewe, B. F., Voigt, F. F., van't Hoff, M. & Helmchen, F. Fast two-layer two-photon imaging of neuronal cell populations using an electrically tunable lens. *Biomedical optics express* **2**, 2035–2046 (2011).

- [71] Shain, W. J., Vickers, N. A., Goldberg, B. B., Bifano, T. & Mertz, J. Extended depth-of-field microscopy with a high-speed deformable mirror. *Optics letters* **42**, 995–998 (2017).
- [72] Olesker, D., Harvey, A. R. & Taylor, J. M. Snapshot volumetric imaging with engineered point-spread functions. *Optics Express* **30**, 33490–33501 (2022).
- [73] Wu, P. *et al.* Large depth-of-field fluorescence microscopy based on deep learning supported by fresnel incoherent correlation holography. *Optics Express* **30**, 5177–5191 (2022).
- [74] Dey, N. *et al.* *3D microscopy deconvolution using Richardson-Lucy algorithm with total variation regularization*. Ph.D. thesis, INRIA (2004).
- [75] Jansson, P. A. *Deconvolution of images and spectra* (Courier Corporation, 2014).
- [76] Temerinac-Ott, M. *et al.* Multiview deblurring for 3-d images from light-sheet-based fluorescence microscopy. *IEEE Transactions on Image Processing* **21**, 1863–1873 (2011).
- [77] Ingaramo, M. *et al.* Richardson–lucy deconvolution as a general tool for combining images with complementary strengths. *ChemPhysChem* **15**, 794–800 (2014).
- [78] Wu, Y. *et al.* Spatially isotropic four-dimensional imaging with dual-view plane illumination microscopy. *Nature biotechnology* **31**, 1032–1038 (2013).
- [79] Preibisch, S. *et al.* Efficient bayesian-based multiview deconvolution. *Nature methods* **11**, 645–648 (2014).
- [80] Andrews, H. C. & Hunt, B. R. *Digital image restoration* (1977).
- [81] Ng, M. K., Chan, R. H. & Tang, W.-C. A fast algorithm for deblurring models with neumann boundary conditions. *SIAM Journal on Scientific Computing* **21**, 851–866 (1999).
- [82] Bertero, M. & Boccacci, P. A simple method for the reduction of boundary effects in the richardson-lucy approach to image deconvolution. *Astronomy & Astrophysics* **437**, 369–374 (2005).
- [83] Cannell, M. B., McMorland, A. & Soeller, C. Image enhancement by deconvolution. In *Handbook of biological confocal microscopy*, 488–500 (Springer, 2006).
- [84] White, R. L. Image restoration using the damped richardson-lucy method. In *Instrumentation in Astronomy VIII*, vol. 2198, 1342–1348 (SPIE, 1994).
- [85] Kempen, V. & Vliet, V. The influence of the regularization parameter and the first estimate on the performance of tikhonov regularized non-linear image restoration algorithms. *Journal of Microscopy* **198**, 63–75 (2000).

- [86] Nicolas, D. *et al.* 3d microscopy deconvolution using richardson-lucy algorithm with total variation regularization. *STITUT NATIONAL de RECHERCHE en INFORMATIQUE et en AUTOMATIQUE* **7**, 3–43 (2004).
- [87] Dey, N. *et al.* Richardson–lucy algorithm with total variation regularization for 3d confocal microscope deconvolution. *Microscopy research and technique* **69**, 260–266 (2006).
- [88] Tzang, O., Feldkhun, D., Agrawal, A., Jesacher, A. & Piestun, R. Two-photon psf-engineered image scanning microscopy. *Optics Letters* **44**, 895–898 (2019).
- [89] Zhang, W., Ye, Z., Zhao, T., Chen, Y. & Yu, F. Point spread function characteristics analysis of the wavefront coding system. *Optics express* **15**, 1543–1552 (2007).
- [90] Gerchberg, R. W. A practical algorithm for the determination of plane from image and diffraction pictures. *Optik* **35**, 237–246 (1972).
- [91] Noll, R. J. Zernike polynomials and atmospheric turbulence. *JOsA* **66**, 207–211 (1976).
- [92] Schmidt, J. D. Numerical simulation of optical wave propagation: With examples in matlab (SPIE, 2010).
- [93] Van der Walt, S. *et al.* scikit-image: image processing in python. *PeerJ* **2**, e453 (2014).
- [94] Tinevez, J.-Y. *et al.* Trackmate: An open and extensible platform for single-particle tracking. *Methods* **115**, 80–90 (2017).
- [95] Hove, J. R. *et al.* Intracardiac fluid forces are an essential epigenetic factor for embryonic cardiogenesis. *Nature* **421**, 172–177 (2003).
- [96] Lu, J., Pereira, F., Fraser, S. E. & Gharib, M. Three-dimensional real-time imaging of cardiac cell motions in living embryos. *Journal of biomedical optics* **13**, 014006 (2008).
- [97] Vennemann, P. *et al.* In vivo micro particle image velocimetry measurements of blood-plasma in the embryonic avian heart. *Journal of biomechanics* **39**, 1191–1200 (2006).
- [98] Lee, J., Leonard, M., Oliver, T., Ishihara, A. & Jacobson, K. Traction forces generated by locomoting keratocytes. *The Journal of cell biology* **127**, 1957–1964 (1994).
- [99] Roy, P., Petroll, W., Cavanagh, H., Chuong, C. & Jester, J. Anin vitroforce measurement assay to study the early mechanical interaction between corneal fibroblasts and collagen matrix. *Experimental cell research* **232**, 106–117 (1997).

