

IMPERIAL

Polarisation Microscopy and Its Application to Optical Data Storage

Design, Informational Analysis, and ML-Based Decoding

Zhonghe Feng

Supervised by
Prof Mark A. A. Neil
Asst Prof Matthew R. Foreman
Dr James Clegg

Imperial College London
Department of Physics

This thesis is submitted in part fulfilment of the requirements
for the degree of Doctor of Philosophy (PhD) in Physics
of Imperial College London, May 2026

Abstract

Optical data storage (ODS) can encode information as birefringent nanostructures in glass, but the resulting polarisation features are weak and can be degraded by photon noise, aberrations, and crosstalk. This thesis develops microscopy and processing techniques for weak polarisation signals, spanning system design and polarisation control, modelling and simulation, practical implementation, calibration, and quantitative evaluation.

Single-shot polarimetry with a polarisation camera is studied first, enabling measurement of the linear components of polarisation. Polarisation-sensitive phase contrast (PPC) microscopy is proposed by combining Zernike’s phase contrast setup with polarisation control to selectively attenuate the surround wave and enhance polarisation contrast. An analytical model is derived under a weak-object assumption and complemented by a partially coherent simulation pipeline based on coherent-mode decomposition of Gaussian Schell-model sources, enabling accurate simulation of the imaging process with a flexible system definition. Experimental implementation is presented, with software for automation, calibration, and real-time visualisation.

Voxel readout is formulated as a communication channel, and mutual information (MI) is used as a decoder-agnostic metric to quantify read performance under noise, aberrations, and intra- and inter-layer crosstalk. A processing pipeline enables MI estimation from experimental and simulated images, and voxel reading is experimentally demonstrated in single-layer and multi-layer settings. By comparing axial MI profiles from simulations and experiments, a thick-voxel model is found to reproduce crosstalk characteristics, and the framework also studies MI degradation from write-side artefacts, intra-layer crosstalk, and photon shot noise. Finally, explicit symbol recovery is demonstrated with a machine-learning decoder, completing an end-to-end demonstration of ODS.

To overcome limitations inherited from spatial filtering in PPC, we introduce a diattenuative module to attenuate the surround wave using a custom partial polariser, while retaining bright-field transfer functions and avoiding azimuthal bias. Beyond voxel imaging, the diattenuative system reveals weak birefringence in other samples such as glass fibres and mycobacteria.

Acknowledgements

The past few years have been both challenging and highly rewarding. Throughout not only the thesis-writing process but my entire PhD journey, I have benefited from generous support from many people, without whom this work would not have been possible.

First and foremost, I would like to thank my supervisors Mark Neil, Matthew Foreman, and James Clegg for their thoughtful guidance, continued support, and insightful ideas throughout my PhD. Working with you, I have not only gained the knowledge and skills needed to complete this thesis, but have also been deeply inspired by your approach to research, work, and teaching, which has had a profound impact on both my professional and personal life.

I would like to thank Hai, Meizhu, Yuheng, and You. It has been a real pleasure working with you in the lab, and I have also enjoyed our thoughtful discussions and the shared interests outside academia. I would like to thank my colleagues at Project Silica as well, who make the wider project possible, especially Rokas and Teodora on the writing team for providing samples for this work. In addition, I would like to give special thanks to Edwin for many inspiring discussions and for the considerable effort he put into preparing the biological samples used in this project. I am grateful to John, Martin, and Simon in the optics workshop, who made many practical aspects of this project possible through their skilled work. I also owe thanks to Judith and Marcia for their consistent support on the administrative side.

I have been fortunate to meet so many inspiring colleagues and friends in the Photonics group, including Paul, Chris, James, Martin, Carl, Kenny, Wenzhi, Huihui, Leo, Nico, Jonny, Sunil, Sara, Connor, Meizhen, Bill, Liuba, Killian, Dan, Mi Qi, Karishma, Hugh, Huaifeng, Xunuo, Meihui, Niall, Yuriy, Nathan, Steven, Abdul, Yuchen, Cheryl, and many others. Beyond working together, the pub drinks, summer trips, Christmas infusions, and film nights have been genuinely unforgettable.

This work was supported by Microsoft Research through its PhD Scholarship Programme.

Finally, I would like to thank my beloved parents for the encouragement, backing, inspiration, and guidance you bring to my life, which have fundamentally shaped who I am. To dear Yichen, thank you for your immense support and companionship through the ups and downs, and for enriching my life, now and in the years to come.

Statement of originality

I declare that this thesis is my own work, except for Figure 1.1 and Figure 2.1(b), which are reproduced from the sources cited in their captions with permission.

The voxel samples used in this work were provided by Project Silica at Microsoft Research.

The glass fibre and mycobacteria samples used in this work were prepared by Edwin Garcia.

I acknowledge the use of ChatGPT from OpenAI for providing language-related suggestions during the preparation of this thesis. No AI-generated text was copied directly into the thesis.

Copyright declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

To my family.

Contents

Abstract	3
Acknowledgements	5
Statement of originality	6
Copyright declaration	7
List of Tables	17
List of Figures	21
List of Abbreviations and Acronyms	25
1 Introduction	27
1.1 Motivation	27
1.2 Optical data storage	28
1.3 Thesis structure and overviews of chapters	31
2 Single-shot polarimetry using pixelated polarisation cameras	34
2.1 Division-of-focal-plane (DoFP) polarimetry with the PolarSens™ sensor	35
2.2 Measuring linear components of polarisation	38
2.3 Mixing of polarisation states	40
2.4 Polarisation camera (PolCam) polarimetry in the circular basis	45
2.4.1 Jones algebra in the circular basis	45
2.4.2 Polarimetric interpretation of the circular-basis formalism	48

2.4.3	False-colour visualisation scheme	52
2.5	PolCam images in the Fourier domain	53
2.5.1	Whittaker-Shannon theorem for 2D sampling	54
2.5.2	Spectral bands for PolCam images	55
2.6	Noise model and characterisation	60
2.6.1	The linear model of PolCam	60
2.6.2	Photon-free characterisation of PolCam	64
2.6.3	Pixel variations and correction	69
2.7	Towards an ideal system for weak birefringence	73
2.8	Conclusions	77
3	Polarisation-sensitive phase contrast (PPC) microscopy	79
3.1	Zernike’s phase contrast	79
3.2	Polarisation control of PPC	81
3.2.1	Objects of birefringence only	81
3.2.2	Birefringent objects with a residual phase	84
3.2.3	The paraxial approximation	86
3.3	Imaging model of PPC upon spatially incoherent illumination	86
3.3.1	Transfer of polarisation signal	88
3.3.2	Transfer of phase signal	95
3.4	Conclusions	98
4	Modal representation of a partially coherent source for image simulation	101
4.1	Second-order statistics of a partially coherent field	103
4.2	Bessel-correlated Schell-model (BSM) sources	107
4.3	Apertured Gaussian Schell-model (GSM) sources	111
4.3.1	Representation in Hermite–Gaussian (HG) modes	114
4.3.2	Representation in Laguerre–Gaussian (LG) modes	118
4.3.3	Numerical modes	122
4.3.4	Comparison of mode efficiency across various representations	126
4.4	Conclusion	127

5	Simulation pipeline for polarisation imaging	130
5.1	Physical model	131
5.1.1	Secondary source at condenser aperture	132
5.1.2	Object space in angular spectrum	133
5.1.3	Voxel arrays	136
5.1.4	Objective and collar correction	137
5.1.5	Tube lens and camera	139
5.2	Implementation of the simulation pipeline	139
5.2.1	System definition	140
5.2.2	Initialisation and array construction	141
5.2.3	Simultaneous mode generation and propagation	142
5.2.4	Output	144
5.2.5	HPC deployment and parallelisation with job arrays	144
5.3	Conclusion	146
6	PPC System Implementation	148
6.1	Hardware assembly	149
6.2	Software for control, visualisation, and automation	151
6.3	Conclusion	153
7	Informational analysis of voxel imaging	155
7.1	A communication model for optical data storage	156
7.2	Entropy, mutual information, and channel capacity	157
7.3	The voxel image processing pipeline	161
7.3.1	Instantaneous field-of-view error (IFOV) correction	162
7.3.2	Voxel localisation	165
7.3.3	Mutual information estimation	172
7.4	Single-layer voxel imaging results	174
7.4.1	Baseline demonstration of voxel reading	174
7.4.2	Impact of voxel pitch	176

7.4.3	Write-side effects and intra-layer crosstalk	177
7.4.4	Camera noise	182
7.5	Axial performance and multi-layer reading	186
7.5.1	Baseline demonstration of multi-layer reading	187
7.5.2	Effect of defocus for single-layer voxel samples	189
7.5.3	Multi-layer simulations and interlayer crosstalk	196
7.5.4	Depth-induced aberrations	200
7.6	Conclusion	202
8	Decoding voxel data using machine learning	207
8.1	Multilayer perceptron decoder	208
8.2	Training and decoding results	210
8.3	Polarisation-based image augmentation	215
8.4	Conclusion and future work	219
9	Clear-aperture polarisation imaging with a diattenuative module	221
9.1	Polarisation mechanism of the diattenuative module	223
9.2	System implementation	227
9.3	Results on birefringent voxel imaging	229
9.3.1	Voxel imaging with a PBS-based diattenuative module	230
9.3.2	Angular dependence of a polarising beam splitter (PBS) and its impact on polarisation contrast	231
9.3.3	Improved voxel imaging with a custom polariser	233
9.4	Results on samples beyond birefringent voxels	235
9.4.1	Glass fibres	236
9.4.2	Mycobacteria	237
9.5	Conclusion and outlook	240

10 Conclusion	242
10.1 Summary of thesis achievements	242
10.2 Future possibilities	246
10.3 Research outputs	248
A Simulating a Mie scatterer	249
A.1 Wave vectors and field components	249
A.2 Scattering in the 1–2 system	250
A.3 Rotation from the lab transverse basis to the 1–2 basis	251
Bibliography	254
Reprint Permissions	289

List of Tables

1.1	Comparison between optical disc formats	28
2.1	Common Jones vectors and matrices in linear and circular basis	47
2.2	Sequence of decoupled quantities from sample birefringence to NLSP observables	50
2.3	Physical meanings of false-colour polarisation maps in the HSV colour space . .	53
2.4	Summary of the three linear regressions used for photon-free noise characterisation	64
2.5	Summary of measured quantities from the photon-free characterisation of a PolCam	69
7.1	Mutual information and 2D data density for single-layer sectors at different voxel pitches	176
7.2	Mutual information between defocused images and the ground truth of every layer in three-layer simulations	197
8.1	Voxel symbols and their corresponding Gray code in a 3-bpv scheme	215
9.1	Summary of different types of polarisers with their extinction ratio and wavefront quality	229

List of Figures

1.1	Different types of laser modification in silica glass	30
2.1	Structure of the Sony PolarSens™ PolCam	36
2.2	Interleaved polarisation channel layout of the Sony PolarSens™ PolCam	37
2.3	PolCam measures the projection of the Stokes vector in the normalised linear Stokes plane	39
2.4	Incoherent and coherent mixing of linearly polarised states	42
2.5	False-colour polarisation maps for PolCam measurements	53
2.6	Fourier spectra of standard sensors and PolCams within the Nyquist region	56
2.7	A model of signal and noise for PolCam	60
2.8	Results of photon-free characterisation of a PolCam	66
2.9	Histograms of pixel values and dark current rates of a PolCam due to the DSNU	70
2.10	Histograms of the polarisation-field gain and dark image pixel values in a PFC calibration	73
3.1	Schematic of ZPC microscopy	80
3.2	Schematic of polarisation modulation introduced by the PPC	84
3.3	Comparison of the STFs for bright-field and PPC setups at the same detection NA	92
3.4	Illustration of the ‘shifting pupil’ picture for the STFs of PPC and bright-field setups	94
4.1	Side lobes in the objective pupil plane when simulating the propagation of a plane-wave component, and the Hann window function to minimise the artefact	102
4.2	Eigenvalues of BSM sources for circular and annular apertures of different sizes	109

4.3	Simulated PPC image of voxels, illuminated by a partially coherent BSM source at the condenser aperture plane	110
4.4	Example HG modes of a GSM source	115
4.5	Eigenvalues of a GSM source in the Hermite–Gaussian representation, and transmission efficiencies through circular and annular apertures	116
4.6	Eigenvalues of a GSM source in the Laguerre–Gaussian representation, and transmission efficiencies through circular and annular apertures	118
4.7	Numerical eigenmodes and eigenvalues of the 2D Fredholm equation	123
4.8	Coherent modes of a circularly apertured GSM source, obtained by solving the 1D Fredholm equation	125
4.9	Eigenvalues of 1D numerical modes for a GSM source upon circular and annular apertures	125
4.10	Comparison of cumulative fractional power error for an apertured GSM source represented by a finite set of coherent modes in various models	126
5.1	Schematic of the simulation pipeline	131
5.2	Modular diagram of the simulation pipeline	140
5.3	Schematic of parallelisation of the simulation pipeline using job arrays on an HPC cluster	145
6.1	Optical and electro-mechanical implementation of the PPC system	149
6.2	Typical workflow for real-time processing of PolCam images	152
7.1	A model of communication for optical data storage with birefringent voxels	156
7.2	Block diagram for the voxel image processing pipeline to obtain MI estimations from raw images	162
7.3	Generation and correction of IFOV errors due to an intensity gradient	163
7.4	Sector isolation with morphological operations and Hough transformation	167
7.5	Estimation of voxel pitch with Fourier analysis	169
7.6	False-colour image and constellation diagram for a single-layer voxel sector	175
7.7	Average traces between adjacent voxel measurements in the NLSP	179
7.8	Mutual information sampled from a biased grid	181
7.9	Effect of PolCam shot noise on mutual information	185
7.10	Ray diagram illustrating the axial shift of focus under refractive-index mismatch	187

7.11	Mutual information plotted against depth into glass for a track of 10 layers	188
7.12	Mutual information from a single layer of voxels plotted against defocus distance	190
7.13	Axial distribution of amplitude PSF for a phase-contrast objective computed with McCutchen’s theorem, compared with bright field	192
7.14	Comparison between experiment and simulation results on the degradation of mutual information of a single-layer sample under defocus	194
7.15	Mutual information from a simulated 3-layer sample against defocus from the middle layer	197
7.16	Mutual information for a simulated 2-layer sample at varying interlayer spacings	198
8.1	False-colour visualisations of example Stokes images provided to the decoder as training dataset	209
8.2	Architecture of the multilayer perceptron decoder	210
8.3	Accuracy of voxel-symbol prediction against number of epochs in an MLP network	212
8.4	First-layer weights of a trained MLP network for decoding of 3-bpv voxels	213
8.5	Polarisation-based image augmentation for training on 4-bpv sector of 114×108 voxels	216
8.6	Prediction accuracy for a 4-bpv sector at a voxel pitch of $0.5 \mu\text{m}$, trained with datasets with and without polarisation-based image augmentation	218
9.1	Schematic of a microscopy system with the diattenuative module	227
9.2	Imaging of a sector of birefringent voxels with the PBS-based diattenuative module	230
9.3	Comparison of the angular dependence of the cross-polarised transmission intensity for a PBS and a custom polariser	231
9.4	Imaging of a sector of birefringent voxels with the diattenuative module implemented with the custom polariser	235
9.5	False-colour images of glass fibres taken with a diattenuative module	236
9.6	False-colour images of <i>Mycobacterium smegmatis</i> observed with a diattenuative module	238
9.7	Simulated false-colour image of a Mie scatterer of a diameter of $2 \mu\text{m}$	239

List of Abbreviations and Acronyms

<i>k</i> -NN	<i>k</i> -nearest neighbour.
<i>M. smegmatis</i>	<i>Mycobacterium smegmatis</i> .
<i>Mtb</i>	<i>Mycobacterium tuberculosis</i> .
nD	n-dimensional.
ACCeNT	automated camera characterization via electron noise tool.
ADC	analogue-to-digital converter.
AoLP	angle of linear polarisation.
AST	absolute sensitivity threshold.
BD	Blu-ray disc.
BER	bit error rate.
bpv	bits per voxel.
BSM	Bessel-correlated Schell-model.
CCD	charge-coupled device.
CD	compact disc.
CMOS	complimentary metal oxide on silicon.
CNN	convolutional neural network.
CPU	central processing unit.
CRLB	Cramér–Rao lower bound.
CSDF	cross-spectral density function.
CTF	coherent transfer function.
CUDA	Compute Unified Device Architecture.
CZT	chirp Z-transform.
DFT	discrete Fourier transform.
DN	digital number.
DoA	division of amplitude.
DoAp	division of aperture.
DoFP	division of focal plane.
DoLP	degree of linear polarisation.
DoT	division of time.

DPC	differential phase contrast.
DR	dynamic range.
DSNU	dark signal non-uniformity.
DVD	digital versatile disc.
EMVA	European Machine Vision Association.
FCL	fully-connected layer.
FFC	flat-field correction.
FFT	fast-Fourier transform.
FoV	field-of-view.
FPN	fixed pattern noise.
FWC	full-well capacity.
GEMM	general matrix multiplication.
GPU	graphics processing unit.
GSM	Gaussian Schell-model.
HDD	hard disk drive.
HG	Hermite–Gaussian.
HPC	high-performance computing.
HSV	hue-saturation-value.
I/O	input/output.
IFoV	instantaneous field-of-view.
IQR	interquartile range.
LED	light-emitting diode.
LG	Laguerre–Gaussian.
MFD	motorised focus drive.
MI	mutual information.
ML	machine learning.
MLP	multilayer perceptron.
MPA	micro-polariser array.
NA	numerical aperture.
NLSP	normalised linear Stokes plane.
ODS	optical data storage.
OTF	optical transfer function.
PBS	Portable Batch System.
PBS	polarising beam splitter.

PCIP	polarisation-camera image processing.
PDF	probability density function.
PFC	polarisation-field correction.
PLM	polarised light microscope.
PolCam	polarisation camera.
PPC	polarisation-sensitive phase contrast.
PRNU	photo-response non-uniformity.
PSF	point-spread function.
PTC	photon-transfer curve.
PVA	polyvinyl alcohol.
QPI	quantitative phase imaging.
QWP	quarter-wave plate.
RDS	Research Data Store.
ReLU	rectified linear unit.
RMS	root-mean-square.
SER	symbol error rate.
SLIM	spatial light interference microscopy.
SLM	spatial light modulator.
SNR	signal-to-noise ratio.
STF	Stokes transfer function.
TCC	transmission cross-coefficient.
VIPP	voxel image processing pipeline.
WGP	wired-grid polariser.
WOTF	weak-object transfer function.
ZPC	Zernike's phase contrast.

Chapter 1

Introduction

1.1 Motivation

With advances of the information technology in the last few decades, humanity is generating data at an unprecedented rate [1]. According to white papers from the International Data Cooperation (IDC) [2, 3], a total of 132.4 ZB (1.324×10^{11} TB) of data was produced in 2023, and is predicted to increase to 393.9 ZB in 2028, driving a seemingly exponentially growing need [4] of large-scale data storage.

Existing data-storage solutions, however, are unlikely to meet this demand. Currently, most of the new information produced are stored in magnetic media [5, 6], notably hard disk drives (HDDs) and magnetic tapes. Conventional HDDs are able to offer a capacity of up to 2 TB per disk [7, 8], but they have a limited lifetime as short as 5 years [6, 9], requiring frequent disk scrubbing [5, 10, 11] to detect disk failures, and data migration as frequent as every 2 years [7, 8], which is both expensive and wasteful. Magnetic tapes, in contrast, offer a more cost-effective solution [12] for large-scale storage of archival data. Nevertheless, they require complex libraries with tape robots and drives that require constant maintenance [13], and has a typical lifetime of 10 years under careful environmental control [6, 12]. Although improvements have been proposed to increase the data density and throughput [14, 15], these technologies are still typically associated with a lifetime of a few decades.

Table 1.1: Comparison between optical disc formats [16].

Format	Wavelength (nm)	Numerical aperture	Data density per layer (Mbit mm ⁻²)	Number of layers
Compact disc (CD)	780	0.45	0.605	1
Digital versatile disc (DVD)	650	0.60	4.29	≤2
Blu-ray disc (BD)	405	0.85	24.6	≤4

In order to preserve the tremendous amount of information and knowledge for future generations to come, a scalable means of data storage is urgently needed. In particular, an ideal technology should offer sufficiently high capacity for the continued growth of global data, while remaining stable over long time periods without frequent maintenance or migration. In addition, the storage medium should be inexpensive and abundant, both for cost-effectiveness and for sustainability at scale. These requirements motivate polarisation-multiplexed optical data storage (ODS), where information can be kept for long-term, high-capacity storage.

1.2 Optical data storage

The first widely adopted form of ODS was compact discs (CDs), which was developed during the 1970s and first released commercially in 1982 [17]. In CDs, data is encoded by varying the thickness of a polycarbonate layer in front of a reflective surface. By scanning a laser across the data pits, information can be retrieved by observing changes in the reflected intensity with a photodiode, offering a data capacity of 680 MB [17, 18] for a disc diameter of 120 mm. Since then, digital versatile discs (DVDs) and Blu-ray discs (BDs) have been developed to improve the data density by employing lasers at shorter wavelengths and objectives with higher numerical apertures (NAs) [16], thereby reducing the diffraction-limited spot size. In addition, the total data capacity can be increased by stacking multiple layers of data pits, and focusing on the layer of interest during reading. Nevertheless, capacity improvements are limited by practical constraints. Further reducing the wavelength poses challenges for the optical media, while dramatic increases of NA are not achievable without immersion. The number of layers is also limited by the transparency of polycarbonate. As a result, practical optical discs have a total

capacity of less than 1 TB, which is below the multi-terabyte capacities available from HDDs, making the optical discs rarely used for large-scale data storage nowadays [7].

Polarisation-multiplexed ODS was proposed to address this limitation by using the polarisation state as additional degrees of freedom for encoding information [19–22]. By introducing anisotropic nanostructures in a transparent medium, local birefringence can modulate the polarisation state of light through its azimuth and retardance, encoding a certain amount of information depending on the number of possible polarisation modulations. The resulting *voxels* serve as the basic storage elements of polarisation-multiplexed ODS. Many voxels can be arranged on a 2D grid to form a *sector*, enabling readout of thousands of voxels within a single field-of-view (FoV). Multiple sectors can then be stacked vertically to form a *track*, enabling multilayer storage and readout, and the concept of 5D-ODS [19, 23, 24] with two polarisation dimensions and three spatial dimensions.

Glass, in particular, has been identified as a medium with great potential due to its low cost, robustness, and a storage lifetime estimated to exceed 10,000 years [23, 25]. Critically, local refractive-index changes can be created by focusing femtosecond laser pulses in glass [26, 27], and the resulting modification can exhibit birefringence that is influenced by polarisation control of the writing laser [28, 29]. Although the modifications are permanent and the medium is therefore read-only after writing, this is not necessarily a limitation for archival data storage, where most data is rarely accessed or modified once written [6].

Figure 1.1 summarises several notable types of modification in silica glass. Type-I modification was first observed as an isotropic increase of refractive index [26, 33]. Although the phase signals introduced by these nanostructures are also valuable in ODS [25, 34], the lack of polarisation multiplexing limits the achievable data density. By incorporating polarisation control of the writing laser, self-assembled nanogratings (type-II modification) [23, 35, 36] can be produced, as shown in Figure 1.1(b). The nanogratings exhibit birefringence depending on the polarisation state of the writing laser, enabling polarisation-multiplexed ODS. Nevertheless, the relatively high transmission loss associated with type-II modification limits the number of layers that can practically be stacked in realistic applications.




Figure 1.1: Different types of laser modification in silica glass. (a) Type I: isotropic refractive index change. (b) Type II: birefringent nanograting. (c) Type X: oblate nanopores. (d) Type S: edge-modified nanovoid. Credits: (a) Reprinted from Miura *et al.*, *Appl. Phys. Lett.* **71**, 3329–3331 (1997), with the permission of AIP Publishing [26]. (b) Reproduced from Shimotsuma *et al.*, *Jpn. J. Appl. Phys.* **44**(7), 4735 (2005), DOI: <https://doi.org/10.1143/JJAP.44.4735>, © The Japan Society of Applied Physics. Reproduced by permission of IOP Publishing Ltd. All rights reserved [30]. (c) Reproduced from Sakakura *et al.* (2020), CC BY 4.0 (<http://creativecommons.org/licenses/by/4.0/>) [31]. (d) Reproduced from Lei *et al.*, *Optica* **8**, 1365 (2021), © 2021 Optical Society of America, under the OSA Open Access Publishing Agreement [32].

To achieve birefringent modification in glass with ultra-low loss [37], type-X [31] and type-S [32] modifications were then demonstrated by precisely controlling writing parameters including pulse energy, number of pulses, repetition rate, pulse duration, and polarisation ellipticity [38]. As shown in Figure 1.1(c), type-X modification appears as a collection of oblate nanopores within the diffraction limit that exhibits an overall birefringence. Type-S modification in Figure 1.1(d), in contrast, is produced when a single round-shaped nanovoid is created with a single energetic pulse, followed by multiple less energetic pulses that modify its edges anisotropically [39, 40], which can be a more energy-efficient process compared to type X.

Both type-S and type-X modifications have ultra-low optical loss, so when they are read with microscopy techniques the out-of-focus layers introduce less scattering and fewer artefacts, even when hundreds of layers are stacked [24, 25], leading to a data capacity as high as 4.8 TB in a 120 mm platter [25]. Nevertheless, the birefringent retardance they introduce is also very small, typically on the order of 10^{-3} rad [38]. A polarisation-resolving microscopy technique

is therefore required that is particularly optimised for imaging weak polarisation modulations, enabling effective readout for ODS.

In this thesis, polarisation-resolving imaging techniques are developed for imaging weak birefringence, enabling readout of birefringent voxels for long-term ODS. The aim is to demonstrate the readout process of polarisation-multiplexed ODS, spanning system design, theoretical and computational modelling, experimental implementation, image processing, and decoding of voxel symbols. This work is carried out in collaboration with Project Silica at Microsoft Research Cambridge [6, 25], where the voxel-writing process is performed and the glass samples used in this thesis are provided. Rather than pushing for an optimal data density that would require broader system-level optimisation, the emphasis is placed on characterising read performance and understanding how it degrades under factors such as noise, aberrations, and crosstalk.

This thesis builds on my MRes work [41], where polarisation-sensitive phase contrast (PPC) was first introduced as a method for imaging weak birefringence using a polarisation camera (PolCam). In the present work, PPC is further developed with improved modelling and simulation under partial coherence, and with an improved experimental implementation for repeatable data acquisition. These developments support quantitative analysis of voxel imaging in an information-theoretic framework, alongside explicit decoding, to evaluate how faithfully symbol information is conveyed. In addition to PPC, a clear-aperture alternative based on diattenuative polarisation modulation is introduced to address limitations inherited from the underlying phase-contrast configuration.

1.3 Thesis structure and overviews of chapters

Chapter 2 starts by establishing a framework for single-shot division of focal plane (DoFP) polarimetry. We carefully identify the capabilities of the PolCam as well as its limitations in terms of accessible polarisation components, spatial resolution, and noise performance. These considerations define the system requirements and motivate the design of the microscopy system.

Chapter 3 then introduces the design and core mechanism of PPC microscopy. An analytical imaging model is presented with polarisation analysis, detailing the imaging process and explaining why polarisation contrast can be enhanced in this configuration.

The analytical model itself, however, relies on an incoherent-illumination assumption, which is fundamentally inconsistent with the partially coherent contrast-generating mechanism of PPC. Chapters 4 and 5 address this problem by developing a numerical modelling approach that accounts for partial coherence using coherent-mode decomposition. Chapter 4 first investigates the coherent-mode structure of apertured Gaussian Schell-model (GSM) sources. Chapter 5 then incorporates this model into a flexible and efficient simulation pipeline, generating results that can be used in complement to experimental results for later informational analysis.

Chapter 6 then focuses on the experimental implementation of the PPC system, including the optical assembly, automation, and software for hardware control and real-time visualisation. Automation and repeatability are prioritised to support the later quantitative analysis and to enable meaningful comparison with simulation results.

With results from both simulations and experiments, Chapter 7 performs quantitative analysis of readout quality from an information-theoretic perspective. An image-analysis pipeline is presented, together with calibration routines, to quantify the amount of information a polarisation image conveys without assuming a particular decoder, whose performance could otherwise confound the system characterisation. This enables analysis of how performance may degrade with respect to factors such as noise, aberrations, and intra- and interlayer crosstalk, through a combination of simulation and experimental results.

Chapter 8 then completes the voxel readout by recovering voxel symbols from polarisation images using a machine learning (ML)-based decoder, which also allows end-to-end ODS performance to be quantified.

Chapter 9 presents an alternative scheme of polarisation modulation that enhances polarisation contrast without introducing the annulus and phase rings required in PPC, allowing clear-aperture imaging of weak birefringent features such as voxels. The chapter also presents

potential applications of the system in areas outside ODS, such as label-free imaging of mycobacteria.

Finally, Chapter 10 summarises the thesis findings and outlines the main limitations and the most direct directions for future work.

Chapter 2

Single-shot polarimetry using pixelated polarisation cameras

Polarisation-resolved imaging enables the extraction of rich structural and material information from samples. While the previous chapter introduced how birefringence can be used to encode information for long-term sustainable optical data storage, it should be recognised that polarisation measurements play a far broader role across the physical and life sciences. In materials science, polarisation-based techniques such as ellipsometry have long been fundamental for determining thin-film thicknesses, refractive indices, and anisotropies [42, 43]. In geology, polarisation microscopy is fundamental for optical mineralogy and petrography, enabling mineral identification and characterisation in rock thin sections [44, 45]. In biological and clinical applications, polarisation contrast enables label-free detection of organised structures from cellular [46] to tissue levels [47], while in fluorescence microscopy, polarimetry can be employed to determine the orientation of fluorophores [48–50].

Starting with a small overview of polarimetry methods, this chapter introduces an approach of single-shot polarimetry based on the division of focal plane (DoFP) architecture, implemented using the Sony PolarSens™ polarisation cameras (PolCams). It will then walk through the physical principles and interpretation, before providing a model of signal and noise for PolCams.

2.1 Division-of-focal-plane (DoFP) polarimetry with the PolarSensTM sensor

To retrieve information encoded within the voxels, it is essential to employ microscopic techniques capable of spatially resolving such densely packed nanostructures. The read-out process, however, cannot be realised in a trivial setup, because measuring the polarisation changes introduced by such birefringent structures requires capturing multiple polarisation components of the optical field. These components may, in the Jones algebra, include amplitude and phase differences for components of the optical field in different directions, or may be interpreted in the Stokes-Müller formalism as various Stokes parameters. Nonetheless, the process is inherently multidimensional because birefringent voxels enable multidimensional data storage.

Conventional imaging sensors, such as charge-coupled device (CCD) or complimentary metal oxide on silicon (CMOS) arrays, record only intensity—a scalar quantity at each spatial location—thereby discarding critical polarisation information. To fully reconstruct the encoded data, some form of polarisation measurement is therefore required.

In imaging polarimetry, a variety of approaches have been developed in which such multiplexed detection is implemented, and they can be broadly categorised into divisions of time, amplitude, aperture, and focal plane.

In the division of time (DoT) methods, multiple images of the specimen are taken, usually with different schemes of polarisation modulation [51, 52], illumination [53], or a combination of both [54, 55]. They are, however, considerably slower, due to not only the multiframe nature of image captures, but also any extra hardware necessary to actively switch between the different schemes, such as liquid crystal retarders, light-emitting diode (LED) arrays, or rotating polarisation optics [56]. In addition, such methods require the specimen to be static, otherwise motion blur and polarisation artefacts can affect the measurements [57]. In the context of optical data storage (ODS), the DoT methods would limit the bandwidth of the read-out process.




Figure 2.1: Structure of the Sony PolarSens™ PolCam. (a) Cross-sectional structure: for each pixel, an on-chip lens focuses the collected light onto the corresponding photodiode through a WGP that analyses the field in a linear polarisation state. (b) Image of WGs in a superpixel. © 2016 IEEE, reprinted from [63].

Division of amplitude (DoA) methods try to tackle these problems by splitting the beams into multiple detection paths, each with different polarisation controls [58, 59]. These methods necessitate careful alignment, extra optics, additional cameras and are generally very sensitive to vibrations [60]. There are also division of aperture (DoAp) techniques, in which light is split at the Fourier planes, typically combines retarder arrays with lens arrays [61] or pyramid prisms [62], to create multiple images on the detection sensors with different polarisation controls. Nevertheless, like the DoA methods, the DoAp ones are also very sensitive to (mis)alignment of the polarisation optics, though they can be more compact and are reported to be more robust against mechanical instabilities.

In recent years, the development of PolCams [63, 64] has offered exciting opportunities in the field of polarimetry. By employing micro-polariser arrays (MPAs) aligned with the sensor pixels, PolCams enable the field to be analysed by polarisers in specific directions before being imaged, enabling single-shot division of focal plane (DoFP) polarimetry, similar to how Bayer filters enable single-shot colour imaging [65]. In particular, the Sony PolarSens™ [66] features a layered design as illustrated in Figure 2.1, where a four-directional array of wired-grid polarisers (WGs) sits between a micro-lens array and the photodiodes of a CMOS sensor. Compared to conventional ‘on-glass’ design where the MPA sits at the other side of the micro-lenses, the improved ‘on-chip’ design features an improved extinction ratio, because the smaller distance between the MPA and the photodiode reduces any optical crosstalk between adjacent pixels [63].

The layout of the PolarSens™ PolCam is illustrated in Figure 2.2. The transmission axes of the

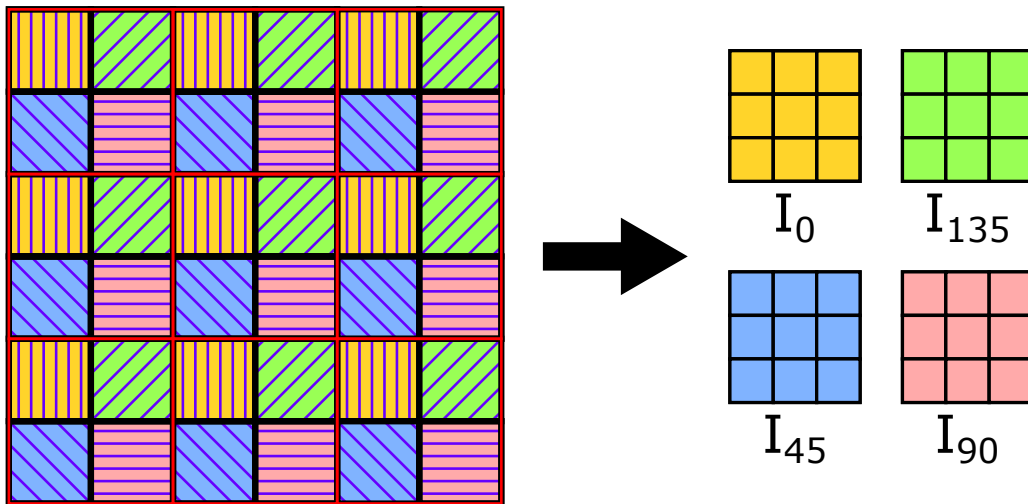


Figure 2.2: Polarisation channel layout of the Sony PolarSens™ PolCam. The micro-polariser array divides the sensor into 2×2 superpixels, each consisting of four pixels analysed at 0° , 45° , 90° , and 135° . Note the lines in each pixel denotes the orientation of the WGP, which is perpendicular to its axis of transmission.

linear polarisers in the array are oriented at four angles (0° , 45° , 90° , and 135°), and the pixels are arranged into 2×2 superpixels (marked with red outlines), each containing one pixel for each polarisation direction. Therefore, the sensor equivalently has four interleaved polarisation channels, and a single shot simultaneously captures four sub-images (I_0 , I_{45} , I_{90} , and I_{135}) where the field is analysed in different linear polarisation states.

One significant advantage of PolCam-enabled DoFP polarimetry is the simplicity of the hardware, which can be as trivial as replacing the regular camera with a PolCam and inserting a polariser for illumination [67]. As part of my MRes thesis [41], we also presented polarisation-sensitive phase contrast (PPC) microscopy which combines PolCam with Zernike's phase contrast (ZPC) [68] to realise simultaneous phase and polarisation imaging in a single shot. In addition to being cost-efficient, the absence of extra optics and splitting of detection paths makes the system more tolerant to misalignment and mechanical instability.

It should be noted that the DoFP comes at its own cost. Firstly, for each polarisation sub-image, pixel pitch is effectively doubled, due to the interleaving scheme of different channels. Therefore, it becomes more challenging to achieve sufficient sampling of the image, where the Nyquist criterion usually needs to be met against the optical transfer function (OTF) of the system [69], or any particular spatial frequency of interest. In our case of voxel imaging, it

is particularly important that the superpixel of the PolCam can be projected by the imaging system to Nyquist-sample the voxel pitch. This will be discussed in more detail in Section 2.5.

Secondly, the interleaving scheme of polarisation channels introduces a spatial offset among them, up to one pixel pitch in each direction. This is fundamental to DoFP methods—multiplexing polarisation information spatially means there is an ambiguity between polarisation signal and intra-superpixel intensity gradient. If left uncorrected, this ambiguity results in instantaneous field-of-view (IFoV) errors [60]. Computational methods have been developed to mitigate these errors [70], as will be discussed in Section 7.3.1.

2.2 Measuring linear components of polarisation

It is critical to understand what polarisation information PolCams can provide, where the Stokes-Müller formalism can be useful. For every superpixel, three out of the four Stokes parameters can be computed [41], by simple linear combination of intensities from different channels, given by

$$S_0 = \frac{1}{2} (I_0 + I_{45} + I_{90} + I_{135}), \quad (2.1)$$

$$S_1 = I_0 - I_{90}, \quad (2.2)$$

$$S_2 = I_{45} - I_{135}, \quad (2.3)$$

where S_0 , S_1 , and S_2 are the first three Stokes parameters. The last Stokes parameter, S_3 , cannot be obtained, because PolCam is agnostic to the polarisation handedness due to the absence of any retarders or circular polarisers.

In the Stokes-Müller formalism, any polarisation state can be represented by a normalised Stokes vector, $\hat{\mathbf{S}} = [\hat{S}_1, \hat{S}_2, \hat{S}_3]^\top$, where \hat{S}_i are normalised Stokes parameters defined as S_i/S_0 . When the field is fully polarised, $\hat{\mathbf{S}}$ sits on the surface of the Poincaré sphere, as illustrated in Figure 2.3(a). For PolCam measurements, the lack of S_3 (and thus the normalised \hat{S}_3) reduces the 3D vector into its 2D projection, $\hat{\mathbf{s}}$, that sits in the \hat{S}_1 - \hat{S}_2 plane, as shown in Figure 2.3(b).

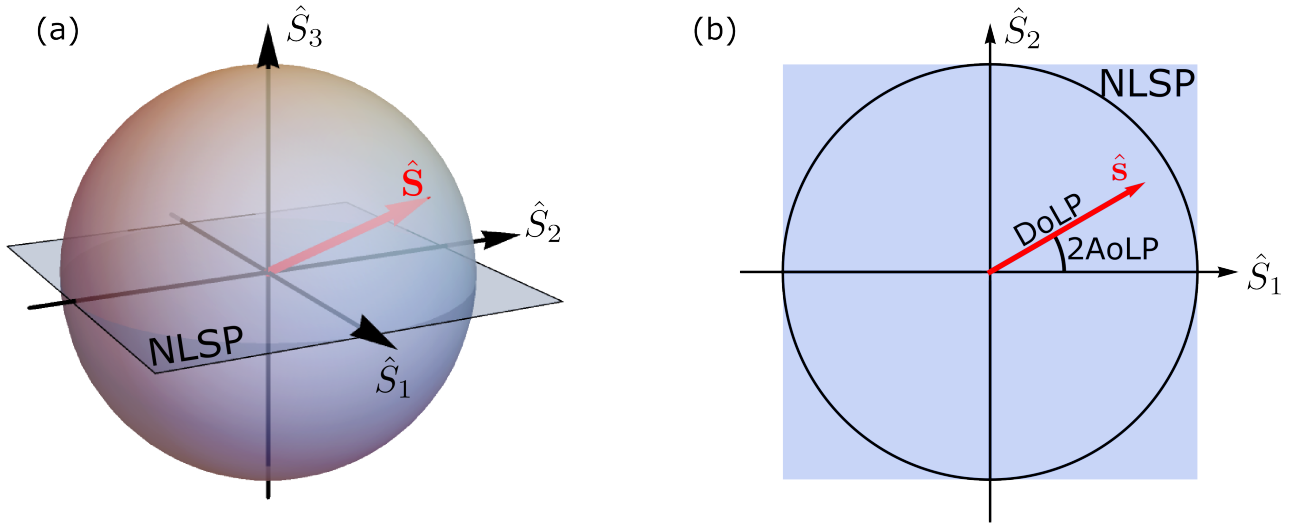


Figure 2.3: PolCam measures the projection of the Stokes vector in the normalised linear Stokes plane (NLSP). (a) A polarisation state to be measured, represented by a normalised Stokes vector $\hat{\mathbf{S}}$ on the Poincaré sphere. (b) PolCam sees the projection of $\hat{\mathbf{S}}$ in the NLSP, $\hat{\mathbf{s}}$, missing any direct measurements in \hat{S}_3 .

The plane contains only linear components of the polarisation states and will therefore be called the normalised linear Stokes plane (NLSP) for the rest of this thesis. Therefore, the NLSP represents the full measurement space captured by the PolCam in response to polarisation.

The Poincaré sphere itself is intersected by the NLSP, forming a unit circle, as shown in Figure 2.3(b). States in this circle represent linear polarisations, and the length of $\hat{\mathbf{s}}$ represents the degree of linear polarisation (DoLP), which can be computed as

$$\text{DoLP} = \frac{\sqrt{S_1^2 + S_2^2}}{S_0}. \quad (2.4)$$

Note that a fully polarised field is not necessarily measured with a DoLP of unity. For example, the circular polarisation can be fully polarised but carries a DoLP of zero. In general, any state with a non-zero ellipticity cannot have a DoLP of unity.

Similarly, the polar azimuth of $\hat{\mathbf{s}}$ also carries important physical meaning, as it is twice the angle of linear polarisation (AoLP) [41], by construction of the Poincaré sphere. The AoLP can thus be computed as

$$\text{AoLP} = \frac{1}{2} \arctan \frac{S_2}{S_1}. \quad (2.5)$$

The extra factor of two comes from the relationship between polarisation azimuth and the

longitude on the Poincaré sphere.

2.3 Mixing of polarisation states

The implication of the Stokes vector projection is far from trivial, especially in the context of polarisation imaging. One may be tempted to assume that, when the illumination is fully polarised and no depolarisation is present in the optical system, the detected field should also be fully polarised. Based on that, it may seem possible to recover the full normalised Stokes vector $\hat{\mathbf{S}}$ from its projection $\hat{\mathbf{s}}$ by extending it along the \hat{S}_3 dimension. In this case, we would get

$$\hat{S}_3 = \pm\sqrt{1 - \text{DoLP}^2}, \quad (2.6)$$

which gives the full polarisation state except for its handedness (as the sign of \hat{S}_3 is unknown).

This idea only works if the detected field is really fully polarised, which is not usually true in practice. Even if the illumination is fully polarised and the optical system has no depolarisation in it, the field at the PolCam may still be only partially polarised due to incoherent mixing between different polarisation states. For example, in fluorescence and transmission imaging, the field is typically spatially incoherent, so light from different spatial locations combines incoherently during image formation in a finite-resolution system. This incoherent superposition can lead to an averaging effect on the polarisation states [71], resulting in measured partial polarisation. A similar depolarising effect can also happen temporally when the polarisation state is varying during the finite exposure time of the sensor, since PolCam integrates intensity instead of the complex field.

To illustrate more clearly how this ambiguity arises, and what role coherence plays, we consider the following example, where eight linear polarisation states (P_0 – P_7) mix with each other. In the context of ODS, the polarisation states can be interpreted as the result of modulation from eight types of birefringent voxels in a 3-bits per voxel (bpv) encoding scheme. The voxels can then be assumed to modulate some circularly polarised illumination by quarter waves, represented by the Jones matrix of a linear retarder $\mathbb{L}\mathbb{R}$, with a retardance $\delta = \pi/2$ and a

fast-axis orientation of $\gamma_\alpha = \pi\alpha/8$ parametrised by the symbol $\alpha \in \{0, \dots, 7\}$ it represents. Therefore, the polarisation state P_α can be represented by its Jones vector \mathbf{E}_α as

$$\mathbf{E}_\alpha = \mathbb{L}\mathbb{R} \left(\gamma_\alpha = \frac{\pi\alpha}{8}, \delta = \frac{\pi}{2} \right) |L\rangle, \quad (2.7)$$

where $|L\rangle = [1, i]^T / \sqrt{2}$ denotes the (normalised) Jones vector for left-handed circular polarisation.

Following the negative phase sign convention [72], electric field in this work is defined as

$$\mathbf{E}(\mathbf{r}, t) = \begin{bmatrix} \tilde{E}_x \\ \tilde{E}_y \end{bmatrix} e^{i(\mathbf{k}\cdot\mathbf{r} - \omega t - \phi)}, \quad (2.8)$$

at a position \mathbf{r} and the time t , where \mathbf{k} denotes the wave vector, ω denotes the frequency, ϕ denotes a phase offset of the wave, and \tilde{E}_x and \tilde{E}_y denotes the complex amplitudes in the x and y-directions, respectively.

In this formalism, the Jones matrix of a linear retarder $\mathbb{L}\mathbb{R}$ can be written as

$$\mathbb{L}\mathbb{R}(\gamma, \delta) = \begin{bmatrix} \cos \frac{\delta}{2} - i \cos 2\gamma \sin \frac{\delta}{2} & -i \sin 2\gamma \sin \frac{\delta}{2} \\ -i \sin 2\gamma \sin \frac{\delta}{2} & \cos \frac{\delta}{2} + i \cos 2\gamma \sin \frac{\delta}{2} \end{bmatrix}. \quad (2.9)$$

By substituting the expression into Equation 2.7, the Jones vector \mathbf{E}_α can be evaluated to be

$$\mathbf{E}_\alpha = e^{i\gamma'_\alpha} \begin{bmatrix} \cos \gamma'_\alpha \\ \sin \gamma'_\alpha \end{bmatrix}, \quad (2.10)$$

where

$$\gamma'_\alpha := \gamma_\alpha - \frac{\pi}{2}, \quad (2.11)$$

by definition.

As expected by construction, the result of Equation 2.10 is in the form of linear polarisation, with an extra global phase of γ'_α which cannot be neglected as one shall see. The offset of

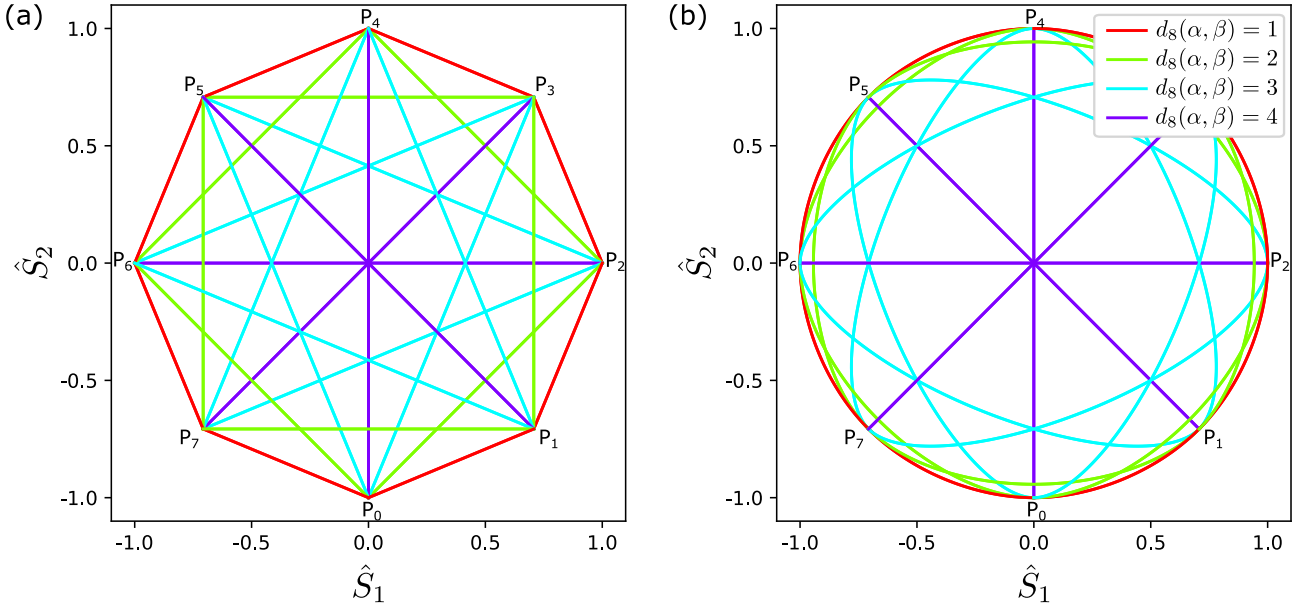


Figure 2.4: Incoherent (a) and coherent (b) mixing of linearly polarised states, P_0 – P_7 . For each pair of states, the projection of the resulting normalised Stokes vector in the NLSP depends on the relative weight $w \in [0, 1]$ between them. The resulting loci traced as w varies are plotted and colour-coded according to the eight-cyclic distance between the symbols, $d_8(\alpha, \beta)$.

γ'_α from γ_α in Equation 2.11 is also expected, and is discussed extensively in Chapter 4 of my MRes thesis [41].

The Stokes parameters can in general be computed from the Jones vector $[\tilde{E}_x, \tilde{E}_y]^T$ [73], according to

$$S_0 = \tilde{E}_x \tilde{E}_x^* + \tilde{E}_y \tilde{E}_y^*, \quad (2.12)$$

$$S_1 = \tilde{E}_x \tilde{E}_x^* - \tilde{E}_y \tilde{E}_y^*, \quad (2.13)$$

$$S_2 = 2\Re(\tilde{E}_x \tilde{E}_y^*), \quad (2.14)$$

$$S_3 = -2\Im(\tilde{E}_x \tilde{E}_y^*). \quad (2.15)$$

As a result, the normalised Stokes vector $\hat{\mathbf{S}}_\alpha$ can be evaluated as $\hat{\mathbf{S}}_\alpha = [\cos 2\gamma'_\alpha, \sin 2\gamma'_\alpha, 0]^T$, which sits evenly around the equator of the Poincaré sphere (i.e. around the unit circle in the NLSP as shown in Figure 2.4).

Now consider P_α to be mixed with another one of the eight states P_β into a new state $P_{\alpha+\beta}$. When they are mixed incoherently, their Stokes vector adds up, resulting in a weighted average

of their normalised Stokes vector as

$$\hat{\mathbf{S}}_{\alpha+\beta}(w) = w\hat{\mathbf{S}}_{\alpha} + (1-w)\hat{\mathbf{S}}_{\beta}, \quad (2.16)$$

where $w \in [0, 1]$ parametrises the relative weight between the two states. The locus of $\hat{\mathbf{S}}_{\alpha+\beta}(w)$ is a straight line connecting P_{α} and P_{β} , as shown in Figure 2.4(a). Since both P_{α} and P_{β} lie in the NLSP, the locus connecting them is also in this plane, with a corresponding \hat{S}_3 component of zero. Note the loci in Figure 2.4(a) are colour-coded according to the eight-cyclic distance, defined as

$$d_8(\alpha, \beta) := \min(|\alpha - \beta|, 8 - |\alpha - \beta|). \quad (2.17)$$

Now consider P_{α} and P_{β} to be mixed coherently (i.e. interference). Instead of the Stokes vectors, their Jones vectors now need to be averaged with weights, viz.,

$$\mathbf{E}_{\alpha+\beta}(w) = w\mathbf{E}_{\alpha} + (1-w)\mathbf{E}_{\beta}. \quad (2.18)$$

The coherently mixed normalised Stokes vector $\hat{\mathbf{S}}_{\alpha+\beta}(w)$ can then be computed again with Equations 2.12–2.15, and the loci traced by its projection $\hat{\mathbf{s}}_{\alpha+\beta}(w)$ in the NLSP are plotted in Figure 2.4(b), for every pair of α and β .

Compared with the incoherent case, the loci for the coherent mixture can be seen to be distinctly different. Although in both cases they connect the states P_{α} and P_{β} , the lines can be much more curved in the coherent case. This is due to the fact that the coherently mixed state $P_{\alpha+\beta}$ is well-described by a Jones vector and therefore is always fully polarised, which forces the corresponding normalised Stokes vector $\hat{\mathbf{S}}_{\alpha+\beta}$ to travel on the surface of the Poincaré sphere, resulting in curved projections of $\hat{\mathbf{s}}_{\alpha+\beta}$ in the NLSP.

It can also be seen how ambiguity can arise from such PolCam detection. For example, for mixtures of orthogonally polarised states (i.e. $d_8(\alpha, \beta) = 4$), in both the lines appear to be straight and the projections in the NLSP appear to be identical. In the coherent case, however, the Stokes vector $\hat{\mathbf{S}}_{\alpha+\beta}$ travels along the great-circle arc connecting P_{α} and P_{β} via the left-

handed circularly polarised state at the ‘South Pole’ of the Poincaré sphere. This can be shown by noting that $\alpha + \beta = \pi$ for orthogonal polarisations,

$$\mathbf{E}_{\alpha+\beta}(w = 0.5) = \frac{1}{2} \left(e^{i\gamma'_\alpha} \begin{bmatrix} \cos \gamma'_\alpha \\ \sin \gamma'_\alpha \end{bmatrix} + e^{i\left(\gamma'_\alpha + \frac{\pi}{2}\right)} \begin{bmatrix} \cos\left(\gamma'_\alpha + \frac{\pi}{2}\right) \\ \sin\left(\gamma'_\alpha + \frac{\pi}{2}\right) \end{bmatrix} \right) \propto |R\rangle, \quad (2.19)$$

where the proportionality (rather than equality) reflects that Equation 2.18 does not conserve total energy.

In contrast, for an equal but incoherent mixture of orthogonal polarisation states, the resultant $\hat{\mathbf{S}}_{\alpha+\beta}$ is simply a zero vector representing the unpolarised state.

Another key observation is that, in the coherent case, the specific arc for $\hat{\mathbf{S}}_{\alpha+\beta}$ is dependent on the relative phase between the two polarisation states, which is the reason why the global phase of γ'_α in Equation 2.10 cannot be ignored. For instance, for the coherent mixture of horizontally and vertically linearly polarised states (P_6 and P_2), if the complex phase is omitted, its Stokes vector would take another great-circle arc through P_6 – P_7 – P_0 – P_1 – P_2 . In fact, depending on their relative phase, P_6 and P_2 can be mixed into any fully polarised state because their Jones vectors $[1, 0]^T$ and $[0, 1]^T$ form an orthonormal basis for the 2D complex vector space \mathbb{C}^2 in which all Jones vectors reside. This is a powerful concept which enables all analysis in the circular basis in the next section.