- [100] Legant, W. R. *et al.* Measurement of mechanical tractions exerted by cells in three-dimensional matrices. *Nature methods* **7**, 969–971 (2010).
- [101] Smith, C. S., Joseph, N., Rieger, B. & Lidke, K. A. Fast, single-molecule localization that achieves theoretically minimum uncertainty. *Nature methods* **7**, 373–375 (2010).
- [102] Nieuwenhuizen, R. P. *et al.* Measuring image resolution in optical nanoscopy. *Nature methods* **10**, 557–562 (2013).
- [103] Steinhauer, C., Jungmann, R., Sobey, T. L., Simmel, F. C. & Tinnefeld, P. Dna origami as a nanoscopic ruler for super-resolution microscopy. *Angewandte Chemie International Edition* **48**, 8870–8873 (2009).
- [104] Kanchanawong, P. *et al.* Nanoscale architecture of integrin-based cell adhesions. *Nature* **468**, 580–584 (2010).
- [105] Bintu, B. *et al.* Super-resolution chromatin tracing reveals domains and cooperative interactions in single cells. *Science* **362**, eaau1783 (2018).
- [106] Xu, K., Zhong, G. & Zhuang, X. Actin, spectrin, and associated proteins form a periodic cytoskeletal structure in axons. *Science* **339**, 452–456 (2013).
- [107] Juette, M. F. *et al.* Three-dimensional sub-100 nm resolution fluorescence microscopy of thick samples. *Nature methods* **5**, 527–529 (2008).
- [108] Song, K.-H., Zhang, Y., Wang, G., Sun, C. & Zhang, H. F. Three-dimensional biplane spectroscopic single-molecule localization microscopy. *Optica* **6**, 709–715 (2019).
- [109] Babcock, H. P. Multiplane and spectrally-resolved single molecule localization microscopy with industrial grade cmos cameras. *Scientific reports* **8**, 1–8 (2018).
- [110] Holtzer, L., Meckel, T. & Schmidt, T. Nanometric three-dimensional tracking of individual quantum dots in cells. *Applied Physics Letters* **90**, 053902 (2007).
- [111] Kao, H. P. & Verkman, A. Tracking of single fluorescent particles in three dimensions: use of cylindrical optics to encode particle position. *Biophysical journal* **67**, 1291–1300 (1994).
- [112] Sims, R. R. *et al.* Single molecule light field microscopy. *Optica* **7**, 1065–1072 (2020).
- [113] Lew, M. D., Lee, S. F., Badieirostami, M. & Moerner, W. Corkscrew point spread function for far-field three-dimensional nanoscale localization of pointlike objects. *Optics letters* **36**, 202–204 (2011).
- [114] Nehme, E. *et al.* Deepstorm3d: dense 3d localization microscopy and psf design by deep learning. *Nature methods* **17**, 734–740 (2020).