It should also be noted that practical microscopy usually involves partial coherence, especially in the spatial domain due to light sources and apertures, such that the reality usually lies in between these two cases. Nonetheless, analysis involving partial coherence would not be as straightforward, because neither amplitudes in Jones vectors nor intensities in Stokes vectors is acted upon linearly by the optical system. Second-order statistical treatment would be necessary, as will be presented in Chapter 4.

2.4 PolCam polarimetry in the circular basis

Since a PolCam directly measures only the linear components of polarisation, rather than the circular component required for full measurements of Stokes vectors, one approach in PolCam polarimetry is to illuminate the sample with circularly polarised light in transmission microscopes [41, 67]. In such cases, the field starts with no DoLP, and any detection of linear components suggests the presence of polarisation signals. The circularly polarised background would also enable the uses of false-colour polarisation mapping schemes that are intuitive and easy to interpret.

In addition, circular states project to the origin of the NLSP. Illumination with circular polarisation therefore provides a maximal range over which the polarisation state can evolve in any direction. In the presence of birefringence, circularly polarised illumination ensures a well-defined polarisation signal in the orthogonal state, as we will see later in this section. In this setting, the circular basis of the Jones calculus is particularly convenient, as it simplifies the algebra and shows how birefringent azimuth and retardance can be decoupled in PolCam measurements.

2.4.1 Jones algebra in the circular basis

As discussed in the previous section, Jones vectors are defined in the \mathbb{C}^2 space. It is therefore equally valid to use any basis that spans this space to describe its polarisation algebra. Consider representing Jones vectors in a circular basis, defined as

$$\mathbf{E}_{\text{circ}} := \begin{bmatrix} \tilde{E}_L \\ \tilde{E}_R \end{bmatrix}_c, \quad (2.20)$$

where \tilde{E}_L and \tilde{E}_R are the complex amplitudes of the field in the left- and right-handed polarisation components, respectively. The subscript c denotes the circular basis, and for distinction, Jones vectors written in their usual linear x - y basis are assigned with the subscript l . The basis vectors in the circular basis are then $|L\rangle$ and $|R\rangle$, where $|L/R\rangle = [1, \pm i]_l^T / \sqrt{2}$ is orthogonal,

i.e.,

$$\langle L|R\rangle = \langle R|L\rangle = 0, \quad (2.21)$$

$$\langle L|L\rangle = \langle R|R\rangle = 1. \quad (2.22)$$

Any Jones vector can be easily transformed from the linear basis into the circular one with a simple matrix multiplication [72], given by

$$\mathbf{E}_{\text{circ}} = \mathbb{U}_{\text{lc}} \mathbf{E}_{\text{lin}}, \quad (2.23)$$

where

$$\mathbb{U}_{\text{lc}} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -i \\ 1 & i \end{bmatrix}. \quad (2.24)$$

Thus, the transformation from the circular basis into the linear one can be represented by its matrix inverse, $\mathbb{U}_{\text{lc}} := \mathbb{U}_{\text{cl}}^{-1}$, and due to unitarity,

$$\mathbb{U}_{\text{cl}} = \mathbb{U}_{\text{lc}}^{-1} = \mathbb{U}_{\text{lc}}^\dagger = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ i & -i \end{bmatrix}, \quad (2.25)$$

where \dagger denotes the Hermitian conjugation (i.e. $\mathbb{A}^\dagger := (\mathbb{A}^*)^T$).

Table 2.1: Common Jones vectors and matrices in linear and circular basis.

	Linear basis	Circular basis
Left-handed circular, $ L\rangle$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}_l$	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}_c$
Right-handed circular, $ R\rangle$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}_l$	$\begin{bmatrix} 0 \\ 1 \end{bmatrix}_c$
Horizontal, $ H\rangle$	$\begin{bmatrix} 1 \\ 0 \end{bmatrix}_l$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}_c$
Vertical, $ V\rangle$	$\begin{bmatrix} 0 \\ 1 \end{bmatrix}_l$	$\frac{-i}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}_c$
Diagonal / Anti-diagonal, $ D/A\rangle$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ \pm 1 \end{bmatrix}_l$	$\frac{1}{2} \begin{bmatrix} 1 \mp i \\ 1 \pm i \end{bmatrix}_c$
Linear at angle θ , $ \theta\rangle$	$\begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix}_l$	$\frac{1}{\sqrt{2}} \begin{bmatrix} e^{-i\theta} \\ e^{i\theta} \end{bmatrix}_c$
Elliptical with azimuth θ and ellipticity ϵ , $ \theta, \epsilon\rangle$	$\begin{bmatrix} \cos \theta \cos \epsilon + i \sin \theta \sin \epsilon \\ \sin \theta \cos \epsilon - i \cos \theta \sin \epsilon \end{bmatrix}_l$	$\begin{bmatrix} \cos(\epsilon + \frac{\pi}{4}) e^{-i\theta} \\ \cos(\epsilon - \frac{\pi}{4}) e^{i\theta} \end{bmatrix}_c$
Linear polariser at angle γ , $\mathbb{P}(\gamma)$	$\begin{bmatrix} \cos^2 \gamma & \cos \gamma \sin \gamma \\ \cos \gamma \sin \gamma & \sin^2 \gamma \end{bmatrix}_l$	$\frac{1}{2} \begin{bmatrix} 1 & e^{-2i\gamma} \\ e^{2i\gamma} & 1 \end{bmatrix}_c$
Linear retarder of retardance δ and fast-axis angle γ , $\mathbb{LR}(\delta, \gamma)$	$\begin{bmatrix} \cos \frac{\delta}{2} - i \cos 2\gamma \sin \frac{\delta}{2} & -i \sin 2\gamma \sin \frac{\delta}{2} \\ -i \sin 2\gamma \sin \frac{\delta}{2} & \cos \frac{\delta}{2} + i \cos 2\gamma \sin \frac{\delta}{2} \end{bmatrix}_l$	$\begin{bmatrix} \cos \frac{\delta}{2} & -ie^{-2i\gamma} \sin \frac{\delta}{2} \\ -ie^{2i\gamma} \sin \frac{\delta}{2} & \cos \frac{\delta}{2} \end{bmatrix}_c$

Similarly, Jones matrices can be rewritten in their circular basis as

$$\mathbb{T}_{\text{circ}} = \mathbb{U}_{\text{cl}}^\dagger \mathbb{T}_{\text{lin}} \mathbb{U}_{\text{cl}}. \quad (2.26)$$

Using Equations 2.23 and 2.26, common Jones vectors and matrices can be written in the circular basis, as summarised in Table 2.1.

2.4.2 Polarimetric interpretation of the circular-basis formalism

Generation of polarisation signal

Although mathematically equivalent, the circular-basis formalism of the Jones algebra carries much physical insight into our polarimetry process. Consider a birefringent sample, acting as a linear retarder described by the following Jones matrix:

$$\mathbb{L}\mathbb{R}(\delta, \gamma) = \begin{bmatrix} \cos \frac{\delta}{2} - i \cos 2\gamma \sin \frac{\delta}{2} & -i \sin 2\gamma \sin \frac{\delta}{2} \\ -i \sin 2\gamma \sin \frac{\delta}{2} & \cos \frac{\delta}{2} + i \cos 2\gamma \sin \frac{\delta}{2} \end{bmatrix}_l = \begin{bmatrix} \cos \frac{\delta}{2} & -ie^{-2i\gamma} \sin \frac{\delta}{2} \\ -ie^{2i\gamma} \sin \frac{\delta}{2} & \cos \frac{\delta}{2} \end{bmatrix}_c \quad (2.27)$$

It can be observed that the Jones matrix is much more compact in the circular basis. The birefringent azimuth γ only appears as complex phases in the off-diagonal terms, while the retardance δ only affects the modulus of each term. In fact, Equation 2.27 can be decomposed neatly as

$$\mathbb{L}\mathbb{R}(\delta, \gamma)_{\text{circ}} = \cos \frac{\delta}{2} \mathbb{I} - i \sin \frac{\delta}{2} \begin{bmatrix} 0 & e^{-2i\gamma} \\ e^{2i\gamma} & 0 \end{bmatrix}, \quad (2.28)$$

where \mathbb{I} denotes the identity matrix. The second term is purely anti-diagonal, and is a rotated form of the Pauli- x matrix [74], representing the generation of a circular polarisation component of the opposite handedness from the illumination. For example, upon left-handed circularly

polarised illumination, the resulting Jones vector is

$$\begin{bmatrix} \tilde{E}_L^{\text{out}} \\ \tilde{E}_R^{\text{out}} \end{bmatrix}_c = \mathbb{L}\mathbb{R}(\delta, \gamma)|L\rangle = \cos \frac{\delta}{2}|L\rangle - ie^{2i\gamma} \sin \frac{\delta}{2}|R\rangle, \quad (2.29)$$

where the second term can be interpreted as a right-handed polarisation signal generated from sample birefringence, which superposes with the left-handed background in the first term. In this picture, a decoupling happens between the two degrees of freedom of birefringence, when the relative amplitude between the signal and the background depends only on the sample retardance, whereas their relative phase depends only on the birefringent azimuth, i.e.,

$$\delta = 2 \arctan \left| \frac{\tilde{E}_R^{\text{out}}}{\tilde{E}_L^{\text{out}}} \right|, \quad (2.30)$$

$$\gamma = \frac{1}{2} \arg \left(\frac{\tilde{E}_R^{\text{out}}}{\tilde{E}_L^{\text{out}}} \right) - \frac{\pi}{4}. \quad (2.31)$$

In particular, in the imaging of weakly birefringent samples such as voxels, Equation 2.29 can be approximated as

$$|E_{\text{out}}\rangle \approx |L\rangle - \frac{i\delta}{2}e^{2i\gamma}|R\rangle, \quad (2.32)$$

and the polarimetric measurement reduces to measuring the phase and amplitude of the orthogonal polarisation component, \tilde{E}_R^{out} , relative to the left-handed circular background. This draws a close resemblance to ZPC, where the sample phase ϕ is usually assumed to be small, such that the output field is considered as the sum the background with a phase-shifted signal whose amplitude is proportional to ϕ , i.e.,

$$e^{i\phi} \approx 1 + i\phi. \quad (2.33)$$

Due to this analogy, which will become more relevant in Chapter 3 in discussions of PPC, we shall call the co-polarised background the ‘surround (S) wave’ and the orthogonally polarised signal the ‘diffracted (D) wave’ which is the commonly used terminology in ZPC [68, 69].

Table 2.2: Sequence of decoupled quantities from sample birefringence to NLSP observables.

Stage	Quantity 1	Quantity 2
Sample birefringence	Birefringent retardance	Birefringent azimuth
Circular polarisation components	Relative amplitude	Relative phase
Polarisation ellipse	Ellipticity	Major-axis orientation
NLSP measurement	DoLP	AoLP

Detection of polarisation signal

Similarly to the Jones matrix $\mathbb{L}\mathbb{R}$, there is also a decoupling in Jones vectors between the polarisation azimuth and ellipticity in the circular basis. Consider a general polarisation state of major-axis angle of θ and ellipticity ϵ , then

$$|\theta, \epsilon\rangle = \begin{bmatrix} \cos \theta \cos \epsilon + i \sin \theta \sin \epsilon \\ \sin \theta \cos \epsilon - i \cos \theta \sin \epsilon \end{bmatrix}_l = \begin{bmatrix} \cos \left(\epsilon + \frac{\pi}{4} \right) e^{-i\theta} \\ \cos \left(\epsilon - \frac{\pi}{4} \right) e^{i\theta} \end{bmatrix}_c. \quad (2.34)$$

It is clear that, in the circular basis, polarisation azimuth and ellipticity corresponds respectively to the relative amplitude and phases between the orthogonal circular components \tilde{E}_L and \tilde{E}_R .

In addition, this sequence of decoupled quantities actually goes further, because the latitude of the state $|\theta, \epsilon\rangle$ on the Poincaré sphere is 2ϵ , and therefore in the NLSP,

$$\text{DoLP} = \cos 2\epsilon, \quad (2.35)$$

$$\text{AoLP} = \theta. \quad (2.36)$$

In summary, in cases when illumination is circular, there is a sequence of decoupled relationships from sample birefringence to NLSP observables, through which retardance and azimuthal information is decoupled, as summarised in Table 2.2. This offers a powerful tool for building an imaging model in later sections of the thesis when specific polarisation microscopy setups are considered.

PolCam as a four-step phase shifting interferometer

The circular-basis analysis also provides another perspective to understand the action of PolCam in polarimetry. Recall that upon circularly polarised illumination, measuring birefringence orientation is equivalent to measuring the phase of an orthogonally polarised signal from the background, which turns it into a problem of quantitative phase imaging (QPI), an active area of research itself [75].

Among various QPI methods, spatial light interference microscopy (SLIM) is particularly relevant, due to its interferometric nature and phase-shifting technique [76]. It operates by spatially separating the sample field (i.e. field whose phase is to be measured) from the reference field, and introduces an extra phase shift in the infinity space (e.g. between an infinity-corrected objective and the tube lens [77]). The amount of phase shift is usually controlled by either a liquid-crystal retarder [76] or a spatial light modulator (SLM). For a phase shift of Δ , the detected intensity would take the form [78]

$$I_{\Delta} = I_r + I_s + 2E_r E_s \cos(\phi + \Delta), \quad (2.37)$$

where ϕ is the sample phase as before, and E_r , E_s , I_r , and I_s are the amplitudes and intensities of the reference and sample fields. By measuring the interference intensity at a set of different phases, a quantitative phase map can be reconstructed. In particular, in the four-step SLIM scheme [79], the phase retrieval algorithm becomes

$$\phi = \arctan \frac{I_{3\pi/2} - I_{\pi/2}}{I_0 - I_{\pi}}. \quad (2.38)$$

This closely resembles Equation 2.5 which described how AoLP is estimated from PolCam measurements. In fact, their equivalence can be seen by considering the measured signal, I_p , for each polarisation channel on PolCams as

$$I_p = |\langle \theta_p | E \rangle|^2, \quad (2.39)$$

where p denotes the polarisation channels, and the analysed field $|E\rangle$ in this case takes the general form of $[E_R, E_L e^{i\phi}]^T$. Note that E_R and E_L are both defined real without loss of generality. Thus, by substituting the expressions for linearly polarised states in Table 2.1, one can compute that

$$I_0 = |\langle H|E\rangle|^2 = \frac{1}{2} (E_R^2 + E_L^2 + 2E_R E_L \cos \phi), \quad (2.40)$$

$$I_{135} = |\langle A|E\rangle|^2 = \frac{1}{2} \left[E_R^2 + E_L^2 + 2E_R E_L \cos \left(\phi + \frac{\pi}{2} \right) \right], \quad (2.41)$$

$$I_{90} = |\langle V|E\rangle|^2 = \frac{1}{2} \left[E_R^2 + E_L^2 + 2E_R E_L \cos (\phi + \pi) \right], \quad (2.42)$$

$$I_{45} = |\langle D|E\rangle|^2 = \frac{1}{2} \left[E_R^2 + E_L^2 + 2E_R E_L \cos \left(\phi + \frac{3\pi}{2} \right) \right]. \quad (2.43)$$

By comparing Equations 2.40–2.43 with the expression for SLIM in Equation 2.37, one can conclude that they are identical up to a factor of two.

In conclusion, when polarimetry is viewed as a form of QPI between orthogonal circularly polarised components, PolCam can be considered as an alternative means of implementing the four-step phase-shifting scheme. The two approaches are mathematically equivalent in both their forward models and the reconstruction algorithms; the distinction is that PolCam introduces the extra phase shift through polarisation modulation at the micro-polariser array (MPA), whereas SLIM uses spatial separation in conjunction with active optics.

2.4.3 False-colour visualisation scheme

In addition, illumination in circular polarisation enables the use of some false-colour polarisation maps to be interpreted intuitively. A polarisation-only (Pol-only) scheme, for example, can be constructed in the hue-saturation-value (HSV) colour space, where the hue represents the AoLP and the value represents the DoLP, as shown in Figure 2.5(a). Upon circularly polarised illumination, the background of any false-colour image in the Pol-only scheme would appear dark, which helps to create a visual contrast.

In cases when the intensity image is relevant, one can also employ a false-colour scheme (Pol+I)

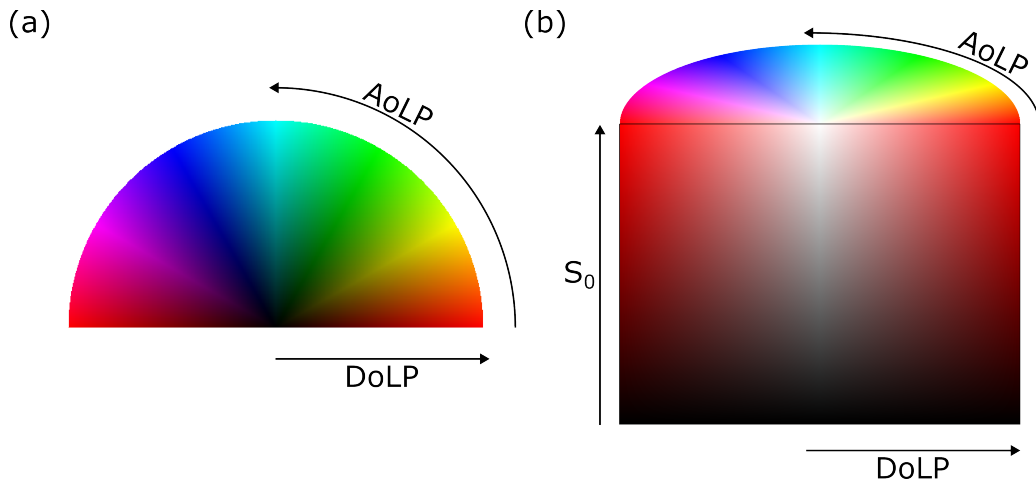


Figure 2.5: False-colour polarisation maps for PolCam measurements, constructed in the HSV colour space. (a) A polarisation-only (Pol-only) scheme where hue corresponds to AoLP, saturation is set to unity, and value corresponds to DoLP. (b) A combined false-colour scheme for polarisation and intensity (Pol+I), where hue represents AoLP, saturation corresponds to DoLP, and value corresponds to intensity.

Table 2.3: Physical meanings of false-colour polarisation maps in the HSV colour space.

HSV channel	Pol-only	Pol+I
Hue	AoLP	AoLP
Saturation	1	DoLP
Value	DoLP	Intensity

that encodes intensity information (S_0) in the value channel while using the saturation channel for DoLP. This colour scheme would create a greyscale background in transmission microscopes, and make any polarisation-neutral structure appear colourless. This particular colour scheme is also used in fluorescence orientation imaging [50], where the intensity channel carries important fluorescence information.

The two false-colour schemes will be used frequently in this thesis, and their HSV channel assignments are summarised in Table 2.3 for quick reference.

2.5 PolCam images in the Fourier domain

In previous Sections 2.2–2.4, discussion focused on the polarisation characteristics of PolCams. Nonetheless, as previously mentioned, one of the most significant drawbacks of PolCam polarimetry is its reduced spatial sampling, which is an intrinsic disadvantage of DoFP methods.

For imaging fine structures, insufficient sampling due to large pitches of (super)pixels can lead to aliasing and thus loss of information, which becomes particularly relevant in the context of ODS. Therefore, the sampling problem has to be properly addressed when designing any microscopy system.

2.5.1 Whittaker-Shannon theorem for 2D sampling

To start with, consider a two-dimensional signal g to be measured by a standard (i.e. non-polarimetry) sensor. Due to the discrete sampling at the pixels, the measured signal becomes

$$g_s(m, n) = \text{III}(m)\text{III}(n)g(m, n), \quad (2.44)$$

where m and n represent the horizontal and vertical indices of pixels, and III denotes the Dirac comb, defined as a ‘train’ of Dirac delta functions at integer positions:

$$\text{III}(m) := \sum_{i \in \mathbb{Z}} \delta(m - i). \quad (2.45)$$

In microscopy systems, the Fourier spectrum G of the signal is band-limited, confined in a finite region which is usually dictated by the OTF or the coherent transfer function (CTF) of the system, as illustrated in the first row of Figure 2.6. As a result, the Fourier transform of the measurements is a powerful tool to determine if sampling is sufficient. It is possible to show that [80], when a signal $g(x, y)$ is sampled at integer positions representing the pixels, the Fourier spectrum of the sampled signal can be expressed as

$$G_s(\xi, \eta) = \sum_{i, j \in \mathbb{Z}} G(\xi - i, \eta - j), \quad (2.46)$$

where ξ and η are the spatial frequencies horizontally and vertically, and i and j represent any integer.

To reconstruct the true signal from the sampled one, it is necessary to be able to infer $G(\xi, \eta)$

from $G_s(\xi, \eta)$. Equation 2.46 suggests that such inference is only possible if G is limited within the so-called Nyquist region of $-\frac{1}{2} < \xi, \eta < \frac{1}{2}$, because i and j take integer values. If the true spectrum G is non-zero beyond this region, ambiguity would occur in the form of aliasing. Conversely, if the OTF (and hence G) is confined in the Nyquist region, perfect reconstruction of the true signal is possible. This is commonly referred to as the Whittaker-Shannon sampling theorem [80, 81]. We also say the field is Nyquist-sampled when this happens, i.e. $W_0 \leq \frac{1}{2}$, where W_0 is the radius of the spectral band.

The Whittaker-Shannon sampling theorem determines the range of spatial frequencies over which alias-free reconstruction is possible, but practical sampled imaging systems can also exhibit additional frequency-response effects associated with the sensor [82]. In particular, Equation 2.44 represents the sensor using Dirac combs and therefore implicitly assumes zero-width detector pixels. Real sensors instead have pixels which are responsive to light over finite areas, which can be described by a pixel response function [83] and incorporated as an additional modulation of the effective OTF. Consequently, although the Nyquist sampling condition determines the alias-free range, spatial frequencies within this range are not necessarily transferred equally, but can be affected by both the OTF of the imaging system and the characteristics of the sensor employed.

2.5.2 Spectral bands for PolCam images

For PolCam images, however, the discussion above would not be directly applicable, as G is not band-limited in the same way due to the MPA. Nevertheless, a similar Fourier analysis is indeed possible [84].

Suppose an intensity signal $I(m, n)$ is sampled as

$$I_s(m, n) = \text{III}(m)\text{III}(n)I(m, n). \quad (2.47)$$

Physically, the signal is only well-defined at integer values of m and n , because of the pixelisation

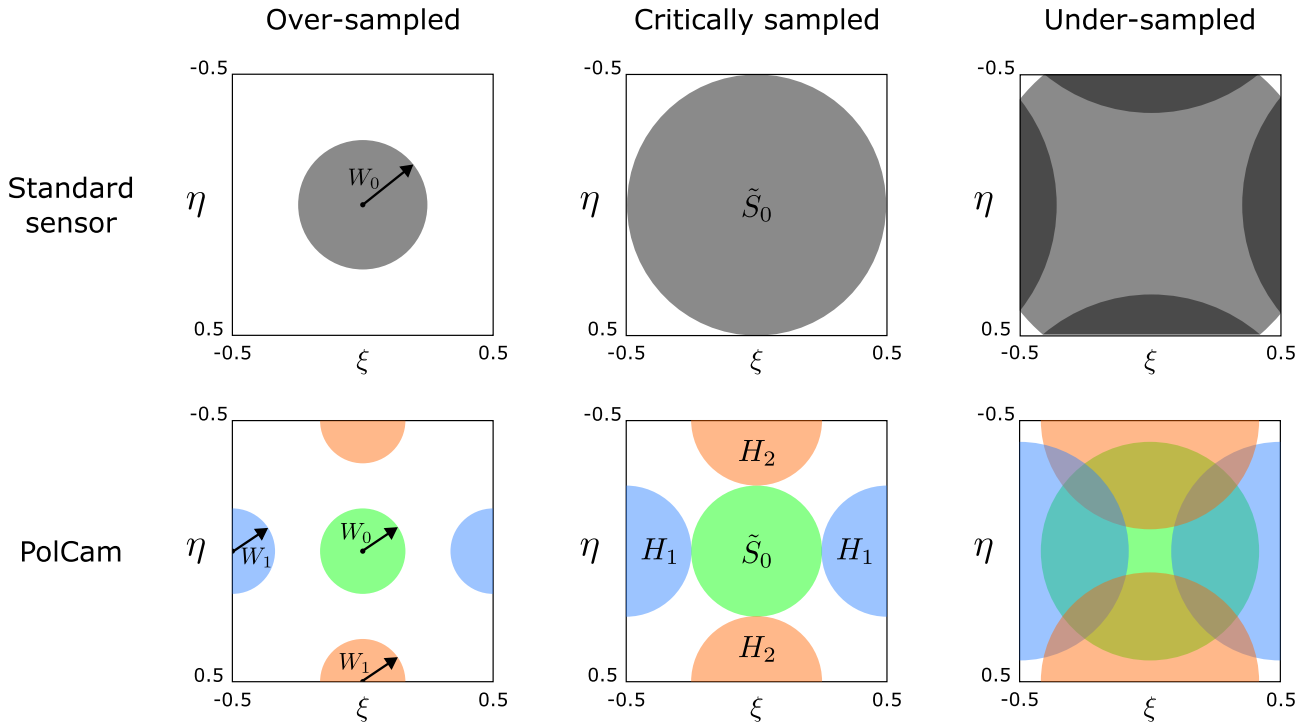


Figure 2.6: Fourier spectra $\tilde{I}(\xi, \eta)$ of images in the Nyquist region. The first row shows the cases for standard (non-polarimetry) sensors, where the Fourier spectrum is only non-zero in a finite region of radius W_0 , which is dictated by the relevant transfer function of the system. The signal can only be reconstructed unambiguously if the \tilde{S}_0 band fits within the Nyquist region $-\frac{1}{2} < \xi, \eta < \frac{1}{2}$. The second row shows that, for PolCam images, there are additional H bands corresponding to the Fourier spectra of the Stokes images. Assuming they are also band-limited to a radius of W_1 , unambiguous reconstruction is only possible if they do not overlap with each other, i.e. $W_0 + W_1 \leq \frac{1}{2}$.

of the MPA. The Müller matrix of the MPA therefore becomes

$$\mathbb{M}_{\text{MPA}}(m, n) = \begin{cases} \mathbb{M}_P(0), & \text{if } m \text{ and } n \text{ are both even,} \\ \mathbb{M}_P\left(\frac{3\pi}{4}\right), & \text{if } m \text{ is odd and } n \text{ is even,} \\ \mathbb{M}_P\left(\frac{\pi}{4}\right), & \text{if } m \text{ and } n \text{ are both odd,} \\ \mathbb{M}_P\left(\frac{\pi}{2}\right), & \text{if } m \text{ is even and } n \text{ is odd,} \end{cases} \quad (2.48)$$

where $\mathbb{M}_P(\gamma)$ denotes the Müller matrix of a linear polariser whose transmission axis sits at an orientation of γ .

To be able to perform a Fourier transform, however, the polarisation modulation has to be modelled as a continuous function with respect to m and n , not just at integer values. To do

that, one can make use of the useful property of

$$\cos(k\pi) = \begin{cases} +1, & k \text{ even,} \\ -1, & k \text{ odd.} \end{cases} \quad (2.49)$$

The choice of cosine also simplifies the algebra in the upcoming Fourier analysis. The Müller matrix can then be constructed as

$$\begin{aligned} \mathbb{M}_{\text{MPA}}(m, n) = \frac{1}{4} & \left[(1 + \cos m\pi)(1 + \cos n\pi) \mathbb{M}_P(0) \right. \\ & + (1 - \cos m\pi)(1 + \cos n\pi) \mathbb{M}_P\left(\frac{3\pi}{4}\right) \\ & + (1 + \cos m\pi)(1 - \cos n\pi) \mathbb{M}_P\left(\frac{\pi}{4}\right) \\ & \left. + (1 - \cos m\pi)(1 - \cos n\pi) \mathbb{M}_P\left(\frac{\pi}{2}\right) \right], \end{aligned} \quad (2.50)$$

which reduces to Equation 2.48 at integer values of m and n , again assuming the pixels to have zero widths.

The Müller matrix of a linear polariser, $\mathbb{M}_P(\gamma)$, can be written as

$$\mathbb{M}_P(\gamma) = \frac{1}{2} \begin{bmatrix} 1 & \cos 2\gamma & \sin 2\gamma & 0 \\ \cos 2\gamma & \cos^2 2\gamma & \cos 2\gamma \sin 2\gamma & 0 \\ \sin 2\gamma & \cos 2\gamma \sin 2\gamma & \sin^2 2\gamma & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (2.51)$$

By substituting the expression, the full Müller matrix for the MPA becomes

$$\mathbb{M}_{\text{MPA}}(m, n) = \frac{1}{4} \begin{bmatrix} 2 & \cos m\pi + \cos n\pi & \cos m\pi - \cos n\pi & 0 \\ \cos m\pi + \cos n\pi & 1 + \cos[(m+n)\pi] & 0 & 0 \\ \cos m\pi - \cos n\pi & 0 & 1 - \cos[(m+n)\pi] & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (2.52)$$

When the PolCam is measuring a field described by a general Stokes vector $\mathbf{S}(m, n)$, the

field after the MPA is then $\mathbb{M}_{\text{MPA}}\mathbf{S}(m, n)$, and the photodiode captures the intensity of the propagated Stokes vector:

$$I(m, n) = [\mathbb{M}_{\text{MPA}}(m, n)\mathbf{S}(m, n)]_0 \quad (2.53)$$

$$= \frac{1}{4} [2S_0 + \cos m\pi (S_1 + S_2) + \cos n\pi (S_1 - S_2)] \quad (2.54)$$

Given that the 2D Fourier transform \mathcal{F} is defined as

$$\mathcal{F}\{f\}(u, v) = \iint_{-\infty}^{\infty} f(x, y) e^{-i2\pi(ux+vy)} dx dy, \quad (2.55)$$

the convolution theorem [85] takes the form of

$$\mathcal{F}\{f \cdot g\}(u, v) = [\mathcal{F}\{f\} * \mathcal{F}\{g\}](u, v), \quad (2.56)$$

where $*$ denotes convolution. Therefore, the Fourier spectrum of the image can be written as the following:

$$\begin{aligned} \tilde{I}(\xi, \eta) &= \frac{1}{2} \tilde{S}_0(\xi, \eta) \\ &+ \frac{1}{8} \left[\delta\left(\xi + \frac{1}{2}\right) + \delta\left(\xi - \frac{1}{2}\right) \right] * \left[\tilde{S}_1(\xi, \eta) + \tilde{S}_2(\xi, \eta) \right] \\ &+ \frac{1}{8} \left[\delta\left(\eta + \frac{1}{2}\right) + \delta\left(\eta - \frac{1}{2}\right) \right] * \left[\tilde{S}_1(\xi, \eta) - \tilde{S}_2(\xi, \eta) \right] \end{aligned} \quad (2.57)$$

Here, $\tilde{S}_i(\xi, \eta)$ is used to denote the Fourier spectrum of $S_i(\xi, \eta)$. As expected, there is no dependence on S_3 . Due to discrete sampling, this Fourier spectrum is again ambiguous upon any integer offset from ξ and η . As a result, $\delta(\xi + \frac{1}{2})$ is equivalent to $\delta(\xi - \frac{1}{2})$, and the spectrum becomes

$$\tilde{I}(\xi, \eta) = \frac{1}{2} \tilde{S}_0(\xi, \eta) + \frac{1}{4} \left[H_1\left(\xi - \frac{1}{2}, \eta\right) \right] + \frac{1}{4} \left[H_2\left(\xi, \eta - \frac{1}{2}\right) \right], \quad (2.58)$$

where

$$H_1(\xi, \eta) := \tilde{S}_1(\xi, \eta) + \tilde{S}_2(\xi, \eta), \quad (2.59)$$

$$H_2(\xi, \eta) := \tilde{S}_1(\xi, \eta) - \tilde{S}_2(\xi, \eta). \quad (2.60)$$

As we will see in Section 3.3, in a microscopy system like the PPC, the Fourier spectra of the Stokes parameters are indeed band-limited, and therefore so are H_1 and H_2 .

Illustrations of such Fourier spectra of PolCam images are shown in the second row of Figure 2.6. As a consequence of Equation 2.58, two sidebands of H_1 and H_2 arise beside the baseband \tilde{S}_0 . The sidebands are displaced to the edges of the Nyquist region, which can also be interpreted as the result of the Fourier shift theorem.

Upon sufficient sampling, these bands are well separated from each other, allowing unmixing of the Stokes spectra by $\tilde{S}_1 = \frac{1}{2}(H_1 + H_2)$ and $\tilde{S}_2 = \frac{1}{2}(H_1 - H_2)$. Then the original signal can be reconstructed due to the Whittaker-Shannon theorem. In contrast, when the bands are large enough to overlap with each other, crosstalk between polarisation components occurs. As a result, unmixing of the original Stokes spectra is not possible, and the field is therefore under-sampled.

Compared with the spectrum of a standard sensor, that of a PolCam needs to fit more bands in the same Nyquist region, suggesting that sufficient sampling is more difficult. Suppose \tilde{S}_0 is band-limited to $\xi^2 + \eta^2 < W_0^2$ and those of \tilde{S}_1, \tilde{S}_2 (and hence H_1 and H_2) are band-limited to W_1 . Critical sampling is achieved when the bands just touch, i.e.,

$$W_0 + W_1 = \frac{1}{2}, \quad (2.61)$$

in contrast to the standard sensor case, where Nyquist sampling requires only $W_0 \leq \frac{1}{2}$. In general, however, W_1 is not necessarily equal to W_0 . As a result, despite the halved spatial sampling pitch per polarisation channel in a PolCam, the Nyquist requirement is not simply doubled, as one might naively expect from the superpixel layout.

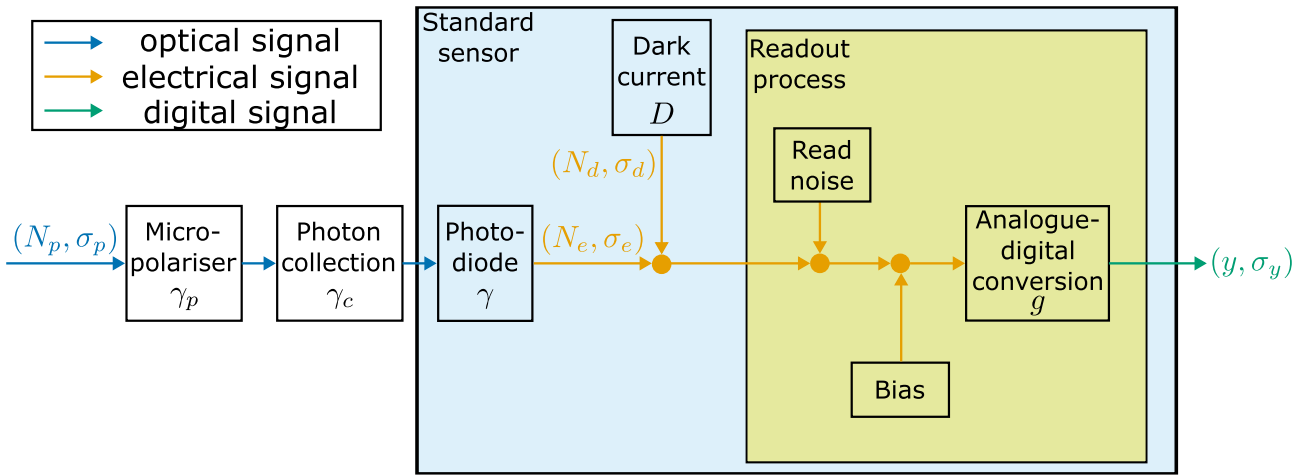


Figure 2.7: A model of signal and noise for PolCam. Photons arriving at a pixel are first collected through the micro-polariser and micro-lens before absorbed photons generate photo-electrons at the photodiode. Together with thermal electrons from the dark current, these charges pass through the analogue readout process, where read noise and a bias offset are introduced before analogue-to-digital conversion into digital numbers (DNs). Quantities are given in pairs of means and standard deviations, separated by commas, where applicable.

2.6 Noise model and characterisation

In addition to sampling capacity, it is also important to understand the noise capability of PolCam, which is especially critical in the context of data storage. In this section, we shall discuss a model of signal and noise, its experimental characterisation, and its limitations.

2.6.1 The linear model of PolCam

Figure 2.7 presents a simplified picture of understanding the process of signal and noise generation. Compared to standard sensors, PolCams feature an on-chip lens and micro-polariser for each pixel. As a result, for a number of N_p photons arriving at the pixel area, the actual number reaching the photodiode should be adjusted. Therefore, we define the transmission ratio, γ_p , as the fraction of photons travelling through the micro-polariser, which largely depends on the input and analysis polarisation states, and the extinction ratio of the polariser.

A collection efficiency, γ_c , is also assigned to the process, to represent the fraction of photons transmitted through the micro-polariser that are optically delivered to the active area of the photodiode. This factor is determined mainly by the geometry of pixels and their optical

coupling, including the fill factor, micro-lens focusing and alignment, angular acceptance, and losses from photons reaching optically inactive regions [86], and is therefore related to the pixel response function of the sensor [83]. It is therefore distinct from the quantum efficiency γ , which describes the subsequent probability of generating collected photo-electrons at the photodiode.

At the photodiode, absorbed photons generate photo-electrons at a ratio parametrised by the quantum efficiency, γ , such that the number of photo-electrons becomes

$$N_e = \gamma\gamma_p\gamma_c N_p. \quad (2.62)$$

Photons, however, are not the only source of electrons, because dark current leads to the generation of thermal electrons (whose mean number is denoted N_d). Both photo- and thermal electrons then go through the readout process to be digitised into digital numbers (DNs), to be sent to a computer. The process is dictated by a set of complex electronics, commonly including charge-to-voltage conversion, analogue amplification, and analogue-to-digital conversion. For the purpose of this simplified mathematical model, these stages are treated as a black box. Read noise and the bias offset are introduced in the analogue readout chain before analogue-to-digital conversion, but their effects are represented in the model as equivalent output quantities in DN.

Read noise comprises random fluctuations introduced during the readout process before digitisation. In this model, it is represented after conversion as a random, zero-mean offset of n_r , with a standard deviation of σ_r , both expressed in DN. The read noise can be assumed to be an intrinsic characteristic of the sensor itself, independent of the optical flux or the exposure time.

A problem arises when the analogue signal after read noise is very negative, which can occur by random chance. In cases when this outweighs the contribution from the thermal and photo-electrons, the digitised output would otherwise be clipped at zero. To prevent that from happening, camera manufacturers typically introduce a bias voltage in the analogue readout chain before digitisation.

Conversion gain, g , in DNs per electron, is used to measure the overall ratio between the digital output and the number of collected electrons, excluding the bias offset. In this model, the corresponding offset after analogue-to-digital conversion is denoted by y_0 in DNs. As a result, the mean grey value of a pixel can be modelled as

$$y = g\gamma\gamma_p\gamma_c N_p + gN_d + y_0. \quad (2.63)$$

In a linear model, both the quantum efficiency γ and the conversion gain g are assumed to be constant, which usually works very well for most CMOS and CCD sensors. The model also complies with the European Machine Vision Association (EMVA) 1288 Standard [87].

In this model, noise is introduced in three components. Firstly, there is the randomness of the light itself, which is intrinsically a stochastic process. Specifically, at constant intensity, the number of photo-electrons is governed by a homogeneous Poisson process [88], giving rise to shot noise with a mean of N_e and a standard deviation $\sqrt{N_e}$. Secondly, the statistics of the thermal electrons can also be modelled as a Poisson process, parametrised by a mean of $N_d = Dt$ where D is the dark current rate and t is the exposure time. The standard deviation of the thermal electrons is therefore $\sigma_d = \sqrt{Dt}$. Thirdly, as previously discussed, the read process introduces a zero-mean random offset n_r , and we shall associate a standard deviation of σ_r with it. Although read noise and bias are physically introduced in the analogue readout chain before the analogue-to-digital conversion, their effects are represented here in equivalent output units after conversion, so that n_r , σ_r , y_0 , and y are all expressed in DNs.

In our linear model, by assuming that all sources of noise are independent of each other, the total noise in the grey value, σ_y^2 in DN^2 , can be expressed as a sum of the variances, given as below:

$$\sigma_y^2 = g^2 (\sigma_e^2 + \sigma_d^2) + \sigma_r^2, \quad (2.64)$$

$$= g^2\gamma\gamma_p\gamma_c N_p + g^2Dt + \sigma_r^2. \quad (2.65)$$

By substituting the expression of the signal in Equation 2.63, the noise can also be written as

$$\sigma_y = \sqrt{g(y - y_0) + \sigma_r^2}, \quad (2.66)$$

eliminating from the equation the explicit appearance of factors such as the quantum efficiency, the rate of dark current, and the exposure time.

Equation 2.66 is particularly useful in characterising camera performance, because it enables measurements of the conversion gain and the read noise by the so-called photon-transfer technique [89]. The technique requires a setup where the sensor is exposed upon illumination with uniform intensity across the field-of-view (FoV) of the sensor, and works when series of images are taken when the sensor is exposed to such flat fields of varying intensity. By computing the mean and standard deviation of the measured grey values, Equation 2.66 can be used to fit to a photon-transfer curve (PTC), that is the plot of σ_y against y [90].

This exact technique, however, can be difficult to realise for the PolCam, because the flat field required for proper illumination would also require a perfectly pure polarisation state across the whole FoV, which is difficult to implement. In addition, this illumination polarisation state should also have no DoLP, otherwise a false spatial nonuniformity would be measured because γ_p becomes polarisation-channel-dependent. In reality, there is always going to be a spread of polarisation states, making such setup infeasible.

By contrast, Diekmann *et al.* [91] presented a photon-free technique for camera characterisation, making use of the thermal electrons instead of the optical ones and therefore eliminating the need of a flat field. Moreover, besides the simple setup it requires, the method provides characterisation of a rich variety of sensor parameters, which will be discussed in greater detail in the next section.

2.6.2 Photon-free characterisation of PolCam

In a photon-free setup which Diekmann *et al.* presented [91], no illumination is required at all—in fact it is critical to block any light from entering the camera. This removes any photo-electrons from our model, such that

$$y = gDt + y_0, \quad (2.67)$$

$$\sigma_y^2 = g^2Dt + \sigma_r^2, \quad (2.68)$$

$$\sigma_y^2 = gy + \sigma_r^2 - gy_0. \quad (2.69)$$

The photon-free model in Equation 2.67–2.69 describes three sets of linear relationships between y , σ_y^2 and t . Therefore, by taking many exposures of dark images at various exposure times, the pixel mean and variance can be computed for every pixel. A linear regression can be carried out for each of the three equations, and their slopes and offsets would yield estimates of various aspects of the camera characteristics including the conversion gain, dark current rate, bias, and read noise, as summarised in Table 2.4.

Table 2.4: Summary of the three linear regressions used for photon-free noise characterisation.

Eqn	Vertical axis	Horizontal axis	Slope	Offset
2.67	Signal y (DN)	Exposure time t (s)	gD (DN/s)	Bias y_0 (DN)
2.68	Variance (DN ²)	Exposure time t (s)	g^2D (DN ² /s)	Read noise σ_r^2 (DN ²)
2.69	Variance (DN ²)	Signal y (DN)	Conversion gain g (DN/e ⁻)	$\sigma_r^2 - gy_0$ (DN ²)

As for implementation, the camera in test is a FLIR Blackfly S BFS-U3-51S5P-C PolCam, integrating a Sony IMX250MZR PolarSens™ sensor. A dark field is easily realised by mounting a thick metal cap to the camera via its C mount.

As for software, Diekmann *et al.* [91] developed the automated camera characterization via electron noise tool (ACCeNT), for the acquisition and analysis of these dark images, based on μ Manager [92] and Fiji [93] respectively. At the acquisition phase, however, ACCeNT stores all raw images, whose average and variance are only calculated later in the analysis phase in Fiji. This becomes very inefficient, especially in the context of noise characterisation when

many exposures are needed for accurate estimate of y and σ_y^2 . Instead, it is identified that, for each exposure time, only a mean image and a variance image is required for the linear regressions, and storing all raw images is unnecessary. Therefore, custom software was written such that upon acquisition, only two images are kept for each level of exposure, the running sum $M_1 = \sum_i I_i$ and the running sum of squares $M_2 = \sum_i I_i^2$. Note that the summing index i represents images at the same exposure level, ranging from 1 up to a total of N . The mean and variance images can then be calculated according to

$$y = \frac{M_1}{N}, \quad \text{and} \quad \sigma_y^2 = \frac{M_2}{N} - y^2. \quad (2.70)$$

Here, temporal averaging is used instead of spatial averaging to avoid the effect of pixel non-uniformity, as we will discuss in Section 2.6.3. As a result, the estimation of y and σ_y is pixel-wise. For an acquisition session involving 89 exposure levels each containing a thousand dark images, the running-sum strategy reduces the need for disk space from 831.1 GB to 1.7 GB, making it much more feasible to use a large N for more accurate estimates.

Care however needs to be taken for acquisition sessions with many exposures, because the photon-free technique relies on the dark current which is very sensitive to any temperature variation. A pre-run of more than 15 minutes is always needed to get the camera into its steady-state temperature, before measurements become representative. Figure 2.8(a) demonstrates how sensitive dark current can be to temperature, as a dual-axis plot of spatial mean (blue) and temperature (orange) against frame index i . It is evident that even changes as small as 2°C have a significant impact on the measured grey values. It is therefore important to interleave frames of different exposure times. In contrast to ACCeNT where all images of the same exposure time are acquired consecutively, our technique acquires them in repeated cycles, each containing one image at each exposure time. This scheme of interleaved acquisition minimises any difference in average temperature for frames with different exposure times.

Custom software was written in Python, making use of the `simple_pyspin` package [94], a wrapper around FLIR's `PySpin`, which in turn interfaces with the Spinnaker SDK which provides APIs for camera control. Computation of the mean and variance maps, and the pixel-wise

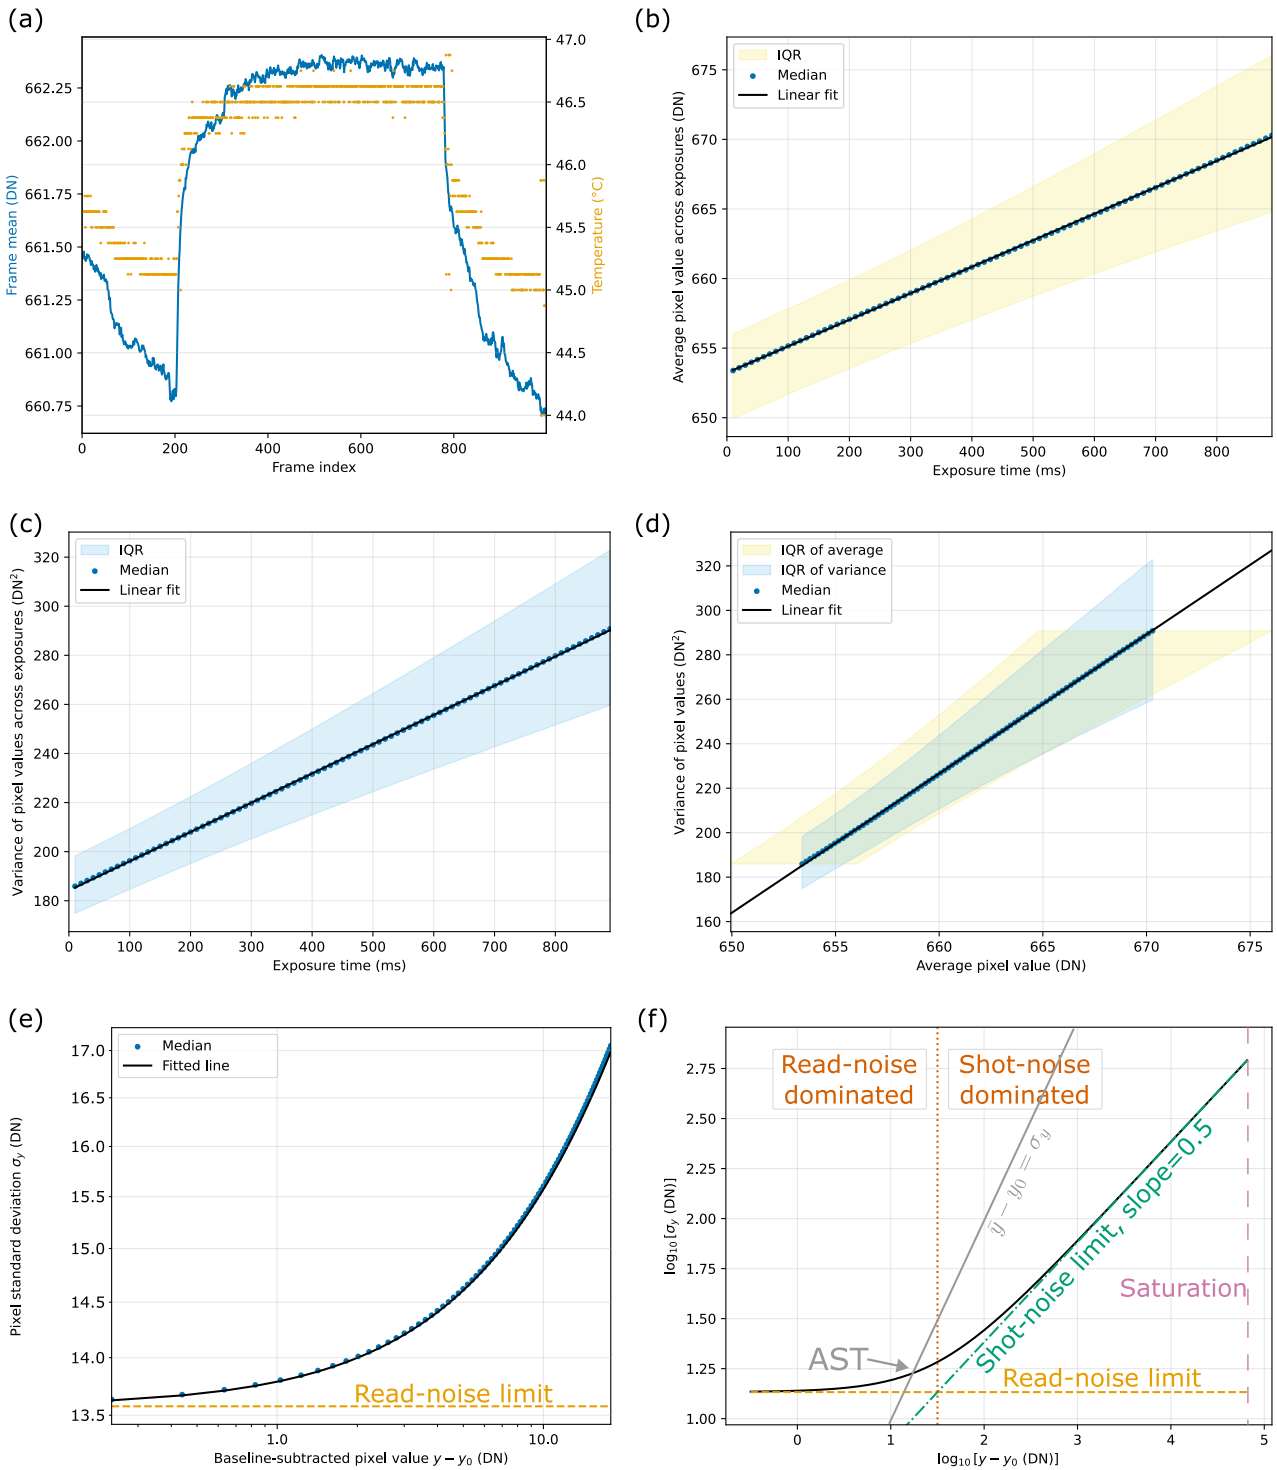


Figure 2.8: Results of photon-free characterisation of a FLIR Blackfly S PolCam. (a) A plot of frame means and temperature in an acquisition session (pre-run not included). (b-d) Linear regressions between temporally averaged pixel grey value y , its temporal variance σ_y^2 , and the exposure time t . (e) The photon-transfer curve (PTC) as a plot of pixel standard deviation against its baseline-adjusted value, converging to a read-noise limit of σ_r as $y - y_0 \rightarrow 0$. (f) The fitted PTC in a log-log plot, extrapolated to the digital saturation of $y = y_{\max}$. Besides the read-noise limit, an asymptote of the shot-noise-dominated regime is drawn in green. The intersection between the two asymptotes marks the read-/shot noise crossover, y_c . The intersection between the PTC and a line of $y - y_0 = \sigma_y$ marks the absolute sensitivity threshold (AST) where the SNR is unity.

linear regressions, were also implemented in Python.

Figure 2.8(b-d) shows the regression results according to Equations 2.67–2.69. Each data point represents the result of a temporal average of 1,000 frames of the same exposure time. Spatially across the FoV, the data point itself represents the median, to minimise any bias due to defect pixels, and the interquartile ranges (IQRs) are also shown in shaded colours where appropriate.

From the figures, it is evident that a linear model works well in the range from 10 ms to 890 ms used in the test, in spite of large IQRs which are implying large pixel non-uniformities. When the standard deviation σ_y is plotted against the baseline-subtracted grey value $y - y_0$, as shown in Figure 2.8(e), a fit (drawn in black) can be carried out using Equation 2.66. If extrapolated to the maximum digital value (y_{\max}) of 65,535, as illustrated in Figure 2.8(f), the fit becomes an estimated PTC for the PolCam. The PTC features two asymptotic behaviours: at small signal $y - y_0 \rightarrow 0$, read noise dominates over shot noise, such that σ_y approaches a flat baseline. In contrast, at $y \gg 0$, the pixel falls in a shot-noise-dominated regime where the read noise is negligible, compared to the sum of Poisson noise from photo- and thermal electrons. As the sum of two Poisson processes is another Poisson process, there is no need to distinguish between them, and the PTC becomes

$$\log \sigma_y \approx \frac{1}{2} \log (y - y_0) + \frac{1}{2} g, \quad (2.71)$$

which is a straight-line asymptote with a gradient of 0.5 in a log-log plot. The cross-over between read- and shot-noise-dominated regimes, y_c , can be defined where the two contributions equal, i.e. $g(y - y_0) = \sigma_r^2$, yielding

$$y_c = \frac{\sigma_r^2}{g} + y_0. \quad (2.72)$$

Using the estimated parameters, it is possible to compute some higher-level metrics. The signal-to-noise ratio (SNR) of a single pixel can be defined as

$$\text{SNR} = \frac{y - y_0}{\sigma_y}, \quad (2.73)$$

which takes its maximum when the grey value is saturated, i.e. $y = y_{\max} = 2^{16} - 1$. The number of electrons needed for such saturation defines the full-well capacity (FWC)¹

$$\text{FWC} = \frac{y_{\max} - y_0}{g}. \quad (2.74)$$

On the other hand, the sensitivity of the pixel is parametrised by the absolute sensitivity threshold (AST), defined by the minimum number of electrons to overpower the read noise. Mathematically, it is defined where SNR is 1, which in our linear model becomes

$$\text{AST} = \frac{1 + \sqrt{1 + 4\sigma_r^2/g^2}}{2}. \quad (2.75)$$

Alternatively, in DN, the AST is adjusted by the conversion gain and the offset and is given by

$$y_{\text{AST}} = g \cdot \text{AST} + y_0, \quad (2.76)$$

$$= \frac{g + \sqrt{g^2 + 4\sigma_r^2}}{2} + y_0. \quad (2.77)$$

As shown in Figure 2.8(f), the AST lies at the intersection of the PTC and $y - y_0 = \sigma_y$. Note it is in general different from the read-/shot-noise crossover, y_c .

Finally, the dynamic range (DR) can be defined as the ratio

$$\text{DR} = \frac{\text{FWC}}{\text{AST}}. \quad (2.78)$$

The measured quantities are summarised and compared with the specifications [95] in Table 2.5. It can be seen that, in entries where the quantity is both measured in our photon-free characterisation and provided by the specification, they agree very well, especially for the inverse conversion gain g^{-1} and the read noise σ_r . The measurements appear slightly worse in terms of SNR_{\max} , the FWC, the AST and the DR, which can be attributed to a slightly larger conversion

¹Some literature (e.g. [87]) distinguishes FWC from saturation capacity because the ‘well’ may not be saturated by electrons when the analogue-to-digital converter (ADC) is saturated, but the distinction is not necessary in our approach of considering the readout process as a black box.

Table 2.5: Summary of measured quantities from the photon-free characterisation of a FLIR Blackfly S PolCam, compared to specifications provided by FLIR by the EMVA 1288 Standard [87].

Quantity	Measurements	Specification[95]
Pixel size (μm)	–	3.45
Quantum efficiency γ (%)	–	25
Conversion gain g (DN/e^-)	5.822	–
Inverse conversion gain g^{-1} (e^-/DN)	0.17	0.18
Read noise σ_r (DN)	13.58	13.59
Dark current D (e^-/s)	3.47	–
Bias y_0 (DN)	653	–
Read-/shot-noise crossover y_c (DN)	685	–
Max signal-to-noise ratio ¹ SNR_{max} (dB)	40.47	40.55
Full-well capacity FWC (e^-)	11145	11359
Absolute sensitivity threshold AST (e^-)	2.89	2.92
Dynamic range ¹ DR (dB)	71.74	71.81

gain. Taking into account the large IQRs in the linear regressions, however, the differences can also be a result of pixel non-uniformities.

2.6.3 Pixel variations and correction

In previous sections, variations from pixel to pixel were mostly neglected. In reality, however, there are such non-uniformities, caused by an inter-pixel difference in any parameter in our model, such as the dark current rate D , the quantum efficiency γ , and the conversion gain g . Such non-uniformities give rise of a spatial variance of pixel values, similar to how various sources of noise discussed above introduce a temporal variance. Therefore, the effect of such non-uniformities is usually called the fixed pattern noise (FPN) [96, 97]. Statistically, the FPN introduces variations in both the means and variances of the pixel values. By estimating the expected offset it introduces to each pixel, the means can be effectively corrected for, although the variance remains due to the intrinsic stochastic nature of noise-generating processes.

In our photon-free characterisation, one particular type of non-uniformity is encountered, namely the dark signal non-uniformity (DSNU). As its name suggests, it is the non-uniformity of signals when no optical flux is present. The DSNU can be clearly seen when plotting a

¹The decibel (dB) is defined as $20 \log_{10}(a/b)$ for SNR and DR.

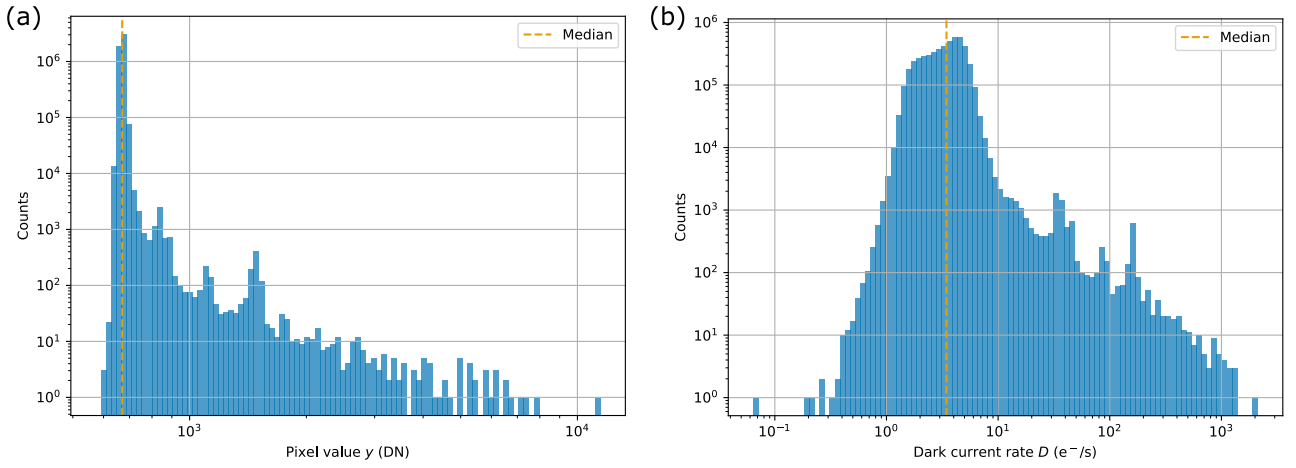


Figure 2.9: Histograms of dark images from a PolCam due to the DSNU. (a) A histogram of the mean pixel values from 1,000 dark images at an exposure time of 890 ms. (b) A histogram of dark current rate D of each pixel in the PolCam, estimated from a photon-free characterisation session of 89 different exposure levels each containing 1,000 images, the same session as discussed in Section 2.6.2. In both cases, the distributions are very positively skewed.

histogram of pixel values y of a dark image, as shown in Figure 2.9(a). As the histogram was produced by averaging 1,000 dark images of the same exposure time, its spread comes mostly from spatial non-uniformities. Another way of seeing the DSNU is to plot a histogram of the estimated dark current rate D , as seen in Figure 2.9(b), enabled by pixel-wise linear regressions in the photon-free characterisation illustrated in the previous section.

In both histograms, the distributions are strongly positively skewed, with a long tail spreading over an order of magnitude. The actual mechanism causing the skew is hardware-dependent and is out of the scope of this thesis, but it is evident that the DSNU contributes to an offset in the mean pixel values in a dark image, which can be corrected for by a simple subtraction, i.e.,

$$Y' = Y - Y_0, \quad (2.79)$$

where Y represents the image in an actual measurement, and Y' denotes the corrected image, while Y_0 is a temporally averaged dark image of the same exposure time. This subtraction removes the mean dark signal of each pixel, although it does not remove its variance, since the number of thermally generated electrons still fluctuates according to the underlying Poisson statistics. This correction becomes particularly important in the read-noise-limited regime, when the signal level is low and any uncorrected offset impacts the apparent SNR significantly.

Another effect of the spatial non-uniformity is the photo-response non-uniformity (PRNU), which accounts for the spatial variance measured when the sensor is illuminated by a flat field. The PRNU can also be affected by the conversion gain like the DSNU, but is mainly due to variations in the quantum efficiency, γ . For a PolCam, the PRNU can also be attributed to variations in the MPA and the micro-lens array, parametrised by γ_p and γ_c , assuming a polarisation-accurate flat-field.

In contrast to the additive nature of the DSNU, the PRNU acts multiplicatively with respect to the signal, as both the quantum efficiency and the conversion gain are multiplicative. Therefore, it becomes very important to correct for the PRNU at high signal levels, because the FPN offset scales linearly with the grey value $y - y_0$, instead of its square root. Actually, in some photon transfer theories when spatial variance is also taken into account [98], there is an FPN regime that overtakes the shot-noise-limited-regime when the signal is near its saturation, featuring a linear asymptote with a gradient of one¹.

Nevertheless, thanks to its multiplicative nature, the PRNU can be effectively corrected for with the aid of a flat field image B , at the same intensity and exposure settings to the image Y . The corrected image can simply be computed by $Y' = Y/B$. By convention, to avoid massive changes to the average signal level, a normalisation factor is added such that $Y' = Y \cdot \langle B \rangle / B$, where $\langle \dots \rangle$ denotes the spatial average across an image.

As a result, by combining the corrections for the DSNU and the PRNU, an overall correction to pixel non-uniformities can be realised by

$$Y' = \frac{\langle B - Y_0 \rangle}{B - Y_0} \cdot (Y - Y_0) = G (Y - Y_0), \quad (2.80)$$

where the flat-field gain, G , is defined as the dimensionless quantity:

$$G := \frac{\langle B - Y_0 \rangle}{B - Y_0}. \quad (2.81)$$

The correction defined by Equation 2.80 is called a flat-field correction (FFC) [99]. Nonetheless,

¹For an example of the so-called ‘classical’ PTC, the reader can refer to Figure 5.4 of Janesick [98].

as previously discussed, it is practically challenging to implement a polarisation-accurate flat field that does not contaminate the polariser transmission ratio γ_p . Moreover, the photon-free method is fundamentally incompatible with the PRNU. Therefore, instead of a flat field, B for PolCam is defined as a bright-field reference that is not necessarily flat or linearly unpolarised. Instead, it is the image from the same optical setup in the same illumination scheme, polarisation setup, and camera settings, only with the object removed. The correction can then no longer be called a FFC due to the dropped flat-field assumption. Alternatively, we shall refer to it as the polarisation-field correction (PFC), and the gain image G becomes the polarisation-field gain.

In contrast to the FFC, the PFC is no longer specific to correcting for the FPN of the sensor. Instead, it depends greatly on the illumination setup as well. A potential intensity gradient across the FoV, for example, would be compensated by a spatially varying gain image G . A polarisation offset of the illumination can also be corrected for, because the measured polarisation-field gain would have a different mean for each polarisation channel.

An example of a PFC calibration is shown in Figure 2.10, with a PPC system, averaged from 100 bright-field and 100 dark-field images, all with the same exposure time of 50 ms. Although the PPC is designed for the left-handed circular polarisation, in reality there is a polarisation offset, which is clearly reflected by the multiple peaks in the histogram of the polarisation-field gain G in Figure 2.10(a). With a standard deviation of 6.43%, G compensates for any variations in the polariser transmission ratio γ_p , the photon collection efficiency γ_c , the quantum efficiency γ , and the conversion gain g from one pixel to another.

The histogram of the dark image Y_0 is shown in Figure 2.10(b). Compared with Figure 2.9(a) where the exposure time is 890 ms, the pixel values in this plot are much closer to the global bias y_0 which was estimated as 653 DN in Table 2.5, because of the much reduced number of thermal electrons. There are, however, still a few significant outliers signalling the existence of hot pixels. During an imaging session, the defect pixels can be visually identified thanks to the real-time polarisation image processing implemented in my MRes project [41]. Ideally, important features (e.g. data sectors) should be placed in the FoV in a way to avoid these defect




Figure 2.10: Pixel histograms from a PFC calibration, with the PPC system in (near-)circular illumination at an exposure time of 50 ms. Both bright-field and dark-field images are averaged from 100 frames. (a) Polarisation-field gain G showing a multiple-peak structure due to a polarisation offset in illumination. (b) Pixel values of the dark image Y_0 . Outliers correspond to hot pixels due to sensor defect.

pixels as much as possible.

It should also be noted that, in theory, the bright-field and dark-field images can be simulated using Equation 2.63 in the linear signal model, eliminating the need to measure them whenever there is a change in the illumination setup, or the exposure time. For ideal performance, however, direct measurements are still preferred whenever possible, to account for potential nonlinearity and temperature variations.

2.7 Towards an ideal system for weak birefringence

With the noise performance of PolCams characterised, it is then important to understand how it impacts polarimetric measurements, and what an ideal system would be to image weakly birefringent features like voxels.

As established in Section 2.6, with the FPN offset corrected for by PFC, a pixel is read-noise limited at a low intensity, and shot-noise limited at a high signal level $y > y_c$. In the context of voxel imaging, there is no need to consider photobleaching or phototoxicity, so it is always preferable to have a higher level of exposure, up to the pixel saturation. For standard sensors, this would imply employing a brighter source or setting a longer integration time such

that $y = y_{\max}$, and the pixels would always be shot-noise limited. For PolCams, however, the saturation of superpixels only implies the best-aligned polarisation channel is saturated, and the orthogonal channel can have a much lower signal, depending on the DoLP. To see this, consider horizontally linearly polarised light being analysed by a PolCam. Adjustment of exposure would ensure I_0 to be saturated, while $I_{90} = 0$ is clearly dominated by the read noise.

To assess the potential impact of this problem, we shall consider a worst-case scenario, when the AoLP aligns perfectly with a polarisation channel, and therefore the orthogonal channel has a minimum polariser transmission at the same time. Calculation from Stokes-Müller algebra yields the corresponding intensities I_{\parallel} and I_{\perp} can be expressed as the following:

$$I_{\parallel} = \frac{1 + \text{DoLP}}{2} S_0, \quad (2.82)$$

$$I_{\perp} = \frac{1 - \text{DoLP}}{2} S_0. \quad (2.83)$$

An apparent extinction ratio ER can be defined as their ratio, I_{\parallel}/I_{\perp} . Assuming the co-polarised channel is filled to digital saturation, the perpendicular channel becomes read-noise limited when

$$\text{ER} = \frac{1 + \text{DoLP}}{1 - \text{DoLP}} > \frac{y_{\max} - y_0}{y_c - y_0}, \quad (2.84)$$

which is solved to be

$$\text{DoLP} > \frac{y_{\max} - y_c}{y_{\max} + y_c - 2y_0}. \quad (2.85)$$

Using the results estimated from the photon-free characterisation, the inequality is evaluated to be $\text{DoLP} > 0.9990$, or $1 - \text{DoLP} < 9.755 \times 10^{-4}$. It becomes clear that the read-noise dominated regime in the NLSP is so limited that it becomes negligible. In realistic situations, all pixels on the PolCam can be considered to be shot-noise limited, which gives us a reduced noise model of

$$\sigma(I_i) = \sqrt{gI_i}. \quad (2.86)$$

Similar considerations were analysed by Foreman and Török [100], when the Cramér–Rao lower bound (CRLB) for AoLP estimation was derived, based on a general stochastic framework for

polarimetry [20]. It was discussed that, in the presence of a non-zero intensity background, shot noise from the background can dominate over the intensity variation due to the AoLP. It is especially significant for polarisation channels oriented nearly perpendicularly, where Malus' law [69] yields a small $\partial I/\partial \text{AoLP}$, leading to 'informational dips' in the Fisher information [100]. Although their polarimeters and assumptions differ (e.g. S_0 to be known *a priori*), the mechanism mirrors our discussion: in our case, the orthogonal channel intensity I_\perp acts as a background term, whose shot noise can exceed read noise and the AoLP-driven intensity change.

We now consider how noise is propagated to estimation of the Stokes parameters. A first-order treatment propagates the uncertainty as per

$$\sigma(\hat{S}_i) = \sqrt{\sum_j \left(\frac{\partial S_i}{\partial I_j}\right)^2 \sigma^2(I_j)}, \quad (2.87)$$

as the intensities at different polarisation channels are independent. By substituting Equations 2.1–2.3, the uncertainties of the normalised Stokes parameters can be estimated as

$$\sigma(\hat{S}_1) = \sqrt{\frac{g}{S_0} \left(1 - \frac{\hat{S}_1^2}{2}\right)}, \quad (2.88)$$

$$\sigma(\hat{S}_2) = \sqrt{\frac{g}{S_0} \left(1 - \frac{\hat{S}_2^2}{2}\right)}. \quad (2.89)$$

For the normalised vector $\hat{\mathbf{s}}$ projected in the NLSP, its uncertainty can be measured by defining a root-mean-square (RMS) spread as

$$\sigma(\hat{\mathbf{s}}) := \sqrt{\frac{1}{2} [\sigma^2(\hat{S}_1) + \sigma^2(\hat{S}_2)]} = \sqrt{\frac{g}{S_0} \left(1 - \frac{1}{4} \text{DoLP}^2\right)}. \quad (2.90)$$

From Equation 2.90, it is evident that for the same level of intensity S_0 , the more linearly polarised the light, the less overall shot noise is propagated to the estimation of the measured Stokes vector in the NLSP. At large DoLP, the noise performance is improved, which is expected because there is more discrimination between polarisation channels. The situation also draws

direct analogy to the nulling technique in polarimetry and ellipsometry [101, 102], when a sharp extinction improves the accuracy by making the optical noise less relevant.

The key assumption of treating S_0 as a constant that is independent of the DoLP is in many cases valid. For example, in PPC, as discussed in Feng [41], phase features from the sample is turned into intensity signals at the camera, which usually decides the exposure settings. This is because the phase features usually introduce larger intensity variations, compared to the relatively minor intensity difference across polarisation channels due to the polarisation. Nonetheless, in cases when S_0 is uniform across the FoV, S_0 may be adjusted such that the co-polarised channel is saturated, i.e. $I_{\parallel} = y_{\max} - y_0$. Using Equation 2.82 for an expression of I_{\parallel} , the optimal setting of S_0 would then depend on the DoLP:

$$S_0 = \frac{2}{1 + \text{DoLP}} (y_{\max} - y_0). \quad (2.91)$$

Recall that the FWC is defined as $(y_{\max} - y_0)/g$, so the RMS spread becomes

$$\sigma(\hat{\mathbf{s}}) = \sqrt{\frac{1}{2\text{FWC}} \left(1 + \text{DoLP} - \frac{1}{4}\text{DoLP}^2 - \frac{1}{4}\text{DoLP}^3 \right)}. \quad (2.92)$$

In this case, besides the DoLP, the noise performance only depends on the FWC limited by the sensor. The overall uncertainty is no longer monotonically decreasing in the range $\text{DoLP} \in [0, 1]$, due to the reduced S_0 needed to avoid over-saturating the co-polarised channel.

It is still usually beneficial, however, to measure states of large DoLP, because the ‘signal’ also improves in the context of a signal-to-noise analysis. The specific definition of the SNR would depend on the specific task, but a meaningful signal is usually the distance from the measured polarisation state to the origin on the NLSP, or the distance between polarisation states of different AoLP for different symbols of voxels, both of which scale linearly with the DoLP. As a result, the SNR scales with

$$\frac{\text{DoLP}}{\sqrt{1 + \text{DoLP} - \frac{1}{4}\text{DoLP}^2 - \frac{1}{4}\text{DoLP}^3}}, \quad (2.93)$$

which is strictly increasing in $[0, 1]$.

In Section 2.4, it was learnt that the Jones vector propagated through a weakly birefringent voxel illuminated in left-handed circular polarisation can be written as

$$|E_{out}\rangle \approx \begin{bmatrix} 1 \\ -\frac{i\delta}{2}e^{2i\gamma} \end{bmatrix}_c, \quad (2.94)$$

expressed in the circular basis. For an optimal SNR performance, the detected polarisation state at the PolCam should be as close to linearly polarised as possible, requiring $\tilde{E}_L^{\text{out}} \approx \tilde{E}_R^{\text{out}}$. Therefore, an ideal microscope system should be able to modulate the amplitude of the left-handed S wave by a further factor of $\frac{\delta}{2}$. Such a polarisation modulation would amplify the DoLP signal on the NLSP to be better detected by the PolCam, leading to the design of the PPC and the diattenuative module in upcoming chapters.

2.8 Conclusions

This chapter has focused on PolCams: how they work, what they provide, and what an ideal microscopy system would need to fully utilise them. Section 2.1 started by presenting them as a way to realise DoFP polarimetry, among many other means of multiplexing polarisation information. The section explained the hardware structure of the PolarSens™ sensor and the arrangement of polarisation channels. In the following analysis, it was shown that three out of the four Stokes parameters can be computed from a PolCam image, restricting polarimetric measurements to a projection of the Stokes vector in the NLSP plane, parametrised by the DoLP and the AoLP. Information about the circular component is not obtainable—it cannot even be inferred due to coherence effects. This was shown with an example of mixing of polarisation states, which will become more relevant in the context of voxel imaging.

To design a microscope system optimised for imaging weak, linear birefringence, illumination with a circularly polarised state was shown to be preferable. With the aid of Jones calculus in the circular basis, it was shown that an orthogonally circularly polarised component emerges

from linear birefringence. Its amplitude and phase relative to the background correspond to not only the retardance and azimuth of sample birefringence, but also the DoLP and the AoLP that can be measured by a PolCam. In addition, the use of circularly polarised illumination enables interpretation of any DoLP as a polarisation signal, with intuitive false-colour visualisation schemes.

In order to address the fundamental drawback of reduced spatial sampling in DoFP techniques, it is crucial to understand how the superpixel structure of PolCams impacts the ability of resolving high-frequency details of a field. The Fourier spectrum of a PolCam image was shown to be divided into multiple bands corresponding to different Stokes parameters. The criterion for sufficient sampling was shown to avoid any crosstalk between these bands, which needs to be considered when designing the microscope.

Furthermore, the noise performance of PolCams was modelled and characterised. As the read-noise-limited regime is very limited and any FPN offset can be effectively removed by a PFC process, the main source of noise was identified to be the shot noise, which is in turn limited by the FWC of the PolCam. From a signal-to-noise perspective, it turns out PolCams favours polarisation states close to linear, which motivates an amplitude modulation (i.e. attenuation) of the S wave in the design of such systems.

Chapter 3

Polarisation-sensitive phase contrast (PPC) microscopy

In the MRes project [41], we developed the polarisation-sensitive phase contrast (PPC) technique. Based on Zernike's phase contrast (ZPC) [68], PPC employs polarisation control and the use of polarisation cameras (PolCams), for simultaneous capture of phase and birefringence information in the sample. In particular, the spatial filtering introduced by the ZPC setup is beneficial for imaging features of weak birefringence, like the voxels. In this chapter, we shall briefly review the working principle of PPC, including its optical setup, and an analytical model of its imaging process.

3.1 Zernike's phase contrast

The key concept of ZPC lies in the field modulation introduced by a matching pair of the condenser annulus and the objective phase ring. As illustrated in Figure 3.1, the condenser aperture takes the shape of a thin annulus, which is projected onto the phase ring of the objective pupil sitting in its conjugate plane. Without scattering at the specimen, all light propagates through the phase ring, where field is modulated in both amplitude and phase. In contrast, in a thin-annulus limit, any light scattered at the specimen would propagate with




Figure 3.1: Schematic of the ZPC microscopy, which features a condenser annular aperture and matching phase ring in the objective pupil plane. At the phase ring, the S wave is modulated by a quarter wave in phase and attenuated in amplitude. PPC employs this setup with extra polarisation controls.

a shifted wave vector, such that it would not be modulated by the phase ring. In ZPC, the former component of the field is called the surround (S) wave, while the latter named as the diffracted (D) wave. The final field (namely, the particle, P, wave) perceived at the detector is a coherent superposition between the two, which is sometimes described as $P = S + D$.

The spatial filtering scheme enables label-free imaging of phase objects, especially in the weak phase approximation of

$$e^{i\phi} \approx 1 + i\phi, \quad (3.1)$$

where ϕ is the object phase. In this picture, the field propagated through the specimen can be interpreted as the sum of an unmodulated constant term corresponding to the S wave, and a phase-shifted perturbation whose amplitude is proportional to the amount of phase, corresponding to the D wave.

In the negative phase contrast setup [69], the S wave is not only attenuated for a better contrast, but also delayed by a quarter-wave compared to the D wave, so that the S and P waves are in phase to be interfered constructively. The phase ring transmission coefficient can be written as $t_p e^{i\pi/2}$, where t_p is a real-valued amplitude transmission coefficient. Additionally, it is assumed

that the entire S wave passes through the phase ring while no D wave does. This approximation is useful for algebraic simplicity, but is only true in the thin-annulus limit, where the condenser annulus, and therefore the phase ring, has infinitesimal width. The P wave thus becomes $i(t_p + \phi)$, encoding phase information in its amplitude.

One can also see how contrast is enhanced by the attenuation of the amplitude. Contrast is usually defined as

$$C = \frac{I_{\text{obj}} - I_{\text{sur}}}{I_{\text{obj}} + I_{\text{sur}}}, \quad (3.2)$$

where I_{obj} and I_{sur} denote the intensity of the object and its surroundings, respectively. In the context of ZPC, I_{obj} can be identified as proportional to $|P|^2$, while I_{sur} is proportional to $|S|^2$. The contrast C can then be evaluated as

$$C = \frac{(t_p + \phi)^2 - t_p^2}{(t_p + \phi)^2 + t_p^2}, \quad (3.3)$$

which is strictly decreasing with respect to t_p for any $\phi > 0$.

In conclusion, we summarise ZPC as a technique that splits the S and D waves by annular illumination, and selectively modulating the S wave at the objective phase ring. The phase modulation brings the S wave in phase with the D wave such that the interference is constructive, and the amplitude modulation attenuates the S wave such that the contrast is enhanced. One can refer to Feng [41] for more detailed discussions on the ZPC, while we shall now move on to how these concepts translate in the PPC setup.

3.2 Polarisation control of PPC

3.2.1 Objects of birefringence only

PPC also employs the basic setup of ZPC in Figure 3.1. Additionally, a PolCam is used instead of a conventional camera as the detector, and the source is controlled to be circularly polarised, as reasoned in Chapter 2. Assuming the object to be birefringent only with a small

retardance, recall Equation 2.32 that the polarisation state $|E_{out}\rangle$ after the sample upon left-handed circularly polarised illumination is

$$|E_{out}\rangle = \mathbb{L}\mathbb{R}(\delta_o, \gamma_o)|L\rangle \approx |L\rangle - \frac{i\delta_o}{2}e^{2i\gamma_o}|R\rangle, \quad (3.4)$$

where δ_o represents the birefringent retardance of the object, and γ_o represents its fast-axis orientation.

It can be observed that the equation is in close analogy to Equation 3.1 for ZPC, in the form of the original field being perturbed by a field proportional to a signal of interest (phase for ZPC and retardance for PPC). By comparing them, one can identify the S and D wave components as

$$|S\rangle = |L\rangle, \quad (3.5)$$

$$|D\rangle = -\frac{i\delta_o}{2}e^{2i\gamma_o}|R\rangle, \quad (3.6)$$

and

$$|P\rangle = |S\rangle + |D\rangle = \begin{bmatrix} 1 \\ -\frac{i\delta_o}{2}e^{2i\gamma_o} \end{bmatrix}_c, \quad (3.7)$$

Using the relations between the normalised linear Stokes plane (NLSP) observables and circular-basis Jones vectors in Equation 2.34–2.36, it can be shown that the degree of linear polarisation (DoLP) and the angle of linear polarisation (AoLP) of the P wave can be expressed as below:

$$\text{DoLP} = \frac{\delta_o}{1 + \delta_o^2/4} \approx \delta_o, \quad (3.8)$$

$$\text{AoLP} = \gamma_o - \pi/4. \quad (3.9)$$

Note a small retardance approximation has been assumed, i.e. $\delta_o \ll 1$.

By employing a phase-contrast objective, the PPC setup modulates the S wave by $t_p e^{i\pi/2}$, corresponding to an amplitude attenuation by t_p and a global phase shift of $\pi/2$ applied to the S wave, just as in ZPC. As in the previous analysis for ZPC, the thin-annulus approximation is retained, so that no D wave is modulated by the phase ring. As a result, the P wave after

phase ring modulation, $|P'\rangle$, becomes

$$|P'\rangle = \begin{bmatrix} it_p \\ -\frac{i\delta_o}{2} e^{2i\gamma_o} \end{bmatrix}_c, \quad (3.10)$$

and the NLSP observables become

$$\text{DoLP} = \frac{t_p \delta_o}{t_p^2 + \delta_o^2/4} \approx \delta_o/t_p, \quad (3.11)$$

$$\text{AoLP} = \gamma_o - \pi/2, \quad (3.12)$$

where the approximation holds for $\delta_o \ll t_p$. Ideally, the DoLP signal is maximised to 1 if the phase ring transmission can be matched to half of the voxel retardance, i.e. $t_p = \delta_o/2$. The P wave would then become linearly polarised, which is making the best use of the NLSP. In practice, however, the phase ring transmission is dictated by the phase-contrast objective that is available to use. For the objective used in this project (Olympus LUCPLFLN 40x/0.60 Ph2), the amplitude transmission t_p was measured as 0.39. In contrast, type-X-modified birefringent voxels are estimated with a retardance δ_o on the order of 10^{-3} rad [38], enabling the small retardance approximation of $\delta_o \ll t_p$.

Comparing the PPC-modulated observables (Equation 3.11–3.12) with the unmodulated case (Equation 3.8–3.9), it can be seen that the DoLP is amplified by approximately $1/t_p$, as illustrated in Figure 3.2. As discussed in the previous chapter, the improved DoLP is very beneficial for the signal-to-noise ratio (SNR) performance for symbol discrimination.

As it can be seen in Figure 3.2, the AoLP is also modulated, by a further $\pi/4$ from the unmodulated case. This is due to the phase modulation of the S wave in the phase ring. For voxel classification, it only adds a systematic offset and therefore plays no role in the performance of symbol discrimination.

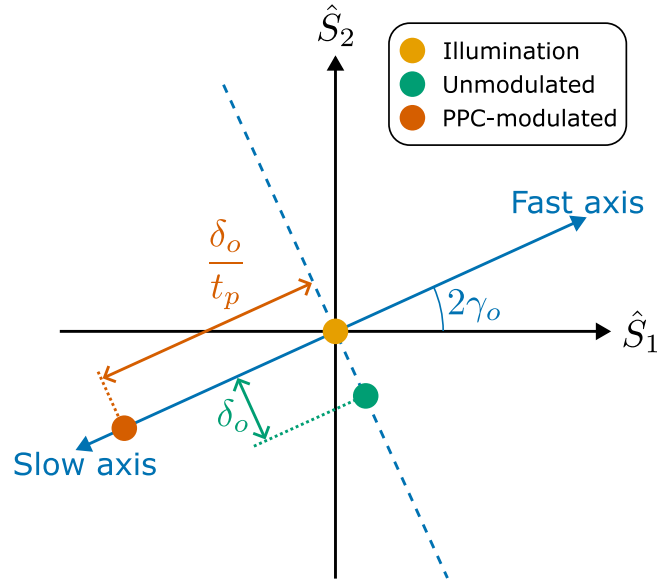


Figure 3.2: Schematic of polarisation modulation introduced by the PPC. Compared with the unmodulated case (green), PPC (shown in orange) modulates the polarisation state of the P wave such that its DoLP signal is improved by $1/t_p$ for weakly birefringent voxels, thanks to the amplitude attenuation of the S wave. The PPC also introduces an azimuthal offset by $-\pi/2$ on the NLSF, due to the phase modulation at the phase ring.

3.2.2 Birefringent objects with a residual phase

The analysis so far has been assuming objects with linear birefringence only, described by the Jones matrix $\mathbb{L}\mathbb{R}(\delta_o, \gamma_o)$. In reality, however, birefringent objects usually associate with a residual phase, due to the objects being made of some different material from the media. For example, laser-induced birefringent voxels carry a different amount of phase compared to the unmodified region of the glass [31].

For a thin object, the residual phase ϕ_o accumulated through a distance d can be modelled [103] as

$$\phi_o = \left(\frac{n_o + n_e}{2} - n_m \right) d, \quad (3.13)$$

where n_o , n_e , and n_m represent the refractive indices of the ordinary and extraordinary axes of the birefringent object, and that of the medium.

In cases when ϕ_o is non-zero, the object introduces a global phase, i.e. $e^{i\phi_o}\mathbb{L}\mathbb{R}(\delta_o, \gamma_o)$. Usually, as the global phase does not affect the polarisation state, it plays no role in the image formation process. Nevertheless, due to ZPC (and therefore PPC) being phase-sensitive, the residual phase can no longer be neglected. Instead, the field E_{out} before phase ring modulation can be

expressed as the following:

$$|E_{out}\rangle = e^{i\phi_o} \mathbb{L}\mathbb{R}(\delta_o, \gamma_o)|L\rangle, \quad (3.14)$$

$$\approx |L\rangle + i\phi_o|L\rangle - \frac{i\delta_o}{2}e^{2i\gamma_o}|R\rangle, \quad (3.15)$$

$$= |S\rangle + |D_L\rangle + |D_R\rangle. \quad (3.16)$$

Note now there are two parts of the D wave—a left-handed component $|D_L\rangle$ proportional to the residual phase, and a right-handed component $|D_R\rangle$ proportional to the birefringence. Such a separation between the phase and birefringence information shows again the advantage of carrying out the polarisation analysis in the circular basis.

Upon the phase ring, any S wave is modulated, resulting in a modulated P wave of

$$|P'\rangle = it_p|S\rangle + |D_L\rangle + |D_R\rangle, \quad (3.17)$$

$$= \begin{bmatrix} i(t_p + \phi_o) \\ -\frac{i\delta_o}{2}e^{2i\gamma_o} \end{bmatrix}_c. \quad (3.18)$$

Compared with Equation 3.10, the new expression retains the same relative phase between the left- and right-handed circular components, so the AoLP measurement is not affected by the residual phase. The DoLP, however, is affected, according to

$$\text{DoLP} = \frac{(t_p + \phi_o)\delta_o}{(t_p + \phi_o)^2 + \delta_o^2/4} \approx \frac{\delta_o}{t_p + \phi_o}. \quad (3.19)$$

Compared with the birefringence-only case described in Equation 3.11, it can be seen that the expression above differs by a substitution of t_p into $t_p + \phi_o$, effectively making the phase ring attenuation weaker. As a result, the measured DoLP is also weaker.

3.2.3 The paraxial approximation

It should be noted that the polarisation analysis above has been making an implicit paraxial approximation, by employing a 2D treatment of the Jones algebra. As the numerical aperture (NA) of the system increases, focusing of the field can introduce a significant z-component, especially for systems with an annular aperture [104]. As focusing changes the wave vector \mathbf{k} , the electric field also changes its direction to remain orthogonal to the wave vector. High-NA treatments like the Debye-Wolf diffraction integral [105] would need to be applied, in combination with the Jones calculus generalised into 3D [72].

According to Sheppard and Matthews [106], the low-NA treatment breaks down when $\text{NA} \gtrsim 1/\sqrt{2} \approx 0.71$. In this project, the PPC system is implemented with an Olympus LUCPLFLN 40x objective with an NA of 0.60. On the illumination side, the NA is even lower, such that the annular aperture matches the phase ring in the objective. For the Ph2 setting, a well-aligned condenser illuminates with a NA estimated to be 0.30, found by imaging the objective pupil through a Bertrand lens. Therefore, the low-NA treatment for the PPC system is deemed appropriate, though a potential extension into higher NA exists.

3.3 Imaging model of PPC upon spatially incoherent illumination

Besides the polarisation analysis presented above, it is important to understand how PPC behaves in the spatial and Fourier domains. It is well known that microscopy systems are band-limited in the Fourier space [77], through the coherent transfer function (CTF) or the optical transfer function (OTF) in the coherent and incoherent regimes respectively. Beyond the transfer functions, information at higher spatial frequencies is irreversibly lost, resulting in a finite optical resolution of the system, usually characterised by the impulse response of the system to a point object, i.e. the point-spread function (PSF). This sets an extra constraint to the ability of the system resolving closely-spaced birefringent voxels or any fine details of bio-

logical specimen, in addition to the detection resolution discussed in Section 2.5. Nevertheless, it requires careful attention to employ the traditional imaging theories to the PPC, not only due to the multidimensional nature of polarimetric detection, but also because of the annular aperture and phase-ring modulations.

For transmission-light illuminated setups like PPC, the sample is usually illuminated with sources with a low degree of spatial coherence, such as halogen lamps and light-emitting diodes (LEDs). Nevertheless, a purely incoherent treatment of the imaging process is deemed insufficient. Using theories of partial coherence, DeVelis and Parrent [107] have shown that an optical system with multiple lenses cannot be described by cascading the OTFs of each individual lens, even if the illumination is perfectly incoherent. This conclusion is of particular relevance for PPC when the Fourier space is modulated in both illumination and detection. From a modern perspective, this can be understood as a consequence of the van Cittert–Zernike theorem [108, 109]—the degree of coherence of a field can increase upon propagation, invalidating the incoherent treatment of further optical elements.

With partial coherence, the optical system is no longer linear with respect to either the field amplitude or the intensity. The second-order statistics of light has to be described, usually with the mutual intensity in the space-time domain [110] or the cross-spectral density function (CSDF) in the space-frequency domain [111]. While the next chapter will provide a more comprehensive picture describing a partially coherent source using the CSDF, this section shall consider the field at the condenser aperture to be a completely incoherent secondary source [112], following a treatment by Zernike [113]. As a result, the sample is illuminated by a set of mutually incoherent plane-wave components at different spatial frequencies, each contributing to a differential amount of intensity (or equivalently Stokes vector) at the sensor. The result is the sum of all sub-images, integrated across the whole condenser aperture. This approach is also known as Abbe’s superposition method [114, 115].

3.3.1 Transfer of polarisation signal

To build a quantitative model, we start by considering the Fraunhofer diffraction integral that propagates the field \mathbf{E} from the front focal plane (x, y) of a lens to its back focal plane (x', y') . The resultant field, \mathbf{E}' , is given by [80] as

$$\mathbf{E}'(x', y') = -\frac{ie^{2ikf}}{\lambda f} \iint_{-\infty}^{\infty} \mathbf{E}(x, y) \exp\left[-\frac{2i\pi}{\lambda f}(xx' + yy')\right] dx dy, \quad (3.20)$$

where f is the focal length of the lens, λ is the wavelength, and $k = 2\pi/\lambda$ is the wave number of the field. By comparing it with the 2D Fourier transform defined with Equation 2.55, it can be recognised that the propagated field \mathbf{E}' is proportional to a scaled Fourier transform $\tilde{\mathbf{E}}(\xi, \eta)$, where

$$\xi = x'/\lambda f, \quad \eta = y'/\lambda f. \quad (3.21)$$

As a result, while the object and image spaces are regarded as real spaces, the condenser and objective pupil spaces can be described in the Fourier domain, parametrised by the spatial frequency vector $\boldsymbol{\nu} = (\xi, \eta)^\top$. Although the scaling between Fourier and real spaces are different depending on the focal length of the lens, the treatment of such infinity space in the Fourier domain is valid as long as the magnification factor is taken into account when propagating into the image space.

Assuming the field at the condenser pupil plane to be a spatially incoherent secondary source, the contribution from each point in the condenser aperture is summed incoherently. Denote the coordinate of a point source of interest $\boldsymbol{\nu}_0$ within the condenser aperture \mathcal{A} . The field from this point source in the object space shall be described by a (polarised) plane wave, written as

$$\mathbf{E}_1(\mathbf{r}_1; \boldsymbol{\nu}_0) = \delta E_0 e^{2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_1} |L\rangle, \quad (3.22)$$

where \mathbf{r}_1 is the physical coordinate in the object space, and δE is an infinitesimal field amplitude.

At the object, the field shall be modulated by the Jones matrix \mathbb{T}_o . Such modulation has to be

spatially confined, otherwise the polarisation modulation is global, and no image is produced. In this case, as we are trying to investigate the impulse response of the system, the modulation is assumed to be infinitesimally small, and therefore described with a Dirac delta function $\delta(\mathbf{r}_1)$.

In analogue to Equations 3.14–3.16, the field immediately after the object, \mathbf{E}'_1 , can be written as

$$\mathbf{E}'_1(\mathbf{r}_1; \boldsymbol{\nu}_0) = \delta E_0 e^{2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_1} |L\rangle + \delta E_0 e^{2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_1} \delta(\mathbf{r}_1) (\mathbb{T}_o - \mathbb{I}) |L\rangle, \quad (3.23)$$

$$\approx \delta E_0 e^{2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_1} |L\rangle + \delta E_0 i \phi_o \delta(\mathbf{r}_1) |L\rangle - \delta E_0 \frac{i \delta_o}{2} e^{2i \gamma_o} \delta(\mathbf{r}_1) |R\rangle, \quad (3.24)$$

upon applying the weak object approximation.

A Fourier transform of \mathbf{E}'_1 then propagates the field into the object pupil space, given by

$$\mathbf{E}_2(\boldsymbol{\nu}; \boldsymbol{\nu}_0) = -\frac{i \delta E_0 e^{2ik f_o}}{\lambda f_o} \left[\delta(\boldsymbol{\nu} - \boldsymbol{\nu}_0) |L\rangle + i \phi_o |L\rangle - \frac{i \delta_o}{2} e^{2i \gamma_o} |R\rangle \right], \quad (3.25)$$

where the S-wave is in turn spatially confined while the D-wave is not. This treatment is valid because the D-wave should spread beyond the objective pupil in the infinity space, and can be commonly found in literature (e.g. Born and Wolf [116]).

Propagation through the objective aperture is simple, as the aperture is modelled to be at the back focal plane of the objective. Assuming the pupil modulation is non-scattering and polarisation-insensitive, the pupil function can be modelled as a scalar multiplier, so the propagated field is simply

$$\mathbf{E}'_2(\boldsymbol{\nu}) = P(\boldsymbol{\nu}) \mathbf{E}_2(\boldsymbol{\nu}), \quad (3.26)$$

where in PPC,

$$P(\boldsymbol{\nu}) = \begin{cases} it_p, & \boldsymbol{\nu} \in \mathcal{A}; \\ 1, & \boldsymbol{\nu} \in \mathcal{P} \setminus \mathcal{A}; \\ 0, & \boldsymbol{\nu} \notin \mathcal{P}. \end{cases} \quad (3.27)$$

In treating the pupil spaces in the Fourier domain, the phase ring sits exactly where the condenser pupil aperture \mathcal{A} is, while the objective pupil \mathcal{P} is dictated by the objective NA.

Here, the thin-annulus approximation is no longer used, as the condenser aperture is assigned a finite width in the pupil function, and the S and D waves are not treated separately. Written explicitly, the field propagated through the objective modulation becomes

$$\mathbf{E}'_2(\boldsymbol{\nu}; \boldsymbol{\nu}_0) = -\frac{i\delta E_0 e^{2ikf_o}}{\lambda f_o} \left[P(\boldsymbol{\nu}_0) \delta(\boldsymbol{\nu} - \boldsymbol{\nu}_0) |L\rangle + i\phi_o P(\boldsymbol{\nu}) |L\rangle - \frac{i\delta_o}{2} e^{2i\gamma_o} P(\boldsymbol{\nu}) |R\rangle \right], \quad (3.28)$$

taking advantage that the S-wave is spatially confined.

To propagate into the image space requires another Fourier transform, although a scaling of $1/m$ is required to account for the difference in scaling of the two pupil spaces. The magnification m can be found as $m = f_t/f_o$, and the field at the sensor becomes

$$\mathbf{E}_3(\mathbf{r}_3; \boldsymbol{\nu}_0) = -\frac{\delta E_0 e^{2ik(f_o+f_t)}}{m} \left[P(\boldsymbol{\nu}_0) e^{-2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_3/m} |L\rangle + i\phi_o h(\mathbf{r}_3) |L\rangle - \frac{i\delta_o}{2} e^{2i\gamma_o} h(\mathbf{r}_3) |R\rangle \right], \quad (3.29)$$

where $h(\mathbf{r}_3)$ is the Fourier transform of the pupil function $P(\boldsymbol{\nu})$, usually known as the amplitude PSF.

To prepare the signal from such point sources to be integrated incoherently, its differential contribution of the Stokes vector shall be found. This is not only because the PolCam can be regarded as a polarimeter imaging the 2D projection of the Stokes vector in the NLSP, as discussed in the previous chapter, but also due to the fact that in the incoherent regime an optical system is linear in Stokes parameters. For mathematical convenience, we shall consider a complex signal $\tilde{S} = S_1 + iS_2$, and it can be shown that

$$\tilde{S} = 2\tilde{E}_L^* \tilde{E}_R. \quad (3.30)$$

Therefore, the differential contribution to this complex Stokes signal, $\delta\tilde{S}$, can be written as

$$\delta\tilde{S}(\mathbf{r}_3; \boldsymbol{\nu}_0) = \frac{2\delta E_0^2}{m^2} \left[P(\boldsymbol{\nu}_0) e^{-2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_3/m} + i\phi_o h(\mathbf{r}_3) \right]^* \cdot \left[-\frac{i\delta_o}{2} e^{2i\gamma_o} h(\mathbf{r}_3) \right], \quad (3.31)$$

$$= -\frac{\delta E_0^2}{m^2} \left[i\delta_o e^{2i\gamma_o} P^*(\boldsymbol{\nu}_0) e^{2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_3/m} h(\mathbf{r}_3) + \phi_o \delta_o e^{2i\gamma_o} |h(\mathbf{r}_3)|^2 \right]. \quad (3.32)$$

Note there are two terms in Equation 3.32. The former term is proportional to the polarisation modulation $\delta_o e^{2i\gamma_o}$ and carries no phase information, while the latter scales with the product $\phi_o \delta_o$. Because in previous derivations a weak object approximation has already been employed, any term higher than the first order becomes invalid, so the latter term shall be ignored. As a result, to the first order, the magnitude of this differential Stokes signal is proportional to the object birefringence, and its complex argument scales linearly with the object retardance, as expected from the polarisation analysis in the previous section.

Now, to account for the whole condenser aperture, the differential signal shall be integrated over the condenser aperture, which is shown as

$$\tilde{S}(\mathbf{r}_3) = \iint_{\boldsymbol{\nu}_0 \in \mathcal{A}} \delta \tilde{S}(\mathbf{r}_3; \boldsymbol{\nu}_0) d\boldsymbol{\nu}_0, \quad (3.33)$$

$$\approx -iI_0 \delta_o e^{2i\gamma_o} h(\mathbf{r}_3) \iint_{\boldsymbol{\nu}_0 \in \mathcal{A}} P^*(\boldsymbol{\nu}_0) e^{2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_3 / m} d\boldsymbol{\nu}_0, \quad (3.34)$$

where $I_0 = \frac{\delta E_0^2}{m^2}$ is the intensity, assuming the source to be uniform within its aperture. By introducing an aperture function

$$A(\boldsymbol{\nu}) = \begin{cases} 1, & \boldsymbol{\nu} \in \mathcal{A}; \\ 0, & \text{otherwise,} \end{cases} \quad (3.35)$$

the domain of integration can be generalised from \mathcal{A} to \mathcal{R}^2 , such that the equation can be recognised as a Fourier transform, given by

$$\tilde{S}(\mathbf{r}_3) = -iI_0 \delta_o e^{2i\gamma_o} h(\mathbf{r}_3) \iint_{\boldsymbol{\nu} \in \mathcal{R}^2} P^*(\boldsymbol{\nu}) A(\boldsymbol{\nu}) e^{2\pi i \boldsymbol{\nu} \cdot \mathbf{r}_3 / m} d\boldsymbol{\nu}, \quad (3.36)$$

$$= -iI_0 \delta_o e^{2i\gamma_o} h(\mathbf{r}_3) \mathcal{F} \{P^*(\boldsymbol{\nu}) A(\boldsymbol{\nu})\}(\mathbf{r}_3). \quad (3.37)$$

By recognising $h(\mathbf{r}_3) = \mathcal{F} \{P(\boldsymbol{\nu})\}(\mathbf{r}_3)$, the equation can be rewritten using the convolution theorem as

$$\tilde{S}(\mathbf{r}_3) = -iI_0 \delta_o e^{2i\gamma_o} \mathcal{F} \{P(\boldsymbol{\nu}) * [P^*(\boldsymbol{\nu}) A(\boldsymbol{\nu})]\}(\mathbf{r}_3). \quad (3.38)$$

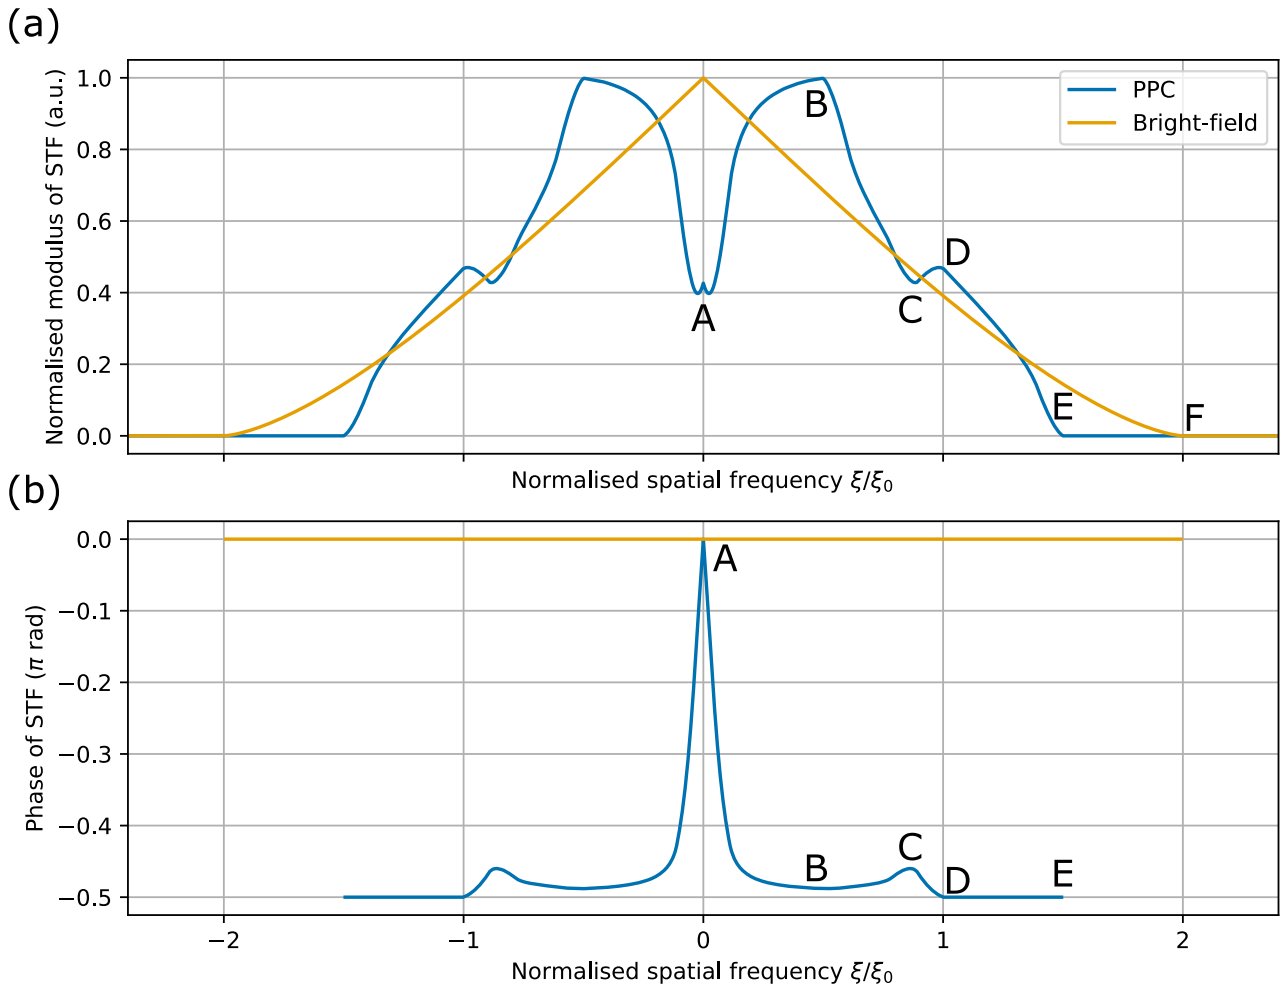


Figure 3.3: The Stokes transfer functions (STFs) of bright-field and PPC setups, shown as line profiles at $\eta = 0$, in magnitudes (a) and phases (b). The horizontal spatial frequency ξ is normalised by $\xi_0 = \text{NA}/\lambda$. The blue lines show the theoretical STF of Olympus LUCPLFLN 40x/0.60 Ph2, while the orange lines show the STF in a setup with bright-field illumination and detection at the same NA.

The result in Equation 3.38 can be broken into three parts. Firstly, it is proportional to the background intensity I_0 , which is expected for a transmission-light illuminated setup. As discussed in the previous chapter, it should be maximised for the best SNR so long as pixels are not over-saturated. Secondly, as expected from the Jones algebra analysis, the signal is proportional to $-i\delta_o e^{2i\gamma_o}$. Thirdly, it can be seen that the signal is band-limited in the Fourier domain. In line with CTF and OTF, we shall call it the Stokes transfer function (STF), identified as

$$\text{STF}(\boldsymbol{\nu}) = P(\boldsymbol{\nu}) * [P^*(\boldsymbol{\nu})A(\boldsymbol{\nu})]. \quad (3.39)$$

Figure 3.3 shows the magnitude and phase of the theoretical STF for Olympus LUCPLFLN 40x/0.60 Ph2 objective, whose size parameters were determined by imaging its pupil through

a Bertrand lens. As Equation 3.3 has been kept general for any aperture and pupil functions, the bright-field case can also be plotted for comparison. The condenser aperture is assumed to be sufficiently large, and the objective aperture is modelled as a circular one at the same NA as the phase-contrast objective, in which case Equation 3.39 recreates the well-known result that its OTF is the autocorrelation of its pupil function [80].

There are many significant features to observe in Figure 3.3. Firstly, the extent in which the STF of PPC is non-zero is reduced due to its annular condenser aperture. In bright-field setup upon coherent illumination, the CTF cuts off at $\xi = \xi_0 = \text{NA}/\lambda$, being the Fourier transform of the pupil function; upon incoherent illumination, the OTF cuts off at $\xi_{\text{cut-off}} = 2\xi_0$ instead, due to the autocorrelation. This cut-off is labelled as Point F in Figure 3.3. In PPC, however, because the illumination has a lower NA than the detection, the cut-off frequency becomes

$$\xi_{\text{cut-off}} = \frac{\text{NA}_{\text{illu}} + \text{NA}_{\text{obj}}}{\lambda}, \quad (3.40)$$

which evaluates to approximately $1.50\xi_0$ for the phase-contrast setting in use (Point E in Figure 3.3). Its diffraction limit is then the reciprocal of the cut-off frequency, reduced to 589 nm from its bright-field counterpart of 442 nm, at the source wavelength of 530 nm.

Secondly, in the STF of PPC, there is a significant dip in the low-spatial-frequency regime, where the STF should have been maximised. The feature, sometimes referred to as the lack of DC components, is classic many phase contrast techniques [69, 117, 118], and can lead to the shade-off artefact in images if not corrected for. Qualitatively, the phenomenon can be understood as a consequence of the finite width of the phase ring. In the simplified picture provided in the previous section, it has been assumed that the phase ring modulates all the S-wave and no D-wave. In reality, however, some parts of the D-wave is not modulated enough to leave the phase ring due to its finite width. As a result, this part of the D-wave arrives at the sensor with a smaller amplitude than it should have, contributing to a smaller DoLP signal, and therefore a smaller Stokes signal in magnitude. This part of the D-wave is also phase-shifted, such that there is no phase difference compared to the S-wave which is phase-shifted in the same way, resulting in a STF phase of zero.




Figure 3.4: Illustration of the ‘shifting pupil’ picture for the STF of PPC (A–E) and bright-field (F) setups. The sample spatial frequency of interest decides the distance by which the objective pupil is shifted. Each case correspond to a position in Figure 3.3 annotated with the same letter.