- [115] Ikoma, H., Peng, Y., Broxton, M. & Wetzstein, G. Snapshot multi-psf 3d single-molecule localization microscopy using deep learning. In *Computational Optical Sensing and Imaging*, CW3B-3 (Optica Publishing Group, 2020).
- [116] Allan, D., Caswell, T., Keim, N., van der Wel, C. M. & Verweij, R. soft-matter/trackpy: Trackpy v0. 5.0. *Genève: Zenodo* (2021).
- [117] Chao, J., Ward, E. S. & Ober, R. J. Fisher information theory for parameter estimation in single molecule microscopy: tutorial. *JOSA A* **33**, B36–B57 (2016).
- [118] Shechtman, Y., Sahl, S. J., Backer, A. S. & Moerner, W. E. Optimal point spread function design for 3d imaging. *Physical review letters* **113**, 133902 (2014).
- [119] Grover, G., DeLuca, K., Quirin, S., DeLuca, J. & Piestun, R. Super-resolution photon-efficient imaging by nanometric double-helix point spread function localization of emitters (spindle). *Optics express* **20**, 26681–26695 (2012).
- [120] Min, J. *et al.* Falcon: fast and unbiased reconstruction of high-density super-resolution microscopy data. *Scientific reports* **4**, 1–9 (2014).
- [121] Nehme, E., Weiss, L. E., Michaeli, T. & Shechtman, Y. Deep-storm: super-resolution single-molecule microscopy by deep learning. *Optica* **5**, 458–464 (2018).
- [122] Sage, D. *et al.* Super-resolution fight club: assessment of 2d and 3d single-molecule localization microscopy software. *Nature methods* **16**, 387–395 (2019).
- [123] Martens, K. J., Jabermoradi, A., Yang, S. & Hohlbein, J. Integrating engineered point spread functions into the phasor-based single-molecule localization microscopy framework. *Methods* **193**, 107–115 (2021).
- [124] Hirsch, M., Wareham, R. J., Martin-Fernandez, M. L., Hobson, M. P. & Rolfe, D. J. A stochastic model for electron multiplication charge-coupled devices—from theory to practice. *PloS one* **8**, e53671 (2013).
- [125] Quirin, S. *et al.* Calcium imaging of neural circuits with extended depth-of-field light-sheet microscopy. *Optics letters* **41**, 855–858 (2016).
- [126] Ji, N. & Lu, R. Video-rate volumetric functional imaging of the brain at synaptic resolution (2020). US Patent 10,578,849.
- [127] Belay, B. *et al.* Optical projection tomography as a quantitative tool for analysis of cell morphology and density in 3d hydrogels. *Scientific reports* **11**, 1–10 (2021).
- [128] Boselli, F. & Vermot, J. Live imaging and modeling for shear stress quantification in the embryonic zebrafish heart. *Methods* **94**, 129–134 (2016).
- [129] Zickus, V. & Taylor, J. M. 3d+ time blood flow mapping using spim-micropiv in the developing zebrafish heart. *Biomedical optics express* **9**, 2418–2435 (2018).

- [130] Hu, Z., Luo, H., Du, Y. & Lu, H. Fluorescent stereo microscopy for 3d surface profilometry and deformation mapping. *Optics express* **21**, 11808–11818 (2013).
- [131] Lacroute, P. & Levoy, M. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, 451–458 (1994).
- [132] Bracewell, R. The sifting property. *The Fourier Transform and Its Applications, 3rd ed. New York, McGraw-Hill* 74–77 (1999).
- [133] Virtanen, P. *et al.* Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods* **17**, 261–272 (2020).
- [134] Botcherby, E. J., Juskaitis, R., Booth, M. J. & Wilson, T. Aberration-free optical refocusing in high numerical aperture microscopy. *Optics letters* **32** (2007).
- [135] Buzug, T. M. *Computed tomography* (Springer, 2011).
- [136] Barrett, H. H. The radon transform and its applications. In *Progress in optics*, vol. 21, 217–286 (Elsevier, 1984).
- [137] Li, M. *et al.* Smoothed norm regularization for sparse-view x-ray ct reconstruction. *BioMed Research International* **2016** (2016).
- [138] Elbakri, I. A. & Fessler, J. A. Statistical image reconstruction for polyenergetic x-ray computed tomography. *IEEE transactions on medical imaging* **21**, 89–99 (2002).
- [139] Kudo, H., Suzuki, T. & Rashed, E. A. Image reconstruction for sparse-view ct and interior ct—introduction to compressed sensing and differentiated backprojection. *Quantitative imaging in medicine and surgery* **3**, 147 (2013).
- [140] Hell, S. & Stelzer, E. H. Properties of a 4pi confocal fluorescence microscope. *JOSA A* **9**, 2159–2166 (1992).
- [141] Verveer, P. J. *et al.* High-resolution three-dimensional imaging of large specimens with light sheet-based microscopy. *Nature methods* **4**, 311–313 (2007).
- [142] Wu, Y. *et al.* Simultaneous multiview capture and fusion improves spatial resolution in wide-field and light-sheet microscopy. *Optica* **3**, 897–910 (2016).