Thirdly, except for the DC regime, the phase introduced by the STF is mostly $-\pi/2$. This is consistent with previous analysis on the quarter-wave-modulation between the S-wave and the D-wave, and can be seen as a generalisation of Equation 3.12. The variation from the expected $-\pi/2$ is physically also due to the crosstalk that some part of the D-wave happens to pass through the phase ring and get accidentally modulated in phase. In the thin-annulus limit, the variation diminishes, such that the STF phase would converge to $-\pi/2$ for all $\xi \neq 0$. This is, however, not practical as the light throughput would be greatly compromised.

This variation in STF phase has important consequences in measuring the birefringence of the sample, because the offset between the measured AoLP and the birefringent azimuth now depends on the spatial frequency of the sample and the width of the phase ring, giving rise to a systematic error. Nevertheless, in the context of voxel imaging for optical data storage (ODS), a systematic error plays minimal role in the classification accuracy.

Furthermore, there is more structure in the STF of PPC compared to the bright-field case. To understand the structure more intuitively, we can interpret the convolution in the expression of STF as if the objective pupil is spatially shifted from the annular illumination in the Fourier space, by a distance corresponding to the sample spatial frequency of interest. Figure 3.4 illustrates a few cases of this ‘shifting pupil’ picture whose labels correspond to the annotations in Figure 3.3. Case A corresponds to the DC component discussed above, where no shift

happens at a spatial frequency of zero. It can be seen that the whole annular aperture coincides with the phase ring, resulting in maximal attenuation. As the shifting amount increases, the overlap between the annular aperture and the phase ring drops, until Case B where the annular aperture touches the edge of the objective pupil. Case B is also an example of the D-wave mostly passing through the objective pupil in regions outside the phase ring, creating a near-quarter-wave phase difference with the S-wave, as can be seen in Figure 3.3(b). After that, some part of the illumination would be blocked by the objective pupil, resulting in a decreasing magnitude in the STF.

Case C then illustrates when the annular aperture and the phase ring are overlapped at the other end, where a locally maximal amount of light is modulated, giving rise to the dip in the STF magnitude and a local maximum of the phase. As the pupil shifts further, the amount of overlap between the two annuli diminishes, until in Case D when they separate. This marks the end of any phase modulation from the phase ring, from which the STF phase is always exactly $-\pi/2$, until the objective pupil is shifted away from the condenser annulus in Case E which is the PPC cut-off. In contrast, the bright-field cut-off (shown as Case F) requires a farther shift due to the larger size of the condenser aperture (i.e. a higher illumination NA).

3.3.2 Transfer of phase signal

The intensity signal, S_0 , can also be derived similarly. The point-source contribution from $\boldsymbol{\nu}_0$ can be written as

$$\delta S_0 = \left| \tilde{E}_L \right|^2 + \left| \tilde{E}_R \right|^2, \quad (3.41)$$

$$= \frac{\delta E_0^2}{m^2} \left[\left| P(\boldsymbol{\nu}_0) e^{-2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_3 / m} + i \phi_o h(\mathbf{r}_3) \right|^2 + \left| -\frac{i \delta_o}{2} e^{2i \gamma_o} h(\mathbf{r}_3) \right|^2 \right], \quad (3.42)$$

$$= I_0 \left\{ \left| P(\boldsymbol{\nu}_0) \right|^2 + 2 \phi_o \Re \left[i P^*(\boldsymbol{\nu}_0) e^{2\pi i \boldsymbol{\nu}_0 \cdot \mathbf{r}_3 / m} h(\mathbf{r}_3) \right] + \phi_o^2 |h(\mathbf{r}_3)|^2 + \frac{\delta_o^2}{4} |h(\mathbf{r}_3)|^2 \right\}. \quad (3.43)$$

Note that the last two terms in Equation 3.43 shall be discarded due to the weak object approximation. Using a similar trick of generalising the domain of integration and identifying

a Fourier transform, the intensity can be found, given as

$$S_0(\mathbf{r}_3) = \iint_{\boldsymbol{\nu}_0 \in \mathcal{A}} \delta S_0(\mathbf{r}_3; \boldsymbol{\nu}_0) d\boldsymbol{\nu}_0, \quad (3.44)$$

$$\approx I_0 |P(\boldsymbol{\nu}_0)|^2 + 2I_0 \phi_o \Re \left[i h(\mathbf{r}_3) \iint_{\boldsymbol{\nu} \in \mathcal{R}^2} e^{2\pi i \boldsymbol{\nu} \cdot \mathbf{r}_3 / m} P^*(\boldsymbol{\nu}) A(\boldsymbol{\nu}) d\boldsymbol{\nu} \right], \quad (3.45)$$

$$= I_0 |P(\boldsymbol{\nu}_0)|^2 + 2I_0 \phi_o \Re [i \mathcal{F}\{P(\boldsymbol{\nu})\}(\mathbf{r}_3) \cdot \mathcal{F}\{P^*(\boldsymbol{\nu}) A(\boldsymbol{\nu})\}(\mathbf{r}_3)], \quad (3.46)$$

$$= I_0 |P(\boldsymbol{\nu}_0)|^2 - 2I_0 \phi_o \Im [\mathcal{F}\{P(\boldsymbol{\nu}) * [P^*(\boldsymbol{\nu}) A(\boldsymbol{\nu})]\}(\mathbf{r}_3)]. \quad (3.47)$$

In contrast to \tilde{S} , S_0 now contains a constant term for background, which is typical for transmission-illuminated setup. In PPC, assuming perfectly aligned condenser annulus and phase ring, the first term becomes $I_0 t_p^2$. The second term is the modulation from the background, whose amount is proportional to the object phase ϕ_o and independent of the object birefringence. In addition, it can be identified that the (imaginary part of) STF also present within the equation.

Denote the modulation term as $M(\mathbf{r}_3)$. Its (inverse) Fourier transform is can then be derived as

$$\mathcal{F}^{-1}\{M(\mathbf{r}_3)\} = \mathcal{F}^{-1}\{-2I_0 \phi_o \Im [\mathcal{F}\{\text{STF}(\boldsymbol{\nu})\}(\mathbf{r}_3)]\}, \quad (3.48)$$

$$= iI_0 \phi_o [\text{STF}(\boldsymbol{\nu}) - \mathcal{F}^{-1}\{(\mathcal{F}\{\text{STF}(\boldsymbol{\nu})\}(\mathbf{r}_3))^*\}], \quad (3.49)$$

$$= iI_0 \phi_o [\text{STF}(\boldsymbol{\nu}) - \text{STF}^*(-\boldsymbol{\nu})], \quad (3.50)$$

making use of the identity

$$\mathcal{F}\{f^*(\mathbf{r})\} = F^*(-\boldsymbol{\nu}). \quad (3.51)$$

From Equation 3.50, it can be observed that the modulation term is always band-limited in the Fourier space. Its inverse Fourier transform can only be non-zero in regions where $\text{STF}(\boldsymbol{\nu})$ or $\text{STF}(-\boldsymbol{\nu})$ is non-zero. In fact, as both the aperture and pupil functions are usually circularly symmetric, the modulation term M is band-limited in the same region as the STF itself in most cases. In bright-field setups, however, the modulation term becomes zero because both the aperture and pupil function are real and circularly symmetric, and thus so is its STF. As a result, Equation 3.50 is evaluated as zero, and no phase information can be transferred. This

confirms bright-field setups can capture no phase information.

For phase information to be transferred to S_0 , the STF of the system has to be non-Hermitian, i.e.,

$$\text{STF}(\boldsymbol{\nu}) \neq \text{STF}^*(-\boldsymbol{\nu}). \quad (3.52)$$

To achieve this, a ZPC/PPC setup utilises the phase ring such that the pupil function P becomes complex and non-Hermitian—in fact, in the thin-annulus limit, the quarter-wave modulation in such a setup maximises the modulation term, because the STF becomes purely imaginary, as discussed before. In other methods such as differential phase contrast (DPC) and Fourier ptychography, the Hermitian symmetry of the STF can be alternatively broken by asymmetric illumination or detection, effectively breaking the central symmetry of P , A , or both [119–121]. In such cases the STF is usually real (except the quarter-phase modulation introduced by the weak phase approximation in Equation 3.1), and the function is commonly referred to as the weak-object transfer function (WOTF) in the field of quantitative phase imaging (QPI). It is also worth noting that, in cases where the weak object assumption has to be abandoned, the object spectrum has to be taken into account. No transfer function would associate with the image formation process [112], because a partially coherent system, in general, is not linear. In those cases, the Hopkins transmission cross-coefficient (TCC) method [122, 123] would be a sensible extension of the STF model.

Additionally, as the imaginary part of the STF plays a critical role in the transfer of phase information, the dip of the STF in the low-spatial-frequency regime also affects phase imaging. In particular, at the DC component where $\boldsymbol{\nu} = 0$, the STF is real. Therefore, the DC component of the object phase is not transferred to the image at all, and cannot be recovered in a conventional ZPC setup. Physically, this is because the low-spatial-frequency content of the object corresponds to only a small change in the wave vector of the D wave. As a result, for a phase ring with finite width, the S wave and this part of the D wave cannot be perfectly separated at the back focal plane of the objective. This part of the D wave is therefore also affected by the phase ring, and does not contribute to the intensity signal in the intended way. Consequently, the resultant phase-contrast image has a high-pass-filtering effect. This effect is

particularly visible for large or slowly varying phase objects, where the broad phase distribution is underestimated. This loss of contrast within an extended phase object is known as the shade-off artefact [69]. The same high-pass-filtering behaviour can also produce bright or dark rims near sharp phase boundaries, which are known as halo artefacts [124–126]. These artefacts can therefore be understood directly from the low-spatial-frequency response of the STF.

3.4 Conclusions

This chapter introduces the design and principle of the PPC microscopy. Starting with the ZPC setup, we introduced how illumination through an annular condenser can be used, in conjunction with the weak phase approximation, to decompose the transmitted field into two components—a phase-encoding, spatially varying D-wave, and a referencing S-wave. It was also identified that the two components are spatially separate in the objective pupil plane, which can be taken advantage of by employing a phase ring to module them differently. It was acknowledged how a quarter-wave modulation in phase encodes phase information into detectable intensity at the sensor, and how an amplitude modulation of the S-wave enhances contrast.

The idea was then taken into the design of PPC. By extending the weak phase approximation into a more general weak object approximation which also assumes a small birefringent retardance in addition to the small residual phase, the idea of considering the detected field as a coherent superposition between the D- and S-waves can be translated into PPC, with the major difference being a vector description is now necessary. Upon circularly polarised illumination, the object now creates two D-waves: a parallelly polarised, phase shifted component with an amplitude proportional to the amount of residual phase (as in ZPC), and an orthogonally polarised component whose amplitude is proportional to the amount of retardance and whose phase encodes the birefringent azimuth. The clean separation of phase and birefringence information in the circular basis of polarisation is what further allows the image model to be built with the two sample features decoupled (in first order of approximation).

To gain a deeper understanding of the image formation process, it was then necessary to understand the impulse response of the optical system. By treating the PPC setup as a 6-f system from the condenser aperture plane to the image plane, the action of all three lenses was described by the Fraunhofer diffraction integral, in the paraxial approximation. Assuming the sensor to be a PolCam, it was found that both the complex signal $\tilde{S} = S_1 + iS_2$ and the intensity signal S_0 are band-limited in the Fourier space by the domain of the STF, encoding birefringence and phase information, respectively.

The STF was found as a Stokes-analogue to the OTF, with its magnitude indicates how the sample retardance translates into the DoLP signal, and its complex phase indicates how birefringent azimuth translates to AoLP at the PolCam. As the STF was found with an expression general to any aperture and pupil function, its features were compared in the bright-field and PPC setups. In the bright-field case, the STF recreates an autocorrelation of the pupil function, which can be found in many non-polarimetric systems. In contrast, for the PPC, the STF has more structure, which can be explained in the ‘shifting-pupil’ picture for the convolution between the condenser aperture and the objective pupil. Notable consequences include:

- (i) the lower illumination NA of the annular aperture reduces maximal resolvable spatial frequency from $2NA/\lambda$ to $(NA_{\text{illu}} + NA_{\text{obj}})/\lambda$;
- (ii) the finite widths of the pupil annuli cause the D-wave in the near-DC regime to be accidentally modulated by the phase ring, leading to shade-off artefacts; and
- (iii) the STF at $\xi < NA/\lambda$ has a varying complex phase, resulting in a systematic error in AoLP depending on the spatial frequency of interest.

These features of the PPC have to be taken into account when designing the system or interpreting any results produced.

Using a similar process, it was found how the intensity signal S_0 is related to the residual phase of the object. In the weak object approximation, the intensity modulation was found independent of any sample birefringence, and is also band-limited by the domain of the STF.

In the bright-field setup, however, the STF is real and circularly symmetric, leading to no phase detection, as one can expect. The employment of phase-contrast objectives solves this problem by introducing a phase modulation, leading to a complex pupil function and thus a non-Hermitian STF, such that the modulation term in the intensity is non-zero. Alternatively, the Hermitian symmetry of the STF can be broken by breaking the circular symmetry of pupils, providing us with an alternative perspective of understanding some QPI techniques such as DPC and Fourier ptychography.

It should be noted that, the imaging model described above is based on a key assumption that the field is spatially incoherent at the start of the 6-f system (i.e. condenser aperture plane). After obtaining the coherent response of the system to a point source within the condenser aperture, the second-order statistics of the field were accounted for by incoherently integrating this response over the entire aperture. This method, however, cannot handle a general case of partial coherence, as there is no degree of freedom for the amount of coherence in the model. As we will discuss in Chapter 4, when implemented numerically for simulating images, the method also suffers from serious spectral leakage.

Chapter 4

Modal representation of a partially coherent source for image simulation

In the previous chapter, an analytical imaging model for polarisation-sensitive phase contrast (PPC) microscopy was introduced, with its Stokes transfer function (STF) derived. Although it is a partially coherent treatment, which was deemed necessary for multi-lens systems [107], this approach [112, 113] necessitate the assumption that the optical field starts with perfect spatial incoherence. In optical systems, coherence can be acquired simply upon propagation of the field [127], as a consequence of the van Cittert–Zernike theorem [108, 116], especially in the presence of apertures and lenses. For the previous chapter, the condenser aperture plane was assumed to be perfectly incoherent, which had little physical reasoning and seemed unnatural in a partially coherent treatment.

In addition, further problems arise when computational simulations are built upon the picture of ‘incoherent points’. As part of my MRes project [41], a simulation framework was developed based on a similar approach. Starting from the object plane, plane-wave components of different wave vectors are propagated numerically through a simulated sample, the objective pupil, and the micro-polariser array (MPA) of the polarisation camera (PolCam). Lens propagations were implemented with the Fraunhofer diffraction integral (Equation 3.20) using discrete Fourier transform (DFT) algorithms [128]. As each plane wave corresponds to a certain




Figure 4.1: (a) Amplitude of a simulated field in the objective pupil plane (after the phase ring), propagated from a single point in the condenser aperture. Significant side lobes arise as the S-wave is locally confined and DFT algorithms introduce a windowing depending on the size of the simulation grid. (b) A Hann function has a narrow Fourier spectrum [129], so it can be multiplied at the object plane to mitigate the side-lobe artefact.

point source in the condenser aperture plane, the intensity contributions can then be summed by sampling the condenser aperture at Gaussian quadrature positions, which is effectively a numerical integration of Equation 3.33.

When a plane-wave component is propagated to the objective pupil plane, however, significant side lobes are usually created, as shown in Figure 4.1(a). Because the objective pupil plane images the condenser aperture, the S-wave produced by a plane wave focuses into a point in the phase ring. The side lobes are a consequence of the finite windowing of DFT algorithms. In the numerical simulation, the field is sampled in a Cartesian grid. In DFT, the finite size of the grid acts as a window, so the result is convolved with the Fourier transform of the window function, leading to the creation of side lobes. The artefact is particularly prominent in this case, as the S-wave is so locally confined in the objective pupil plane. In fact, the side-lobes can outweigh the D-waves as the object is typically weak in the simulations, ruining the image formation process.

To reduce the amount of side lobes, a multiplication with a Hann function [129] (shown in Figure 4.1(b)) was introduced at the object plane, such that the Fourier transform of the window function is narrower, reducing the side lobes. Nevertheless, the Hann windowing is merely

a computational trick instead of a representation of any meaningful physical process, and its action depends on the size of the simulated grid. In addition, it effectively changes the illumination into windowed plane waves, which cannot be produced by a point source in the condenser aperture plane, and therefore not only changes the illumination of the system, but also breaks the assumption of perfect incoherence. As a result, although the simulation produced qualitatively sensible images [41], the method is self-contradictory and any quantitative results are not well-supported.

In this chapter, coherent-mode decomposition is offered as an alternative to the incoherent-point description of the source. It shall be seen that, a non-zero amount of spatial coherence can be specified for the source, which can also be set independently of the intensity profile of the source. By decomposing the source into mutually incoherent modes, each mode can still be propagated coherently through the system. In contrast to the ‘incoherent points’, the coherent modes are not locally confined, and therefore no side lobes would be generated, removing the need of Hann windowing. As there are different ways in which the same sources can be decomposed mathematically, a few choices shall be discussed and compared.

4.1 Second-order statistics of a partially coherent field

To set context for following discussions, some theoretical background needs to be provided. In coherence theory [116], light is regarded as a stochastic process, with the electromagnetic field treated as a random variable. In the scalar theory [88, 111], the field is usually parametrised with the complex-valued analytic signal, $V(\mathbf{x}, t)$ in a domain $\mathbf{x} \in \mathcal{D}$.

Correlation between two space-time points (\mathbf{x}_1, t_1) and (\mathbf{x}_2, t_2) can therefore be quantified with the mutual coherence function, Γ , defined as

$$\Gamma(\mathbf{x}_1, \mathbf{x}_2, t_1, t_2) = \mathbb{E} [V^*(\mathbf{x}_1, t_1)V(\mathbf{x}_2, t_2)], \quad (4.1)$$

where \mathbb{E} indicates the expected value across all possible realisations. Assuming the field to be

statistically stationary, the mutual coherence depends only on the time difference $\tau := t_2 - t_1$, such that

$$\Gamma(\mathbf{x}_1, \mathbf{x}_2, \tau) = \mathbb{E} [V^*(\mathbf{x}_1, 0)V(\mathbf{x}_2, \tau)]. \quad (4.2)$$

While the mutual coherence can be regarded as a representation of the second-order statistics in the space-time domain, it is often useful to express the correlation in the space-frequency domain, which can be obtained by a Fourier transform with respect to time. The resultant quantity is the cross-spectral density function (CSDF) [130, 131], defined as

$$W(\mathbf{x}_1, \mathbf{x}_2, \nu) = \int_{-\infty}^{\infty} \Gamma(\mathbf{x}_1, \mathbf{x}_2, \tau) \exp(-i2\pi\nu\tau) d\tau, \quad (4.3)$$

where ν denotes the frequency. The spectral density $S(\mathbf{x}, \nu)$ can be found by

$$S(\mathbf{x}, \nu) = W(\mathbf{x}, \mathbf{x}, \nu). \quad (4.4)$$

The spectral degree of coherence is then defined as

$$\mu(\mathbf{x}_1, \mathbf{x}_2, \nu) = \frac{W(\mathbf{x}_1, \mathbf{x}_2, \nu)}{\sqrt{S(\mathbf{x}_1, \nu)}\sqrt{S(\mathbf{x}_2, \nu)}}. \quad (4.5)$$

The quantity is a useful measure of the correlation as Born and Wolf [116] proved it is bounded between unity for perfect coherence and zero for incoherence, such that

$$0 \leq \mu(\mathbf{x}_1, \mathbf{x}_2, \nu) \leq 1. \quad (4.6)$$

The linear system theory [80, 132] provides an effective framework for describing the first-order statistics of light. In this theory, the input $f(\mathbf{x})$ and the output $g(\mathbf{x}')$ of an optical system can be related in the general form of

$$g(\mathbf{x}') = \int_{-\infty}^{\infty} f(\mathbf{x})h_1(\mathbf{x}', \mathbf{x}) d\mathbf{x}, \quad (4.7)$$

where $h_1(\mathbf{x}', \mathbf{x})$ is the linear impulse response of the system. When light is completely coherent, f and g takes the meaning of field amplitudes, and h_1 represents the coherent transfer function (CTF). In contrast, in the incoherent regime, f and g represents the light intensity, and the optical transfer function (OTF) is used.

When second-order effects are considered [133], a linear impulse response is no longer sufficient. Instead, it is replaced with the second-order Volterra kernel $h_2(\mathbf{x}'; \mathbf{x}_1, \mathbf{x}_2)$, representing the double-impulse response of the system. In this formulation, most optical systems can be considered to be bilinear [111, 134], with the spectral density $S'(\mathbf{x}', \nu)$ regarded as the output and the CSDF $W(\mathbf{x}_1, \mathbf{x}_2, \nu)$ regarded as the input. For simplicity, the explicit dependence on frequency ν will be omitted in discussions below. Consequently, the resultant spectral density as the output of a system can be related to the CSDF as the input, given as

$$S'(\mathbf{x}') = \iint_{\mathcal{D}^2} W(\mathbf{x}_1, \mathbf{x}_2) h_2(\mathbf{x}'; \mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_1 d\mathbf{x}_2. \quad (4.8)$$

Computational implementation of Equation 4.8 is, however, rather difficult, due to the high dimensional nature. For a Cartesian grid of the size $N \times N$, the double-impulse response $h_2(\mathbf{x}'; \mathbf{x}_1, \mathbf{x}_2)$ scales with N^6 , posing significant challenges to both memory allocation and processing power. Alternatively, the CSDF can be propagated by itself [135–137], although it still requires arrays of size N^4 . To address such challenges, modal expansion of the CSDF becomes useful, such that a partially coherent system can be represented by a sum of coherent systems that are mutually incoherent with each other. While it is also possible to carry out eigenanalysis [138] on the Hopkins transmission cross-coefficient (TCC) [112, 122, 123], this work primarily focuses on the decomposition of the source while trying to keep the optical system general.

Coherent mode decomposition [131, 139] takes advantage of the fact that the CSDF is square-integrable, Hermitian symmetric and non-negative definite [88, 111]. As a consequence of Mercer's theorem [140], a CSDF can therefore always be expanded as a sum of mutually incoherent

modes, in the form of a Mercer's expansion, given as

$$W(\mathbf{x}_1, \mathbf{x}_2) = \sum_{n \in \mathcal{N}} \lambda_n \phi_n^*(\mathbf{x}_1) \phi_n(\mathbf{x}_2), \quad (4.9)$$

where λ_n and ϕ_n are the n -th eigenvalue and eigenfunctions of the homogeneous Fredholm integral:

$$\int_{\mathcal{D}} W(\mathbf{x}_1, \mathbf{x}_2) \phi_n(\mathbf{x}_1) d\mathbf{x}_1 = \lambda_n \phi_n(\mathbf{x}_2), \quad (4.10)$$

and \mathcal{N} represents the set of all coherent modes. Mercer's theorem states that all eigenvalues are real and non-negative, and the eigenfunctions form a complete and orthonormal basis in \mathcal{D} , such that

$$\int_{\mathcal{D}} \phi_n^*(\mathbf{x}) \phi_m(\mathbf{x}) d\mathbf{x} = \delta_{nm}. \quad (4.11)$$

In coherence theory, it can then be shown that the eigenfunctions $\phi_n(\mathbf{x})$ satisfies the Helmholtz equation [111] and can propagate linearly through the optical system. It was also shown by Wolf [131] that each eigenfunction is completely coherent. Therefore, the coherent mode decomposition illustrates a picture where a partially coherent source can be represented by a sum of coherent modes $\phi_n(\mathbf{x})$ that are mutually incoherent with each other. Together, the modal expansion, the eigenvalues λ_n , and the modes ϕ_n represent the coherence structure of the field.

To simulate the partially coherent field propagating through an optical system, we can thus treat each mode independently. Suppose a coherent mode $\phi_n(\mathbf{x})$ propagates through a linear system \mathcal{L} and produces an output $\phi'_n(\mathbf{x}') = \mathcal{L}\phi_n(\mathbf{x})$. The output CSDF can then be reassembled as

$$W'(\mathbf{x}'_1, \mathbf{x}'_2) = \sum_n \lambda_n \phi_n'^*(\mathbf{x}'_1) \phi_n'(\mathbf{x}'_2), \quad (4.12)$$

such that the output spectral density would be

$$S'(\mathbf{x}') = W'(\mathbf{x}', \mathbf{x}') = \sum_n \lambda_n |\phi_n'(\mathbf{x}')|^2. \quad (4.13)$$

The coherent-mode structure of the field simplifies image simulation, because only the coher-

ent response \mathcal{L} of the system needs to be found, instead of the higher-dimensional Volterra kernel h_2 . As the modes are mutually uncorrelated, the simulation can treat each propagation independently, and only record a cumulative sum of their contributions to the output spectral density S' . In cases where the coherence structure can be expressed in closed form, numerical simulations can start with the coherent modes without the need of constructing a CSDF, reducing the memory need to N^2 .

It is worth noting that, more generally, the coherence of electromagnetic waves can be described by a vector theory [88, 141, 142], where partial polarisation can also be taken into account. In the vector coherence theory, a similar approach is also possible, with the aid of coherency matrices, cross-spectral density matrices, and vector coherent modes [20, 143]. In this work, however, discussion is limited to fully polarised fields whose polarisation state is uniform across the source. Therefore, when a vectorised representation of the coherent modes is needed for propagating through the polarisation-sensitive optical system, we simply write [144]

$$\phi_n(\mathbf{r}) = \phi_n(\mathbf{r})|E\rangle, \quad (4.14)$$

where $|E\rangle$ represents the normalised Jones vector of the polarisation state of the field. In addition, it is assumed that the spectral band of the source is sufficiently narrow that the coherence length is much greater than any optical path difference in the system, so a quasi-monochromatic treatment of the source will be employed throughout the discussion.

4.2 Bessel-correlated Schell-model (BSM) sources

To simulate a partially coherent source, its CSDF first has to be defined, such that the coherence structure can be found by solving the Fredholm integral in Equation 4.10. In particular, a useful type of source to consider is the Schell-model class [139, 145], which assumes the spectral degree of coherence μ depends only on the difference $\mathbf{x}_2 - \mathbf{x}_1$, not on them individually. Following

Equation 4.5, the CSDF can be written in the form of

$$W(\mathbf{x}_1, \mathbf{x}_2) = \sqrt{S(\mathbf{x}_1)}\sqrt{S(\mathbf{x}_2)}\mu(\mathbf{x}_2 - \mathbf{x}_1). \quad (4.15)$$

Using the van Cittert–Zernike theorem [108, 109], it was shown that the Schell-model represents well the CSDF of the far-field diffraction from an extended, incoherent source [114, 145, 146]. It is therefore ideal to represent secondary sources, for example, the condenser aperture plane in transmission-light setups, Köhler-illuminated with incoherent sources such as halogen lamps or light-emitting diodes (LEDs), under the paraxial approximation.

Among Schell-model sources, we consider the Bessel-correlated Schell-model (BSM) source [147], whose CSDF is given in the form of

$$W(\mathbf{x}_1, \mathbf{x}_2) = \sqrt{S(|\mathbf{x}_1|)}\sqrt{S(|\mathbf{x}_2|)}J_0(\beta|\mathbf{x}_1 - \mathbf{x}_2|), \quad (4.16)$$

where J_0 represents the Bessel function of the first kind and zeroth order.

Comparing the definition of the BSM to the general form of a Schell-model source in Equation 4.15, it can be seen that the Bessel function acts as the spectral degree of coherence. It is also assumed to depend only on the scalar distance between the two points, independent of direction. Moreover, the spectral density S is restricted to circularly symmetric functions that depend only on the polar radius. Under these assumptions, the coherence structure of a BSM source can be given by [147]:

$$W(\mathbf{x}_1, \mathbf{x}_2) = \sum_0^{\infty} \lambda_n \phi_n^*(\mathbf{x}_1) \phi_n(\mathbf{x}_2), \quad (4.17)$$

$$\lambda_n = 2\pi \int_0^{\infty} S(r) J_n^2(\beta r) r dr, \quad (4.18)$$

$$\phi_n(r, \theta) = \frac{1}{\sqrt{\lambda_n}} \sqrt{S(r)} [a_n J_n(\beta r) \exp(-in\theta) + b_n J_{-n}(\beta r) \exp(in\theta)], \quad (4.19)$$

in the polar coordinate system $\mathbf{x} = (r, \theta)$. The degeneracy constants a_n and b_n satisfy $a_n^2 + b_n^2 = 1$, while their ratio is otherwise arbitrary according to Gori et al. [147].




Figure 4.2: Comparisons of eigenvalues of BSM sources for apertures of different sizes. (a) Circular apertures of different radii. (b) Annular apertures of the same outer radius $R_o = 250/\beta$ but different sizes of inner radii. All eigenvalues are normalised to $[0,1]$ for comparison.

An important advantage of this model is that the spectral density function can be arbitrary as long as it is circularly symmetric with a finite norm. Therefore, the BSM is well suited to modelling circular or annular apertures that are commonly used in microscopy by choosing the spectral density function as

$$S(r) = \begin{cases} 1, & R_i \leq r < R_o, \\ 0, & \text{otherwise,} \end{cases} \quad (4.20)$$

where R_i and R_o represent the inner and outer radii of an annular aperture, and $R_i = 0$ for circular apertures. As a result, the absence of azimuthal dependence enables the coherence structure to be reduced to 1D, reflected in both its modal expansion and the expression of the eigenvalues.

Another key advantage is that, although the coherence structure in Equation 4.18 takes the form of an infinite sum, in practice a finite set of modes can represent the source very well, due to a characteristic ‘cut-off’ in the eigenvalues λ_n . Figure 4.2(a) shows the eigenvalues λ_n against the mode index n , for BSM sources with circular apertures of different sizes. Cut-offs can be observed at $n \sim \beta R_o$, due to the asymptotic behaviour of Bessel functions at high orders [147, 148]. As a result, in numerical simulations, it would not be necessary to go far beyond this cut-off. The dimensionless quantity βR_o is a usual measure of the degree of global coherence for BSM sources [111, 147].




Figure 4.3: Simulated PPC image of voxels, illuminated by a partially coherent BSM source at the condenser aperture plane, in the (a) Pol+I and (b) Pol-Only false-colour polarisation maps as defined in Section 2.4.3. Illumination from the BSM source can be seen to be restricted in an annular region, dictated by the Bessel form of its spectral degree of coherence.

Figure 4.2(b) shows how the (normalised) eigenvalues varies for annular apertures with the same outer radius but different inner radii. It can be seen that there is still the cut-off at $n \sim \beta R_o$ in all four cases. Depending on the size of the inner radius, however, there are different levels of suppression at $n \lesssim \beta R_i$, due to the inner stop of the aperture.

When simulating a BSM source as the secondary source at the condenser aperture plane in a microscope, however, significant problem arises as a non-uniform illumination in the object plane, as illustrated in Figure 4.3. Although the details of the Fourier simulation of PPC will be discussed in the next chapter, it can be seen that a BSM source only illuminates an approximate annular region, as can be seen in the Pol+I false-colour image in Panel (a). In Panel (b), despite the Pol-Only false-colour map not including the intensity signal, the degree of linear polarisation (DoLP) signal also comes in an annular shape, which is not ideal for representing a realistic imaging. This phenomenon can be understood as a consequence of the generalised van Cittert–Zernike theorem, which describes the relationship between the intensity distribution of an incoherent field and the mutual intensity of its far-field diffraction [88, 108]. Specifically, De Santis et al. [149] proved that any Schell-model source with a spectral degree of coherence $\mu(\mathbf{x})$ can be synthesised at the back focal plane of a single converging lens, by putting an incoherent extended source at the front focal plane. The intensity distribution of the incoherent primary

source should be proportional to $\mathcal{F}\{\mu(\mathbf{x})\}$. In the case of BSM sources where $\mu(\mathbf{x}) = J_0(\beta r)$, since the function is circularly symmetric, the 2D Fourier transform can be reduced into a zeroth-order Hankel transform (also known as a Fourier–Bessel transform) [80, 148], where a Bessel function J_0 is transformed into a Dirac delta function. As a result, if the BSM is used to model a secondary source focused through the field lens, it would correspond to an incoherent primary source that takes the shape of an infinitely thin annulus, which is inconsistent with the practical system. Furthermore, in this picture, the annularly shaped primary source is critically illuminating the object plane, and therefore casting an uneven contrast.

4.3 Apertured Gaussian Schell-model (GSM) sources

To solve the problem of annular illumination as in the BSM sources, it is necessary to change the form of μ in the CSDF. A relatively well-studied type is the 2D Gaussian Schell-model (GSM) source (also known as a Collet–Wolf source) [150], which is widely considered to be a good representation of realistic fields, gaining popularity in both theoretical and experimental aspects [114, 151–155]. For a GSM source, both the intensity and the spectral degree of coherence are Gaussian, such that

$$W(\mathbf{x}_1, \mathbf{x}_2) = I_0 \exp \left[-\frac{|\mathbf{x}_1|^2 + |\mathbf{x}_2|^2}{4\sigma_I^2} \right] \exp \left[-\frac{|\mathbf{x}_1 - \mathbf{x}_2|^2}{2\sigma_g^2} \right], \quad (4.21)$$

where σ_I and σ_g are the RMS widths of the Gaussian distributions for intensity and correlation respectively, while I_0 being a positive constant. Their ratio,

$$\gamma := \frac{\sigma_g}{\sigma_I}, \quad (4.22)$$

is defined as the degree of global coherence for GSM sources. In contrast to $\mu(\mathbf{x}_1, \mathbf{x}_2)$ which is defined between two points, γ measures the amount of coherence for the whole source, and varies from zero to infinity.

Compared to BSM, a critical advantage of GSM is that the correlation is now Gaussian, whose

Hankel transform is also Gaussian, creating a much more uniform illumination profile. Nevertheless, the GSM demands a Gaussian intensity profile that is infinitely extensive, in contradiction to realistic optical systems. For the simulation of imaging systems, in particular, a well-defined numerical aperture (NA) is vital for an accurate representation of the resolving ability. It is therefore necessary to investigate apertured GSM sources.

Consider an on-axis, circularly symmetric aperture $\mathcal{A} \subset \mathcal{D}$, sitting immediately after a GSM source. The apertured source is then described by a CSDF of

$$W'(\mathbf{x}_1, \mathbf{x}_2) = \begin{cases} I_0 \exp \left[-\frac{|\mathbf{x}_1|^2 + |\mathbf{x}_2|^2}{4\sigma_I^2} \right] \exp \left[-\frac{|\mathbf{x}_1 - \mathbf{x}_2|^2}{2\sigma_g^2} \right], & \mathbf{x}_1, \mathbf{x}_2 \in \mathcal{A}, \\ 0, & \text{otherwise.} \end{cases} \quad (4.23)$$

The CSDF can also be reconstructed by aperturing the coherent modes of the original GSM source according to Equation 4.12, i.e.

$$W'(\mathbf{x}_1, \mathbf{x}_2) = \sum_n \lambda_n [A(\mathbf{x}_1)\phi_n(\mathbf{x}_1)]^* [A(\mathbf{x}_2)\phi_n(\mathbf{x}_2)], \quad (4.24)$$

where A denotes the masking from the aperture, given as

$$A(\mathbf{x}) = \begin{cases} 1, & \mathbf{x} \in \mathcal{A}, \\ 0, & \text{otherwise.} \end{cases} \quad (4.25)$$

Upon further propagation, the apertured modes $A(\mathbf{x})\phi_n(\mathbf{x})$ can then be regarded as the coherent modes of the apertured GSM source, although they are no longer necessarily orthonormal. From a physics perspective, light from a coherent mode is partly blocked by the aperture, and the fractional power transmitted can be quantified by the transmission efficiency, given as

$$\eta_n^2 = \iint_{\mathcal{A}} |\phi_n(\mathbf{x})|^2 d\mathbf{x}. \quad (4.26)$$

Therefore, a normalised coherent mode of the apertured GSM source can be expressed as

$$\phi'_n(\mathbf{x}) = \frac{1}{\eta_n} A(\mathbf{x}) \phi_n(\mathbf{x}), \quad (4.27)$$

with a corresponding weight of

$$\lambda'_n = \eta_n^2 \lambda_n, \quad (4.28)$$

such that

$$W'(\mathbf{x}_1, \mathbf{x}_2) = \sum_n \lambda'_n \phi_n'^*(\mathbf{x}_1) \phi'_n(\mathbf{x}_2). \quad (4.29)$$

Note that, even after normalisation, the apertured coherent modes do not necessarily form an orthonormal set, because the aperture effectively changes the integration domain from \mathcal{D} to \mathcal{A} , breaking the original orthogonality relation. Nevertheless, they can still be treated as coherent modes to be propagated independently through optical systems, with the final intensity obtained by summing their contributions with the updated weights λ'_n .

In computational simulations, a finite set of modes \mathcal{S} must be chosen to be generated and propagated. To ensure accuracy, a meaningful criterion can be the fraction f of cumulative weights, i.e.

$$f := \frac{\sum_{n \in \mathcal{S}} \lambda'_n}{\sum_{n \in \mathcal{N}} \lambda'_n}. \quad (4.30)$$

As a coherent mode ϕ'_n is normalised, its weight λ'_n represents the power being carried in a particular mode n . The denominator physically represents the total power of the apertured source, and thus can be calculated by integrating the spectral density within the aperture, i.e.

$$\sum_{n \in \mathcal{N}} \lambda'_n = \int_{\mathcal{A}} S(\mathbf{x}) d\mathbf{x}. \quad (4.31)$$

As we will see in later sections, the choice of ϕ' for the apertured GSM source is not unique. To achieve a given threshold of f , the minimum number of modes in \mathcal{S} can act as a measure of efficiency of a particular coherent-mode representation for computational simulations. The more modes a simulation requires, the more expensive it becomes to coherently propagate

them computationally, and the less efficient the representation is. This is especially relevant for quasi-homogeneous sources [156] whose intensity varies slowly compared to coherence, i.e. $\sigma_I \gg \sigma_g$ such that $\gamma \ll 1$. The following sections will present and compare such representations in different bases.

4.3.1 Representation in Hermite–Gaussian (HG) modes

In laser theory, Hermite–Gaussian (HG) modes are well established and commonly used to represent the transverse modes of laser beams [157, 158]. Gori [151] showed that the HG modes can be employed as the coherent modes of a GSM source, whose coherence structure is then given by

$$W(\mathbf{x}_1, \mathbf{x}_2) = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \lambda_{mn} \phi_{mn}^*(\mathbf{x}_1) \phi_{mn}(\mathbf{x}_2), \quad (4.32)$$

$$\lambda_{mn} = \lambda_0 q^{m+n}, \quad (4.33)$$

$$\phi_{mn}(x, y) = \sqrt{\frac{2}{\pi}} \sqrt{\frac{1}{2^{m+n} m! n!}} \frac{1}{w} H_m \left(\frac{x\sqrt{2}}{w} \right) H_n \left(\frac{y\sqrt{2}}{w} \right) \exp \left(-\frac{x^2 + y^2}{w^2} \right), \quad (4.34)$$

in the Cartesian coordinate system $\mathbf{r} = (x, y)$. In the equations, (m, n) denotes the mode indices in the x and y directions respectively, H_n denotes the Hermite polynomial of the n -th order [148]. Parameters are defined as

$$\lambda_0 := \frac{\pi I_0}{(a+b)^2}, \quad q := \frac{a-b}{a+b}, \quad w := \frac{1}{\sqrt{2ab}}, \quad (4.35)$$

where

$$a := \sqrt{\frac{1}{8\sigma_I^2} + \frac{1}{2\sigma_g^2}}, \quad b := \frac{1}{\sqrt{8\sigma_I^2}}. \quad (4.36)$$

Figure 4.4 illustrates a few HG coherent modes as examples, for a GSM source with a circular aperture of radius R_o . Note that the lateral extension of an HG mode increases with its mode index in that direction. In the x -direction, for instance, an HG mode is proportional




Figure 4.4: Example HG modes of a GSM source whose degree of global coherence $\gamma = 0.05$, upon a circular aperture of the size $R_o/\sigma_I = 1.1$. Mode indices m, n are denoted at the top of each mode. Modes are shown in their complex magnitudes.

to the expression $H_m(x')e^{-x'^2/2}$, where $x' := \sqrt{2}x/w$. Plancherel and Rotach [159] studied the asymptotic behaviour of the expression for large m , and concluded that it is oscillatory in the region $|x'| < \sqrt{2m+1}$, beyond which it transitions to an exponential decay [160, 161]¹. The so-called Plancherel–Rotach asymptotic of Hermite polynomials becomes useful in estimating the number of coherent modes required to represent an apertured GSM source, as an HG mode is significantly blocked by a circular aperture if

$$\sqrt{2m+1} > \frac{R_o\sqrt{2}}{w}, \quad (4.37)$$

which rearranges to

$$m > m_c := \left(\frac{R_o}{w}\right)^2 - \frac{1}{2}, \quad (4.38)$$

where m_c is used to denote the expected cut-off. Equations 4.35–4.36 can be used to find m_c for a given source as

$$m_c = \frac{R_o^2}{4\sigma_I^2} \sqrt{1 + \frac{4}{\gamma^2}} - \frac{1}{2}. \quad (4.39)$$

From Equation 4.39, it is evident that, for the same aperture and intensity distribution, the lower the coherence, the more coherent modes are needed to effectively represent the source. For quasi-homogeneous sources where $\gamma \ll 1$, m_c becomes large and can be approximated by

$$m_c \approx \frac{1}{2\gamma} \left(\frac{R_o}{\sigma_I}\right)^2, \quad (4.40)$$

which becomes useful in estimating the modes required in practical simulations.

¹For full asymptotic expansions of the expression, the reader can refer to Theorem 8.22.9, Page 201 of Szegő [161].

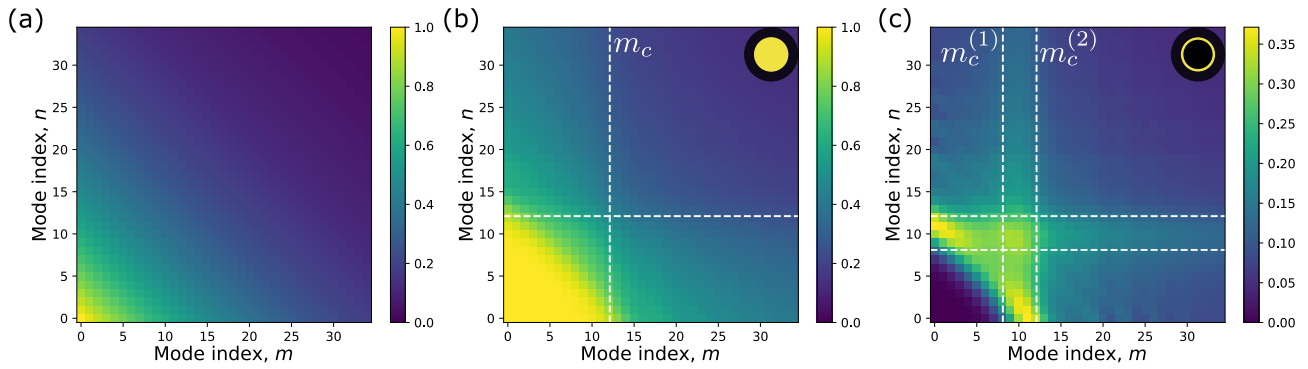


Figure 4.5: Eigenvalues of a GSM source in the Hermite–Gaussian (HG) representation, and transmission efficiencies through circular and annular apertures. (a) Normalised eigenvalues, λ_{mn}/λ_0 , of a GSM source whose degree of global coherence $\gamma = 0.05$. (b) Transmission efficiencies η_{mn}^2 of HG modes through a circular aperture of radius $R_o/\sigma_I = 1.1$. White dashed lines denote the expected cut-off m_c calculated according to Equation 4.40. (c) Transmission efficiencies η_{mn}^2 of HG modes through an annular aperture of inner radius $R_i/\sigma_I = 0.9$ and outer radius $R_o/\sigma_I = 1.1$. White lines denote the expected cut-offs m_c^1 and m_c^2 , corresponding to the inner and outer radii, respectively.

To illustrate how the cut-off affects the weights of the HG modes, Figure 4.5 shows the eigenvalues of a GSM source, and the transmission efficiencies of these modes through circular and annular apertures. As discussed in the previous section, the weight of a coherent mode λ'_{mn} of an apertured GSM source is the product of its original eigenvalue λ_{mn} and the transmission efficiency η_{mn}^2 through a specified aperture. Panel (a) shows the normalised eigenvalue λ_{mn}/λ_0 follows a simple geometric progression with respect to the sum $m+n$ according to Equation 4.33. In contrast, the transmission efficiencies η_{mn}^2 have much more structures. In particular, cut-offs can be seen in both directions, at mode indices close to m_c predicted by Equation 4.40, illustrated with the white lines. Through a circular aperture, as shown in Panel (b), the transmission efficiency is almost at unity for modes $m+n \lesssim m_c$ whose lateral extent lies completely within the circular aperture. It also has a considerable value if the mode is at least confined by the aperture in one direction, i.e. $m < m_c$ or $n < m_c$. At the top-right region where both indices go beyond the cut-off, the transmission efficiency is significantly compromised.

A similar pattern can be found for annular apertures as shown in Panel (c), although complicated by the inner stop of radius R_i of the annular aperture. Two cut-offs can be calculated, given as

$$m_c^{(1)} \approx \frac{1}{2\gamma} \left(\frac{R_i}{\sigma_I} \right)^2, \quad (4.41)$$

and

$$m_c^{(2)} \approx \frac{1}{2\gamma} \left(\frac{R_o}{\sigma_I} \right)^2, \quad (4.42)$$

corresponding to the inner and outer radii, respectively. It can be seen that, a low-order mode at $m + n \lesssim m_c^{(1)}$ has a transmission efficiency close to zero, because hardly any fraction of the mode extends beyond the inner stop of the aperture. A bright band of modes exists beyond $m_c^{(1)}$ and before the cut-off at $m_c^{(2)}$, which are therefore the most important modes to include in simulations. Beyond $m_c^{(2)}$, again, the transmission efficiency becomes suppressed due to the finite radius of the aperture.

It can however be quite challenging to come up with a systematic strategy of selecting which HG modes are the best to be included in \mathcal{S} for a simulation, because the weights of coherent modes should be sorted in a decreasing order for maximal efficiency. For a generic GSM source, this can be achieved easily by sorting the modes in an increasing order of $m + n$. In contrast, for an apertured GSM source in the HG representation, the complicated structure of η_{mn}^2 breaks this order, as the weight is now also dependent on how well the geometries match between an HG mode and the specified aperture. Moreover, as the apertures are circular while the modes are Cartesian, these transmission efficiencies have to be calculated numerically according to Equation 4.26, which requires the numerical generation of the modes themselves. The mode generations, in general, can be computationally expensive. Although the computational implementation of simulations will be discussed in more detail in the next chapter, it can be appreciated that the numerical computation of Hermite polynomials is difficult at high orders, due to its normalisation [148],

$$\int_{-\infty}^{\infty} H_m(x) H_n(x) e^{-x^2} dx = \sqrt{\pi} 2^n n! \delta_{mn}, \quad (4.43)$$

scales with $2^n n!$. At orders where $m + n \gtrsim 100$, standard double-precision floating-point arithmetic is prone to overflow problems, and mode generation would require higher-precision computations. It would therefore be beneficial to find a representation whose transmission efficiencies η^2 have a simpler structure and can ideally be calculated analytically, such that the

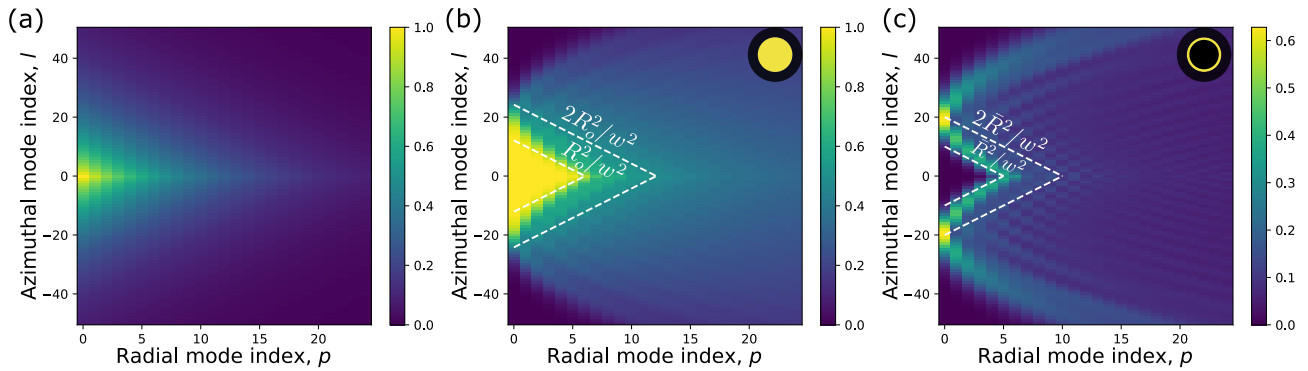


Figure 4.6: Eigenvalues of a GSM source in the Laguerre–Gaussian (LG) representation, and transmission efficiencies through circular and annular apertures. (a) Normalised eigenvalues, λ_{pl}/λ_0 , of a GSM source. (b) Transmission efficiencies η_{pl}^2 of LG modes through a circular aperture. White dashed lines denote the expected cut-offs at $2p + l = R_o^2/w^2$ and $2p + l = 2R_o^2/w^2$. (c) Transmission efficiencies η_{pl}^2 of LG modes through an annular aperture. White lines denote the approximate cut-offs at $2p + l = \bar{R}^2/w^2$ and $2p + l = 2\bar{R}^2/w^2$, corresponding to the mid-radius $\bar{R} := (R_i + R_o)/2$. All source and aperture parameters are consistent with Figure 4.5.

modes can be selected before they are generated numerically.

4.3.2 Representation in Laguerre–Gaussian (LG) modes

As an alternative to the HG modes, a logical basis to consider is the Laguerre–Gaussian (LG) modes [157, 158], which is defined in a polar geometry $\mathbf{x} = (r, \theta)$. Zhang and Lü [162] showed that, in the LG representation, the coherence structure of a GSM source can be written as:

$$W(\mathbf{x}_1, \mathbf{x}_2) = \sum_{p=0}^{\infty} \sum_{l=-\infty}^{\infty} \lambda_{pl} \phi_{pl}^*(\mathbf{x}_1) \phi_{pl}(\mathbf{x}_2), \quad (4.44)$$

$$\lambda_{pl} = \lambda_0 q^{2p+|l|}, \quad (4.45)$$

$$\phi_{pl}(r, \theta) = \sqrt{\frac{2p!}{\pi (p + |l|)!}} \frac{1}{w} \left(\frac{\sqrt{2}r}{w} \right)^{|l|} \exp\left(-\frac{r^2}{w^2}\right) L_p^{|l|}\left(\frac{2r^2}{w^2}\right) e^{-il\theta}, \quad (4.46)$$

where p and l are the radial and azimuthal mode indices respectively, and L_p^l denotes the associated Laguerre polynomial [148]. Note that expressions from Zhang and Lü [162] has been rewritten to be consistent with prior notations and generalised to take negative l into account.

Figure 4.6(a) shows the eigenvalues λ_{pl} for the same GSM source as in the previous section. Similar to the LG modes, it follows a geometric progression with the same λ_0 and ratio q . The

only difference is that the progression is now a function of $2p + |l|$ instead of $m + n$.

Now consider the GSM source placed immediately before a circular aperture. The transmission efficiency of LG modes can be calculated according to Equation 4.26, and for a circular aperture, it becomes

$$\eta_{pl}^2 = \int_{r=0}^{R_o} \int_{\theta=0}^{2\pi} |\phi_{pl}^2(r, \theta)| r dr d\theta, \quad (4.47)$$

$$= \frac{p!}{(p + |l|)!} \int_{x=0}^{2R_o^2/w^2} x^{|l|} e^{-x} [L_p^{|l|}(x)]^2 dx, \quad (4.48)$$

using a substitution of $x := 2R_o^2/w^2$. Thanks to the polar separability of LG modes, the expression reduces to a single integral.

Panel (b) shows the transmission efficiency η_{pl}^2 for the same circular aperture as in Figure 4.6(b). It is evident that, in a roughly triangular region of $p + |l| \lesssim R_o^2/w^2$, the coherent modes experience negligible energy loss. This can be seen as a consequence of the asymptotic formulae of Plancherel—Rotach type [163], which describes $e^{-x/2} L_p^{|l|}$ as oscillatory at $x < 2p + |l| + 1$ and exponentially decaying beyond, for associated Laguerre polynomials of high orders [161, Theorem 8.22.8, Page 200]. This approximate cut-off, however, gets worse as $|l|/p$ increases, as can be seen in the figure. In the extreme case of $p = 0$, the Laguerre polynomial $L_p^{|l|} = 1$, such that the integrand in Equation 4.48 becomes proportional to $x^{|l|} e^{-x}$, which maximises at $x = 2R_o^2/w^2 = |l|$. Therefore, for a circular aperture, cut-off in the transmission efficiency can be expected in the rough region of

$$\frac{R_o^2}{w^2} \lesssim 2p + |l| \lesssim \frac{2R_o^2}{w^2}, \quad (4.49)$$

as shown between the dashed white lines in Figure 4.6(b).

For an annular aperture, as shown in Figure 4.6(c), the low-order LG modes have a transmission efficiency of nearly zero, as their oscillatory regime is blocked by the inner stop. When the annulus is thin, only modes near the cut-off is efficiently transmitted, in the region of $\bar{R}^2/w^2 \lesssim 2p + |l| \lesssim 2\bar{R}^2/w^2$, where the mid-radius $\bar{R} := (R_i + R_o)/2$. The region is also shown between

the white-dashed lines in the figure. Compared to the HG case in Figure 4.5(c), the LG representation seems to contain some more efficient modes whose $\eta_{pl}^2 > 0.6$, thanks to its polar geometry.

When selecting the coherent modes to be simulated, it is usually reasonable to search within this region, as modes beyond have both low eigenvalues and low transmission efficiencies. Nevertheless, it is still preferable to find an efficient way of predicting η_{pl}^2 to determine which modes carry the most energy from the apertured GSM source, preferably without numerically generating the modes themselves. We shall therefore show in the following discussion that a closed-form expression for η_{pl}^2 can be obtained, which greatly facilitates mode selection, whereas an analogous closed-form result is not available for HG modes because the circular symmetry of the aperture conflicts with the Cartesian symmetry of the HG basis.

Recall that the Laguerre polynomial can be written as a series expansion [161, Page 101], given as

$$L_p^{|l|}(x) = \sum_{c=0}^p a_c x^c, \quad (4.50)$$

where

$$a_c = \frac{(-1)^c}{c!} \binom{p+|l|}{p-c}, \quad (4.51)$$

using the binomial coefficient. Its square can thus be written as

$$[L_p^{|l|}(x)]^2 = \sum_{c=0}^{2p} \sum_{d=0}^c a_d a_{c-d} x^c. \quad (4.52)$$

To include only non-zero terms, the range of second summation can be changed, such that

$$[L_p^{|l|}(x)]^2 = \sum_{c=0}^{2p} \sum_{d=\max(c-p,0)}^{\min(c,p)} a_d a_{c-d} x^c. \quad (4.53)$$

Substituting the expression into Equation 4.48, we have

$$\eta_{pl}^2 = \frac{p!}{(p + |l|)!} \int_{x=0}^{2R_o^2/w^2} x^{|l|} e^{-x} \left(\sum_{c=0}^{2p} \sum_{d=\max(c-p,0)}^{\min(c,p)} a_d a_{c-d} x^c \right) dx, \quad (4.54)$$

$$= \frac{p!}{(p + |l|)!} \sum_{c=0}^{2p} \sum_{d=\max(c-p,0)}^{\min(c,p)} a_d a_{c-d} \int_{x=0}^{2R_o^2/w^2} x^{|l|+c} e^{-x} dx. \quad (4.55)$$

The integral

$$\int_{x=0}^{2R_o^2/w^2} x^{|l|+c} e^{-x} dx, \quad (4.56)$$

is, by definition, the lower incomplete gamma function [148, §8.2]

$$\gamma \left(|l| + c + 1, \frac{2R_o^2}{w^2} \right). \quad (4.57)$$

Here, we employ the regularised lower incomplete gamma function, $P(x, y) := \gamma(x, y)/\Gamma(x)$, where Γ denotes the gamma function. A closed-form expression for η_{pl}^2 can then be found as

$$\eta_{pl}^2 = \sum_{c=0}^{2p} (-1)^c \sum_{d=\max(c-p,0)}^{\min(c,p)} A_{pl}^{cd} P \left(|l| + c + 1, \frac{2R_o^2}{w^2} \right), \quad (4.58)$$

where

$$A_{pl}^{cd} := \frac{p}{(p + |l|)!} \times \frac{1}{d!} \binom{p + |l|}{p - d} \times \frac{1}{(c - d)!} \binom{p + |l|}{p - c + d} \times (|l| + c)!, \quad (4.59)$$

$$= \frac{p! (p + |l|)! (|l| + c)!}{(p - d)! (p - c + d)! (|l| + d)! (|l| - d + c)! d! (c - d)!}, \quad (4.60)$$

$$= \binom{p}{d} \binom{p + |l|}{p - c + d} \binom{|l| + c}{c - d}. \quad (4.61)$$

For an open aperture as $R_o \rightarrow \infty$, the regularised lower incomplete gamma function becomes unity, and it can be shown $\eta_{pl}^2 \rightarrow 1$ as expected. For an annular aperture, the transmission efficiency takes the difference between two gamma functions, such that

$$\eta_{pl}^2 = \sum_{c=0}^{2p} (-1)^c \sum_{d=\max(c-p,0)}^{\min(c,p)} A_{pl}^{cd} \left[P \left(|l| + c + 1, \frac{2R_o^2}{w^2} \right) - P \left(|l| + c + 1, \frac{2R_i^2}{w^2} \right) \right]. \quad (4.62)$$

The closed-form expressions of η_{pl}^2 enables the transmission efficiency of any given mode to be calculated without the need to generate the mode itself. As the original eigenvalue λ_{pl} follows a simple geometric progression, the overall weight $\lambda'_{pl} = \eta_{pl}^2 \lambda_{pl}$ can therefore be computed and sorted relatively quickly for a large number of mode candidates.

4.3.3 Numerical modes

In previous sections, the coherent modes of apertured GSM sources were found by exploiting existing solutions of the coherence structure (λ_n, ϕ_n) of unapertured GSM sources. With the introduction of transmission efficiencies η_n^2 , the appropriate weights and apertured modes (λ'_n, ϕ'_n) can be found to represent the apertured sources, because aperturing is a linear process. Although this approach is totally valid, it should be noted that (λ'_n, ϕ'_n) are not themselves eigenvalues and eigenmodes of the Fredholm integral in Equation 4.10. In addition, although the normalised individual modes ϕ'_n remain fully coherent individually and mutually incoherent, they are not necessarily orthogonal, due to the change of domain from \mathcal{D} to \mathcal{A} .

As an alternative to the HG and LG modes, a viable approach would be to directly consider the eigenvalue problem in \mathcal{A} , i.e.

$$\int_{\mathcal{A}} W'(\mathbf{x}_1, \mathbf{x}_2) \phi'_n(\mathbf{x}_1) d\mathbf{x}_1 = \lambda'_n \phi'_n(\mathbf{x}_2), \quad (4.63)$$

which can be solved numerically [164], for example, with the help of NumPy [165]. As an indication of the computational scale, one representative run generated 5000 2D numerical modes on a 512×512 sampling grid, for an apertured GSM source with a global degree of coherence $\gamma = 5 \times 10^{-2}$. This run took approximately 493 s on a workstation with an Intel Core i9-10900K central processing unit (CPU) at 3.70 GHz and 32 GB of installed memory, using a CPU-only NumPy 1.26.0 implementation.

Figure 4.7 shows some example modes obtained using this approach, for a circular aperture. Modes are sorted in a decreasing order of eigenvalue, such that $\lambda'_{n+1} \leq \lambda'_n$. It can be noted that, especially in lower orders, some eigenmodes exhibit levels of polar separability or degeneracy,

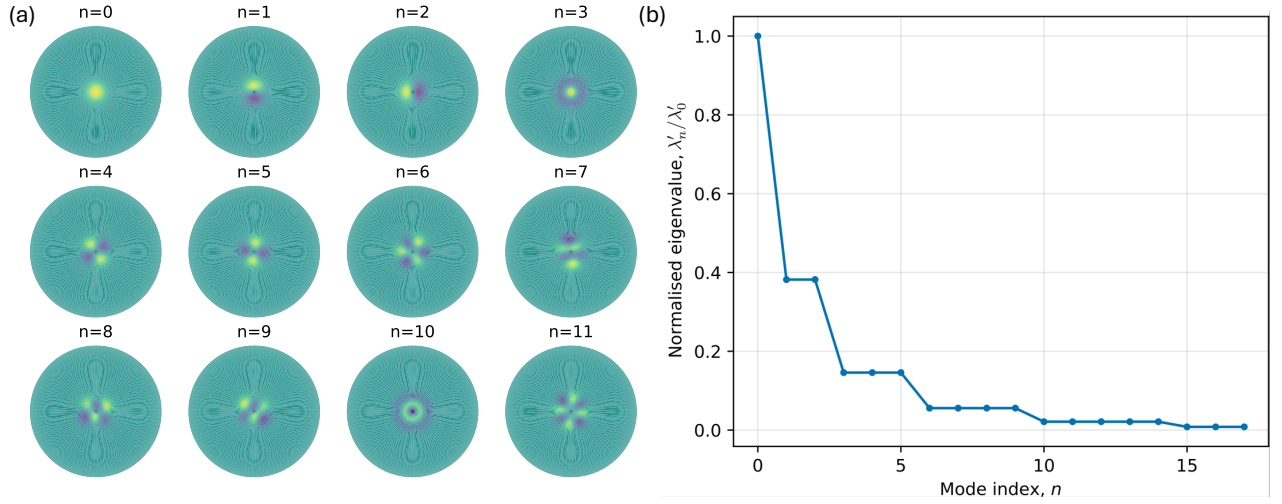


Figure 4.7: Eigenmodes and eigenvalues λ'_n of the apertured GSM source, obtained by solving the 2D Fredholm equation numerically. (a) Eigenmodes ϕ'_n of a circularly apertured GSM source. (b) The normalised eigenvalues λ'_n/λ'_0 of the 2D numerical modes.

although this polar symmetry is not generally guaranteed through the 2D numerical method.

Nevertheless, since the aperture \mathcal{A} is indeed circularly symmetric, we can attempt to solve the Fredholm equation by assuming the following polar form of modes, given as

$$\phi'_{pl}(r, \theta) = \psi_{pl}(r)e^{il\theta}, \quad (4.64)$$

where $\psi_{pl}(r)$ denotes the radial part, and the azimuthal dependence is in the form of a complex exponential, as in the LG modes. The Fredholm equation can then be rearranged as

$$\int_{R_i}^{R_o} \exp\left(-\frac{r_1^2 + r_2^2}{4\sigma_I^2} - \frac{r_1^2 + r_2^2}{2\sigma_g}\right) F_l(r_1, r_2)\psi_{pl}(r_1)r_1 dr_1 = \lambda'_{pl}\psi_{pl}(r_2)e^{il\theta_2}, \quad (4.65)$$

where

$$F_l(r_1, r_2) = \int_0^{2\pi} \exp\left[\frac{r_1 r_2 \cos(\theta_1 - \theta_2)}{\sigma_g^2}\right] \times e^{il\theta_1} d\theta_1, \quad (4.66)$$

$$= e^{il\theta_2} \int_0^{2\pi} \exp(il\Delta + \alpha \cos \Delta) d\Delta, \quad (4.67)$$

with

$$\Delta := \theta_1 - \theta_2, \quad \alpha := r_1 r_2 / \sigma_g^2. \quad (4.68)$$

Employing the Jacobi–Anger identity [148, §10.12], we can write

$$\exp(\alpha \cos \Delta) = \sum_{j=-\infty}^{\infty} i^j J_j(-i\alpha) e^{ij\Delta}. \quad (4.69)$$

Therefore,

$$F_l(r_1, r_2) = 2\pi e^{i\theta_2} \sum_{j=-\infty}^{\infty} i^j J_j(-i\alpha) \delta_{l,-j}, \quad (4.70)$$

$$= 2\pi e^{i\theta_2} i^{-l} J_{-l}(-i\alpha), \quad (4.71)$$

using the orthogonality of complex exponentials over $[0, 2\pi]$.

Using the Bessel identities $J_{-l}(z) = (-1)^l J_l(z)$ and making use of the modified Bessel function $I_l(\alpha) = i^{-l} J_l(i\alpha)$ [148], we can then find that

$$F_l(r_1, r_2) = 2\pi e^{i\theta_2} I_l(\alpha). \quad (4.72)$$

The Fredholm equation can then be written in a 1D form, given as

$$2\pi \int_{R_i}^{R_o} \exp\left(-\frac{r_1^2 + r_2^2}{4\sigma_I^2} - \frac{r_1^2 + r_2^2}{2\sigma_g}\right) I_l(\alpha) \psi_{pl}(r_1) r_1 dr_1 = \lambda'_{pl} \psi_{pl}(r_2). \quad (4.73)$$

Therefore, by solving Equation 4.73 for $\psi_{pl}(r)$ as a 1D eigenproblem for every value of l , one could obtain a valid set of coherent modes. In contrast to solving the original Fredholm integral in Equation 4.63, this approach features the advantage of reduced dimensionality (r_1, r_2) compared to the case of the full CSDF ($\mathbf{x}_1, \mathbf{x}_2$), benefiting both memory allocation and computational speed in a simulation. For the same apertured GSM source with $\gamma = 5 \times 10^{-2}$ and the 512×512 sampling grid, a representative run solving the 1D eigenproblems for 81 azimuthal indices generated 6601 numerical modes in approximately 0.84 s, on the same workstation and CPU-only NumPy 1.26.0 implementation.

Figure 4.8 shows some examples of coherent modes generated by this method, compared with LG modes of the same order. In the case of a circular aperture, features similar to LG modes




Figure 4.8: Coherent modes $\phi'_{pl}(r, \theta)$ of a circularly apertured GSM source, obtained by solving the 1D Fredholm equation in the radial direction, compared with LG modes of the same orders p, l .




Figure 4.9: Normalised eigenvalues $\lambda'_{pl}/\lambda'_{00}$ of 1D numerical modes for a GSM source, upon (a) a circular aperture and (b) an annular aperture. All parameters of the source and the apertures are consistent with Figure 4.5 and 4.6.

are recreated in low orders, due to the same form of polar separability between the two. In contrast, at higher orders, as the LG modes get closer to the cut-off described by Equation 4.49, they start to behave differently from the 1D numerical modes in the radial direction, which is particularly apparent for the mode (10, 0) in the figure.

Figure 4.9 shows the distribution of eigenvalues λ'_{pl} in the 1D numerical model, upon a circular and an annular aperture respectively, where the parameters of both the source and the apertures are consistent with previous models. As expected, for the circular aperture, the distribution is very similar to that of the LG modes as the mode themselves are similar. In contrast, in the annular case, the eigenvalue distribution in the 1D numerical model looks much simpler compared to the LG counterpart, as there is no complication introduced by the transmission

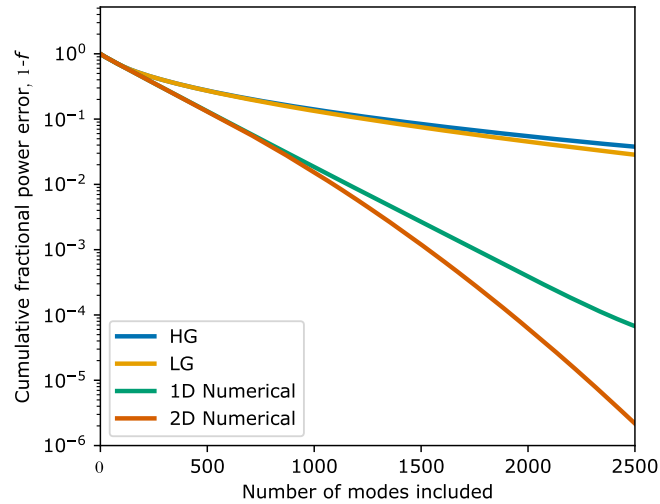


Figure 4.10: Comparison of cumulative fractional power error for an apertured GSM source represented by a finite set of coherent modes in various models, viz. Hermite–Gaussian (HG), Laguerre–Gaussian (LG), 1D numerical and 2D numerical modes.

efficiencies η_{pl}^2 . For the 1D numerical modes, the eigenvalues are significant only over a limited range of p , which can be explained by the thin width of the annulus (R_i of 9 mm and R_o of 11 mm).

4.3.4 Comparison of mode efficiency across various representations

To evaluate how efficiently a GSM source could be decomposed into a finite set of coherent modes in various models, the cumulative fractional power f is considered at varying number of modes being included in the set \mathcal{S} . By considering the error $1 - f$ for a varying number of modes, the four models discussed can be compared, as shown in Figure 4.10. As expected, in all cases, the error diminishes as the set of modes \mathcal{S} gets larger, as required by the completeness of all coherent-mode representations. The error in various models, however, decreases at different rates. Both analytical methods (i.e. HG and LG) exhibit compromised efficiencies, where coherent modes are originally eigenmodes of the Fredholm equation within a larger domain \mathcal{D} . As a result, the fractional power being carried by each individual mode is reduced by a factor of η_n^2 due to the mismatch between \mathcal{D} and the aperture \mathcal{A} . In particular, due to the mismatch between Cartesian geometry in HG modes and the polar geometry of the aperture, more HG modes are needed to achieve the same level of fractional power compared to the LG modes.

In contrast, both 1D and 2D numerical models are superior at using a finite number of modes to efficiently represent the GSM source, because the modes are generated within \mathcal{A} from the first place and no propagation through the aperture is required. As can be seen from the figure, the fractional error $1 - f$, decays roughly exponentially with respect to the number of modes included in \mathcal{S} . The 1D numerical model requires a larger number of modes because of the sacrifice of efficiency for mode separability, as Equation 4.64 is put as an extra constraint on the form of the coherent modes.

For these numerical approaches, the smaller number of modes has a significant benefit when using the models in computational simulations, as less resources are required to generate and propagate all modes in \mathcal{S} through the optical system. Nevertheless, the lack of analytical form limits the number of coherent modes that can be obtained, because a maximum of N eigenvectors exist for an $N \times N$ matrix. As a consequence, in cases where coherence is very low (i.e. $\gamma \ll 1$), a greater demand of the number of coherent modes requires a CSDF whose size scales quadratically and quartically for the 1D and 2D numerical models, respectively. For the 1D modes, the polar symmetry further requires an interpolation if the propagating field is defined in a Cartesian grid, which can reduce the accuracy.

4.4 Conclusion

In this chapter, the coherent-mode representation of a partially coherent field was introduced, which can be employed to model the secondary source in a transmission-light microscope, e.g. the condenser aperture plane for PPC. In contrast to the STF treatment which assumes a particular optical setup and sufficiently weak objects, this approach focuses on the second-order statistics of the source, and therefore holds for any subsequent optical system that behaves linearly with respect to the field amplitude. When implemented computationally, the coherent-mode method is superior to the ‘incoherent-point’ picture, producing no significant Fourier artefacts without the introduction of a windowing function. It is also more physical, because in a realistic optical system, the amount of spatial coherence has to be non-zero, as a consequence

of the van Cittert–Zernike theorem.

Besides being physically accurate, the coherent-mode representation also fits well the need for efficient computations. For a simulated grid of the size $N \times N$, it only requires an array of N^2 to be propagated each time, if the coherence structure is known in closed form, in contrast to a general double-impulse response (N^6) or a CSDF (N^4). Furthermore, due to the circular or annular aperture immediately after the secondary source, a cut-off is usually observed at a certain order of the coherent modes, beyond which the associated power becomes negligible and can be omitted from simulations. The BSM model was discussed to show an example of this cut-off.

In a microscopy system like the PPC, however, a Bessel-type correlation was deemed unreasonable. The GSM was then introduced, whose intensity and correlation are both described in Gaussian forms. With a circularly symmetric aperture placed on the GSM source, four modal representations were investigated in detail, namely the HG modes, the LG modes, the 2D numerical modes, and the 1D numerical modes. In the first two models, the HG and LG modes are masked from a larger domain \mathcal{D} to the aperture \mathcal{A} , and therefore lose a fraction of their power, quantified by $1 - \eta_n^2$. This loss of power not only requires a larger number of modes to adequately represent the source, but also introduces an extra structure depending on how well a particular mode may fit within the aperture. Using the Plancherel–Rotach asymptotics of Hermite and Laguerre polynomials at high orders, empirical guides can be found for where the cut-off in η_n^2 approximately sits. For the LG modes which have a polar geometry, we found that the transmission efficiency η_{pl}^2 can be expressed in closed form, allowing the modes with the largest weights to be selected without being generated computationally.

In contrast, coherent modes can be found directly within the aperture \mathcal{A} , by solving the eigenvalue problem in 2D, or assuming polar separability, in 1D. Compared to the HG and LG representations, the numerical models are found to be more efficient at representing the apertured GSM source with a minimal number of modes, as the modes are themselves eigenfunctions within the aperture. Nevertheless, the number of modes required is set by the global degree of coherence of the source, with lower γ requiring more modes for an accurate representation up

to a given fractional power f , regardless of the model in use. If the coherent modes are solved numerically, the number of modes is then limited by the sampling of the CSDF. Therefore, for very low-coherence sources (i.e. $\gamma \ll 1$), the required sampling can become so large that constructing and solving the CSDF becomes impractical under memory constraints.

In Chapter 5, we will consider how transmitted-light setups such as PPC can be simulated, building on the modal-decomposition framework developed in this chapter. Using a modal representation of a partially-coherent secondary source, we develop a simulation pipeline that is physically accurate, suitable for subsequent quantitative, informational analysis of voxel imaging.

Chapter 5

Simulation pipeline for polarisation imaging

Numerical simulations play an important role in modern optical imaging and are widely used alongside experiments because they offer several advantages. Widely used alongside experiments [50, 166–169], they provide a perfectly controlled and repeatable environment to characterise system performance over parameters that are otherwise difficult to control. Moreover, they can impose ideal imaging conditions beyond practical physical limitations, or be scaled up to generate large datasets for applications such as machine learning.

The versatility of simulations becomes especially relevant in the context of voxel imaging for optical data storage. For example, one could assess the read quality for voxels sitting in air instead of glass, to eliminate aberrations due to the glass thickness, which is of course impossible in practice. One can also systematically vary the voxel depth within the glass, which would be too time-consuming and impractical to do experimentally. With a simulation pipeline, images can be obtained for a range of situations that are impractical or impossible to realise experimentally. With the help of high-performance computing (HPC) [170], high-throughput simulations are possible for parameter sweeps as well, which is enabling and provides valuable insights into system performance.

For a simulation pipeline to be effective, high fidelity is crucial so that the results are consistent




Figure 5.1: Schematic of the simulation pipeline.

with physical reality. Achieving high-quality simulations is particularly challenging for systems under partially coherent illumination [171]. In the previous chapter, we introduced how partially coherent sources can be represented by a set of mutually uncorrelated modes, such that any linear optical system can be simulated by propagating the coherent modes individually. In contrast to the Stokes transfer function (STF) method in Chapter 3, the coherent-mode method makes no assumptions about the optical system, providing maximal flexibility in how the system is defined.

In this chapter, we present how the coherent-mode method can be used to build a versatile and efficient simulation pipeline for polarisation imaging. We will start by defining a physical model of the system, specifying the secondary source, the object space, the objective (including a depth-correction collar), the polarisation camera (PolCam), and how a coherent field is propagated between them. We will then describe the implementation aspects, including the software structure, parallelised coherent-mode generation, graphics processing unit (GPU) acceleration, and the use of job arrays on HPC clusters for parameter sweeps. Results of the simulations will be presented in Chapter 7 alongside experimental results.

5.1 Physical model

Figure 5.1 shows the schematic of the physical model for the simulations. The setup in general draws similarity to that introduced in Chapter 2, where coherent fields are propagated through various optical elements in the framework of paraxial Fourier optics, except the distinction

between S and D waves is no longer necessary in numerical simulations.

The model starts at the condenser aperture plane, which is regarded as a secondary source of the system. For improved versatility, the object plane is generalised into an object space where the back focal plane of the condenser lens ($z = 0$) is no longer assumed to necessarily coincide with the front focal plane of the objective lens sitting at $z = z_f$, enabling the effect of defocus to be simulated. In addition, there can be multiple layers of features (L_1, \dots, L_M) sitting at different z -positions, and the refractive indices between different layers may not be unity either, such that depth-dependent aberrations can be introduced. The objective pupil function P then needs to be generalised as well, to include a tunable depth correction that is commonly found in objectives. Finally, after the tube lens, the sensor captures the field intensities analysed in four directions (0° , 45° , 90° , and 135°). If the intensities are mosaicked according to the micro-polariser array (MPA) layout, PolCam images can then be obtained.

5.1.1 Secondary source at condenser aperture

As discussed in the previous chapter, the field at the front focal plane of the condenser (\mathbf{r}_0) can be described by a Gaussian Schell-model (GSM) source, masked by a circular or annular aperture $A(\mathbf{r}_0)$. In realistic situations, the intensity variation within the aperture is usually small, such that its width σ_I is much larger than the aperture radius R_o . As the choice of R_o has to reflect physical reality for a correct illumination numerical aperture (NA), σ_I needs to take a large value. Furthermore, due to the van Cittert–Zernike theorem, the more coherent the field is at the condenser aperture plane, the less uniform illumination becomes in the object space. In the extreme case of a completely coherent field, for example, illumination in the object space would become an Airy disk for a circular aperture or a Bessel beam for a thin annular aperture, in contrast to reality where illumination is approximately uniform across the field-of-view (FoV). As a result, the GSM source should be in the low-coherence regime, i.e. the degree of global coherence $\gamma \ll 1$.

Overall this description fits the definition of quasi-homogeneous sources by Carter and Wolf [156], and such sources intrinsically require a large number of coherent modes to be included in

the simulation, regardless of the coherent-mode representation. In this pipeline, we employ the Laguerre–Gaussian (LG) representation of apertured GSM sources, as discussed in the previous chapter, which comes with the advantage that the mode weights λ'_{pl} can be written in closed form. Therefore, the set of simulated modes \mathcal{S} can be selected without being generated, as computation of all potential modes can be expensive in the low-coherence regime. In addition, any LG mode ϕ'_{pl} can be written analytically, in contrast to the numerical models where the number of orders obtainable is dependent on the sampling within the aperture. Computation of LG and their eigenvalues also require no matrix construction of the cross-spectral density function (CSDF), enabling the use of large simulation grids given limited computer memory.

As an appropriate model is selected for the apertured source, modes can then be generated and propagated independently. At each iteration, the simulation starts with a coherent field

$$\mathbf{E}_0(\mathbf{r}_0) = \phi'_{pl}(\mathbf{r}_0)|L\rangle, \quad (5.1)$$

where ϕ'_{pl} can be calculated according to Equation 4.46 from the previous chapter.

5.1.2 Object space in angular spectrum

After the condenser lens, the field at its back focal plane $\mathbf{E}_1(x_1, y_1, z = 0)$ can be described with Fraunhofer diffraction in the paraxial approximation, given by Equation 3.20. In contrast to Chapter 3, the z -coordinate of the object space is now defined to allow multiple layers of features in the object space, including voxel arrays, material interfaces, and general thin scatterers. To model the wave propagation between these layers, we employ the angular-spectrum description, where a coherent field can be regarded as the sum of plane waves travelling in different directions. Parametrised by the 2D wave vector $(k_x, k_y)^T$, the angular spectrum can be written as a 2D Fourier transform of the field [80], given as

$$\mathbf{U}(k_x, k_y, z) = \iint_{-\infty}^{\infty} \mathbf{E}(x, y, z) \exp[-i(k_x x + k_y y)] dx dy. \quad (5.2)$$

As the Fraunhofer diffraction integral is also in the form of a 2D Fourier transform, we therefore have

$$\mathbf{U}(k_x, k_y, 0) \propto \mathbf{E}_0 \left(x_0 = -\frac{f_c k_x}{k_0}, y_0 = -\frac{f_c k_y}{k_0} \right), \quad (5.3)$$

where f_c is the focal length of the condenser lens, and $k_0 := 2\pi/\lambda$ is the wave number in vacuum.

In this description, propagation of light acts as a simple transfer function, given as

$$\mathbf{U}(k_x, k_y, z) = \exp(ik_z \Delta z) \mathbf{U}(k_x, k_y, 0), \quad (5.4)$$

where Δz denotes the distance of propagation, and k_z is given by

$$k_z := \sqrt{n^2 k_0^2 - k_x^2 - k_y^2}. \quad (5.5)$$

Note we keep the expression general for any refractive index n , which is critical for modelling wave propagation in dielectric materials like glass.

Refraction at interfaces can also be included in the framework. For a planar interface of mismatched refractive indices normal to the z -axis, field continuity at the boundary ensures the component of the wave vector parallel to the interface is unchanged, i.e. the 2D vector $(k_x, k_y)^T$ is conserved, which is consistent with the Snell's law. As a result, the angular spectrum of the field immediately before and after the interface is identical. Propagation through an interface should be dealt in parts, with a different k_z on each side. It is worth noting that this approach, however, does not take the Fresnel coefficients for transmission into account.

Scattering at objects, however, is more complicated. For an incident plane wave, scattering in general depends on the incident wave vector \mathbf{k}^i , the scattered wave vector \mathbf{k}^s , and the polarisation. The scattered field \mathbf{U}' and the incident field \mathbf{U} are usually related by a 2×2 scattering amplitude matrix [172], given as

$$\begin{bmatrix} U'_1(\mathbf{k}^s) \\ U'_2(\mathbf{k}^s) \end{bmatrix} = \begin{bmatrix} S_{11}(\mathbf{k}^s, \mathbf{k}^i) & S_{12}(\mathbf{k}^s, \mathbf{k}^i) \\ S_{21}(\mathbf{k}^s, \mathbf{k}^i) & S_{22}(\mathbf{k}^s, \mathbf{k}^i) \end{bmatrix} \begin{bmatrix} U_1(\mathbf{k}^i) \\ U_2(\mathbf{k}^i) \end{bmatrix}, \quad (5.6)$$

where the subscripts 1 and 2 denotes orthogonal axes perpendicular to their respective wave vectors. In the paraxial approximation they can be defined as the x and y -axes. Integrating over the incident wave vector, the output angular spectrum is then given as

$$U'_i(k_x, k_y, z) = \sum_j \iint S_{ij}(k_x, k_y, k'_x, k'_y) U_j(k'_x, k'_y, z) dk'_x dk'_y. \quad (5.7)$$

For thin objects, however, the expression can be further simplified, because its action through a Jones matrix \mathbb{T} is local, i.e.

$$E'_i(x, y, z) = \sum_j T_{ij}(x, y) E_j(x, y, z), \quad (5.8)$$

in physical coordinates. In the angular-spectrum description, the transmitted angular spectrum can be found by Fourier transforming the original Jones matrix, such that

$$\tilde{T}_{ij}(k_x, k_y) := \iint T_{ij}(x, y) \exp[-i(k_x x + k_y y)] dx dy. \quad (5.9)$$

The output angular spectrum can then be found using the convolution theorem, given as

$$U'_i(k_x, k_y, z) = \sum_j \tilde{T}_{ij}(k_x, k_y) * U_j(k_x, k_y, z). \quad (5.10)$$

In summary, the pipeline propagates a coherent field through the object space, starting with the angular spectrum at the condenser back focal plane, i.e.

$$\mathbf{U}^0(k_x, k_y, 0) = \mathbf{E}_0(x_0 = -f_c k_x/k_0, y_0 = -f_c k_y/k_0), \quad (5.11)$$

omitting global factors for simplicity. In the object space, there can be an arbitrary number of layers, $L_1(z_1), L_2(z_2), \dots, L_M(z_M)$. Each layer can be a dielectric interface, a thin feature described by the transmitted angular spectrum \tilde{T}_{ij} , or a general scatterer described by the scattering amplitude matrix S_{ij} . Denote the angular spectrum after the m -th layer as $\mathbf{U}^m(k_x, k_y, z)$ and the refractive index as n_m , for $m = 1$ up to M . For each layer, the pipeline first needs to

propagate the field from the previous layer, using

$$\mathbf{U}^{m-1}(k_x, k_y, z_m) = \mathbf{U}^{m-1}(k_x, k_y, z_{m-1}) \exp \left[i(z_m - z_{m-1}) \sqrt{n_{m-1}^2 k_0^2 - k_x^2 - k_y^2} \right]. \quad (5.12)$$

After that, the field is propagated through the layer depending on its type, i.e.

$$U_i^m(k_x, k_y, z_m) = \begin{cases} U_i^{m-1}(k_x, k_y, z_m), & \text{for interfaces,} \\ \sum_j \tilde{T}_{ij}^m(k_x, k_y) * U_j^{m-1}(k_x, k_y, z_m), & \text{for thin features,} \\ \sum_j \iint S_{ij}^m(k_x, k_y, k'_x, k'_y) U_j^{m-1}(k'_x, k'_y, z_m) dk'_x dk'_y, & \text{for general scatterers.} \end{cases} \quad (5.13)$$

After all layers, the field needs to be propagated to the objective front focal plane at $z = z_f$.

For an air objective, its immersion refractive index is unity (i.e. $n_M = 1$), such that

$$\mathbf{U}^M(k_x, k_y, z_f) = \mathbf{U}^M(k_x, k_y, z_M) \exp \left[i(z_f - z_M) \sqrt{k_0^2 - k_x^2 - k_y^2} \right]. \quad (5.14)$$

Note that $z_f - z_M$ is non-positive for a system properly focused on the layers, so this step is usually a virtual propagation in the $-z$ direction, and is independent of the actual refractive index between $z = z_f$ and $z = z_M$.

Finally, the field can be propagated to the objective back focal plane \mathbf{r}_2 . Similar to the condenser, the field is related to the angular spectrum in the object space by a simple scaling in the paraxial approximation, given as

$$\mathbf{E}_2(x_2, y_2) = \mathbf{U}^M(k_x = k_0 x_2 / f_o, k_y = k_0 y_2 / f_o, z = z_f), \quad (5.15)$$

where f_o is the focal length of the objective.

5.1.3 Voxel arrays

Among all possible objects, one type of particular interest is the voxel array, which can be used to represent a data sector in the context of optical data storage (ODS). Although in real sectors

the voxels may not be arranged in perfect grids depending on the scanning mechanism in the writing process [25], a Cartesian grid of voxels is representative in most cases. In such models, there are $N_v^x \times N_v^y$ voxels in a sector, and a voxel indexed (m, n) sits at the position of

$$x_{mn} = \left(m - \frac{N_v^x - 1}{2}\right) p_v^x, \quad y_{mn} = \left(n - \frac{N_v^y - 1}{2}\right) p_v^y, \quad (5.16)$$

where p_v^x and p_v^y represent the voxel pitch in the x and y-directions, respectively. Each voxel is assumed to be a birefringent disk of radius r_v , so its extension lies within

$$\chi_{mn}(x_1, y_1) = \begin{cases} 1, & \sqrt{(x_1 - x_{mn})^2 + (y_1 - y_{mn})^2} \leq r_v, \\ 0, & \text{otherwise.} \end{cases} \quad (5.17)$$

The polarisation action is then described by the Jones matrix of a linear retarder:

$$\mathbb{T}_{mn} = \text{LR} \left(\delta = \delta_v, \gamma = \frac{\pi g_{mn}}{N_g} \right), \quad (5.18)$$

where $g_{mn} \in \{0, 1, \dots, N_g - 1\}$ is the symbol the voxel represents, and δ_v denotes the birefringent retardance of the voxel. A data sector therefore carries $N_v^x N_v^y \log_2 N_g$ bits of data.

Therefore, the overall Jones matrix of the voxel array becomes

$$\mathbb{T}(x_1, y_1) = \mathbb{I} + \sum_{m,n} (\mathbb{T}_{mn} - \mathbb{I}) \chi_{mn}(x_1, y_1), \quad (5.19)$$

and the corresponding transmitted angular spectrum can be found through Equation 5.9.

5.1.4 Objective and collar correction

For infinity-corrected objectives, the objective pupil is modelled at the back focal plane of the lens. Similar to Equation 3.26, the output field \mathbf{E}'_2 can be simply written as

$$\mathbf{E}'_2(x_2, y_2) = P(x_2, y_2) \mathbf{E}_2(x_2, y_2), \quad (5.20)$$

where P is a polarisation-insensitive pupil function.

The parameters of the objective should be set to correctly reflect physical realities. A microscope objective is usually marked with magnification \mathcal{M} and NA, designed to be used in conjunction with tube lenses of a particular tube lens f_u . Thus, the objective focal length is given by $f_o = f_u/\mathcal{M}$, and the radius of its pupil can then be found as $R = \text{NA} \times f_o$.

Nevertheless, in practice, the microscope is usually focused some distance into the glass. This can be due to the sector depth for optical data storage, or caused by coverslips in biological applications. As a result, aberration is introduced by the refractive-index mismatch between glass and air. Accordingly, many objectives incorporate a correction collar that compensates these aberrations, and it is vital to include it in the simulations.

Consider a slab of thickness d and refractive index n is inserted in the object space. The wavefront aberration $W(d, n, \rho)$ it introduces can then be written as [173]

$$W(d, n, \rho) = d\sqrt{n^2 - \rho^2\text{NA}^2}, \quad (5.21)$$

by considering the extra optical path assuming the Abbe sine condition [174].

Assuming the correction collar to perfectly remove the aberration caused by the refractive index mismatch, it should introduce a negative wavefront of

$$\psi = -[W(d_c, n_c, \rho) - W(d_c, 1, \rho)], \quad (5.22)$$

where n_c is the refractive index for which the correction is designed for, and d_c is the correction collar setting. The pupil function therefore becomes

$$P(\rho) = \begin{cases} \exp(ik_0\psi), & \rho \leq 1, \\ 0, & \rho > 1. \end{cases} \quad (5.23)$$

Additionally, for phase-contrast objectives, an extra factor of it_p should be multiplied within the phase ring, as discussed previously in Chapter 3.

5.1.5 Tube lens and camera

The tube lens, parametrised by its focal length f_{tl} , focuses the field onto the sensor. In the absence of camera defocus, the process can be described by the Fraunhofer diffraction integral, given as

$$\mathbf{E}_3(x_3, y_3) = -\frac{ie^{2ik_0f_{tl}}}{\lambda f_{tl}} \iint_{-\infty}^{\infty} \mathbf{E}'_2(x_2, y_2) \exp\left[-\frac{2i\pi}{\lambda f_{tl}}(x_2x_3 + y_2y_3)\right] dx_2 dy_2. \quad (5.24)$$

In simulations, this can be implemented with discrete Fourier transform (DFT) algorithms.

As the sensor is a PolCam that analyses the field in four linear states, four intensity images of the systems are recorded, i.e.

$$I_q(x_3, y_3) = \sum_{pl} \lambda'_{pl} |\mathbb{P}(q)\mathbf{E}_3(x_3, y_3)|^2, \quad (5.25)$$

where q represents one of the four polarisation channels at 0° , 45° , 90° , and 135° , and \mathbb{P} is the Jones matrix of linear polarisers in Table 2.1. Contributions from all coherent modes (p, l) are summed over, following the coherent-mode method in Equation 4.13.

In taking this approach, the field at every position is analysed in all four linear states. This offers perfect registration between polarisation channels, which, though not strictly realistic, is advantageous in situations where the instantaneous field-of-view (IFoV) error is preferred to be avoided. If desired, realistic PolCam images can be obtained by mosaicking the four intensity images according to the superpixel layout of PolCam.

5.2 Implementation of the simulation pipeline

The modular structure of the simulation pipeline is summarised in Figure 5.2. The simulator accepts a configuration file at a time, and produces the four-channel polarisation image corresponding to the system specified, utilising the physical model described in the previous section. It is optimised for deployment on the HPC clusters provided by Imperial College Research

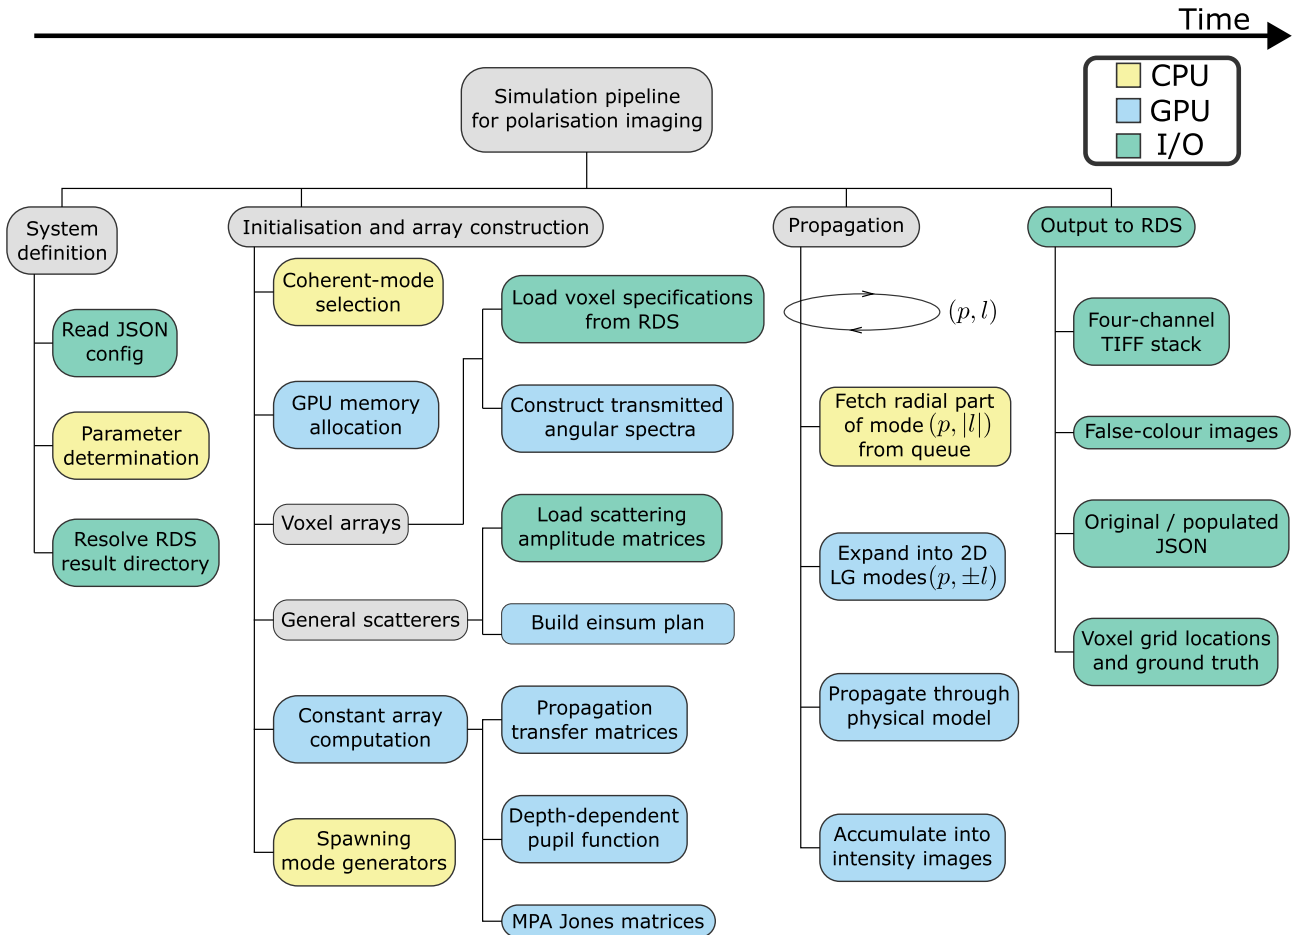


Figure 5.2: Modular diagram of the polarisation image simulation pipeline, which consists of system definition, initialisation and array construction, propagation, and output to the Research Data Store (RDS). Each module can be further decomposed into lower-level submodules. Colours indicate the primary resource used by each module (CPU, GPU, or I/O), featuring the CPU-based mode generation and the GPU-accelerated mode propagation.

Computing Service [170], with external input/output (I/O) handled with the Research Data Store (RDS), although it can also run on any local machine as long as Compute Unified Device Architecture (CUDA) Toolkit [175] is available for GPU parallelisation.

5.2.1 System definition

A simulation starts by reading a single configuration file defining the system to be simulated. Written in the JSON format, the file specifies the states of all system components in the physical model. A determination and validation step then checks that the configuration is physically meaningful, and computes any additional parameters required by the simulator. This step produces a post-determined parameter object that is then passed around the parent and sub-

processes of the pipeline.

At this stage, the pipeline also tries to identify the machine it is running on, by inspecting the local environment such as the hostname and the hardware address, which helps to determine the I/O paths the pipeline works with. This is essential for placing the results in the correct directory and, when required, locating the cache files and voxel specifications.

5.2.2 Initialisation and array construction

The pipeline then needs to prepare any prerequisites for the coherent mode propagation. First, it needs to identify which modes (p, l) needs to be propagated. As discussed in the previous chapter, the mode selection process is to find a set of simulated modes \mathcal{S} such that the cumulative fractional power f satisfies a certain threshold. Ideally, the modes should be of the greatest weights λ'_{pl} such that \mathcal{S} is as small as possible. The selection process is carried out by searching the modes near the asymptotic cut-offs in Equation 4.49. With the help of the closed-form expression of the mode transmission efficiency (Equation 4.62), the mode-selection process can be carried out without actually generating any LG mode, which greatly accelerates the computation. Mode-selection results are saved as a Python pickle file (`.pkl`) for potential future reuses.

The pipeline then prepares for GPU computations for mode propagations. GPU acceleration is realised with the aid of `cupy` [176], an open-source array library based on the CUDA Toolkit. To prevent repeated initialisation of GPU arrays, they are preallocated with the expected size and data type on the GPU memory, which can then be provided to the mode propagator and reused over iterations.

The action of all layers in the object space must also be specified. Layers of thin features are described by their transmitted angular spectra, which can be obtained from Equation 5.9 when a Jones matrix is given. Alternatively, when the object space contains one or more voxel arrays, the pipeline reads the relevant sector information from a library of voxel specifications on the RDS, referred to as the glass library. The glass library stores the specifications according

to which all glass samples are written. Each sample is defined by a JSON file that contains the writing parameters, the sector definitions, and a reference to the corresponding ground-truth data. For samples that were imaged experimentally, the JSON file also specifies the corresponding sector configuration, enabling straightforward comparison with the experiments. By specifying the glass sample name and sector name in the simulation configuration, the pipeline can simulate any sector in the glass library together with the correct ground truth it is designed to represent.

In cases when general scatterers are to be simulated, the pipeline instead requires a scattering amplitude matrix to be provided as a `.pk1` file. Taking into account the simulation grid, the scattering amplitude matrix is a rank-6 tensor acting on a rank-3 field tensor, contracting over three dimensions to produce a rank-3 output field tensor, which can be computationally expensive. To optimise the general matrix multiplication (GEMM) process, an open-source package `opt_einsum` [177] is employed to determine an optimal GEMM path prior to the propagation phase, with the constant arrays cached in GPU memory. It is also found that the CUDA-based `cuBLAS` library [178] greatly speeds up the GEMM computation, when used as the backend for `opt_einsum`.

In addition, the pipeline constructs several other constant arrays on the GPU during the initialisation process, including the interlayer propagation phases $\exp(ik_z\Delta z)$, the depth-corrected pupil functions $P(\mathbf{r}_2, d_c)$, and the MPA Jones matrices.

5.2.3 Simultaneous mode generation and propagation

In the propagation phase, all coherent modes (p, l) need to be efficiently generated and propagated through the physical model, for which GPU-acceleration is greatly helpful. Nevertheless, the GPU operations employ a fixed precision floating-point arithmetic, which is problematic for generating high-order LG modes.

Recall the radial part, $\psi_{pl}(r)$, of the LG mode, given as

$$\psi_{pl}(r) = \sqrt{\frac{2p!}{\pi (p + |l|)!}} \frac{1}{w} \left(\frac{\sqrt{2}r}{w} \right)^{|l|} \exp\left(-\frac{r^2}{w^2}\right) L_p^{|l|}\left(\frac{2r^2}{w^2}\right). \quad (5.26)$$

Note that, at very high orders of p and $|l|$, the associated Laguerre polynomial is prone to numerical overflow, while the other terms are prone to underflow. As a result, although the value of $\psi_{pl}(r)$ itself may be well-represented by a floating-point number, the computation of it requires arithmetic of higher precisions.

To overcome the problem, the `mpmath` package [179] is employed, to enable floating-point arithmetic with arbitrarily high precision using the central processing unit (CPU). To compensate the lack of parallelisation functionality offered by `mpmath`, the simulation pipeline is designed with an asynchronous producer-consumer structure, for which multiple subprocesses are spawned as mode generators at the end of initialisation phase. At the propagation phase, the mode generators repeatedly take the next available mode index and computes the radial mode $\psi_{pl}(r)$ at high-precision on CPU, while the parent process acts as the propagator on GPU, that expands the radial modes into full 2D modes, propagates through the physical model, and accumulates into the result intensity image. An inter-process queue is used to connect between the mode generators and the propagator, enabling asynchronous computation.

The mode generators also take advantage of symmetries to avoid redundant computations. Notably, as $\psi_{p,l}(r) = \psi_{p,-l}(r)$, each radial part actually corresponds to a mode pair $(p, \pm l)$ for the propagator to process. In addition, as the simulation grid takes a Cartesian geometry, many points on the grid share the same radius: for instance, grid points $(\pm i, \pm j)$ and $(\pm j, \pm i)$ all share the same radial distance and only need to be considered once. At the initialisation phase, the pipeline considers a set of unique radial distances according to the definition of the simulation grid to take advantage of this symmetry, and constructs a routine such that the propagator can expand the 1D radial distances back to the 2D grid.

It is also worth noting that, for modes of higher orders, the mode generators get slower due to the higher precision arithmetic they require, while the time it takes for the mode propagator

to process a mode pair stays the same. Therefore, the mode indices (p, l) are shuffled during initialisation, to help the load balancing between the mode generators and the propagator. Ideally, the number of generator subprocesses (and therefore the number of CPU cores requested from the cluster) should be set such that the production rate of modes matches the consumption rate of the propagator, maximising resource utilisation.

5.2.4 Output

After all coherent modes are propagated and the four-channel polarisation image is obtained, the final results are transferred from the GPU and written to the result directory on the RDS as a TIFF stack. In addition, false-colour images are generated as a quick visual preview and saved as separate TIFF files.

The pipeline records all configuration for reproducibility. Both the original and post-determined system parameters are saved as JSON files in the result directory, together with log files, GPU status, and the environment variables. The voxel grid locations and associated ground truths for voxel layers are also saved, which can be useful for subsequent informational analysis of the system.

5.2.5 HPC deployment and parallelisation with job arrays

Although the pipeline can run on any machine where the CUDA context and the I/O paths to the glass library and the result directories are defined, it is primarily optimised for HPC deployment and is typically executed on the CX3 or HX1 clusters at Imperial College London [170]. With the code base cloned to a RDS location, the simulation can be launched via the Portable Batch System (PBS).

To start a simulation session, a PBS script is submitted to the job scheduler, specifying the required computational resources, a list of module dependencies (e.g. CUDA), as well as the path to the JSON configuration file. Once resources become available, the scheduler assigns

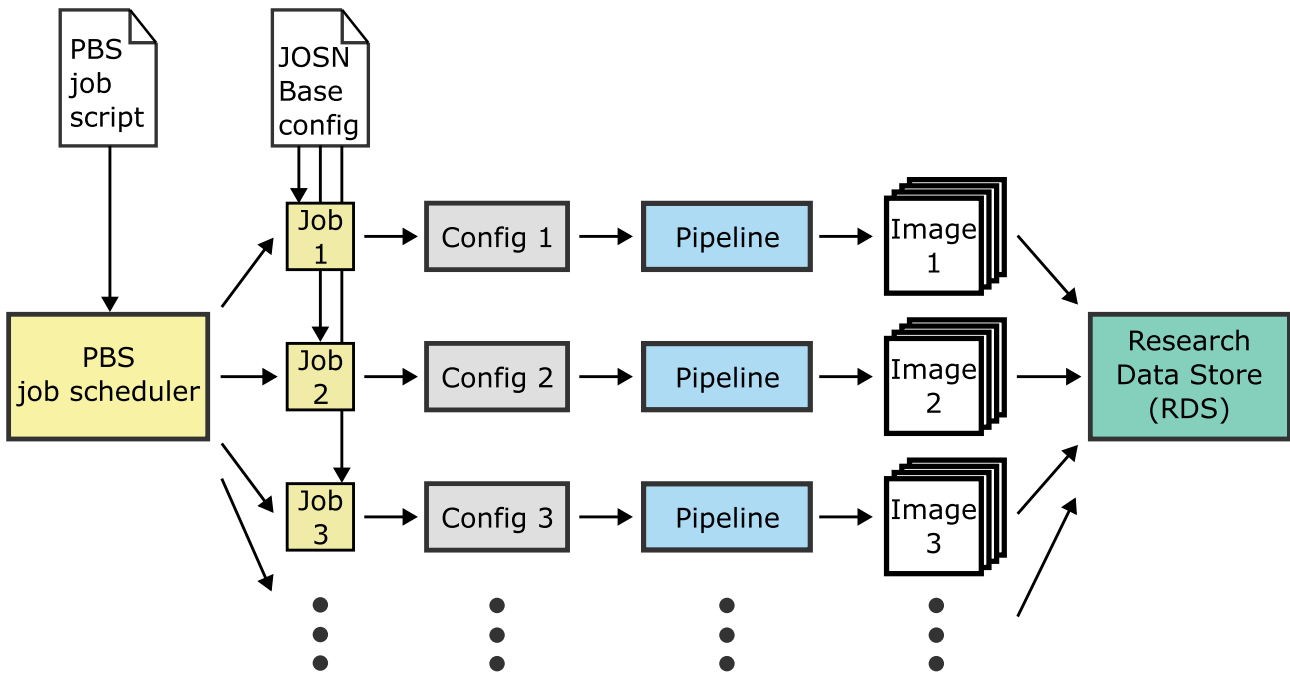


Figure 5.3: Schematic of the parallelisation of the simulation pipeline using job arrays on an HPC cluster. A single Portable Batch System (PBS) job script is submitted to the scheduler, which spawns multiple jobs that each read the same base JSON configuration and derive their own set of parameters. Each job runs an independent instance of the simulation pipeline, and the resulting images are written to structured result directories on the RDS.

nodes to execute the computation, and the resulting data are written back to the RDS. This method of software deployment provides not only a consistent software environment, but also access to high-end hardware. For instance, the NVIDIA A100 GPU with 80 GB of memory is particularly advantageous for simulations involving a large number of layers in the object space.

Another major advantage of HPC deployment is the ease of parallelisation. In cases like parameter sweeps where many simulation sessions are needed, multiple instances of the simulation pipeline can be launched simultaneously using PBS job arrays, which is illustrated in Figure 5.3. When a PBS script is submitted to the scheduler as a job array, multiple jobs can run at the same time on the cluster, subject to hardware availabilities. Each job reads the same JSON file as a base configuration, and applies necessary modifications according to its array index. The modified configuration is then passed to a separate instance of the simulation pipeline, with the resulting images written to the RDS. In this way, the high level of automation makes efficient use of a rich set of computational resources provided by the HPC clusters, enabling exploration of large parameter spaces that is otherwise impractical.

As an indication of the computational scale, one representative single-layer simulation of birefringent disks used an apertured GSM source with global degree of coherence $\gamma = 2 \times 10^{-3}$, involving 233 019 mode pairs on a 2048×2048 spatial grid with two polarisation components. This run was performed on the Imperial College London CX3 HPC cluster [170] using 32 CPU processes on Intel Xeon Silver 4110 CPUs and an NVIDIA RTX 6000 GPU. The job completed in approximately 183 min, using 6 GB of memory at peak, with an average CPU utilisation of 31.99 cores, indicating that the simulation was CPU-bound. When simulating optical systems with a more complex object space, jobs can become GPU-bound instead, requiring less CPUs. In general, job sizing is required to determine the appropriate amount of resources to request from the PBS for any given set of parameters.

5.3 Conclusion

In this chapter, we described both the physical model and the software implementation of a simulation pipeline for transmitted-light polarisation imaging. We first presented a physical model, with a flexible framework to define the microscopy system, consisting of a partially coherent source, an extended object space, a depth-correcting objective, the tube lens and a PolCam. The LG representation of an apertured GSM source from the previous chapter was employed to model the secondary source at the condenser front focal plane. The object space was treated in the angular spectrum, containing an arbitrary number of layers at different z-positions. Each layer can be a thin object represented by their transmitted angular spectra, a general scatterer defined by a scattering amplitude matrix, or a dielectric material interface for refractive-index changes, while interlayer propagation acts as a simple transfer function in the angular spectrum. In particular, voxel arrays were modelled as birefringent disks sitting in a Cartesian grid, with their azimuths corresponding to the ground truth for ODS applications. The objective pupil function was then generalised to incorporate depth-correction for various thicknesses of glass, and the focusing by the tube lens was described by the Fraunhofer diffraction.

We then described how the physical model is used to implement an efficient simulation pipeline that accepts a single configuration file and produces the polarisation images for the system defined. After determining a consistent set of parameters, the pipeline finds a set of coherent modes to be included in the simulation, according to the asymptotic analysis of LG modes and the closed-form expressions for the transmitted efficiencies from the previous chapter. At the propagation phase, the pipeline employs an asynchronous producer-consumer structure, with multiple subprocess spawned for high-precision computation of radial modes in CPU. Simultaneously, the parent process fetches the radial modes from an interprocess queue, expands it into 2D LG mode pairs, and propagates them through the physical model. The propagator was further accelerated with CUDA-based GPU operations, including faster GEMM operations via `cupy` and `opt_einsum`. We also described how the pipeline was deployed on HPC clusters, and can be efficiently parallelised with PBS job arrays.

In the next chapter, we move from the numerical simulation to the practical implementation. We will describe the hardware realisation of the polarisation-sensitive phase contrast (PPC) system, including the optical setup, hardware controls, measurement automations, and real-time visualisation with the polarisation-camera image processing (PCIP) module. Together, the development of a physical system will provide the experimental counterpart to the simulations, such that their results in voxel imaging can be compared using informational analysis in Chapter 7.

Chapter 6

PPC System Implementation

In previous chapters, we discussed how polarisation-sensitive phase contrast (PPC) is designed for imaging weak polarisation features in a sample, and developed both analytical and numerical models of the imaging process. In this chapter, we instead focus on how the system is implemented in practice. We describe the optical and electro-mechanical hardware configuration, including sample positioning, illumination control and system stability, and then introduce the software components used for control, automation and real-time visualisation of polarisation images.

The motivation for this chapter is that imaging weak polarisation signals in practice requires stability, repeatability, and usability of the system. In particular, both biological samples and birefringent voxels can produce weak contrast that is difficult to locate and focus on from single raw frames, and optical data storage (ODS) experiments would also benefit from long acquisition sessions and large datasets. In addition, the demonstration of high-density ODS requires the capability of positioning the sample in 3D and navigating within the glass. A robust implementation is therefore essential, both to ensure consistent data quality and to support the quantitative analysis in later chapters.

The PPC system was initially implemented during my MRes project [41], which introduced the initial optical layout and basic acquisition process. Since then, the system has been extended and refined in several aspects, including the illumination scheme, sample positioning, hardware

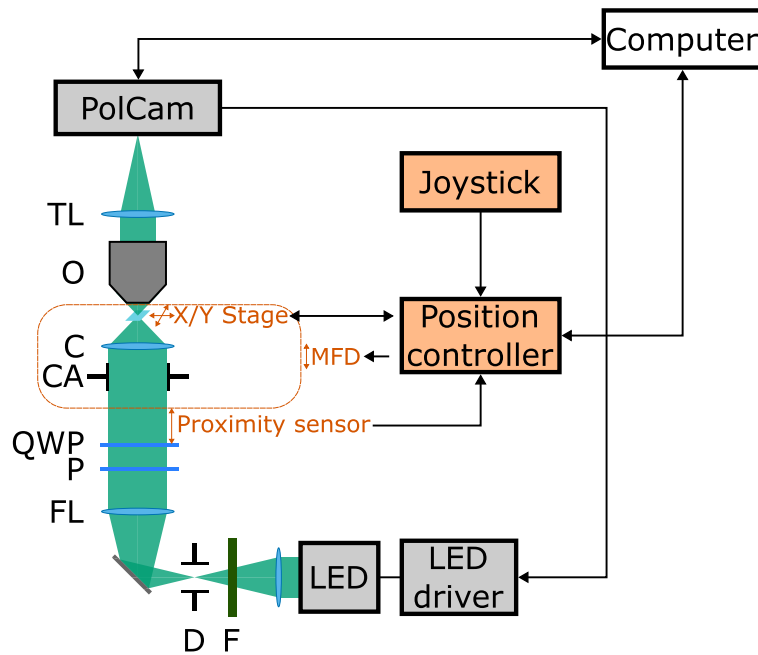


Figure 6.1: Implementation of the PPC system, including the optical setup and its integration with the electro-mechanical components for controls and measurements. F, spectral filter; D, field diaphragm; FL, field lens; P, linear polariser; QWP, quarter-wave plate; CA, condenser aperture; C, condenser lens; O, objective lens; TL, tube lens; MFD, motorised focus drive.

integration and software automation. In this chapter, we therefore only summarise the elements that are essential for understanding related experiments, and focus on the upgrades and new control capabilities that were developed beyond the original MRes work.

6.1 Hardware assembly

Figure 6.1 summarises the optical and electro-mechanical setup for implementing the PPC system. An external light-emitting diode (LED) (Thorlabs, SOLIS-565C) is used as the light source, replacing the original halogen lamp at the transmitted illumination port of the Olympus BX41 frame. Combined with a spectral filter (Comar, 547 IB 50), the source spectrum has a bandwidth of 50 nm and peaks at 530 nm. The LED collimator and the field lens project an image of the LED chip at the condenser aperture plane, achieving a Köhler-illumination setup. Between the field lens and the condenser aperture, a linear polariser (Thorlabs, LPVISE200-A) and a quarter-wave plate (QWP) (Thorlabs, WPQ10E-532) are placed with their orientation offset by 45° to produce left-handed circular polarisation. A phase-contrast condenser (Olym-

pus, U-PCD-2) is employed to match the objective (Olympus, LUCPlanFLN 40x/0.60 Ph2). Finally, light is focused by the 180 mm tube lens offered by the microscope frame onto the polarisation camera (PolCam) (FLIR, Blackfly S BFS-U3-51S5P-C) connected to a computer. Note the trinocular component is removed from the Olympus frame to provide a mirrorless light path between the polariser and the PolCam, to prevent the polarisation being changed upon reflections.

The upgrade of the light source from the halogen lamp to the LED provides a much higher optical power up to 6.1 W, massively reducing the exposure time required to make full use of the full-well capacity (FWC) offered by the PolCam. It is particularly helpful in a PPC setup where the light throughput is inherently poor due to both the condenser annulus and the phase ring attenuation. With a high optical power, however, thermal management is required to prevent damages to the optical components, especially the polariser which dissipates roughly half of the incident power. To limit the heat load, the LED is placed in pulsed operation, triggered by the PolCam exposure signal. Therefore, the LED is only switched on during the exposure window of the PolCam, minimising thermal effects to the system in long acquisition sessions.

Another significant upgrade from the previous iteration of the system is the motorised positioning of the sample, which enables controlled and repeatable sample placement in 3D. This is particularly useful for imaging glass samples for ODS. As a glass sample can contain millions of unique sectors, accurate positioning is essential to place a sector of interest within the camera field-of-view (FoV) and to achieve proper focus. In the transverse directions, this is achieved by a scanning stage (Märzhäuser Wetzlar, SCAN 75 x 50) with a custom insert that can hold either a 75 mm × 25 mm microscope slide or a 25 mm square glass sample at the centre. Positioning in the axial direction is controlled by a motorised focus drive (MFD) (Märzhäuser Wetzlar, MFD for Olympus) attached to the focusing knob. Compared with a piezoelectric stage, the MFD provides a long axial travel range limited only by the microscope frame. Both transverse and axial motion are commanded via the position controller (Märzhäuser Wetzlar, TANGO 3 DT) using either a three-axis joystick or the computer. A proximity sensor (Omron, EE-SX4009-P1) provides a repeatable axial reference when the condenser assembly approaches the QWP, while the limit switches of the scanning stage define the transverse reference position, ensuring

consistent sample placement in 3D. Together, these features allow reliable navigation to any glass location containing a data sector and support acquisition of z -stacks, relevant to both ODS and biological applications.

Mechanical stability is also critical, because high-quality images may be obtained by averaging repeated exposures, during which the sample must remain stationary. During the MRes project [41], the sample stage was observed to slowly drift downwards, and significant millisecond-scale vibrations were present. These issues are effectively mitigated by mounting the system on a floated optical bench and using the MFD, which suppresses axial drift. Together, the enhanced mechanical stability and accurate motorised sample positioning enable long, automated acquisition sessions, which are essential for both statistical characterisation of imaging quality and high-throughput data reading.

6.2 Software for control, visualisation, and automation

The computer plays a central role in hardware control, real-time visualisation, and automated acquisition of images, all of which were realised with Python based on the open-source ScopeFoundry platform [180].

Camera control and image acquisition was achieved through the `simple_pyspin` package [94] that interfaces with FLIR's `PySpin` package. For stage control, we wrote `TangoPy`, a Python wrapper over the vendor DLL that issues commands and queries status via an emulated COM port. In both cases, the hardware control is integrated to the ScopeFoundry app as hardware classes.

During an imaging session, raw images are continuously acquired from the PolCam. They are written to a circular buffer that maintains the most recent frames and computes a running average. The averaged image is then passed to polarisation-camera image processing (PCIP) to generate false-colour visualisations. Figure 6.2 shows a typical workflow of this process, demonstrated with an example image of closely packed voxels whose birefringent fast axes orientates at 0° , 45° , 90° , and 135° , arranged in a chequerboard pattern. From the raw image,

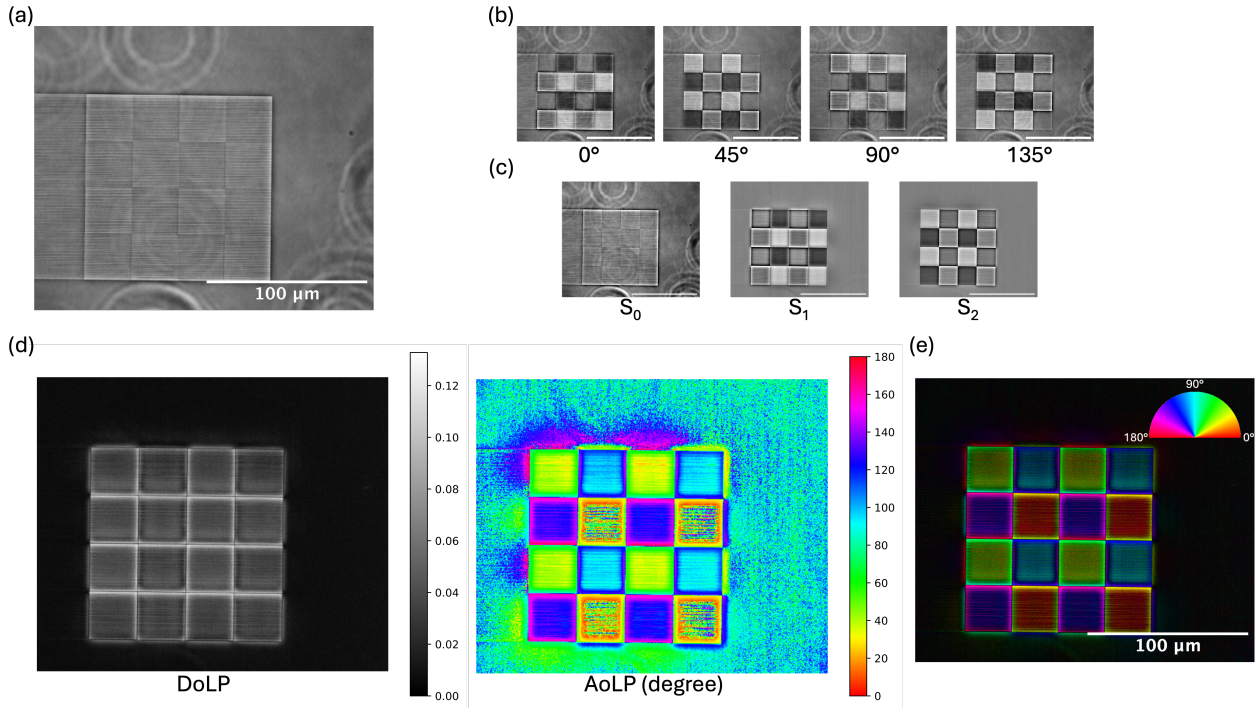


Figure 6.2: Typical workflow for real-time processing of PolCam images. Example shows a sector of closely packed voxels with fast-axes orientations 0° , 45° , 90° , and 135° , arranged in a chequerboard pattern. The raw image (a) is split into the four polarisation channels (b); from these the three Stokes images are computed (c), yielding the DoLP and AoLP maps (d). A false-colour image (e), described in Section 2.4.3, is then displayed for visualisation. Hue, AoLP; value, DoLP. Scale bars for all images in the figure represents 100 μm .

the PCIP module splits it into four polarisation channels, before computing the Stokes images (S_0 , S_1 , S_2) using Equations 2.1–2.3. It then computes the degree of linear polarisation (DoLP) and angle of linear polarisation (AoLP) maps with Equations 2.4–2.5, before displaying a false-colour visualisation using the hue-saturation-value (HSV) scheme described in Section 2.4.3.

Optionally, the PCIP module estimates and removes a polarisation background, assuming the background contains only low spatial frequencies. The process is carried out on the Stokes images $S_i(m, n)$ which are band-limited due to the Stokes transfer functions (STFs). The background-removed Stokes images can then be expressed as

$$S'_i(m, n) = S_i(m, n) - (S_i * g_\sigma)(m, n), \quad (6.1)$$

where

$$g_\sigma(m, n) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{m^2 + n^2}{2\sigma^2}\right), \quad (6.2)$$

represents a 2D Gaussian kernel with a width of σ . Additionally, the module may attempt to correct the instantaneous field-of-view (IFoV) error caused by the spatial offset between polarisation channels. This issue is discussed in Chapter 2. The correction uses a method based on the Fourier–shift theorem introduced in Feng [41], which is evaluated in detail in Section 7.3.1. After all processing, the false-colour image is passed back to ScopeFoundry for display. Thanks to GPU-acceleration with `cupy` [176], the image processing can be carried out in real time as images are constantly acquired.

Finally, the ScopeFoundry application acquires single frames, time series, and z -stacks, saving them in a structured layout at a specified location (typically with cloud storage due to the large size) as TIFF images with accompanying JSON files that record hardware attributes. This automation is critical, enabling statistical analysis on large datasets for read-quality characterisation or decoding, as will be presented in the following chapters.

6.3 Conclusion

In this chapter, the practical implementation of the PPC system has been described. While the preceding chapters established how polarisation contrast is generated in PPC and how the imaging process can be modelled analytically and numerically, reliable imaging in practice requires a system that is sufficiently stable, repeatable and automated, especially in the context of ODS where large datasets are required for the upcoming informational analysis. The hardware upgrades reported in this chapter, including high-power illumination with pulsed operation for thermal management, the mirrorless implementation, and motorised 3D positioning with repeatable referencing, were motivated by these requirements. In particular, the 3D positioning capability enables accurate navigation to a sector of interest, and long acquisition sessions can then be carried out without significant drift.

On the software side, the system needs not only to acquire raw frames, but also to provide real-time feedback and be sufficiently automated for scripted acquisition. Because PPC is optimised for samples with weak polarisation contrast, the features of interest are often difficult to locate

and to focus on directly from single raw frames, whether the sample is biological or birefringent voxels. The real-time visualisation capability provided by the PCIP module is therefore of great value. Together with averaging raw images using a circular buffer, the software enables the user to rapidly identify weak features and achieve proper focus during an imaging session. In addition, based on the ScopeFoundry platform, the software integrates camera and stage operations, enabling acquisition of not only single frames, but also time series and z -stacks, with structured storage of image data and hardware metadata.

In the next chapter, we build on this experimental foundation and move from implementation and visualisation to quantitative evaluation of voxel imaging performance. Rather than treating the PCIP output as a visualisation tool, Chapter 7 assumes polarisation images have been acquired, and applies an image processing and decoding pipeline, including sector masking, voxel localisation, and symbol classification. The polarisation measurements are then analysed using mutual information. By making use of the experimental system and the simulation pipeline described in the previous chapter, the informational analysis will provide a framework for understanding how different sources of error affect the readout quality.

Chapter 7

Informational analysis of voxel imaging

In the previous chapters, we introduced polarisation-sensitive phase contrast (PPC) as an imaging technique optimised for weakly birefringent samples such as voxels, and showed how it can be modelled numerically and implemented experimentally. In the application of optical data storage (ODS), however, it is not sufficient to obtain a polarisation image of voxels. Instead, the polarisation states from each voxel must be measured quantitatively and compared to its ground truth in order to assess how well the original data can be retrieved. Furthermore, it is critical to understand how the optical system affects, and potentially limits, the performance of ODS through the underlying physical processes.

In this chapter, we take an information-theoretic perspective on how well a microscopy system can serve as the read head for ODS. We first introduce a framework for informational analysis in which mutual information (MI) is used as a metric for quantifying how much information is carried by birefringent voxels. By viewing ODS as a communication channel, the estimated MI between the ground-truth voxel symbols and the measured polarisation states provides an estimate, and in particular a lower bound, on the channel capacity according to Shannon's noisy-channel coding theorem.

We then present a practical image-processing pipeline that automatically extracts a polarisation measurement for each voxel in a polarisation camera (PolCam) image, from which MI can be estimated. Together, the theoretical framework and the processing pipeline allow us to

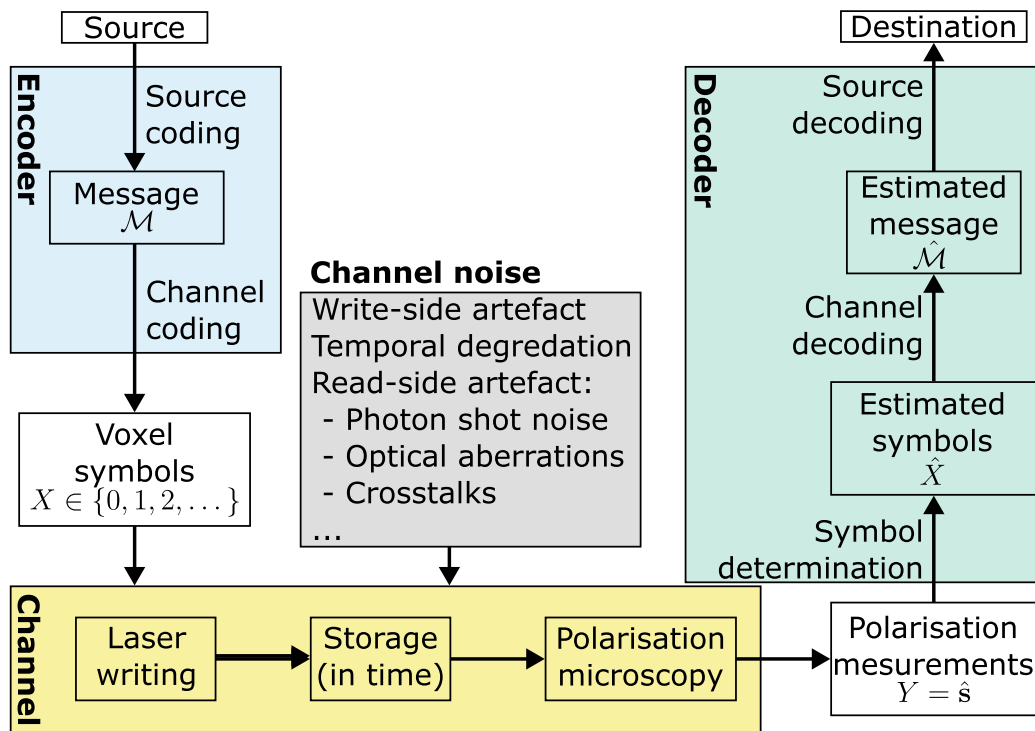


Figure 7.1: A model of communication for optical data storage with birefringent voxels. To store a source-coded message \mathcal{M} , it is firstly encoded into a sequence of voxel symbols, X , as the input of a communication channel. The channel consists of voxel writing with femtosecond lasers, storage of glass, and the voxel reading process with polarisation microscopy, each of which introduces noises to the channel. The output of the channel is the polarisation measurement of voxels, \hat{s} , from which the sequence of voxel symbols is estimated and decoded into the estimated message $\hat{\mathcal{M}}$.

investigate how the channel capacity is influenced by physically meaningful parameters in the optical system, such as photon shot noise, defocus, and aberrations, for which experimental results will be presented and compared to simulations.

7.1 A communication model for optical data storage

From an informational perspective, data storage is considered as a communication process, for which the process can be abstracted using Shannon's linear model [181], as illustrated in Figure 7.1. Consider a source-coded message \mathcal{M} that needs to be stored. A deterministic channel-coding algorithm determines a sequence of symbols to be written as voxels in glass. A (noisy) communication channel, \mathcal{C} , can be defined for the process of voxel writing, the storage of glass, and the polarimetry of voxels using a read head like the PPC. The decoder then try to determine the voxel symbols, from which the original message can be estimated.

From the communication model, it is apparent that the whole process includes many elements beyond the scope of this thesis, such as the coding strategy and the writing and storage of voxels. Instead, we are interested in how well a method of polarisation imaging serves as the read head of ODS, and how that performance may be affected in various situations. Therefore, rather than investigating the input and output messages \mathcal{M} and $\hat{\mathcal{M}}$, or focusing primarily on how well \hat{X} estimates the original voxel symbols X , we mostly limit our discussion to the communication channel \mathcal{C} itself. As illustrated in the figure, the input of the channel is the ground truth of the voxel symbol, treated as a discrete random variable. In contrast, the output of the channel can be continuous, as it represents the polarisation measurements made on the voxels. In the case of single-shot polarimetry with a PolCam, the measurement is a 2D projection of the Stokes parameter on the normalised linear Stokes plane (NLSP), $\hat{\mathbf{s}} := [\hat{S}_1, \hat{S}_2]^\top$, as discussed in Chapter 2. As a result, the statistical relationship between the input and the output can be used to quantify the performance of the channel itself, largely independently of the particular encoder and decoder.

At the same time, to demonstrate the feasibility of voxel-based ODS in a concrete setting, we will later introduce a symbol-determination step that maps measurements Y to estimated symbols \hat{X} , enabling direct comparisons between X and \hat{X} . With an assumed Gray-code mapping [182], this also permits illustrative comparisons between messages \mathcal{M} and $\hat{\mathcal{M}}$. Nevertheless, the pair (X, Y) remains the most relevant basis for evaluating the intrinsic performance of the PPC read head, since in practical applications the encoder and decoder can be designed around a well-characterised channel, for example by introducing an appropriate amount of data redundancy depending on the channel capacity, which is in turn affected by channel impairments such as photon shot noise and optical aberrations in the microscopy system.

7.2 Entropy, mutual information, and channel capacity

The statistical relationship between the channel input X and the corresponding output Y is ultimately governed by the joint distribution $p_{X,Y}(x, y)$ between the two random variables. In the

context of communication, we are particularly interested in how much information is conveyed about the input when we measure the output, which is the mutual information (MI) [183]. In cases where X is discrete while Y is continuous, the MI is given as¹

$$\text{MI}(X, Y) = \sum_x \int p_{X,Y}(x, y) \log \frac{p_{X,Y}(x, y)}{p_X(x) p_Y(y)} dy, \quad (7.1)$$

where $p_X(x)$ denotes the marginal probability

$$p_X(x) = \int p_{X,Y}(x, y) dy, \quad (7.2)$$

and $p_Y(y)$ denotes the marginal probability density function (PDF), given as

$$p_Y(y) = \sum_x p_{X,Y}(x, y). \quad (7.3)$$

MI is a particularly powerful metric, because it quantifies how much information is actually received at the channel output Y for any given input distribution $p_X(x)$, regardless of how the measurements may be classified into symbols. In other words, it measures how much the uncertainty in X is reduced given the knowledge of Y , which can be seen by expressing MI using Shannon entropies, given as

$$\text{MI}(X, Y) = H(X) - H(X | Y). \quad (7.4)$$

In this equation, $H(X)$ represents the Shannon entropy of X , defined as

$$H(X) := - \sum_x p_X(x) \log p_X(x), \quad (7.5)$$

¹Throughout the chapter, the base of the logarithm is omitted for simplicity, as it can be arbitrarily chosen depending on the unit of information in use. In particular, a quantity is given in bits if base 2 is used, or in nats if base e is used instead.

and the conditional entropy $H(X | Y)$ can be written as

$$H(X | Y) = - \sum_x \int p_{X,Y}(x, y) \log \left[\frac{p_{X,Y}(x, y)}{p_Y(y)} \right] dy. \quad (7.6)$$

The conditional entropy $H(X | Y)$ is the average uncertainty that remains about X after the channel output Y is observed. In perfect communication, the channel input X can be determined from Y without ambiguity, in which case $H(X | Y) = 0$. In contrast, the noisier the communication, the greater the ambiguity that remains after observing Y , leading to a larger $H(X | Y)$. Therefore, the difference between the input entropy $H(X)$ and the conditional entropy $H(X | Y)$ quantifies the expected reduction in uncertainty about X by observing Y , and thus represents the amount of information conveyed through the communication channel.

From Equation 7.4, it can be seen that the maximum possible value for MI is the input entropy $H(X)$. This happens only in the absence of any channel imperfections, so that perfect inference of input X from the output Y is possible, resulting in the conditional entropy $H(X | Y)$ being zero. If the input X is one of N_X possible symbols denoted with integers from 0 to $N_X - 1$, the entropy $H(X)$ is maximised when X follows a discrete uniform distribution, i.e.

$$X \sim \mathcal{U}\{0, N_X - 1\}. \quad (7.7)$$

In a perfect channel, the maximum MI takes the maximum possible value of $\log N_X$.

Nevertheless, realistically, some information is inevitably lost in a noisy channel, as the channel noise introduces uncertainty in the process through $H(X | Y)$. In fact, an informational limit can be set for a noisy channel \mathcal{C} beyond which no further information may be transmitted even with the increase of the input entropy $H(X)$. The limit is quantified as the channel capacity C [183], given as

$$C(\mathcal{C}) = \max_{p_X} \text{MI}(X, Y), \quad (7.8)$$

and Shannon's noisy channel coding theorem states that error-free communication is possible through a noisy channel at a rate C [181, 183]. The theorem thus enables the performance of the channel to be characterised without the specification of a coding strategy or a classification

algorithm. In fact, the coding strategy and classification algorithm should instead be designed according to the channel characteristics, in order to make the best use of the channel capacity.

It should be noted that, in the expression of the channel capacity, the input probability p_X ought to be determined in a way that maximises the mutual information given the joint distribution $p_{X,Y}(x, y)$, which, in general, may not be a uniform distribution. In our particular channel, there is a rotational symmetry that can be taken advantage of. Consider a writing scheme where the birefringent azimuths of the voxel symbols spreads evenly within $[0, \pi)$, i.e.

$$\gamma_x = \frac{\pi x}{N_X}. \quad (7.9)$$

Upon imaging, PPC introduces an enhancement in degree of linear polarisation (DoLP) and a constant offset in angle of linear polarisation (AoLP), according to our analysis in Chapter 3. Nevertheless, given the illumination polarisation state of PPC is circular, the polarisation contrast is uniform for any birefringent azimuth of the sample. As a result, a different input symbol X should only result in a rotation of the measured Stokes vector $\hat{\mathbf{s}}$ in the NLSP, i.e.

$$x \mapsto x + 1 \quad \Longrightarrow \quad \hat{\mathbf{s}} \mapsto \mathbb{R}(\pi/N_X) \hat{\mathbf{s}}, \quad (7.10)$$

where \mathbb{R} denotes the 2D rotation matrix, given as

$$\mathbb{R}(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (7.11)$$

Assuming the channel noises to be also independent of the azimuth, this creates a symmetry in the communication channel, as one should expect

$$p(\hat{\mathbf{s}} \mid X = x) = p(\mathbb{R}(\pi/N_X) \hat{\mathbf{s}} \mid X = x + 1) \quad (7.12)$$

to hold for any x . As a result, the channel is said to be symmetric, for which the MI is maximised when the input symbol follows a uniform distribution [184].

In the process of characterising the channel capacity, we therefore use a discrete uniform distribution for the voxel symbols, and use MI to set a lower bound for the channel capacity, which is usually also a good estimate as long as $H(X) > \text{MI}(X, Y)$ and the imaging process is reasonably rotationally symmetric on the NLSP. This provides a practical and decoder-agnostic way to quantify the intrinsic performance of PPC for voxel reading, and motivates the remainder of this chapter where $\text{MI}(X, Y)$ is used to characterise how factors such as noise, defocus, and aberrations impact voxel readout.

7.3 The voxel image processing pipeline

To estimate the MI between the ground-truth voxel symbols X and the polarisation measurements $Y = \hat{s}$, it is necessary to obtain a dataset of paired samples $\{(x_v, y_v)\}_v$ where v denotes the voxel index. The polarisation-camera image processing (PCIP) module introduced in Section 6.2 is helpful to compute a map of \hat{s} across the whole field-of-view (FoV), but the polarisation measurements need to be extracted at the voxel level and be associated with the correct ground truth X for that symbol. To achieve this, it is necessary to develop a pipeline that not only performs the low-level polarisation calibration and image processing, but also automatically localises the voxels in the image. The corresponding Stokes measurements can then be resampled at these voxel positions and paired with the relevant ground-truth symbols, enabling statistical analysis and MI estimation.

The voxel image processing pipeline (VIPPP) is developed to be responsible for this process, for which a schematic diagram is shown in Figure 7.2. Receiving raw PolCam images as inputs, the pipeline firstly carries out the polarisation-field correction (PFC) when appropriate bright-field and dark-field reference images are provided. As described in Section 2.6.3, the PFC calibration is particularly effective in removing any fixed pattern noise (FPN) in the PolCam. This is especially critical for imaging birefringent voxels, because any polarisation modulation introduced by voxels is inferred from small differences between the raw polarisation channels. Without the PFC, these differences would be dominated by the FPN of the PolCam. In

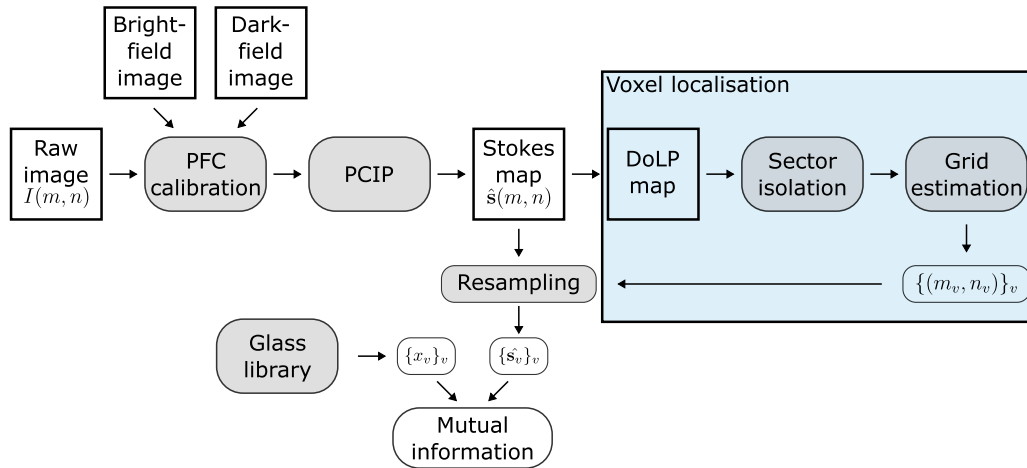


Figure 7.2: Block diagram for the voxel image processing pipeline (VIPPP) to obtain MI estimations from raw images. The pipeline starts with low-level image processing including polarisation-field correction (PFC) and polarisation camera image processing (PCIP) that obtains a map of normalised 2D Stokes vectors $\hat{\mathbf{s}}$ across the image (m, n) . To resample the Stokes measurements at the voxel locations, a voxel localisation process is carried out on the degree of linear polarisation (DoLP) map of the image that estimates the pixel coordinates (m_v, n_v) of any voxel v . Finally, by comparing the Stokes measurements at voxel locations $\hat{\mathbf{s}}_v$ with their corresponding ground-truth symbols x_v , mutual information (MI) can be estimated.

addition, the PFC corrects for any polarisation offset in the illumination of the system, by taking the difference between the raw image and the bright-field image that would have the same offset.

The second step of the VIPPP is to compute the polarimetric parameters at the image level, by making use of the PCIP module introduced in Section 6.2. Within the PCIP processing, the raw image is split into polarisation channels, from which the Stokes images S_0 – S_2 can be computed. From the Stokes images, the $\hat{\mathbf{s}}$ vectors in the NLSP can therefore be estimated for every superpixel of the PolCam. In addition, any residual polarisation background is estimated and removed using a 2D Gaussian kernel described in Equation 6.1. Together with the PFC correction, this step eliminates from the image any DoLP background as much as possible, enabling the subsequent voxel localisation to be performed on the DoLP signal.

7.3.1 Instantaneous field-of-view error (IFOV) correction

As discussed in Chapter 2, a consequence of PolCam polarimetry as a division of focal plane (DoFP) method is the generation of instantaneous field-of-view (IFoV) errors. This ultimately




Figure 7.3: Generation and correction of instantaneous field-of-view (IFoV) errors in the presence of an intensity gradient. A 2D Gaussian intensity profile (a) is provided to the polarisation-camera image processing (PCIP) module without IFOV correction. The resulting false-colour output (b) exhibits false polarisation signal as IFOV errors. With correction enabled, the false-colour output (c) shows no visually perceivable artefact. False-colour scheme in (b) and (c): Hue, AoLP; value, DoLP under identical normalisation.

stems from the ambiguity between an intensity gradient and a polarisation signal, because the four polarisation channels are sampled at different pixel locations and are therefore misaligned by one pixel with each other on the sensor. To illustrate the generation of IFOV errors, a computationally generated 2D Gaussian intensity profile is provided to the PCIP module, as an example of a pure intensity signal with no polarisation signal, as shown in Figure 7.3(a). Without any IFOV correction, a spurious polarisation contrast can be seen in the false-colour output in Figure 7.3(b), where annular artefacts arise with a varying estimated AoLP.

To correct for the error, polarisation channels must be co-registered with each other, which requires some form of interpolation or resampling [70]. Within the PCIP module, a number of methods have been implemented and can be selected by specifying it in the options. Here, we present one interpolation-based method and one resampling method that generally performs well in voxel imaging.

In applications where no prior bandwidth model is available, the underlying polarimetric images may not be well approximated as band-limited. In contrast, in a well-sampled microscopy system like PPC, the true signal is band-limited by the Stokes transfer function (STF) and Nyquist-sampled by the PolCam, as discussed in Section 3.3 and 2.5, respectively. Therefore, the Whittaker–Shannon interpolation [81] (also known as the sinc interpolation) is deemed ideal, because it interpolates without introducing additional spatial-frequency content in the Fourier

spectrum. For PolCam images, this can be implemented through zero-padding, by first computing the Fourier spectra $\tilde{I}_p(\xi, \eta)$ of the polarisation channels using a discrete Fourier transform (DFT) algorithm, zero-padding them by a factor of two in both x and y -directions, and inverse Fourier transforming the images. The resultant images are oversampled by a factor of two in each direction, and the misalignment between channels then becomes an offset by one pixel. As a result, channel registration is realised by shifting the arrays by one column and/or row according to the micro-polariser array (MPA) layout, and the registered polarisation-channel images are then used to compute Stokes images and any subsequent measurements in the pipeline. The Whittaker–Shannon interpolation outperforms commonly used methods such as bilinear [185] and bicubic spline [186] in voxel imaging, by taking advantage of the prior knowledge that the Fourier spectrum of the true signal should be negligible outside the Nyquist region.

Alternatively, one could perform a translation in $I_p(m, n)$ itself to correct for the misalignment between polarisation channels. Nevertheless, as each polarisation channel contains one sample from each superpixel, sub-pixel shifts on $I_p(m, n)$ are required for channel registration. To preserve the Fourier spectrum, we employ the Fourier shift theorem [85], which states that the translation of a function $f(x)$ by an amount of x_0 corresponds to a multiplication with a tilted phase in its Fourier spectrum, i.e.

$$\mathcal{F}\{f(x - x_0)\}(\xi) = e^{-i2\pi\xi x_0} \mathcal{F}\{f\}(\xi), \quad (7.13)$$

which enables sub-pixel shifts by an arbitrary offset.

For channel registration in PolCam, the desired shift is $(\pm 0.5, \pm 0.5)$ in pixels, so the corrected polarisation-channel images can be computed as

$$I_p^{\text{corrected}}(m, n) = \mathcal{F}^{-1} \left\{ e^{i\pi(\pm\xi \pm \eta)} \mathcal{F} \{ I_p(m, n) \} (\xi, \eta) \right\}. \quad (7.14)$$

Similar to the Whittaker–Shannon interpolation method, the Fourier shift method also introduces no additional spatial-frequency content, as it is a multiplication in the Fourier spectrum. This method is usually preferred when PCIP is supposed to be running in real-time (for example

within ScopeFoundry), as it does not result in a larger image like the interpolation method that slows further processing down. In contrast, the Whittaker–Shannon interpolation method usually yields a higher estimated MI, due to a larger, accurately reconstructed image from which the subsequent voxel resampling may benefit. For the example of a simulated 2D Gaussian, both methods are very effective and yield a false-colour image with no visible IFoV error, as shown in Figure 7.3(c).

7.3.2 Voxel localisation

After image-level processing through the PCIP module, the next step in the VIPP pipeline is to determine the voxel locations (m_v, n_v) in image coordinates, such that the Stokes parameter \hat{s} can be resampled at these locations. The voxel pitch within a sector is typically smaller than one micrometer, in contrast to the accuracy of $\pm 3\mu\text{m}$ for the translation stage, so it becomes infeasible to determine the voxel locations from stage coordinates alone. Instead, an image-based procedure is developed.

In voxel localisation, rather than detecting each voxel individually, the voxels are assumed to lie on a finite affine grid [187], i.e.

$$\{(m_v, n_v)\}_v = \{\mathbf{v}_0 + p\mathbf{v}_1 + q\mathbf{v}_2\}_{p,q}, \quad (7.15)$$

where \mathbf{v}_1 and \mathbf{v}_2 represent the two basis vectors of the affine grid, \mathbf{v}_0 represents a translational offset, and p, q are integers used for indexing voxel coordinates. Such an affine grid can be obtained by applying an affine transformation to an underlying rectangular grid. The transformation may include rotation, scaling, and shear, which maps the fundamental unit element to a parallelogram. Without loss of generality, \mathbf{v}_0 can be chosen so that it lies in the fundamental parallelogram of the affine grid, i.e.

$$\mathbf{v}_0 = a\mathbf{v}_1 + b\mathbf{v}_2, \quad (7.16)$$

where

$$-\frac{1}{2} \leq a, b < \frac{1}{2}. \quad (7.17)$$

In an ideal situation, the basis vectors correspond to the horizontal and vertical voxel pitches specified for writing, adjusted by the microscope magnification and the pixel pitch of the camera. In reality, however, such an estimate is not usually accurate, due to a variety of factors such as magnification discrepancy, sample placement, writing error, or misalignment in the system. The task of voxel localisation is therefore to estimate the affine parameters $(\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2)$ and the ranges of p, q such that the voxel coordinates can be computed using Equation 7.15.

As the voxel localisation algorithm should be agnostic to the symbols, the determination process should not rely on the AoLP information in the image. Instead, we use the DoLP image, which is expected to be largely free from any polarisation background after PFC calibration and Gaussian-filter-based background subtraction. The voxel localisation algorithm therefore takes DoLP as its input, and provides robust localisation in practice.

Sector isolation

The voxel localisation process is broken down into two steps, sector isolation and grid estimation, because there are usually multiple voxel sectors within the FoV, as can be seen from the DoLP map in Figure 7.4(a). Sector isolation aims to compute a Boolean mask that covers only the sector of interest. This is not only helpful for determining the valid range of grid indices (p, q) , but is also vital for estimating the affine grid parameters $(\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2)$, which can take different values across different sectors. Rather than providing a theoretically optimal segmentation, this procedure is primarily pragmatic: a simple and effective method that works robustly across the datasets considered in this project, enabling grid estimation and subsequent informational analysis.

Starting with the DoLP image (a), a maximum filtering (b) is first applied such that the spacing between voxels is connected while that between sectors is not. After that, a simple thresholding operation produces a Boolean mask of all data sectors (c), from which different




Figure 7.4: Sector isolation with morphological operations and Hough transformation. (a) DoLP image for the whole FoV that contains multiple voxel sectors. (b) Maximum-filtered DoLP image. (c) Boolean mask for all voxel sectors. (d) Labelled masks for voxel sectors. (e) Largest voxel sector. (f) Largest voxel sector from a noisy image. (g) Hole-filled mask for the largest detected voxel sector. (h) Edge of the sector. (i) Detected sector edge in Hough space (r_l, θ_l) . (j) Boolean mask for the voxel sector with straight-line edges, constructed from estimated parameters in the Hough space.

connected regions can be labelled (d). The desired sector (e) is then chosen by assuming it is the largest connected component in the FoV. If the desired sector is not the largest one by area, cropping should be specified such that it is. In cases where the original image is noisy, a ‘broken’ sector mask $M(m, n)$ is sometimes generated (f), where holes exist within the sector region and the boundary may be fragmented.

To overcome this problem, a morphological hole-filling is carried out (g). In practice, this step is implemented by iterative binary erosion [188] under a constraint mask, which fills holes while preserving the outer boundary of the sector. Erosion is a morphological operation that ‘shrinks’ a binary mask, upon which a pixel remains True only if a chosen kernel (the structuring element) fits entirely inside the True region of the mask. Symbolically, the hole-filling process can be described as

$$A^{(k+1)} = A^{(k)} \ominus B \vee M, \quad (7.18)$$

where \ominus denotes erosion and \vee denotes the binary OR. A 3×3 squarely-connected kernel is

used as the structuring element B for the erosion, i.e.

$$B = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad (7.19)$$

and $A^{(k)}$ is a Boolean mask in the k -th iteration, starting with A^0 being a full matrix of the same dimension as $M(m, n)$, with every entry being True except at the edge where it is False. Through iterations, the mask A is gradually eroded from its boundary until it reaches the outermost edge of M . At that point erosion no longer proceeds and $A^{(k+1)} = A^{(k)}$. The converged mask $A^{(k)}$ is then regarded as the hole-filled mask M^{filled} .

After hole filling, an edge map (h) is extracted using erosion followed by an exclusive OR operation, given as

$$M^{filled} \ominus B \vee M^{filled}, \quad (7.20)$$

where \vee denotes the binary exclusive OR (XOR).

To obtain straight edges, a Hough transformation [189] is then carried out on the edge map. Expressing straight lines in the form of

$$r_l = n \cos \theta_l + m \sin \theta_l, \quad (7.21)$$

each point in the Hough space (r_l, θ_l) represents a straight line in the image. The Hough transform is therefore computed by letting each True pixel in the Boolean edge map vote for all parameters (r_l, θ_l) corresponding to lines passing through that pixel, as illustrated in Figure 7.4(i). By summing votes from all edge pixels, four peaks in the Hough space correspond to the dominant straight edges in the mask boundary. Finally, the four vertices are obtained from the intersections of the four detected lines, and used to produce a reliable Boolean mask of a single data sector (j). The mask can then be applied to the original DoLP image for grid estimation.




Figure 7.5: Estimation of voxel pitch vectors \mathbf{v}_1 and \mathbf{v}_2 with Fourier analysis. The Fourier spectrum of a masked DoLP image contains discrete peaks due to the underlying affine grid of voxels. Using a chirp z-transform (CZT), the spectrum near the zeroth-order (black) and first-order (red, green) peaks is evaluated with a spectral zoom by a factor of 100, enabling high-precision estimation of both peak locations and phases.

Grid estimation

To estimate the affine parameters $(\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2)$, the masked DoLP image is passed to a Fourier-based grid estimation algorithm. As the voxels lie on an affine grid, the Fourier transform of the masked DoLP image exhibits strong discrete peaks, as shown in Figure 7.5. The zeroth-order peak (circled in black) sits at the centre of the spectrum and corresponds to the DC component of the masked DoLP image. The first-order peaks (circled in red and green) are the strongest non-zero peaks closest to the centre, and their positions encode the two fundamental reciprocal vectors of the voxel grid.

The vector positions of the first-order peaks relative to the zeroth-order are denoted $\boldsymbol{\nu}_1$ and $\boldsymbol{\nu}_2$, in cycles per pixel, and their complex phases are denoted ϕ_1 and ϕ_2 , respectively. With these quantities, several useful relationships with the affine parameters can be identified. Firstly, the

Fourier transform demands a reciprocal relationship, given as

$$\boldsymbol{\nu}_i \cdot \mathbf{v}_i = 1, \quad (7.22)$$

for $i \in \{1, 2\}$. Nevertheless, the reciprocal vectors and basis vectors are, in general, not necessarily parallel, so the element-wise division $\boldsymbol{\nu}_i = 1/\mathbf{v}_i$ may not hold. Instead, a reciprocal vector is orthogonal to the other basis vector, i.e.

$$\boldsymbol{\nu}_i \cdot \mathbf{v}_j = 0, \quad (7.23)$$

for $j \neq i$. To see this, $\boldsymbol{\nu}_1$ corresponds to a set of constant-phase lines normal to $\boldsymbol{\nu}_1$. Each such line contains voxels with the same index p , and these lines are therefore parallel to the other basis vector \mathbf{v}_2 . Combining the two constraints yields

$$\boldsymbol{\nu}_i \cdot \mathbf{v}_j = \delta_{ij}, \quad (7.24)$$

which is a standard relationship in lattice geometry [190].

In addition, the phases at the first-order peaks, ϕ_1 and ϕ_2 , can be related to the translation vector \mathbf{v}_0 of the affine grid, given as

$$\boldsymbol{\nu}_i \cdot \mathbf{v}_0 = -\frac{\phi_i}{2\pi}, \quad (7.25)$$

as a consequence of the Fourier shift theorem in Equation 7.13.

In summary, these relationships can be written in the following matrix equation:

$$\begin{bmatrix} \boldsymbol{\nu}_1^T \\ \boldsymbol{\nu}_2^T \end{bmatrix} \begin{bmatrix} \mathbf{v}_0 & \mathbf{v}_1 & \mathbf{v}_2 \end{bmatrix} = \begin{bmatrix} -\frac{\phi_1}{2\pi} & 1 & 0 \\ -\frac{\phi_2}{2\pi} & 0 & 1 \end{bmatrix}. \quad (7.26)$$

Therefore, the affine parameters \mathbf{v}_0 , \mathbf{v}_1 , and \mathbf{v}_2 can be estimated by matrix inversion if the reciprocal vectors $\boldsymbol{\nu}_1, \boldsymbol{\nu}_2$ and phases ϕ_1, ϕ_2 of the first-order peaks can be determined accurately from the Fourier spectrum.

In practice, the localisation accuracy is limited because the DFT of a discrete signal is itself sampled on a discrete frequency grid. When a conventional algorithm such as the fast-Fourier transform (FFT) is employed, the localisation of a peak is therefore limited to a resolution of $(\frac{1}{M_x}, \frac{1}{M_y})$, where M_x and M_y are the pixel dimensions of the masked DoLP image in the x and y -directions, respectively. This also limits the accuracy of the phase estimation.

To overcome this limitation, the chirp Z-transform (CZT) is used [85, 191] instead of conventional DFT algorithms, enabling spectral zooming by an arbitrary factor¹ around any specified spatial frequency of interest. By using the FFT result as an initial guess, the Fourier spectrum around each peak is evaluated with CZT, as shown on the right of Figure 7.5. A 2D sinc-like pattern is observed for each peak, which arises from the finite size of the Boolean mask acting as a windowing function. In addition, because the mask is not centred in the image, the Fourier transform of the windowing function carries a tilted phase, which is most visible in the phase map around the zeroth-order peak. With CZT, the peak locations and the corresponding phases can be determined with dramatically improved accuracy, enabling robust construction of the affine parameters.

With additional knowledge of the sector size in terms of the number of voxels in each direction, a finite affine grid can then be constructed according to Equation 7.15. The Stokes vector $\hat{\mathbf{s}}(m, n)$ can then be resampled at these locations $\{m_v, n_v\}_v$ using interpolation, such that $\{\hat{\mathbf{s}}_v\}_v$ may be provided for subsequently MI estimation. Empirically, the choice of the interpolation method at this stage was found to have negligible practical influence, especially when IFoV error has been corrected using the Whittaker–Shannon interpolation method, because the voxel pitch is much larger than the pixel pitch. In VIPP, bicubic spline interpolation is employed in this step.

Finally, as the estimated sector mask is typically slightly larger than the true extent of the sector due to maximum filtering, the valid index ranges for (p, q) in the finite affine grid need be determined. This can usually be achieved by optimising a small translation of the grid to maximise the total resampled DoLP.

¹The CZT implementation used for spectral zooming is adapted from MATLAB [192], ported by Mark A. A. Neil who incorporated the Fourier shift theorem to enable zooming around an arbitrary spatial frequency, and was subsequently revised and integrated into the present work by the author.

7.3.3 Mutual information estimation

Now that the Stokes measurements at the voxel positions $\{\hat{\mathbf{s}}_v\}_v$ have been obtained, they can be compared with the corresponding ground-truth symbols $\{x_v\}_v$ such that the MI can be estimated. The ground-truth symbols are read from the glass library, a database containing the specifications according to which the voxel samples are written, as introduced in Section 5.2.2. The goal of this section is therefore to estimate $\text{MI}(X, Y)$ between the discrete input random variable X and the continuous output random variable Y of the communication channel, from the measurement pairs $\{(x_v, \hat{\mathbf{s}}_v)\}_v$ within a voxel sector.

Binning estimator

A conceptually simple approach is to estimate the MI according to its definition in Equation 7.1, by first estimating the joint distribution $p_{X,Y}$. The corresponding marginals p_X and p_Y can then be obtained by summation or integration. A practical challenge arises, however, due to the continuous nature of the channel output $Y = \hat{\mathbf{s}} \in \mathbb{R}^2$. A common approach is therefore to discretise Y by binning, and treat the binned measurement as a discrete variable $Y_{\mathcal{B}}$. The MI can then be estimated by a finite sum, given as

$$\text{MI}(X, Y) \approx \text{MI}(X, Y_{\mathcal{B}}) = \sum_{x,j} p(x, j) \log \frac{p(x, j)}{p_X(x)p_{Y_{\mathcal{B}}}(j)}, \quad (7.27)$$

where j denotes the bin index. The binned marginal probability can then be defined as

$$p_{Y_{\mathcal{B}}}(j) := \int_j p_Y(y) dy, \quad (7.28)$$

and can be estimated in practice by counting the number of measurements in bin j , denoted n_j , and dividing it by the total number of measurements N in the dataset.

In the limit where the bin size tends to zero and the number of voxels tends to infinity, the binning estimator converges to the true MI. Nevertheless, in realistic settings the estimator is known to be biased [193, 194]. When the number of bins is too large relative to the number

of voxels, many bins are poorly populated and the estimations of $p_{Y_B}(j)$ and $p(x, j)$ become unreliable. As each bin contributes non-negatively to the sum, these fluctuations tend to yield a spurious positive contribution on average, leading to an overestimate. Conversely, when the bins are too wide, information within each bin is discarded and lost, leading to an underestimate. Between the two regimes there is an optimal binning choice, but Ross [194] showed that the determination of this optimum is not straightforward, and is in general not possible from simple statistics of the data.

***k*-nearest neighbour (*k*-NN) estimator**

To avoid the bias associated with binning, or an arbitrary discretisation of the continuous measurement Y , the VIPP employs a k -nearest neighbour (k -NN)-based estimator [194] that does not require estimating $p_{X,Y}$ explicitly. The approach builds on the relationship between MI and Shannon entropies [183], given by

$$\text{MI}(X, Y) = H(X) + H(Y) - H(X, Y), \quad (7.29)$$

where $H(X, Y)$ is the joint entropy, representing the overall uncertainty of the pair (X, Y) . The joint entropy thus relates to the conditional entropy according to

$$H(X, Y) = H(Y) + H(X | Y), \quad (7.30)$$

so that Equation 7.29 follows directly from Equation 7.4.

As the entropy of a continuous variable can be estimated using a k -NN algorithm [195], MI can be computed without discretising Y . In the present case where X is discrete and $Y = \hat{s}$ is continuous, the algorithm iterates through each measurement \hat{s}_v and finds the distance d_v to its k -th nearest neighbour on the NLSF, only counting measurements that share the same label with \hat{s}_v . Here, the Euclidean distance is adopted for simplicity, rather than, for example, the angular distance on the Poincaré sphere. For each voxel v , the algorithm then counts the number of measurements of any label μ_v that lie within distance d_v of \hat{s}_v . The MI can then be

estimated as

$$\text{MI}(X, Y) = \langle \psi(N) - \psi(N_v) + \psi(k) - \psi(\mu_v) \rangle_v, \quad (7.31)$$

where N_v denotes the number of voxels in the dataset that share the same input symbol as voxel v , ψ is the digamma function [148], and $\langle \dots \rangle_v$ denotes averaging over all voxels.

The k -NN estimator has been shown to be reasonably stable with respect to the choice of k [194] when it takes a small integer. Unless otherwise stated, $k = 3$ is used for all results reported in this work.

7.4 Single-layer voxel imaging results

7.4.1 Baseline demonstration of voxel reading

To demonstrate the capability of the PPC system in resolving voxels, the results from an example imaging session are presented in Figure 7.6. The sector of interest contains 162×162 voxels, separated by $0.7 \mu\text{m}$ with each other both horizontally and vertically. In the 3-bits per voxel (bpv) scheme, there are 8 possible symbols (labelled as A–H) for which each voxel may represent, and the step in birefringent azimuth is 22.5° between each of them. The sector sits at a depth of $150 \mu\text{m}$ from the top surface of the glass, and the spherical aberrations are corrected for by manually setting the correction collar of the objective.

Figure 7.6(a,b) shows the false-colour image computed by the PCIP module. As expected, point-like features of various colours can be observed for the birefringent voxels, and the background is black due to successful polarisation background removal by the PFC and Gaussian filtering introduced in the previous section.

It is, however, difficult to judge from visual inspection alone whether individual voxels are truly well resolved. In some cases, adjacent voxels appear to visually connect, particularly when they encode symbols with similar birefringent azimuths. In such situations, the polarisation responses from neighbouring voxels partially overlap and mix, which can lead to a non-zero

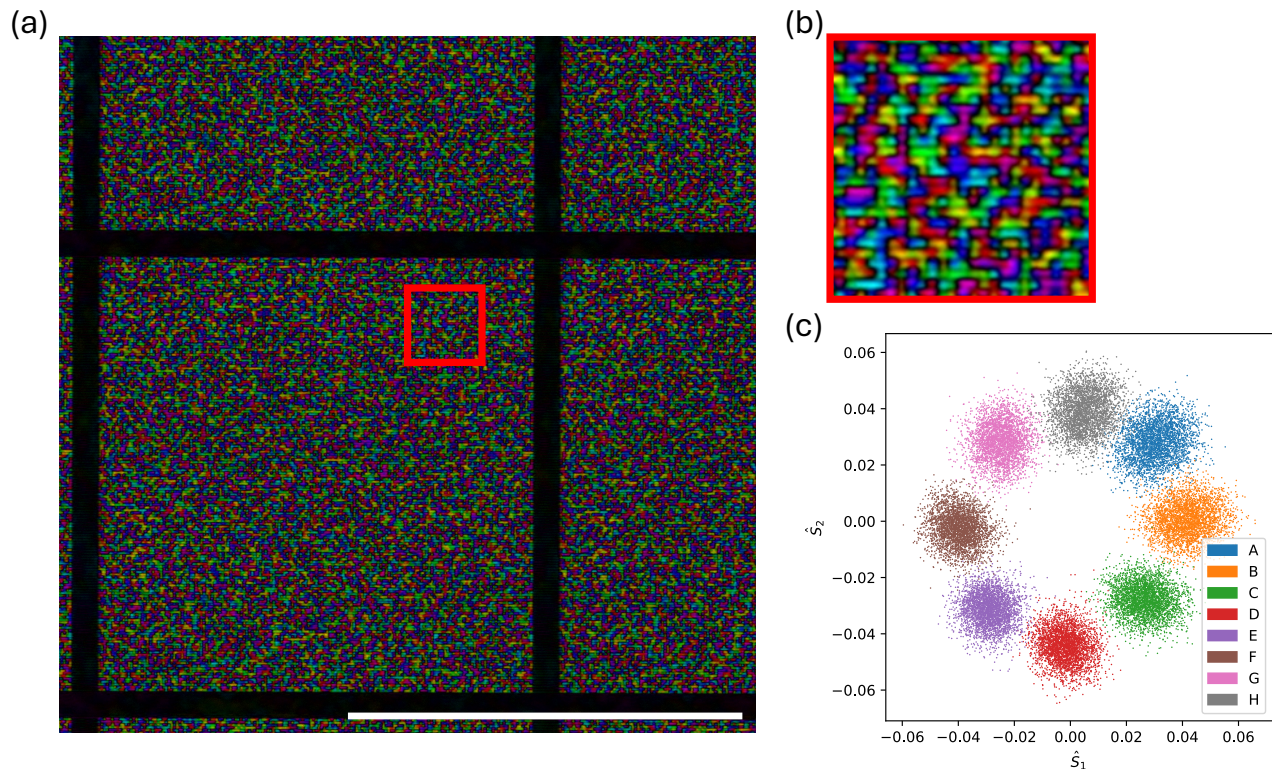


Figure 7.6: False-colour image and constellation diagram for a single-layer voxel sector. (a) False-colour image as the output from the PCIP module. The sector contains 162×162 voxels, with a voxel pitch of $0.7 \mu\text{m}$ in both directions, sitting at a depth of $150 \mu\text{m}$ from the top surface of the glass. (b) An enlarged image of the region-of-interest, marked with a red rectangle in (a). (c) Scatter plot of Stokes measurements $\hat{\mathbf{s}}$ resampled at estimated voxel locations, in the normalised linear Stokes plane (NLSP). Voxel symbols are labelled as alphabets from A to H for a 3-bit-per-voxel scheme. False-colour scheme in (a) and (b): Hue, AoLP; value, DoLP. Bar= $100 \mu\text{m}$.

DoLP and an apparently continuous false-colour region, as discussed in Section 2.3. It may, of course, be a result of the underlying voxel morphology. Conversely, neighbouring voxels may appear visually well separated even if the polarisation signals they produce overlap with each other, as long as the polarisation states are of a nearly orthogonal AoLP. It should be emphasised that, in general, the DoLP (and hence the false-colour rendering) is not band-limited, because the STF acts on the Stokes parameters. As a result, the apparent connectivity between voxels depends on multiple factors, including the extent of the STF, the sampling of the PolCam, and the voxel morphology and the ground truth.

A scatter plot in the NLSP of Stokes measurements $\hat{\mathbf{s}}_v$ after resampling at estimated voxel locations is shown in Figure 7.6(c). The data points are plotted in different colours, depending on their corresponding ground-truth symbol. It is obvious that eight well-separated clusters are present, as expected in a 3-bpv scheme. In digital communications, the plot is known as a

Table 7.1: Mutual information and 2D data density for single-layer sectors at different voxel pitches.

Input entropy (bpv)	Voxel pitch (μm)	MI (bpv)	Data density per layer (Mbit mm^{-2})
4	0.9	3.754	4.63
4	0.8	3.795	5.93
3	0.7	2.964	6.05
4	0.5	1.541	6.16
3	0.45	0.376	1.86

constellation diagram [196], representing how a digitally modulated signal is received.

The eight clusters are mostly well separated with little overlap, and therefore leaving little chance that a voxel may be misinterpreted. That is to say, given the knowledge of a voxel measurement, the uncertainty of the corresponding input is minimal, i.e. $H(X | Y) \approx 0$, such that one can expect

$$\text{MI}(X, Y) \approx H(X) \approx 3 \text{ bpv}. \quad (7.32)$$

In fact, the k -NN method estimated an MI of 2.964 bpv, which is remarkably close to the maximal possible MI dictated by the 3-bpv scheme. If allowed to be extrapolated, the MI also translates to a data density of 6.05 Mbit mm^{-2} per layer, which shows great potential as a single-layer characterisation. In addition, the fact that the estimated MI is close to the Shannon entropy of the input implies that the channel capacity is very likely to be higher.

In addition, the eight clusters exhibit similar spreads and comparable mean radial distances from the origin. Their angular separations are also approximately uniform. Any residual non-uniformity is small, and is potentially due to a weak background polarisation due to the circular polariser. Overall, the constellation remains close to rotationally symmetric, supporting the symmetric-channel assumption discussed earlier and making a uniform input distribution a reasonable choice for channel coding.

7.4.2 Impact of voxel pitch

Among the physical parameters related to voxel writing, the voxel pitch is of obvious importance, as a smaller pitch apparently imposes challenges to any imaging system due to a finite

resolution. Table 7.1 summarises the results of imaged voxel sectors at several voxel pitches. When the input entropy increases to 4 bpv, the MI is no longer capped at 3 bpv and can increase accordingly. At a pitch of 0.8 μm , an MI as high as 3.795 bpv was achieved, beyond which further increasing the pitch brings no noticeable improvement in this dataset. Conversely, decreasing the pitch increases crosstalk between neighbouring voxels, as polarisation signals from them increasingly overlap. When the pitch drops below the diffraction limit of 589 nm estimated in Section 3.3, the MI decreases substantially. A smaller pitch also makes voxel localisation more challenging, which directly impacts the resampling of Stokes measurements. At 0.45 μm , localisation was no longer reliable, and the estimated MI drops to 0.376 bpv.

Even when the MI per voxel is strongly compromised, error-free storage can in principle still be achieved by introducing redundancy through channel coding, due to Shannon's noisy-channel coding theorem. From a storage perspective, there is therefore a trade-off between the information conveyed per voxel and the area occupied by each voxel, and the relevant metric becomes the data density. In Table 7.1, the highest per-layer data density, 6.16 Mbit mm^{-2} , occurs at a pitch of 0.5 μm , even though this pitch lies below the diffraction limit and yields a reduced MI of 1.541 bpv. This illustrates that sacrificing MI per voxel can still increase the overall data density, at the cost of a more demanding coding strategy and a more challenging symbol determination.

7.4.3 Write-side effects and intra-layer crosstalk

In building a physical model for numerical simulations in Section 5.1, the voxels were modelled as birefringent disks of a definite radius. In reality, this is never strictly true. For both type-X [31] and type-S modifications [32], voxels arise from nanovoid structures whose typical sizes are far below the diffraction limit. Moreover, their actual morphology depends strongly on the writing parameters [38], including the number, duration, energy, and polarisation state of the writing pulses. It is therefore unrealistic, within the scope of this work, to build a fully physics-accurate model for individual voxels.

Instead, we focus on the typical characteristics of the signals produced by voxels, and how

they affect the performance of ODS. Since PPC does not resolve features below the diffraction limit, we perform a statistical analysis over all voxels in a sector to study how the polarisation measurement $\hat{\mathbf{s}}$ transitions from one voxel to its neighbour, either horizontally or vertically. Because voxel locations are provided by the VIPP, the corresponding image coordinates (m_v, n_v) and (m_{v+1}, n_{v+1}) on the Stokes map $\hat{\mathbf{s}}(m, n)$ are known. We can therefore interpolate along the straight line segment between these two locations to obtain a trace that describes how the measured Stokes vector varies between adjacent voxels. Repeating this procedure for every adjacent voxel pair in the sector, and averaging the traces according to their corresponding ground-truth symbols α and β , an average trace for each ordered pair (α, β) can be obtained, given by

$$\hat{\mathbf{s}}_{\alpha\beta}(w) = \left\langle \hat{\mathbf{s}}(m = (1-w)m_v + wm_{v+1}, n = (1-w)n_v + wn_{v+1}) \right\rangle_{x_v=\alpha, x_{v+1}=\beta}, \quad (7.33)$$

where $w \in [0, 1]$ parametrises the spatial position along the segment. Plotting these averaged traces in the NLSP reveals how clusters in the constellation diagram (Figure 7.6) are typically connected. This is conceptually close to the example on mixing of polarisation states in Section 2.3, but here w represents a spatial interpolation parameter rather than a mixing ratio between two pure states.

Figure 7.7 shows these traces in the NLSP, for horizontally (a) and vertically (b) adjacent voxels, obtained from the image discussed in Section 7.4.1. The traces are colour-coded using the 8-cyclic symbol difference between the two ground-truth symbols, defined as

$$\Delta_8(\alpha, \beta) := (\beta - \alpha) \pmod{8}. \quad (7.34)$$

For both horizontal and vertical cases, when the birefringent axes of adjacent voxels are perpendicular to each other, the connected trace are roughly straight lines going through the origin. As discussed previously, the absence of DoLP signal between the corresponding voxel images may come purely from the mixing of polarisation signals, instead of indicating a physical separation of the voxels. Nevertheless, when examining the vertical traces for adjacent symbols

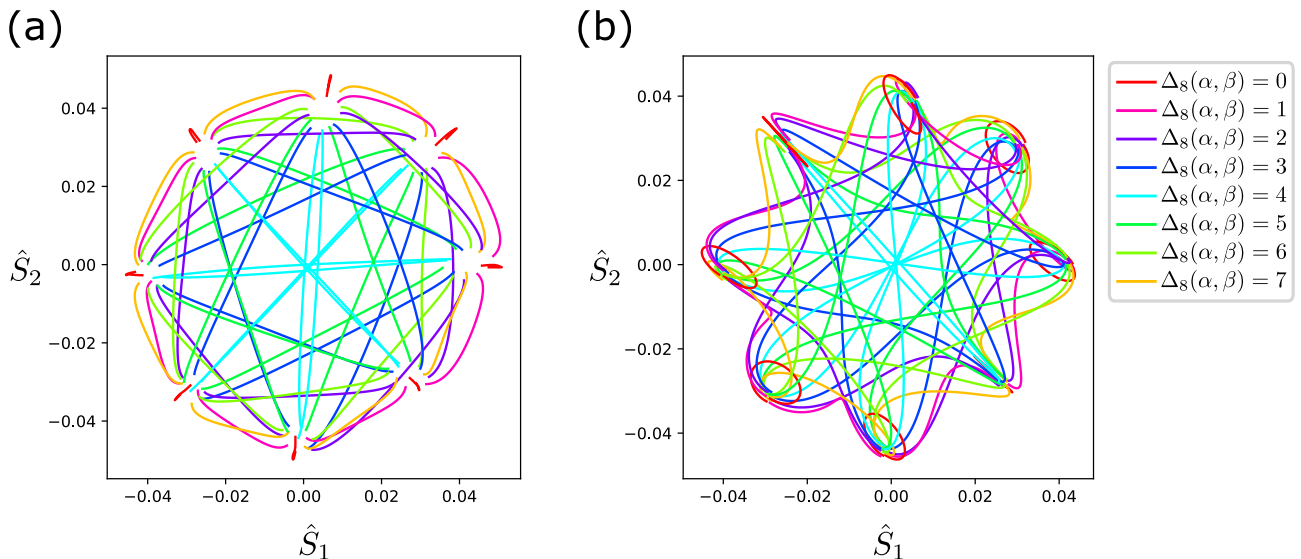


Figure 7.7: Average traces of the measured Stokes vector between adjacent voxels in a sector separated by $0.7\ \mu\text{m}$, plotted in the NLSP. For each neighbouring voxel pair $(v, v + 1)$, the Stokes map $\hat{\mathbf{s}}(m, n)$ is sampled along the straight line segment connecting their localised coordinates (m_v, n_v) and (m_{v+1}, n_{v+1}) obtained from the VIPP, and the resulting traces are averaged according to the ground-truth symbol pair (α, β) for the voxels. (a) Horizontal traces, where $v + 1$ is the voxel immediately to the right of v , corresponding to an increment in the affine-grid index p . (b) Vertical traces, where $v + 1$ is the voxel immediately below v , corresponding to an increment in the affine-grid index q . Traces are colour-coded by the cyclic symbol difference $\Delta_8(\alpha, \beta)$ to highlight how spatial mixing depends on symbol separation.

(i.e. $|\Delta_8(\alpha, \beta)| = 1$), there is clearly a strong reduction in the DoLP signal. The mixing of polarisation states in itself cannot explain the shape of the trace, because that should produce traces that either bend outwards for a coherent mixture, or straight-line traces for an incoherent mixture. The fact that these $|\Delta_8(\alpha, \beta)| = 1$ lines bends inwards is clear evidence that the voxels at the pitch of $0.7\ \mu\text{m}$ can be resolved.

The horizontal traces, in contrast, are distinctively different for $|\Delta_8(\alpha, \beta)| = 1$ that no reduction in DoLP can be observed. Every two clusters in the diagram of horizontal traces are connected roughly by straight lines, which is consistent with the mixing of polarisation signals. This difference between the horizontal and vertical traces cannot be explained in terms of the imaging process, as the PPC should have a uniform resolving power in all lateral directions. Instead, this strongly implies a write-side effect, and likely arises due to the direction of laser scanning. It is likely that the voxels are in fact continuous in the horizontal direction, only with a periodic modulation in the birefringent azimuth corresponding to the ground truth.

It should be noted that the measured traces are affected by the finite pixel size of the sensor,

or more generally by its pixel response function [83]. As a result, each sampled pixel value represents a spatially averaged signal rather than a perfect point, which can reduce the measured polarisation contrast and potentially lead to an underestimation of the DoLP. Nonetheless, this averaging effect is expected to act similarly in the horizontal and vertical directions, and therefore does not explain the strong difference between the two sets of traces.

Another important feature of interest is where a trace starts and ends. In both cases, all traces starts and ends approximately at a maximal DoLP, indicating a successful voxel localisation in both directions. Nevertheless, the endpoints of the traces for horizontally adjacent voxels separate from each other, which cannot be observed for the vertically adjacent voxels. As the endpoints represent the actual polarisation measurement sampled for MI estimation, this indicates an intra-layer crosstalk in the horizontal direction, because the polarisation measurement depends on the symbol of the adjacent voxel. This is, again, a write-side artefact, as it happens preferably in the horizontal direction, and contributes to the spread of clusters in the constellation diagram in Figure 7.6(c).

To investigate the intra-layer crosstalk further, an analysis of the ‘spread’ of voxel information is carried out, using biased grids. Starting with the best-estimate grid from the VIPP, a translational bias of (Δ_1, Δ_2) is applied to shift the grid to a slightly different location, where Δ_i lies in the direction of \mathbf{v}_i in the affine grid and is measured in the physical distance in the object space.

For each bias configuration, the Stokes measurement $\hat{\mathbf{s}}$ can then be sampled at the corresponding biased grid, from which a MI with the ground truth can then be estimated. By plotting the estimated MI with respect to the translational bias, a map can be obtained showing how a typical voxel crosstalks into its neighbouring region.

Figure 7.8(a) shows such a map, computed from the baseline dataset discussed in Section 7.4.1. As anticipated from the traces, the map shows significant difference in the horizontal and vertical directions. It can be seen that, an increased level of crosstalk happens preferably in the horizontal direction, which goes as far as the position for the neighbouring voxel (marked with white dashed lines). Conversely, in the vertical direction, the amount of information spread is

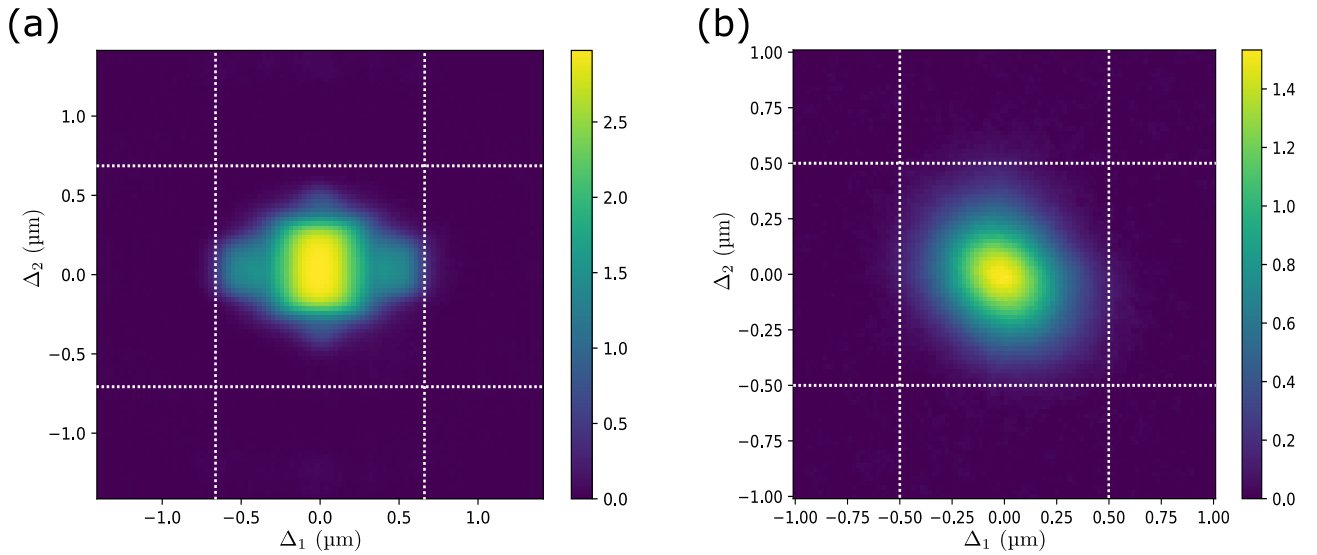


Figure 7.8: Mutual information (in bits) sampled from a biased grid. The estimated grid from VIPP was shifted in the \mathbf{v}_1 and \mathbf{v}_2 directions, by a physical distance of Δ_1 and Δ_2 , respectively. MI is then estimated by sampling the Stokes measurement $\hat{\mathbf{s}}$ with the biased grid. (a) Sector of type-X-modified 3-bpv voxels at a pitch of $0.7 \mu\text{m}$, with a maximum MI of 2.964 bpv. (b) Sector of type-S-modified 4-bpv voxels at a pitch of $0.5 \mu\text{m}$, with a maximum MI of 1.541 bpv. White lines sit one voxel pitch away from the origin.

much reduced, which is consistent with the traces shown in Figure 7.7(b). This further confirms that, due to the direction in which writing of voxels took place, an increased level of crosstalk happens between neighbour voxels, even if the read head is capable of resolving at the voxel pitch.

In contrast, Figure 7.8(b) shows the result for another dataset, which differs in both the voxel pitch ($0.5 \mu\text{m}$ instead of $0.7 \mu\text{m}$) and the type of voxel (type-S instead of type-X). The MI spread in this case has much less directionality compared to the previous example. The spread is not in a perfect round shape, which may be potentially caused by an imperfect grid estimation as the voxel pitch goes below the diffraction limit, but its width in the horizontal and the vertical directions are approximately equal. Conceptually, this can be explained with the qualitative knowledge of the sub-diffractive nature of the voxels, as introduced in Section 1.2. Instead of being made from a collection of nanopores in type-X modifications, type-S voxels are formed from a single round-shaped nanovoid [197] before being modified at its edge [37] for the anisotropic nanostructure it needs for birefringence. As a result, the MI peak obtained for this dataset sees no evident write-side artefact, and the broadening of the MI peak may be mostly due to the finite resolving power of PPC at the read-side.

Overall, the two analyses presented in Figures 7.7 and 7.8 attempts to investigate the sub-diffractive morphology of the voxels by carrying out a statistical analysis for the whole sector. Results suggest that intra-layer crosstalk may not always be determined by the optical performance of PPC at the read-side, but may also be affected by the writing process. While the readout appears capable of resolving the voxel pitch in one direction in Figure 7.7(b), the inconsistency between the traces for horizontal and vertical neighbouring voxels indicates an additional write-side contribution that breaks the expected lateral isotropy of the imaging system.

The biased-grid analysis in Figure 7.8 takes an informational perspective, visualising this effect as a spatial ‘information spread’ around the neighbouring region of each voxel. An increased level of crosstalk between neighbour voxels is observed for type-X voxels along the writing direction. In contrast, type-S voxels do not show such directionality, and the corresponding MI peak appears mostly circular, consistent with the laterally isotropic resolving power of PPC and the underlying voxel morphology.

In conclusion, it should be remembered that the MI characterisation presented in this work is not a measure of the read side alone. Though it serves a convenient metric of the statistical correlation between the ground truth X and the polarisation measurement Y , the MI is set by the full end-to-end communication channel, including write-side effects and, in theory, any degradation of the glass sample during the storage period. In [25], a lifetime of more than 10,000 years was estimated for the voxels, so the span of a few years of this PhD project is not likely to introduce any measurable effects. Nevertheless, our analysis shows that the write-side effects are not negligible, and they can contribute to MI degradation meaningfully.

7.4.4 Camera noise

On the read side, an apparent source of noise is that introduced by the camera. As discussed in Section 2.6, the PolCam operates well in the shot-noise-dominated regime, as the illumination is circularly polarised, and the birefringence introduced at the sample is weak. The use of a high-power light-emitting diode (LED) also ensures the full-well capacity (FWC) is well utilised

with a minimal exposure time, such that the amount of dark current produced is small. With these assumptions, the standard deviation of the noise at each pixel can simply be written as

$$\sigma_y = g\sqrt{N_e}, \quad (7.35)$$

where σ_y denotes the pixel standard deviation in digital numbers (DNs), g representing the gain, and N_e being the number of photo-electrons. The fewer photo-electrons there are, the lower the signal-to-noise ratio (SNR) becomes for each pixel, such that the computed polarisation measurement $\hat{\mathbf{s}}$ becomes unreliable. The propagation of this uncertainty has been discussed in detail in Section 2.7.

In this section, instead, we take an informational perspective, to investigate how the camera noise (in particular, the shot noise), affects the MI. The experiment was carried out by taking 100 frames of the same sector, and averaging a varying number of frames to provide to VIPP for informational analysis. The sector consists of 3-bpv voxels, and the dataset in use is actually the same as 7.4.1. The more frames from which an image is averaged from, the more photo-electrons are included to produce the image, which can be calculated using the measurement of camera gain in Table 2.5. To prevent voxel-localisation errors from contaminating the MI estimation, the voxel coordinates were determined once from the average of all 100 frames and then reused for every VIPP session, regardless of how many frames were averaged to form the raw image in that session.

In addition, simulations at different noise level were carried out to provide a comparison to the experimental results. In this case, the only channel noise considered was the camera shot noise, so an image-based simulation using the pipeline in Chapter 5 was deemed not necessary because the imaging process itself is irrelevant. Instead, a Monte–Carlo simulation was performed, by generating a noisy constellation diagram by adding shot noise to ideal polarisation channel measurements.

In a 3-bpv scheme, the Stokes measurements in a noise-free constellation diagram are expected

to be

$$\hat{\mathbf{s}} = \begin{bmatrix} \text{DoLP} \times \cos(2\text{AoLP}) \\ \text{DoLP} \times \sin(2\text{AoLP}) \end{bmatrix}, \quad (7.36)$$

where a typical value from the experiment can be used for DoLP, and AoLP can be modelled as $\text{AoLP} = \frac{\pi x}{4}$. As a global phase of the constellation diagram is irrelevant to MI estimation, any offset of the AoLP with respect to the birefringent azimuth can be safely ignored. For any polarisation channel whose transmission axis sits at an orientation of γ , the intensity at each polarisation channel can then be obtained by considering the first row of the Müller matrix of a linear polariser, resulting in

$$I_p(\gamma) = \frac{S_0}{2} (1 + \hat{s}_1 \cos 2\gamma + \hat{s}_2 \sin 2\gamma), \quad (7.37)$$

where S_0 represents the intensity of light before the MPA on the PolCam. The average number of photo-electrons in each polarisation channel, $N_p(\gamma)$, can then be obtained as

$$N_p(\gamma) = \frac{N_e}{2} \{1 + \text{DoLP} \cos [2(\gamma - \text{AoLP})]\}. \quad (7.38)$$

To introduce shot noise to the simulation, Poisson statistics were applied to the number of photo-electrons in each polarisation channel with a standard deviation of $\sqrt{N_p(\gamma)}$, and their corresponding grey value could then be obtained with a multiplication with the camera gain. Finally, a noisy Stokes measurement $\hat{\mathbf{s}}$ could be obtained using the PCIP module. With a large number of realisations, a noisy dataset of $\hat{\mathbf{s}}$ could be obtained, whose MI with their ground truth could be estimated.

Figure 7.9 summarises how the estimated MI depends on the mean number of photo-electrons N_e per pixel, using both experimental data and Monte-Carlo simulations. For the experimental dataset, four estimators are shown for comparison of both MI estimating method (binning or k -NN), and the representation of the channel output (i.e. using either AoLP alone or the 2D Stokes measurement $\hat{\mathbf{s}}$). All estimators exhibit the same qualitative behaviour, with MI increasing rapidly at low N_e and then approaching an asymptote as shot noise is reduced.




Figure 7.9: Effect of PolCam shot noise on the estimated mutual information for a sector of 3 bpv voxels. Experimental results are compared with Monte–Carlo simulations assuming the camera is operating at the shot-noise-dominated regime, estimated using the k -NN method from simulated $\hat{\mathbf{s}}$ measurements. Four estimators are shown for the experimental data, corresponding to binning and k -NN estimators, each applied to either the scalar AoLP or the 2D Stokes measurement $\hat{\mathbf{s}}$ as the channel output.

A clear discrepancy is observed for the binning estimator when $Y = \hat{\mathbf{s}}$, which produces consistently larger MI values than the other three estimators, and also converges to a higher value at large N_e . This behaviour is consistent with the known positive bias of the estimator, especially when many bins are poorly populated. The effect is especially relevant when the 2D vector $\hat{\mathbf{s}}$ is regarded as the channel output, because a 2D histogram contains substantially more poorly populated bins for the same number of samples. For this reason, the binning estimator applied to $\hat{\mathbf{s}}$ is deemed not reliable in the present analysis.

The remaining three estimators agree closely across the full range of N_e . Among them, the k -NN estimator using $\hat{\mathbf{s}}$ yields a slightly higher value than those based on AoLP alone. This is expected if the Stokes vector carries weak additional information beyond the angle, for example through a small residual dependence of DoLP on the ground-truth symbol. In this dataset, it is very possible that this dependence arises from a small background polarisation introduced

by the circular polariser, which breaks the rotational symmetry in the NLSP. In practice, this suggests that using $\hat{\mathbf{s}}$ as the channel output is advantageous when paired with the k -NN estimator that is not significantly biased at 2D. Therefore, the k -NN estimator on $\hat{\mathbf{s}}$ is adopted as the default estimator in this work.

Comparing experiments with Monte–Carlo simulations, the curves follow the same overall trend, while the simulations consistently predict a higher MI. This difference is expected, because the simulations isolate shot noise out by construction, whereas the experimental measurements include additional channel noises through the writing and the imaging process. In the limit of sufficient photo-electrons, the Monte–Carlo result converges at 3 bpv as expected, while the experimental curves flatten at a slightly smaller value due to other channel noises.

Finally, Figure 7.9 also illustrates a diminishing return when averaging multiple frames from decreasing gradients of the curves with respect to N_e . In addition, it may not make sense to always obtain enough frames and try to eliminate the effect of shot noise as much as possible, which takes a longer time. On the other side, if the level of shot noise is high in a working system, a higher amount of redundancy is then required at channel coding, to compensate for a lower channel capacity. This implies a trade-off between the data density and the read throughput, which is ultimately a system-design problem. For the purpose of characterising the intrinsic performance of the read head, however, it is beneficial to always capture a sufficient number of frames so that the contribution of shot noise is minimised, such that other factors can be isolated more clearly.

7.5 Axial performance and multi-layer reading

For efficient ODS, it is critical that the sectors can be stacked in the vertical direction into ‘tracks’ of many layers. In a well-optimised system, a glass of a few millimetres can contain hundreds of layers of voxel arrays, enabling a much higher data density of more than 1 Gbit mm⁻³ [24, 25]. Therefore, it is critical to understand the axial performance of using PPC as a read head, and to identify what potential challenges there are.

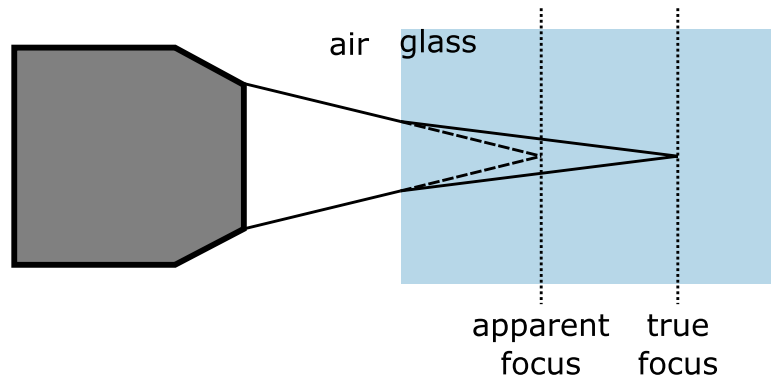


Figure 7.10: Ray diagram illustrating the axial shift of focus under a mismatch between the refractive indices of the glass (n_{glass}) and the immersion medium ($n_{\text{air}} = 1$). As a result, upon a stage movement, the shift of the true focus is scaled by a factor of n_{glass} .

7.5.1 Baseline demonstration of multi-layer reading

As a baseline, we attempt to read a track of 10 layers at an interlayer spacing of $20\ \mu\text{m}$, starting from $150\ \mu\text{m}$ below the glass surface. In each layer, the sector contains 162×162 3-bpv voxels at a pitch of $0.7\ \mu\text{m}$. From Section 7.4.1, it has been demonstrated that such a sector configuration can be read almost perfectly (at an MI of 2.964 bpv) in a single-layer setting, so any degradation from that should be a result of axial effects. In this section, the same set of ground-truth symbols is shared across all layers, due to the sample availabilities. Although it is not ideal for demonstrating an increased volumetric data density, it is convenient for analysis, since it allows a single MI to be estimated using the same ground truth.

Thanks to the motorised focus drive (MFD) built in the PPC system, the glass sample can be moved reliably in the vertical direction at a sub-micron precision. A z -stack for a range of $240\ \mu\text{m}$ was imaged at a step of $0.5\ \mu\text{m}$ for the stage.

It should be noted that, due to the refraction at the glass-air interface when light exits, a focus-shifting effect exists in the axial direction, which is illustrated in Figure 7.10. Using the Snell's law, it can be shown that any axial movement of the stage is amplified by a factor determined by the ratio of the refractive indices of the sample (n_{glass}) and the immersion medium ($n_{\text{air}} = 1$). As a result, the range of $240\ \mu\text{m}$ by the MFD translates to the true focus travelling by $360\ \mu\text{m}$, at a step of about $0.75\ \mu\text{m}$.

At each z -position, 100 frames were captured to produce an average that is practically free from




Figure 7.11: Mutual information (in bits) plotted against the depth of true focus from the top surface of the glass. The track contains 10 layers of voxel sectors, from a depth of 150 μm to 330 μm , with an interlayer spacing of 20 μm .

camera shot noise, before the resulting image is provided to VIPP for informational analysis. All VIPP sessions uses the same voxel localisation result obtained from the nearest in-focus layer. Otherwise, many VIPP sessions would fail due to the difficulty in estimating a voxel grid. The in-focus z -positions were accurately determined by finding the frames where the sum of coefficients of variation [198] across polarisation channels is maximised, given as

$$\sum_p \frac{\sigma_p^2}{\mu_p^2}, \quad (7.39)$$

where σ_p and μ_p denote the standard deviation and mean of a polarisation channel p over all pixels.

Figure 7.11 shows the result of the estimated MI from each z -position, plotted against the depth of the true focus from the top of the glass. Ten dominant peaks of more than 1.5 bpv are clearly shown, with a corrected separation of 20 μm between each other, showing the viability of using PPC for reading multi-layer voxel samples in this configuration.

Across the layers, the top layer at a depth of 150 μm provides the highest MI of 2.554 bpv. Beyond this depth, the MI generally decreases with increasing depth until the 8th layer at 270 μm , where the MI is 1.766 bpv. This overall reduction is consistent with increasing uncorrected aberrations when imaging deeper into the glass. The current implementation of PPC cannot automate the objective correction collar, so it was fixed during acquisition. In this dataset, the collar was set for a depth of 150 μm . As the focus moves deeper, depth-dependent aberrations, primarily the spherical aberration, reduce the resolving power of the readout and therefore reduce the estimated MI. The MI peaks also broaden and become increasingly asymmetric with depth, which is consistent with an increasing separation between the paraxial and marginal foci along the axial direction.

Spherical aberration alone does not explain all features in Figure 7.11. Interlayer crosstalk is another significant factor, because the reduction in MI near the top layer cannot be attributed to depth-induced aberrations when the collar is correctly set for that depth. Interlayer crosstalk is also suggested by the increase in MI for the last two layers. In the absence of crosstalk, one would expect the deepest layer to have the lowest MI because it experiences the strongest spherical aberration.

A further feature is the presence of smaller MI side lobes, which are most apparent before the first layer and after the last layer. This behaviour is likely to be a read-side effect as well, which prompted the experiments and modelling of how MI side lobes may be produced from a nearby, defocused layer.

7.5.2 Effect of defocus for single-layer voxel samples

Experimental results

To gain a deeper understanding of the axial performance of PPC, especially the emergence of side lobes in the MI plot, it is helpful to first investigate what happens when a single voxel layer is imaged under defocus. Therefore, an experiment was carried out in which a single layer of voxels was imaged while scanning in the axial direction. Similar to the multi-layer experiment




Figure 7.12: Mutual information from a single layer of voxels plotted against defocus distance. Significant side lobes can be seen on both sides. The voxels are type-X modifications using a 4 bpv scheme at a pitch of $0.5\ \mu\text{m}$.

discussed above, frames at each z -position were averaged to minimise the effect of camera noise. The in-focus position was then determined, from which voxel locations were estimated with VIPP. The averaged image at each z -position was then processed through VIPP using the same set of voxel coordinates. This enabled estimation of MI under low-SNR conditions that would otherwise be impractical when the defocus is large.

Figure 7.12 shows the estimated MI as a function of the stage defocus for a single layer of type-X voxels at a pitch of $0.5\ \mu\text{m}$ and a symbol entropy of 4 bpv. Significant side lobes are clearly visible on both sides of the in-focus position. These side lobes plausibly correspond to the smaller peaks above the first layer and below the last layer in the multi-layer results.

There is also a degree of asymmetry between the two side lobes, which is likely due to residual spherical aberration in the system. In practice, manual setting of the correction collar is challenging, because the collar is so sensitive that even a small rotation produces a large change in the amount of spherical-aberration correction. Through repeated experiments, it was found that the shapes, positions, and relative heights of these side lobes are all very sensitive to the exact correction-collar setting.

McCUTCHEEN'S THEOREM FOR AXIAL RESPONSE

To understand the origin of these side lobes, it is helpful to relate them to the axial characteristics of the imaging process. In conventional imaging methods where the system is either coherent or incoherent and the measurements are non-polarimetric, point-spread functions (PSFs) are well-defined to describe the response of the system to a point source. In the axial direction, the field and intensity distributions can be described analytically using McCutchen's theorem [199]. For a system with rotationally symmetric pupils like the PPC, the pupil function P becomes a one-dimensional function of the radial coordinate r , and McCutchen's theorem takes the form of a one-dimensional Fourier transform. For a bright-field objective with a pupil radius R , the pupil function is a top-hat function, and the normalised axial field distribution takes the form of a sinc function, given by

$$U_{bf}(d) = -i \exp \left[-\frac{i\pi l(R)d}{\lambda} \right] \operatorname{sinc} \left[\frac{\pi l(R)d}{\lambda} \right], \quad (7.40)$$

where $l(r)$ denotes the cosine component of a pupil coordinate r projected to the spherical principal plane, which can be expressed as

$$l(r) = \sqrt{1 - \frac{r^2}{f_o^2}}, \quad (7.41)$$

where f_o denotes the focal length of the objective.

A similar expression can be obtained for an annular pupil function. If the annulus has an inner radius R_i and an outer radius R_o , the axial field distribution is given by

$$U_a(d) = -i \exp \left[-\frac{i\pi [l(R_i) + l(R_o)] d}{\lambda} \right] \operatorname{sinc} \left[\frac{\pi [l(R_o) - l(R_i)] d}{\lambda} \right]. \quad (7.42)$$

The equations indicate a Fourier relation between the cosine-adjusted width of the pupil function and the axial extent of the amplitude PSF. As the annular aperture becomes thinner, the axial PSF broadens, which extends the axial range over which the system has an optical response.




Figure 7.13: Axial distribution of amplitude PSF for a phase-contrast objective computed with McCutchen’s theorem, compared with bright field. The field amplitude is logarithmically scaled and normalised for visualisation.

This provides a plausible semi-quantitative explanation for the side lobes in PPC due to the thin phase ring in the phase-contrast objective. Effectively, the pupil function in PPC can be written as the difference between a bright-field pupil function P_{bf} and that of a thin annulus P_a , given by

$$P_{ppc} = P_{bf} - (1 - it_p)P_a, \quad (7.43)$$

where it_p denotes the transmission coefficient of the phase ring that includes both phase and amplitude modulation. By linearity of the Fourier transform, the amplitude response of a phase-contrast objective can then be written as

$$U_{ppc} = U_{bf} - (1 - it_p)U_a. \quad (7.44)$$

Figure 7.13 compares the axial amplitude PSF for a phase-contrast objective shown in blue and a bright-field objective shown in orange. At small defocus, the two objectives behave similarly, due to the same numerical aperture (NA) of 0.6 for both objectives. This is consistent with

the well-known relationship between depth of focus and NA [80], given by

$$\text{Depth of focus} \approx \frac{\lambda}{\text{NA}^2} \approx 1.5 \mu\text{m}. \quad (7.45)$$

At larger defocus, the two axial profiles start to deviate. Both profiles remain oscillatory with a similar period. The phase-contrast amplitude PSF becomes considerably higher than the bright-field one over a range of tens of micrometres.

Strictly speaking, this analysis based on McCutchen's theorem is not directly applicable to PPC. As discussed in Chapter 3, a more complete imaging model of PPC depends on both the illumination side and the detection side, with the resulting STF given by a convolution of the two. McCutchen's theorem takes the Fourier transform of the objective pupil only, and it does not include illumination-side optics, which indicates that this approach is not the full picture for a partially coherent imaging method like PPC. In addition, the vector nature of the polarimetric measurements in PPC complicates the interpretation, as the concept of a scalar PSF cannot be trivially translated. Nevertheless, McCutchen's theorem provides useful intuition and points to the phase ring as a physical origin of the side lobes. It also predicts a higher optical response for a phase-contrast objective over a range of tens of micrometres, which is sufficient for a semi-quantitative explanation.

Simulation result and the thick-voxel model

To recreate the side lobes observed in the experiments, it is necessary to consider a full imaging model rather than the pupil function alone. The simulation pipeline introduced in Chapter 5 is suitable for this purpose, because it models the full optical system and properly represents partial coherence. The simulation also accounts for the finite extent of voxels, in contrast to a point-source response. Finally, the simulation pipeline supports informational analysis of the generated images using VIPP, which enables direct comparison with experimental MI results.

Batch simulations described in Section 5.2.5 were performed using the high-performance computing (HPC) facility at Imperial College London's Research Computing Service [170]. The




Figure 7.14: Comparison between experiment and simulation results on the degradation of mutual information of a single-layer sample under defocus. The single-slice simulation shows oscillatory behaviour similar to the field distribution from McCutchen’s theorem, while a multi-slice model gives a result closer to experiments. All three curves are obtained from a 3-bpv voxel sector at a pitch of $0.7\ \mu\text{m}$ and show significant side lobes at a similar relative height.

base configuration includes a single layer of voxels modelled as an array of thin, birefringent disks, following the description in Section 5.1.3. The layer sits at a depth of $150\ \mu\text{m}$ from the top surface of the glass with a refractive index of 1.5. It also lies $1850\ \mu\text{m}$ from the bottom surface, recreating the experimental conditions as closely as possible. Objective collar correction is set to cancel the depth-induced aberrations at focus, and it is kept constant for all simulation sessions. At focus, the objective front focal plane is set to a depth of $225\ \mu\text{m}$ to compensate for the optical path between the voxel layer and the glass-air interface, and it is then shifted axially to simulate defocus across the simulation sessions. For processing, each simulation session provides four polarisation-channel images that are perfectly registered without DoFP mosaicking, so that no IFOV is introduced. In addition, the true voxel coordinates are provided from the simulation pipeline to VIPP, eliminating voxel-localisation error.

Figure 7.14 shows the estimated MI from these simulations shown in blue, compared with the experimental results shown in green. The side lobes observed in the experiment were successfully

recreated by the single-layer simulation, which supports the interpretation that this is a read-side effect of the imaging process. The single-slice simulation also shows oscillatory behaviour similar to the field distribution U_{ppc} in Figure 7.13. Such high-frequency oscillations in MI were not observed experimentally.

Physically, the absence of this high-frequency feature in the experiment was likely due to the finite axial extent of the voxels. Lei et al. [32] reported that the axial extension of a voxel can range from $3.5\ \mu\text{m}$ to over $5\ \mu\text{m}$. This scale is comparable to, or larger than, the period of these axial fluctuations, which makes the thin-voxel model invalid in this regime. To include the axial extension of a voxel, a multi-slice model is introduced, where many slices of thin birefringent disks are stacked with an axial spacing much smaller than the depth of focus of the system.

The orange curve in Figure 7.14 shows the result from simulations using this multi-slice model. In this model, 20 identical slices of birefringent disks are stacked within a thickness of $4\ \mu\text{m}$. The multi-slice simulations show a low-pass filtering effect compared with the single-slice model. The overall trend of MI remains similar, while the high-frequency oscillations are suppressed.

Compared with the experiment, the multi-slice simulations agree well in general. For small defocus, the main peak has a similar width, and this width is primarily determined by the imaging process because it changes little between the single-slice and multi-slice models. The side lobes in the multi-slice simulations and the experiment are, however, centred at different axial positions. This mismatch can be caused by multiple factors. The emergence of side lobes is a consequence of phase-ring modulation at the objective, so it depends on phase-ring parameters including the radius, the width, and the phase modulation. Although the size parameters were measured experimentally using a Bertrand lens, systematic errors can arise from misalignment or magnification error. In the simulations, the phase modulation introduced by the phase ring is assumed to be an ideal quarter-wave, which is not guaranteed in practice. Any phase deviation changes the relative phase between U_{bf} and U_a , which can shift the side lobes in the z -direction. In addition, the axial MI profile in the experiment is highly sensitive to the correction-collar setting, which affects both the positions and the heights of the side lobes.

Despite the mismatch in the exact side-lobe positions, the results support the conclusion that

the side lobes observed in the multi-layer experiment arise from the phase ring extending the axial response of the imaging process. This behaviour can be recreated by the simulation pipeline. The results also show that a multi-slice model is necessary to account for the axial extension of the voxels, which enables the informational analysis of multi-layer samples using simulations in the following sections.

7.5.3 Multi-layer simulations and interlayer crosstalk

Now that the multi-slice model has been validated against experiments in reproducing the axial MI profile, including the main-peak degradation and the side lobes, the simulations could be extended to multi-layer samples. Although it is possible to simulate an arbitrary number of layers, such as the 10-layer configuration discussed in Section 7.5.1, this would be conceptually unnecessary and computationally expensive, due to the hundreds of slices needed in each simulation and the extended defocus range required.

Instead, to investigate how interlayer crosstalk affects voxel reading, a three-layer configuration was adopted, where the middle layer was expected to experience a higher level of crosstalk than the other two. Each of the three layers was a 3-bpv sector at a pitch of $0.7\ \mu\text{m}$, at a depth of $140\ \mu\text{m}$ (top), $150\ \mu\text{m}$ (middle), and $160\ \mu\text{m}$ (bottom) below the glass-air interface. In each layer, 20 slices of birefringent disks were evenly spaced across a thickness of $4\ \mu\text{m}$. In contrast to the experimental demonstration, the three layers in this case had different sets of ground-truth symbols, so three estimates of MI could be computed for each image, each corresponding to the ground truth of one layer.

Figure 7.15 plots these MI estimates against a simulated amount of defocus, with a step size of $0.2\ \mu\text{m}$. Three peaks can be clearly identified for the different layers, with the top and bottom peaks centred at $\pm 6.4\ \mu\text{m}$ of defocus from the middle peak, which can be explained by the shift of focus due to the refractive-index mismatch. In contrast to the single-layer simulations where the MI maximises close to 3 bpv, all three layers in this simulation show a significantly compromised MI, which indicates the presence of interlayer crosstalk in this configuration. The middle layer, in particular, was estimated with a MI as low as 2.634 bpv, due to the increased




Figure 7.15: Mutual information from a simulated 3-layer sample against defocus from the middle layer, with an interlayer spacing of $10\ \mu\text{m}$. Each layer represents a 3-bpv voxel sector with a different ground truth from the other two. A multi-slice model is employed for each layer, which consists of 20 slices of thin birefringent disks across an axial range of $4\ \mu\text{m}$. Significant reduction of MI can be observed for the central peak due to interlayer crosstalk.

Table 7.2: Mutual information between defocused images and the ground truth of every layer in three-layer simulations.

Interlayer spacing (μm)	MI (bpv) bottom layer	MI (bpv) middle layer	MI (bpv) top layer
10	2.769	2.634	2.776
20	2.955	2.911	2.937

amount of crosstalk from both sides. In fact, the middle layer shows an increased level of DoLP at voxel positions compared with the other two layers, which suggests a stronger polarisation signal. The fact that the MI is reduced in spite of a stronger signal further supports the interpretation that interlayer crosstalk is the dominant factor.

For comparison, another set of simulations was carried out at an interlayer spacing of $20\ \mu\text{m}$, consistent with the configuration used in the multi-layer experiment in Section 7.5.1, and the peak MI values for both sets of simulations are summarised in Table 7.2. As expected, increasing the interlayer spacing reduces crosstalk and therefore increases the MI of the middle layer to 2.911 bpv. It is worth noting that, in all simulations, the objective correction collar was set to a depth of $150\ \mu\text{m}$ for the middle layer, so the top and bottom layers in the larger-spacing setting




Figure 7.16: Mutual information for a simulated 2-layer sample at varying interlayer spacings. Both layers contain 3-bpv voxel sectors at a pitch of $0.7\ \mu\text{m}$, with a thickness of $4\ \mu\text{m}$ simulated using 20 evenly spaced slices of birefringent disks. In each session, the stationary layer at $150\ \mu\text{m}$ is in focus with perfect correction of depth-induced aberrations, while a moving layer contaminates the image of the stationary layer from a varying axial distance. Each simulated image is compared against the ground truth of both (a) the stationary layer and (b) the moving layer.

experienced a higher level of spherical aberration. The fact that their MI increases despite a higher level of aberration indicates that crosstalk is more important at this distance scale.

To further understand how crosstalk is affected by the interlayer spacing, another set of simulations was carried out involving two layers of sectors in the multi-slice model. A stationary layer sat at a depth of $150\ \mu\text{m}$ and was in focus in all sessions, while a moving layer sat at a different distance from the stationary layer in each session. As the moving layer got closer, the image of the stationary layer became contaminated, and the MI could be estimated with respect to the ground truth of both the stationary layer and the moving layer. For comparison, the bright-field setting was also simulated under the same conditions, where a clear aperture was used instead of the annular aperture on the illumination side, and a normal objective was used instead of a phase-contrast objective.

Figure 7.16 shows the results of VIPP processing on the simulated images, evaluated against the ground truth of (a) the stationary layer and (b) the moving layer. As the moving layer got closer, the MI for the stationary layer decreased substantially from 3 bpv, while that for the

moving layer increased sharply, which is direct evidence of interlayer crosstalk. The PPC curve for the moving layer followed its single-layer profile in Figure 7.14, apart from an axial scaling factor of 1.5 due to the shift of focus caused by the refractive-index mismatch.

Although a side lobe is observed for the moving layer under PPC, no significant features were found for the stationary layer when the interlayer spacing increased beyond 20 μm . From an imaging perspective, this could be understood as a reduced contribution of the polarisation signal from the moving layer when it is considerably out of focus, which was outweighed by the signal from the in-focus stationary layer. From an informational perspective, it is a consequence of underutilisation of the communication channel, since the source entropy $H(X)$ is likely smaller than the true channel capacity C at large interlayer spacing, resulting in the lack of features for the stationary layer. In practice, other sources of channel noise would limit the channel capacity, and the capacity itself may decrease under optimisation of channel parameters such as the voxel pitch. It is therefore unnecessary to increase the source entropy dramatically to the true channel capacity of the simulations, since such a high SNR is practically impossible. Accordingly, the MI degradation caused by long-range interlayer crosstalk associated with side lobes is likely less significant than other sources of channel noise in realistic conditions.

Nevertheless, the short-range interlayer crosstalk is significant, as seen in Figure 7.16(a). As the MI drops well below the source entropy of 3 bpv, short-range interlayer crosstalk clearly limits the channel capacity, which ultimately limits the achievable data density. Compared with the bright-field configuration, PPC appears to suffer a greater level of short-range interlayer crosstalk at the same NA, which is likely due to inferior optical sectioning characteristic to Zernike's phase contrast (ZPC) methods, as reported by Munro and Török [200]. Nevertheless, the enhancement of polarisation signal provided by PPC remains critical, enabling voxel reading in the first place. In practice, when a bright-field configuration was used, voxel localisation became unreliable or failed due to the low SNR.

This evaluation also highlights the advantage of using a simulation pipeline whose results are reliable and validated against experiment in a single-layer setting, since these conditions are difficult to create experimentally. Firstly, in practice, a voxel layer cannot physically move, so

such an experiment would require a different track for each interlayer spacing. This required a specific sample to be written and was also prone to writing inconsistencies across tracks, for example due to depth-induced aberrations or thermal effects, and it required high precision of the written sector depth. Secondly, on the read side, it was difficult to focus on the stationary layer with a sub-micrometre accuracy, especially when the two layers were close to each other. Thirdly, experimental images could be too noisy for voxel localisation, especially in the bright-field configuration where the polarisation signal is much weaker than PPC. Even for PPC images, voxel localisation could become unreliable when layers were this close, which could lead to failed MI estimation or an uncalibrated amount of MI degradation. By providing the set of true voxel coordinates, the simulation pipeline removed voxel-localisation error from the analysis and isolated the impact of interlayer crosstalk and defocus on the estimated MI.

7.5.4 Depth-induced aberrations

As observed in the baseline demonstration of multi-layer voxel reading in Section 7.5.1, another observable source of channel noise is the depth-induced optical aberrations. These aberrations can lead to both an asymmetric axial profile and a reduced maximum MI. The wavefront aberration introduced through this process is detailed in Section 5.1.4, from which the uncorrected aberration can be written as

$$W_{uncorrected}(\rho) = W(d_g, n_g, \rho) - [W(d_c, n_c, \rho) - W(d_c, 1, \rho)], \quad (7.46)$$

where d_g denotes the sector depth in glass, n_g denotes the refractive index of glass, d_c represents the collar setting for depth correction, n_c represents the design refractive index for which the correction is specified, ρ is the radial coordinate of the normalised pupil, and W denotes the wavefront aberration due to refractive-index mismatch, given in Equation 5.21. In general, n_g and n_c may not perfectly match, which leaves a residual aberration that cannot be removed by the collar setting alone and can in principle limit the channel capacity.

In practice, this residual aberration can be further reduced by refocusing. Introducing an

additional defocus d is equivalent to adding an extra axial distance in air, and the expression becomes

$$W_{uncorrected}(\rho) = W(d_g, n_g, \rho) - [W(d_c, n_c, \rho) - W(d_c + d, 1, \rho)], \quad (7.47)$$

which can be optimised with respect to both d and d_c . Using $n_g = 1.461$ for fused silica at a wavelength of 530 nm [201, 202] and $n_c = 1.515$ as a representative design refractive index for cover-glass correction [77, 203], the optimisation predicts a Strehl ratio above 0.9994 for depths up to 2 mm when both the correction collar and refocusing are set optimally. Here the Strehl ratio is defined as [204]

$$S := \frac{1}{\pi^2} \left| \int_0^1 e^{iW(\rho)} 2\pi\rho d\rho \right|^2. \quad (7.48)$$

These results indicate that, under an idealised model where the collar setting and refocus can be adjusted accurately, depth-induced aberrations can be reduced to a negligible level over the whole thickness of the glass. With motorised collar correction [25], depth-induced aberrations can in principle be compensated with the collar setting and the stage moving in coordination, and the residual amount of depth-induced aberration is mostly limited by engineering constraints such as the precision of the collar setting and its calibration with respect to the vertical position of the stage, rather than posing a fundamental informational limit in the system. A detailed experimental characterisation of MI as a function of uncorrected depth was not pursued here, because it would require single-layer sectors written at multiple depths, and because manual collar adjustment is highly sensitive and difficult to reproduce consistently across sessions, as discussed before. A similarly comprehensive simulation study is also possible within the existing pipeline, but it was not prioritised because the results in the previous subsection indicate that interlayer crosstalk is a more important limitation in the configurations considered. In future work, a combined study using motorised collar control and corresponding simulations could quantify the residual impact of depth-induced aberrations on MI under realistic calibration errors and practical constraints. Alternatively, depth-induced optical aberrations may be reduced using a remote-refocusing approach, where an aberration-free optical copy of the sample is formed in a remote space and then reimaged [205]. This provides another potential route for imaging at depth into the glass without relying on the

correction collar of the objective, which can be of interest for future work.

7.6 Conclusion

This chapter studied the performance of PPC on voxel imaging from an informational perspective, motivated by the need to evaluate voxel readout for the application of ODS. In contrast to image-based analysis, this chapter aims to quantify how much information the PPC read head preserves from the written voxel symbols. Using Shannon’s linear model, a communication channel is identified where the input is the ground-truth voxel symbol X and the output is the polarisation measurement \hat{s} in the NLSP. The estimated MI between them can then serve as a practical, decoder-agnostic metric that captures the information available for reliable decoding thanks to Shannon’s noisy-channel coding theorem, which also sets a lower bound for the channel capacity.

The chapter then established the VIPP, a dedicated image-processing pipeline for informational analysis on voxel images, consisting of both image-level processing and voxel-level characterisation. Specifically, the PolCam images were firstly calibrated for PFC and handled by PCIP, where IFoV error is effectively corrected for using a Fourier-based technique. The resultant DoLP image is then provided to a voxel localisation algorithm, which isolates a voxel sector using morphological operations and affine transform, before a finite affine grid is found to a sub-pixel level using CZT, according to which polarisation measurements are resampled. For MI estimation, both binning and k -NN-based estimators were considered, and the default choice in this work is the k -NN estimator operating on the 2D \hat{s} measurement, as the binning estimator was found to be significantly biased in certain situations.

With this framework and pipeline in place, the chapter then demonstrated that voxel reading with PPC is already strong in both single-layer and multi-layer settings. In the single-layer demonstration, a sector of 162×162 voxels at a pitch of $0.7 \mu\text{m}$ and a depth of $150 \mu\text{m}$ produced well-separated clusters for a 3-bpv scheme. The estimated MI for this dataset is 2.964 bpv, which is close to the input entropy of 3 bpv, indicating that most of the information encoded in the

voxel symbols are effectively preserved throughout the communication process.

The single-layer results also quantified trade-offs that govern achievable data density. Across voxel pitches from $0.9\ \mu\text{m}$ to $0.45\ \mu\text{m}$, the estimated MI decreases as the pitch is reduced, consistent with increasing crosstalk and increasing difficulty in voxel localisation at small pitches. At a pitch of $0.8\ \mu\text{m}$ with a 4-bpv scheme, a high MI of 3.795 bpv was achieved. At a pitch of $0.5\ \mu\text{m}$, the MI is reduced to 1.541 bpv for a 4-bpv scheme, while the data density reaches $6.16\ \text{Mbit mm}^{-2}$ per layer. This illustrates that maximum per-layer density does not necessarily coincide with maximum MI per voxel, and that higher density can be achieved at the cost of the more challenging coding and decoding.

The analyses of adjacent-voxel traces and biased-grid MI maps further showed that intra-layer crosstalk can be anisotropic and can depend strongly on write-side effects. For the baseline 3-bpv dataset at $0.7\ \mu\text{m}$ pitch, the traces and the MI spread indicate stronger crosstalk along the writing direction. In contrast, a type-S dataset at $0.5\ \mu\text{m}$ pitch showed a more laterally symmetric MI spread, consistent with a reduced write-side directional artefact. These results reinforce that $\text{MI}(X, Y)$ characterises the full end-to-end channel rather than the read head alone, which is essential when interpreting information loss in experimental data.

The shot-noise study quantified how the available information depends on the photon budget. An experimental series was obtained by capturing 100 frames of the same sector and averaging a varying number of frames, while reusing voxel coordinates determined from the 100-frame average to avoid localisation failures at low SNR. Across the range of mean photo-electrons per pixel, MI increases rapidly at low photon budgets and then approaches an asymptote, showing diminishing returns for further frame averaging once shot noise is sufficiently suppressed. Monte-Carlo simulations that include shot noise alone converge to 3 bpv for a 3-bpv scheme, while the experimental curves flatten below 3 bpv, consistent with additional sources of channel noise besides camera noise.

The chapter then extended the analysis to axial performance and multi-layer reading, motivated by the requirement that practical ODS should operate with many stacked layers. As a baseline, a track of 10 layers at an interlayer spacing of $20\ \mu\text{m}$ was imaged from a nominal depth of $150\ \mu\text{m}$

down to $330\ \mu\text{m}$. A z -stack was acquired over $240\ \mu\text{m}$ of stage travel with a step of $0.5\ \mu\text{m}$, which corresponds to a $0.75\ \mu\text{m}$ of true focus due to refractive-index mismatch. Distinct MI peaks were observed for all layers, demonstrating that stacked readout is feasible within the same pipeline used for single-layer evaluation. The highest MI peak in this dataset is $2.554\ \text{bpv}$ at the top layer near $150\ \mu\text{m}$, and the MI generally decreases with depth to $1.766\ \text{bpv}$ at $270\ \mu\text{m}$. The effects of side lobes, interlayer crosstalk, and depth-induced optical aberrations were all observable from the multi-layer profile, prompting detailed analysis and modelling.

Firstly, defocus experiments on single layers showed a main MI peak together with pronounced side lobes, consistent with the multi-layer profile. McCutchen's theorem was used to compute the axial field distributions from pupil functions, which points to the phase ring as a plausible origin of an extended axial response, although the theorem was not directly applicable to PPC due to PPC being a partially coherent method. Full simulations, in contrast, take the whole partially coherent imaging process into account, and successfully reproduced the main features of the measured MI profiles and showed that a multi-slice voxel model is needed to represent finite voxel thickness. In the simulations, 20 slices over $4\ \mu\text{m}$ suppress high-frequency oscillations present in thin-voxel simulations and yield MI profiles that agree better with experiments.

Secondly, as for the interlayer crosstalk, three-layer simulations of thick voxels show MI of $2.634\ \text{bpv}$ for the middle layer at an interlayer spacing of $10\ \mu\text{m}$, which is, despite less depth-induced optical aberrations and a higher measured DoLP, lower than the MI for the top and bottom layers, indicating the information loss due to interlayer crosstalk. At a spacing of $20\ \mu\text{m}$, all three layers see an increased amount of MI, indicating that the increased spacing can recover information even when depth-induced aberrations are increased for the top and bottom layers, placing interlayer crosstalk as a major source of channel noise. A two-layer spacing sweep further showed that, as a moving layer approaches an in-focus stationary layer axially, the MI evaluated against the stationary ground truth decreases while the MI evaluated against the moving ground truth increases, providing direct evidence of crosstalk between these layers. The two-layer simulations also reveal that the long-range features like the side lobes are unlikely to cause an information loss as significantly as the short-range interlayer crosstalk does, and the poor optical-sectioning capability of ZPC methods puts PPC at a worse position

than bright-field configurations, although the enhancement of polarisation signal enables voxel reading in PPC in the first place.

Lastly, depth-induced aberrations such as the spherical aberration were discussed as another potential sources of channel noise. Nevertheless, analysis with wavefront aberration shows that any depth-induced optical aberrations can be effectively corrected with a combination of objective collar adjustment and refocusing, predicting a Strehl ratio of over 0.9994 at a depth up to 2 mm where the system is effectively aberration-free. It was therefore concluded that it is possible to practically eliminate any MI degradation with a well-calibrated motorised collar correction, and depth-induced optical aberration does not place a fundamental informational limit in voxel reading. Potential future work includes the implementation of such automation, which also enables precise control of collar coordinate that can be compared with simulation results using the pipelines. In addition, remote-refocusing offers a potential alternative where such aberrations may be removed without the need of a depth-correcting objective.

Overall, this chapter established an informational framework for characterising voxel imaging as part of a communication process, together with a practical pipeline for extracting per-voxel measurements from experimental images and estimating MI for the communication channel. Importantly, the voxel-reading capability of PPC and the VIPP pipeline was demonstrated for samples with both single and multiple layers, while MI provided a decoder-agnostic metric for understanding how factors such as voxel pitch, noise, defocus, and crosstalk limit the information available at the measurement stage. This perspective is valuable because it separates readout quality from the particular choice of symbol determination and coding, and it provides a principled basis for system-level characterisation with respect to these potential sources of channel noise. Nevertheless, an end-to-end storage demonstration ultimately requires a decoding process that maps from measurements $Y = \hat{s}$ to estimated symbols \hat{X} , so that the recovered data can be evaluated directly against the written ground truth. The next chapter therefore takes a more practical approach by proposing a machine-learning decoder based on a multilayer perceptron (MLP) network, together with data preparation and polarisation augmentation strategies. By extending the present channel characterisation to explicit symbol decoding, the thesis provides a complete end-to-end demonstration from voxel symbols to po-

larisation measurements and back to recovered symbols, supporting the viability of ODS based on birefringent voxels and PPC readout.

Chapter 8

Decoding voxel data using machine learning

In the previous chapter, an informational framework was established to evaluate the performance of polarisation imaging for voxel readout. By using mutual information (MI) as a decoder-agnostic metric, factors such as camera noise and defocus were studied to characterise polarisation-sensitive phase contrast (PPC) as a read head for optical data storage (ODS). Shannon’s noisy-channel coding theorem implies that the information quantified by the estimated MI can, in principle, be recovered with an appropriate coding and decoding strategy. This chapter goes beyond that existence statement by introducing a practical decoder and demonstrating that symbol information is indeed recovered from experimental images.

With the decoder added to the readout pipeline, the aim of this chapter is therefore an end-to-end demonstration of ODS with birefringent voxels. Within Project Silica, sector-level decoding has been demonstrated using a multi-scale convolutional neural network (CNN) [6, 25], where the input is images of an entire sector and the output is an array of symbol predictions for all voxels within it. This approach can be robust because the network can learn to account for factors such as intra-layer crosstalk and sector geometry, and it does not require explicit voxel localisation. At the same time, it typically requires a large amount of training data because the input is high dimensional and the model complexity is correspondingly higher.

In this work, we take a different approach and focus on voxel-wise decoding using a shallow neural network enabled by voxel image processing pipeline (VIPP). Given voxel coordinates and polarimetric measurements estimated by VIPP, the decoder operates at a voxel level, predicting the symbol of a single voxel from local pixels around its estimated position. This formulation reduces the input dimensionality and allows training with a substantially smaller dataset. The results in this chapter therefore test whether a shallow neural network is sufficient for voxel decoding when the polarisation imaging and voxel localisation are reliable, establishing a foundation from which more advanced decoders can be developed in future work.

8.1 Multilayer perceptron decoder

From the perspective of machine learning (ML), decoding can be interpreted as a supervised classification task, where polarimetric measurements are provided as inputs and a probability distribution over symbols is predicted. In contrast to the informational analysis where the Stokes images are resampled at voxel positions, the neural-network-based decoder can accept a higher-dimensional input than typical MI estimators, enabling the use of polarimetric measurements extracted from multiple pixels in the local neighbourhood of a voxel. For each estimated voxel coordinate, a 16×16 pixel region is extracted from the vicinity of the voxel for all three Stokes parameters. By using Stokes measurements at multiple nearby pixels, the decoder is expected to be more robust to voxel localisation errors and intra-layer crosstalk. The input to the network is therefore a $(3, 16, 16)$ array, and the output is a set of probabilities over all symbols.

By imaging a sector of interest, the Stokes images $S_0(m, n)$ to $S_2(m, n)$ and the voxel coordinates (m_v, n_v) are provided by VIPP, as introduced in the previous chapter. For the example sector used as a baseline demonstration in Section 7.4.1, 26 244 voxels are available from one imaging session. With 8 possible symbols, voxels from a single sector are usually sufficient for both training and testing a shallow neural network using an 80-20 split between training and test sets. The ground-truth symbol x_v for each voxel is used as the target on the training set, or




Figure 8.1: False-colour visualisations of example Stokes images provided to the decoder as training dataset. Labelled by the ground-truth symbols (A–H), each false-colour image contains a voxel at the centre, with different crosstalk from adjacent voxels. Hue, AoLP; value, DoLP.

provides the reference for evaluating predictions on the test set. False-colour representation of some example images in the training dataset is shown in Figure 8.1, labelled against their ground truth (A–H). From the images, it can be seen that a voxel of a similar colour corresponding to their ground truth is visible at the centre of each image, while adjacent voxels crosstalk differently in these images. The Stokes images, therefore, provide information both at the voxel location itself and the surrounding context.

Unlike many image classification tasks, voxel classification in this work has a relatively small dimensionality and does not require translational invariance, due to the voxel localisation provided by VIPP. This motivates the use of a small and shallow model such as the multilayer perceptron (MLP) [206], whose architecture is illustrated in Figure 8.2. The $(3, 16, 16)$ input array is flattened into a vector before being passed as input. Two hidden layers are used in this feedforward neural network, and nodes in adjacent layers are connected by a fully-connected layer (FCL) [206] followed by a rectified linear unit (ReLU) activation function [207]. Therefore,




Figure 8.2: Architecture of the multilayer perceptron decoder. The feedforward network contains three fully-connected layers with learnable weights, taking flattened Stokes images around each voxel as the input, and gives probabilities for every symbol as the output.

a node h'_j in the next layer can be computed from those in the previous layer h_i , according to

$$h'_j = \max \left(\sum_{i=1} w_{ij} h_i + w_{0j}, 0 \right), \quad (8.1)$$

where w_{ij} denotes the weights in the FCL. Through training, the model learns these weights from the ground-truth labels. Nonlinearity is introduced by the ReLU activation functions to the affine transforms at the FCLs, which is critical in enabling non-trivial decision boundaries of the decoder.

Feedforward networks of this form have been well studied [208–210], showing that such networks can represent an arbitrary decision boundary to an arbitrary accuracy [206], given sufficient network width. The so-called universal approximators are therefore, in principle, capable of serving as the voxel decoder, in spite of the simplicity of the architecture.

8.2 Training and decoding results

The decoder network was implemented and trained with PyTorch [211], with layer sizes of 768 for the input, 512 and 128 for the two hidden layers, and an output size of 8 for 3-bits per voxel (bpv) voxels or 16 for 4-bpv voxels. For each input voxel image, the model outputs a vector

of logits z_x [212], where x denotes the voxel symbol. The logits are related to the predicted probabilities $q(x)$ by a softmax function, given by

$$q(x) = \frac{e^{z_x}}{\sum_{x'} e^{z_{x'}}}. \quad (8.2)$$

Training proceeded in epochs, where each epoch corresponds to a full pass through the training set. Within an epoch, the training data was iterated in mini-batches, and the model performed a forward pass to produce logits for each voxel. Using cross-entropy loss, the loss for a mini-batch of B samples with ground-truth labels x_b is

$$L = -\frac{1}{B} \sum_{b=1}^B \log q_b(x_b), \quad (8.3)$$

where $q_b(x_b)$ denotes the predicted probability of the correct class for sample b . A backward propagation was then performed to compute gradients of the loss with respect to the weights in the FCLs, and a momentum stochastic gradient descent (SGD) optimiser [213] updated the parameters accordingly. Through multiple epochs, the model parameters were gradually adjusted to minimise the average loss over the training set. Prediction accuracy was recorded after each epoch by assigning each voxel to the symbol with the largest predicted probability $q(x)$, and computing the fraction of voxels whose predicted symbols matched the ground truth.

As a proof of viability, the network was trained on a sector image of 162×162 voxels in a 3-bpv scheme, with a voxel pitch of $0.7 \mu\text{m}$ in both directions. As demonstrated in Section 7.4.1, such a sector was estimated with a MI of 2.964 bpv, implying near-perfect decoding should be possible. For this sector, each voxel was assigned one of eight possible symbols drawn independently from a discrete uniform distribution. No line coding, run-length-limited coding, or error-correction coding was applied before writing the sample. Although no significant local clustering was observed in this particular realisation, such clusters or long runs of similar symbols could occur by random chance, in a sequence of sufficient length containing independently generated symbols. In a practical storage system, sequence-level coding schemes could be important, because local clusters of similar symbols can cause the local average of polarisation measurements to




Figure 8.3: Accuracy of voxel-symbol prediction against number of epochs in an MLP network, trained on the sector image of a 3 bpv voxel sector containing 26 244 voxels at a pitch of 0.7 μm . A final test accuracy of 0.9996 was reached for the test set.

significantly deviate from zero, introducing systematic errors to background estimations. In addition, it may bias the polarisation measurements of nearby voxels and neighbouring layers. Coding schemes such as the run-length-limited coding [214, 215] could therefore be helpful, by suppressing long runs or clusters of similar symbols, although imposing such constraints may reduce the effective information density by restricting the set of allowed symbol sequences. The analysis reported in this thesis should therefore be interpreted for this uncoded, independently and uniformly distributed symbol sequences.

Figure 8.3 shows the prediction accuracy against the number of epochs, for both the training and test datasets. From the plot, it is evident that the MLP learns the decoding process quickly, and within 5 epochs, a prediction accuracy of 0.9996 was reached. This proves the capability of a shallow MLP network in decoding on a voxel level.

Besides prediction accuracy, the decoding performance can be assessed with MI, offering comparison with the VIPP-estimated results. Instead of computing the MI between the input symbol X and the polarisation measurement $Y = \hat{\mathbf{s}}$, MI is now computed between X and the estimated symbol, \hat{X} , given as

$$\text{MI}(X, \hat{X}) = \sum_x \sum_{\hat{x}} p_{X, \hat{X}}(x, \hat{x}) \log \frac{p_{X, \hat{X}}(x, \hat{x})}{p_X(x) p_{\hat{X}}(\hat{x})}. \quad (8.4)$$




Figure 8.4: First-layer weights of a trained MLP network for decoding of 3-bpv voxels. For the i -th input node and the j -th node in the first FCL, the quantity $\max_j w_{ij}$ is computed and reshaped into a $3 \times 16 \times 16$ array, corresponding to the original Stokes images for a voxel.

For the test set after 5 epochs, out of the ground-truth entropy of 2.999 bpv, the MI was evaluated to be 2.995 bpv. Note that this number is even higher than the 2.964 bpv quoted for $\text{MI}(X, Y = \hat{s})$, which would not be expected if \hat{X} were obtained solely by classifying the single-point measurement \hat{s} , since no additional information can be created by post-processing. Instead, because \hat{X} is decoded from $3 \times 16 \times 16$ pixels from the Stokes images, more information is provided to the MLP network than a single Stokes measurement in the case of VIPP.

To visualise which input pixels the trained model relied on most strongly, the weights of the first FCL were inspected. For each input node i , the maximum weight over the first hidden layer, $\max_j w_{ij}$, was computed and then reshaped into a $3 \times 16 \times 16$ array corresponding to the original Stokes images.

The resulting maps are shown in Figure 8.4. The trained network relies mainly on pixels in the central regions of the S_1 and S_2 images. In contrast, the S_0 image carries limited information about the voxel symbol, since its variation is dominated by the phase-contrast signal arising from residual phase changes of the voxels. Intensity variations can also arise from transmission loss introduced by voxel modifications, but this effect is expected to be very small for the type-X voxels used here [37]. The central regions of S_1 and S_2 , by comparison, contain the polarimetric features needed to distinguish symbols. As discussed in Section 7.4.3, the polarisation signal extends over a finite area, due to a combination of the finite modified volume from the write side, and the resolving power of the microscope from the read side. By using multiple pixels

within this area, the model is expected to be more robust to camera noise, voxel-localisation error, and crosstalk from neighbouring voxels, offering a plausible explanation to the increase in MI.

In digital communications, the reliability of a noisy channel is often measured by the bit error rate (BER) [216], defined as the number of bit errors divided by the total number of received bits. It should be noted that the BER is not simply the complement of the symbol accuracy, because it depends on the mapping between symbols and bits. In this work with birefringent voxels and PPC reading, voxel symbols are primarily distinguished through angle of linear polarisation (AoLP) when the degree of linear polarisation (DoLP) signal is sufficiently high, so symbol errors are most likely to occur between neighbouring symbols. It is therefore natural, when assigning bit labels to these symbols, to minimise the bit difference between neighbouring states. Therefore, the Gray code [182], as shown in Table 8.1, is used as the symbol-to-bit mapping in this work.

In the 3-bpv scheme, the Gray code ensures that adjacent symbols differ by only one binary digit (i.e. a Hamming distance [217] of one). Therefore, for the present system in which symbol errors are most likely to occur between neighbouring AoLP-encoded states, a symbol error typically corresponds to one bit error under this labelling. In general, the specific mapping between symbols and bits could be optimised according to the joint probability $p_{X,\hat{X}}$, but for symmetric channels and with the number of possible symbols being an exponential of 2, the Gray code is usually a sensible choice.

It is important to emphasise that this symbol labelling does not change the underlying symbol accuracies, the symbol confusion matrix, or the information throughput of the communication channel. The apparent reduction from symbol error rate (SER) to BER should therefore not be interpreted as an improvement in the readout process or as an increase in the amount of information conveyed. This is the reason why the information analysis in Chapter 7 was based on MI, which is independent of the actual bit labels assigned to the symbols.

Assuming Gray coding, the BER was estimated to be $(1.30 \pm 0.92) \times 10^{-4}$ from the test set, which is approximately one third of the original SER, since in this dataset symbol confusion typically

Table 8.1: Voxel symbols and their corresponding Gray code in a 3-bpv scheme. Adjacent symbols are assigned bit labels with a Hamming distance of 1.

Voxel symbol	3-bit Gray code
A	000
B	001
C	011
D	010
E	110
F	111
G	101
H	100

occurs at a Hamming distance of 1. This value is therefore a useful bit-level description of the measured symbol errors under the specified labelling, rather than a separate measure of the information throughput. In a 3-bpv scheme and a voxel pitch of $0.7\ \mu\text{m}$, the BER corresponds to a decoded data density of $6.12\ \text{Mbit mm}^{-2}$ per layer.

8.3 Polarisation-based image augmentation

Now that successful decoding had been demonstrated using a 3-bpv image, a natural next step was to attempt decoding a more challenging sector, such as the $0.5\ \mu\text{m}$ -pitch sector presented in Figure 7.8(b). As the voxel pitch decreases below the diffraction limit, there is inevitably crosstalk between neighbouring voxels, and the ML-based decoder can potentially help by taking into account the surrounding context. A practical challenge, however, was the insufficiency of training data from a single sector, due to the smaller number of voxels present in the sector, and the increased number of possible symbols each voxel represents.

To see how the insufficiency may arise, it is helpful to consider a simple argument based on local symbol configurations. If the task is defined as classifying the centre voxel while being robust to nearest-neighbour crosstalk from the four adjacent voxels, the local configuration contains five voxels in total, as shown in Figure 8.5(a). Results presented in the previous section were obtained from a 3-bpv sector, for which the total number of possible local configurations is $8^5 = 32\,768$. The available training set was 80% of a 162×162 voxel sector, resulting in

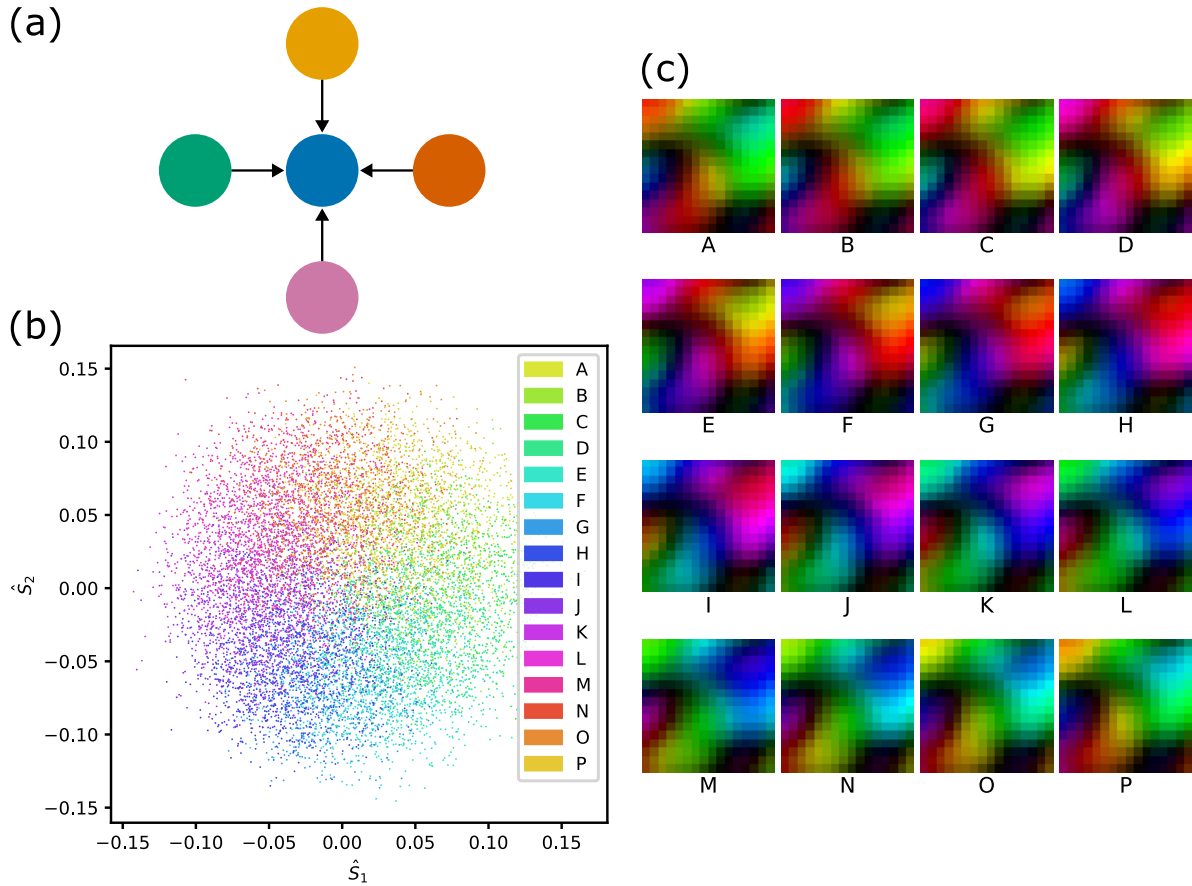


Figure 8.5: Polarisation-based image augmentation for training on 4-bpv sector of 114×108 voxels, with a voxel pitch of $0.5 \mu\text{m}$. (a) Schematic showing intra-layer crosstalk comes mostly from the four closest neighbouring voxels. (b) A nearly-circular constellation diagram produced by the VIPP showing the channel is highly symmetric. (c) False-colour images of the original image (A) and the 15 augmented images (B–P) by rotating the Stokes vectors in the normalised linear Stokes plane (NLSP) by 22.5° each time. Hue, AoLP; value, DoLP.

20995 training voxels. In contrast, the more challenging sector uses a 4-bpv scheme with $16^5 = 1\,048\,576$ possible local configurations, and the sector contains only 114×108 voxels. After an 80 to 20 split, this leads to a smaller training set containing 9850 voxels. As a result, the training set can cover at most 0.94% of all possible configurations, whereas in the previous section it could cover about 64%. This makes the training set for the challenging sector inevitably sparse, which can limit the generalisability of the model to unseen configurations.

It should be noted that this argument is only approximate, to illustrate why the training dataset is likely to be insufficient, as a model can generalise without observing every configuration during training. Nevertheless, the sharp increase in the number of possible local configurations, together with the more limited experimental data available for the sector, makes data insuffi-

ciency a real concern. This motivates the use of image augmentation techniques to increase the effective size of the training set.

In Section 7.2, it was noted that the communication channel is symmetric, due to the uniform polarisation contrast of PPC across all AoLP. In particular, it was noted that a rotation in the normalised linear Stokes plane (NLSP) corresponds to a shift in voxel symbol, i.e.

$$x \mapsto x + 1 \quad \implies \quad \hat{\mathbf{s}} \mapsto \mathbb{R}(\pi/N_X) \hat{\mathbf{s}}, \quad (8.5)$$

where \mathbb{R} denotes a 2D rotator matrix, and N_X denotes the number of possible symbols for a voxel. Implicit assumption of this symmetry has been made through the uniform distribution of the ground-truth symbols, and is supported with the rotational symmetry observed in the constellation diagram provided by the VIPP, as shown in Figure 8.5(b). By exploiting this symmetry, images can be augmented by rotating the Stokes measurement $\hat{\mathbf{s}}$ in the NLSP for all pixels by 22.5° while labelling the ground truth as the next symbol, as shown in the false-colour images in Figure 8.5(c). Repeating this operation for each training image therefore increases the effective training set size by a factor of 16.

For comparison, the same MLP model was trained with and without image augmentation, and the resulting accuracies are plotted against the number of epochs in Figure 8.6. In both cases, the test accuracy improves rapidly over the first tens of epochs and then increases more slowly towards an asymptote well below unity. Meanwhile, the training accuracy continues to increase, suggesting that the model fits characteristics of the training set that do not generalise to the test set.

Relative to the 3-bpv result in Figure 8.3 where training and test accuracies converge close to unity within five epochs, training on the 4-bpv sector is slower and converges to a lower final accuracy. The slower convergence and degraded final performance on the 4-bpv sector illustrate the increased difficulty of the decoding task and the limited amount of training data available for this dataset.

Despite these challenges, image augmentation provides a clear improvement. At any given




Figure 8.6: Prediction accuracy for a 4-bpv sector at a voxel pitch of $0.5 \mu\text{m}$, trained with datasets with and without polarisation-based image augmentation.

number of epochs, the model trained with augmentation achieves a higher test accuracy, and it also converges to a higher final accuracy compared to the same decoder trained without augmentation. At 89 epochs, the augmented model reaches a test accuracy of 0.565, corresponding to a symbol MI of 2.15 bpv, a significant improvement from the 1.34 bpv evaluated by VIPP with the Stokes measurement at a single point.

The improvement obtained after decoding is a clear indication of the effectiveness of a ML-based decoder in exploiting the local polarisation context for symbol classification. This advantage is particularly relevant as the voxel pitch decreases below the diffraction limit, where crosstalk becomes more severe. At the same time, the MI computed from decoded symbols, $\text{MI}(X, \hat{X})$, is intrinsically dependent on the choice and performance of the decoder. Small changes in the model, the training strategy, or dataset can change the estimation significantly even when the PPC measurements remain unchanged. This sensitivity further justifies the use of $\text{MI}(X, \hat{s})$ in the previous chapter, where the informational analysis was performed on the Stokes measurements themselves for a decoder-agnostic characterisation of PPC as a read head.

In addition, assuming 4-bit Gray code, the corresponding BER is 0.135 ± 0.004 , giving a decoded data density of $13.85(6) \text{ Mbit mm}^{-2}$ per layer. This decoded data density is higher than that obtained for the $0.7 \mu\text{m}$ 3-bpv sector despite the worse BER, which shows that a higher data density can be achieved even when decoding is less reliable.

8.4 Conclusion and future work

This chapter completes an end-to-end demonstration of ODS with birefringent voxels spanning voxel writing, PPC readout, VIPP-based image processing and voxel localisation, and ML-based decoding. Together with the decoder-agnostic informational analysis in the previous chapter, these results establish both that information is recoverable in principle and that symbol information is recovered in practice from experimental images.

For the 3-bpv sector at a voxel pitch of $0.7\ \mu\text{m}$, a shallow MLP network achieved classification of single-voxel images with a test accuracy of 0.9996 after 5 epochs. The decoded MI reached 2.995 bpv out of a ground-truth entropy of 2.999 bpv, indicating near-lossless recovery of symbol information. Assuming Gray coding, the corresponding BER was evaluated as $(1.30 \pm 0.92) \times 10^{-4}$, giving a decoded data density of $6.12\ \text{Mbit mm}^{-2}$ per layer. This decoded data density can be regarded as a system-level metric for the whole ODS process.

Decoding was further demonstrated in a more challenging regime using the 4-bpv sector at a voxel pitch of $0.5\ \mu\text{m}$, where intra-layer crosstalk is stronger and the experimental training set from a single sector is inevitably sparse. By exploiting the rotational symmetry of the channel, polarisation-based image augmentation was developed to increase the size of the training set and hence improve the decoding performance. Assuming 4-bit Gray coding, the augmented MLP achieved a decoded data density of $(13.85 \pm 0.06)\ \text{Mbit mm}^{-2}$ per layer, despite a worse BER than the 3-bpv experiment.

As future possibilities, the voxel-wise decoding strategy can be extended in several ways. Firstly, the current model is trained and tested with data extracted from a single polarisation camera (PolCam) image, which limits generalisation. By including images across different sectors, depths, and imaging sessions, the model is expected to be more robust to variations in these factors. Alternatively, training data can be synthesised with the simulation pipeline presented in Chapter 5. Practical factors can then be introduced to the simulated images through augmentation, for example by adding camera noise consistent with the noise model illustrated in Section 2.6, or by applying a varying amount of translation to model voxel localisation errors.

A second direction is to incorporate a larger spatial context for the model by expanding the neighbourhood used by the decoder beyond the 16×16 pixel region, which can potentially help when intra-layer crosstalk extends beyond the four nearest neighbours. Nevertheless, this also requires a much larger training set, either acquired experimentally or through simulations, due to the exponentially increasing number of possible configurations. In addition, a much larger context is likely to provide diminishing benefit in crosstalk correction and will also require a more complex network.

Thirdly, there are possibilities in exploring alternative architectures for the decoder. In a preliminary comparison on the 4-bpv dataset, a voxel-wise CNN reduced the BER from 0.135 ± 0.004 to 0.123 ± 0.004 , indicating that further optimisation of the network architecture or width may provide improved performance.

In addition to predicting voxel symbols, a decoder can alternatively be trained to predict the bits represented by each voxel directly, as a more direct objective for optimising the BER [25]. Bit prediction requires the definition of a coding strategy, which can also go beyond the Gray code and be optimised for the particular communication channel.

Overall, this chapter establishes ML-based voxel-wise decoder as a viable constituent of an end-to-end ODS process, and it motivates work on generalisation across sectors and imaging conditions, robustness to practical factors, and the use of a larger spatial context or alternative architectures as future possibilities.

Chapter 9

Clear-aperture polarisation imaging with a diattenuative module

In the previous chapters, polarisation-sensitive phase contrast (PPC) microscopy was developed and employed for readout of birefringent voxels for optical data storage (ODS). As a read head, PPC has been shown to provide reliable reading at voxel pitches close to the diffraction limit, before read-side crosstalk becomes dominant. The implementation of PPC can be realised relatively simply through modification of a standard Zernike's phase contrast (ZPC) system, by introducing a circular polariser at the illumination side, and replacing the camera with a polarisation camera (PolCam), providing a major advantage in practical systems.

Nevertheless, closer examination of the system is required to make further improvements on the PPC setup, in areas such as enhancing the resolution, minimising the interlayer crosstalk, and biological applications. The underlying ZPC configuration was originally developed to image phase signals rather than polarisation signals. As a result, some components that are critical for phase imaging are not necessarily ideal when the primary goal becomes the detection of weak polarisation signals. In PPC, the key role of the phase-contrast objective is to attenuate the S wave relative to the D wave, thereby enhancing polarisation contrast. Nonetheless, the associated phase modulation does not provide additional benefit and can introduce artefacts such as halo and shade-off [69, 124]. This motivates an alternative implementation in which

the required S-wave attenuation is achieved without a phase ring, which also removes the need for an annular aperture.

Firstly, lateral resolution is affected by the limited illumination numerical aperture (NA) due to the annular aperture, as the diffraction limit [80] is given by $\lambda/(\text{NA}_{\text{illu}} + \text{NA}_{\text{obj}})$. In addition, a ZPC setup requires objectives with an integrated phase ring, which reduces the range of suitable stock lenses. When combined with practical constraints such as working distance, depth correction range, immersion media, and NA, identifying an appropriate objective can become challenging.

Secondly, the ZPC configuration can lead to a degraded axial performance. While this thesis does not include experimental measurements on voxel tracks with sufficiently small interlayer spacings to directly observe the effect, simulations in Section 7.5.3 predict increased interlayer crosstalk for PPC compared with a bright-field configuration for interlayer spacings below approximately 15 μm . For practical ODS applications, degraded axial performance would limit the optimal interlayer spacing and therefore the achievable data density [25].

In addition, the imaging of weak polarisation signal is not limited to ODS applications, but also of interest in biological imaging of anisotropic structures which can also exhibit birefringence, such as collagen fibres, actin filaments, and mitotic microtubules [52, 102, 218, 219]. In biological imaging, the mixing between phase and polarisation contrast can be problematic, because a much stronger phase signal can be found in many biological samples [75, 220] compared to the weak residual phase of type-X and type-S modifications. While PPC is sensitive to both phase and polarisation signals, the former is translated to the average across polarisation channels in a PolCam, while the latter translates to the difference between the channels. As a result, the phase signal usually overpowers the polarisation features, making PPC less ideal for observing weak polarisation signals in biological contexts.

Furthermore, the ZPC setup is known to be susceptible to image artefacts such as halo and shade-off [69, 124], due to the crosstalk between the S and D waves with a finite width of the phase ring. The problem can become especially severe if the condenser annulus and the phase ring are slightly misaligned.

A further issue for polarisation microscopy with PPC arises from the non-zero complex phase of the Stokes transfer function (STF), as formulated in Section 3.3.1. The STF phase introduces a systematic offset in the measured angle of linear polarisation (AoLP), which, in the thin-annulus limit, corresponds to the phase modulation at the objective. In practice, however, the offset can deviate from the expected $\pi/2$ radians due to the difference between the design wavelength and the illumination, as well as the finite spectral bandwidth of the source. The finite width of the phase ring also introduces crosstalk between S and D waves, producing a spatial-frequency-dependent AoLP offset that makes calibration difficult. In ODS applications, this offset is much less consequential as voxel reading is a classification problem in nature. In biological polarisation microscopy, however, an accurate AoLP measurement is often of interest because it can carry orientational information about the underlying biological structure. Lastly, removing the annular aperture can greatly improve light throughput, which is beneficial in a commercial setting.

These considerations, therefore, motivate the development of an alternative system that enhances polarisation contrast without introducing any phase modulation with a phase-contrast objective. In the following sections, a polarisation-based system will be presented that aims to image weak polarisation features in a specimen without complication from its phase. The polarisation mechanism will be described, followed by the practical implementation of the system, and the imaging results on a range of samples.

9.1 Polarisation mechanism of the diattenuative module

To start with, we shall first recall the polarisation mechanism of PPC. From Section 3.2, it was found that the unmodulated P wave can be expressed as

$$|P\rangle = |S\rangle + |D_L\rangle + |D_R\rangle, \quad (9.1)$$

in the presence of weak birefringence and a weak residual phase of the object. Besides the S wave component $|S\rangle$ for the left-handed circularly polarised background, the emergence of

the left- and right-handed D waves are caused by the phase and birefringence at the sample, respectively, given as

$$|S\rangle = |L\rangle, \quad (9.2)$$

$$|D_L\rangle = i\phi_o |L\rangle, \quad (9.3)$$

$$|D_R\rangle = -\frac{i\delta_o}{2} e^{2i\gamma_o} |R\rangle, \quad (9.4)$$

where ϕ_o denotes the residual phase of the object, δ_o denotes the birefringent retardance, and γ_o denotes its fast-axis orientation.

In ZPC, the S wave is then modulated by spatial filtering using a matched pair of condenser annulus and phase ring. As discussed in Section 3.3, this spatial filtering also results in a narrower STF, poor optical sectioning, and image artefacts such as shade-off and halo. Over the years, improvements to this scheme of spatial filtering have been developed, including the use of apodised phase plates [124] and spatial light modulators (SLMs) [125, 126], in efforts of reducing crosstalk between the S and D waves and increasing the effective illumination NA.

For polarisation imaging, however, polarisation modulation can be used in place of spatial filtering. This is because the signal of our interest is the right-handed circularly polarised component of the D wave, $|D_R\rangle$, only, as the orthogonal component $|D_L\rangle$ encodes phase instead of birefringent information. A module attenuating the left-handed circularly polarised component selectively would therefore reduce the polarisation background from the S wave while retaining the polarisation signal in $|D_R\rangle$, enhancing the polarisation contrast. Such optical behaviour is called circular diattenuation [72, 221], and can be described by a Jones matrix of

$$\mathbb{T}_{dia} = \begin{bmatrix} t_s & 0 \\ 0 & 1 \end{bmatrix}_c, \quad (9.5)$$

where $t_s \in \mathbb{R}$ denotes the amplitude transmittance for the left-circular component (corresponding to the S wave).

With a diattenuative module installed at the detection side, the Jones vector of the modulated

P wave would therefore become

$$|P'\rangle = \mathbb{T}_{dia}|P\rangle = \begin{bmatrix} t_s(1 + i\phi_o) \\ -\frac{i\delta_o}{2}e^{2i\gamma_o} \end{bmatrix}_c, \quad (9.6)$$

and the degree of linear polarisation (DoLP) therefore becomes

$$\text{DoLP} = \frac{t_s\delta_o\sqrt{1 + \phi_o^2}}{t_s^2(1 + \phi_o^2) + \delta_o^2/4}, \quad (9.7)$$

$$\approx \frac{t_s\delta_o}{t_s^2 + \delta_o^2/4}. \quad (9.8)$$

Compared to the expression of DoLP for PPC in Equation 3.19, it can be observed that the two expressions coincide in the limit of $\phi_o \rightarrow 0$ with a substitution of t_s with the phase ring attenuation t_p . In this limit, DoLP is optimised to unity when t_s (or t_p) is set to $\delta_o/2$.

Nevertheless, as discussed in Chapter 3, this optimal t_p is not realistic in PPC, because the transmittance is fixed by the design of the phase-contrast objective. Even if t_p could be adjusted, its optimum would in general depend on both ϕ_o and δ_o , because the weak phase assumption does not necessarily imply $\phi_o/t_p \ll 1$. In contrast, it is always reasonable to assume $\phi_o^2 \ll 1$, such that the DoLP for the diattenuative module becomes independent of ϕ_o , including in the regime where t_s is close to its optimum at $\delta_o/2$. This provides the diattenuative module with the potential to be optimised for imaging of weak birefringent features beyond what is achievable with PPC.

As for the practical design of the diattenuative module, it is useful to write its Jones matrix in the linear basis, because this makes it easier to identify how the module can be constructed from standard polarisation components. As discussed in Section 2.4.1, the Jones matrices in the linear and circular bases are related by

$$\mathbb{T}_{\text{lin}} = \mathbb{U}_{\text{lc}}^\dagger \mathbb{T}_{\text{circ}} \mathbb{U}_{\text{lc}}, \quad (9.9)$$

where

$$\mathbb{U}_{lc} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -i \\ 1 & i \end{bmatrix}. \quad (9.10)$$

Therefore, in the linear basis, the Jones matrix of the diattenuative module can be written as

$$\mathbb{T}_{dia} = \mathbb{U}_{lc}^\dagger \begin{bmatrix} t_s & 0 \\ 0 & 1 \end{bmatrix} \mathbb{U}_{lc}. \quad (9.11)$$

In the linear basis, the diagonal matrix $\text{diag}(t_s, 1)$ corresponds to a partial polariser. The unitary matrix \mathbb{U}_{lc} can be implemented using quarter-wave plates (QWPs) by noting that

$$\mathbb{U}_{lc} = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} \mathbb{LR} \left(\gamma = \frac{\pi}{4}, \delta = \frac{\pi}{2} \right), \quad (9.12)$$

where $\mathbb{LR}(\gamma = \pi/4, \delta = \pi/2)$ denotes a QWP with its fast axis oriented at 45° . A QWP at 45° maps the linear basis to the circular basis, up to the additional phase difference of a quarter wave.

The Jones matrix of the diattenuative module therefore becomes

$$\mathbb{T}_{dia} = \mathbb{LR} \left(\gamma = -\frac{\pi}{4}, \delta = \frac{\pi}{2} \right) \begin{bmatrix} 1 & 0 \\ 0 & -i \end{bmatrix} \begin{bmatrix} t_s & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} \mathbb{LR} \left(\gamma = \frac{\pi}{4}, \delta = \frac{\pi}{2} \right), \quad (9.13)$$

$$= \mathbb{LR} \left(\gamma = -\frac{\pi}{4}, \delta = \frac{\pi}{2} \right) \begin{bmatrix} t_s & 0 \\ 0 & 1 \end{bmatrix} \mathbb{LR} \left(\gamma = \frac{\pi}{4}, \delta = \frac{\pi}{2} \right), \quad (9.14)$$

because the additional phases from the two QWPs cancel out since the Jones matrix of the partial polariser is diagonal. As a result, the diattenuative module can be built by cascading a QWP at 45° , a partial polariser, and another QWP at -45° .

Regarding the imaging model, the analytical model presented in Chapter 3 is still applicable. In the point-source model, the diattenuative configuration shares the same STF as a bright-




Figure 9.1: Schematic of a microscopy system with the diattenuative module installed between the objective and the tube lens. The diattenuative module consists of a partial linear polariser (P2) placed between two quarter-wave plates (QWP2 and QWP3), at an offset of 45° and -45° , respectively. At the illumination side, the combination of a linear polariser (P1) and another quarter-wave plate (QWP1) is employed for circularly polarised light.

field system, up to a normalisation factor, as neither an annular aperture nor a phase ring is required. As presented in Section 3.3, the bright-field STF has a cut-off at a spatial frequency of $2NA/\lambda$, and therefore a finer diffraction limit. In addition, in the absence of phase modulation, the STF is real over all spatial frequencies, so the AoLP measurements are not offset as they are in PPC. Finally, since both illumination and detection are circularly symmetric, the STF exhibits Hermitian symmetry, leading to no transfer of phase signal to any Stokes parameter under the weak object assumption.

9.2 System implementation

The hardware implementation of the diattenuative system is a relatively simple upgrade from the PPC system discussed in Chapter 6. Based on the same microscope frame (Olympus, BX41), bright-field imaging can be achieved by simply disengaging the annular aperture of the phase-contrast condenser (Olympus, U-PCD-2). The only modification required is the installation of the diattenuative module between the objective and the tube lens, as illustrated in Figure 9.1. All other components remain identical to PPC, including high-power light-emitting diode (LED) illumination, 3D sample positioning, and the control software.

In the system, three identical QWPs (Thorlabs, WPQ10E-532) were employed, one at the illumination side to produce circularly polarised light, and the other two being part of the diattenuative module. In the ideal case, the relative orientation between the illumination side and the detection side should have no impact on the imaging process, because the illumination polarisation in the object space would be perfectly circular and therefore rotationally symmetric. Nevertheless, due to the finite bandwidth of the source, not all spectral components are modulated by an exact quarter-wave by QWP1, resulting in a small spread of the illumination polarisation. By employing identical QWPs and orienting QWP2 perpendicularly to QWP1, such spectral error can be corrected for all wavelengths in the S wave component. Using Jones algebra, it is a consequence of the following identity, given as

$$\mathbb{L}\mathbb{R}\left(\gamma = \frac{\pi}{4}, \delta\right) \mathbb{L}\mathbb{R}\left(\gamma = -\frac{\pi}{4}, \delta\right) = \mathbb{I}, \quad (9.15)$$

which holds for all values of δ .

It should also be noted that, in the limit of P2 becoming a full polariser instead of a partial one, the P1–QWP1–QWP2–P2 setup resembles a classic setup for a polarised light microscope (PLM) [222, 223] for imaging birefringent structures. Nevertheless, such a setup would completely remove the S wave from being imaged. As discussed in Section 2.4, it is the relative phase between the S and D waves that encodes AoLP information, so complete removal of the S wave would make it impossible to measure the birefringent azimuth.

A practical setup requires careful consideration on the implementation of P2. Although theoretical analysis implies its extinction ratio can be as high as possible up to $4/\delta_o^2$ for a maximal DoLP signal, practical considerations have to be taken. Firstly, the extinction ratio has to be much smaller than that of P1 (Thorlabs, LPVISE200-A), otherwise a non-negligible right-handed S wave would become comparable to the attenuated left-handed S wave. Secondly, the use of a high-extinction-ratio polariser as a partial polariser would make the system very difficult to align. For a given angular misalignment, a higher extinction ratio leads to a larger percentage increase in the transmitted S wave relative to the aligned case. Thirdly, as it is placed in the infinity space, it needs to be flat enough to minimise the wavefront error it produces.

Table 9.1: Summary of different types of polarisers with their extinction ratio and wavefront quality. Values are either cited from manufacturer specifications or measured experimentally in this work.

Type	Thorlabs Item #	Extinction ratio	RMS wavefront error (nm)	RMS phase error ($\lambda = 530$ nm)
Film polariser	LPVISE2X2	1000 [224]	215	0.406
Crystal polariser	GT10	100 000 [225]	—	0.250 [225]
WGP	WP25M-VIS	683 [226]	27	0.051
PBS (in reflection)	PBS251	20–100 [227]	16	0.031
Custom polariser	—	89	40	0.075

The combination of a low extinction ratio and a good wavefront performance makes the selection of P2 a practically difficult task. A few common types of polarisers are summarised in Table 9.1, with their quoted extinction ratio, and root-mean-square (RMS) wavefront error measured with a Zygo interferometer. Film polarisers typically quote an extinction ratio of over 1000 at 530 nm, and usually associate with a large transmitted wavefront error that introduces significant aberrations to the imaging process. Crystal polarisers have a better wavefront performance, with a distortion well below a wavelength, but their extinction ratio is usually on the order of 10^4 to 10^5 , too high for our purpose. Wired-grid polarisers (WGPs) usually have a lower extinction ratio and a transmitted wavefront error, but the reflection they produce in the $-z$ -direction can lead to a ghost image superimposed onto the intended image. Polarising beam splitters (PBSs), in contrast, can be used in the reflection mode that has an extinction ratio less than 100 [227], and it usually has a wavefront error much less than a quarter wavelength, making it a potential candidate for the partial polariser.

9.3 Results on birefringent voxel imaging

Birefringent voxel samples were used to assess the practical performance of the diattenuative system. A PBS in reflection was first used as P2, and its limitations will motivate the use of a custom polariser for an improved uniformity.

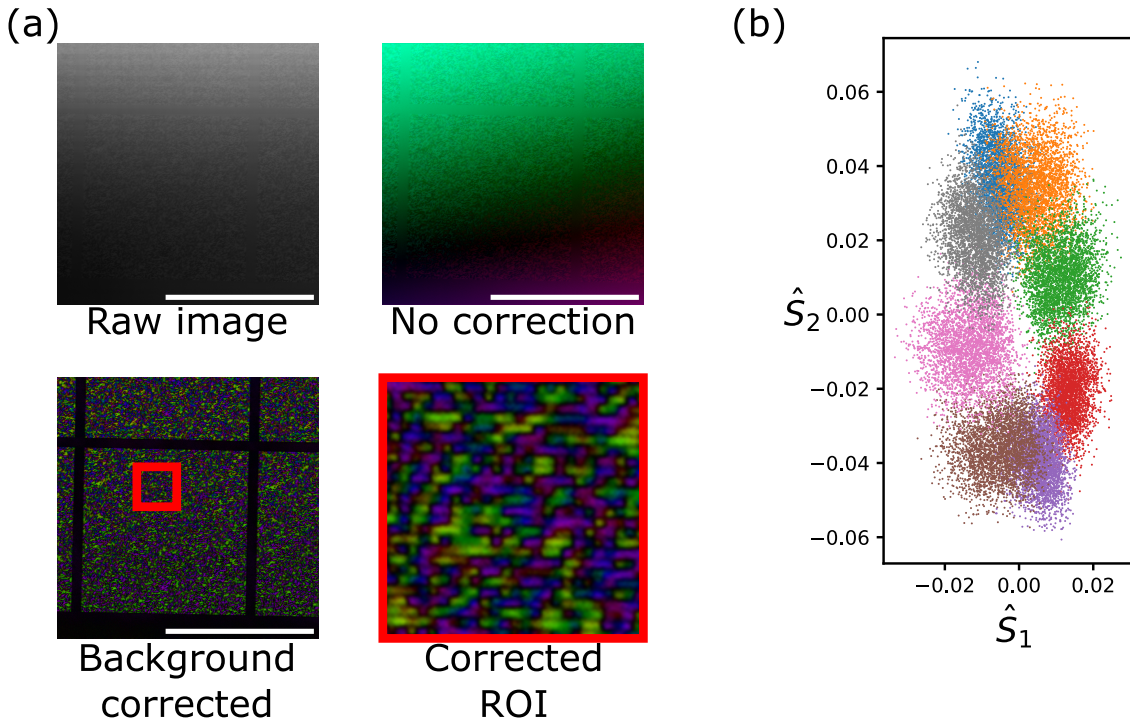


Figure 9.2: Imaging of a sector of birefringent voxels with the PBS-based diattenuative module. (a) A significant intensity gradient can be observed on the raw image, which translates to severe instantaneous field-of-view (IFoV) errors in the false-colour image without correction. With correction, the IFoV error was effectively removed, but the voxels in false colour are shown only in purple and green. Hue, AoLP; value, DoLP. Bar=100 μm . (b) Constellation diagram corresponding to the background-corrected image, showing an anisotropic polarisation contrast in the normalised linear Stokes plane.

9.3.1 Voxel imaging with a PBS-based diattenuative module

To characterise the capability of a PBS-based diattenuative module, a sector of birefringent voxels were imaged and processed through the voxel image processing pipeline (VIPP). The sector contains 162×162 voxels at a pitch of $0.7 \mu\text{m}$ and a source entropy of 3 bpv, a configuration that has been shown to be imaged and decoded reliably with PPC. Nevertheless, as can be seen from Figure 9.2(a), the raw image contains a large intensity gradient across the field-of-view (FoV), leading to a severe level of instantaneous field-of-view (IFoV) error shown as an apparent polarisation gradient in the false-colour image. With the help of polarisation-field correction (PFC), the IFoV-error correction, and the Gaussian filtering in VIPP, the IFoV error can be seen effectively corrected in the lower-left image in Figure 9.2(a). Nonetheless, the voxels are shown almost exclusively in green and purple, as opposed to the more varied colour palette one would expect.

As the DoLP signal was reliable enough to enable voxel localisation, the local Stokes measure-

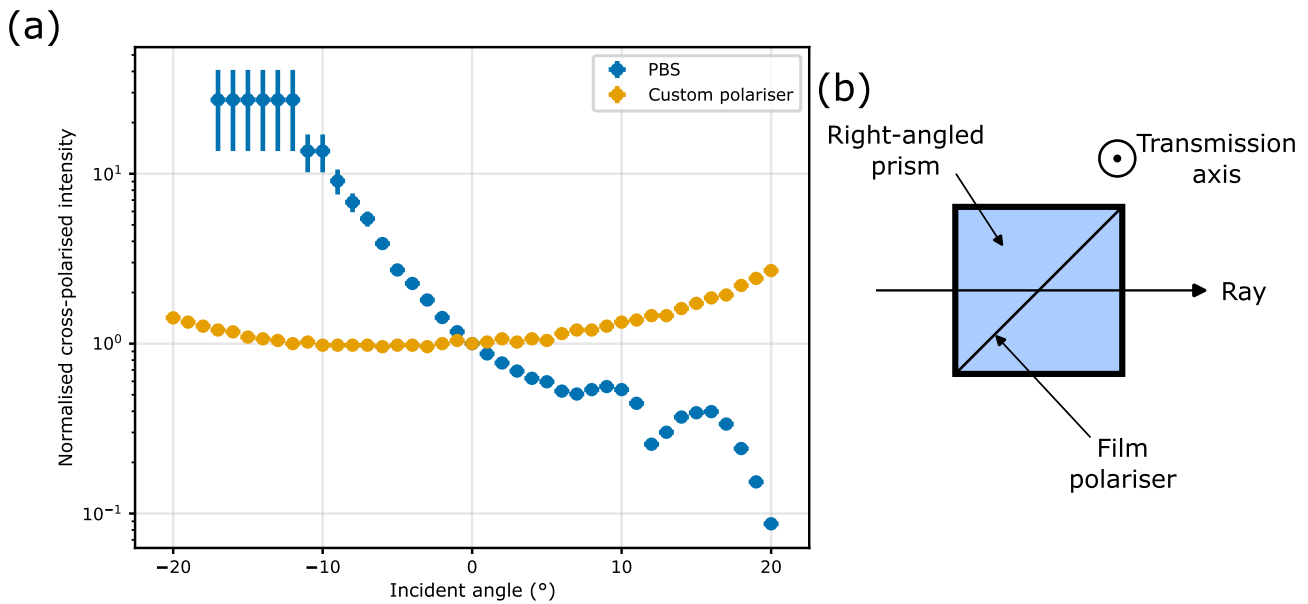


Figure 9.3: Comparison of the angular dependence of the cross-polarised transmission intensity for a PBS and a custom polariser. (a) Normalised cross-polarised intensity as a function of the incident angle. (b) Schematic of the custom polariser, built by sandwiching a film polariser between two right-angled prisms. The transmission axis is perpendicular to the page, and the configuration reduces the extinction ratio.

ments \hat{s} were plotted in the normalised linear Stokes plane (NLSP) as a constellation diagram, as shown in Figure 9.2(b). It is evident that the polarisation contrast became anisotropic, as the constellation became elliptical with a major axis roughly in the \hat{S}_2 direction.

9.3.2 Angular dependence of a PBS and its impact on polarisation contrast

A possible explanation for the anisotropic contrast enhancement from the PBS-based diattenuation module is the angular dependence of the PBS itself. As the diattenuative module sits within the infinity space between the objective and the tube lens, any angular dependence translates into non-uniformity across the FoV. Characterisation was carried out using a diode laser (Thorlabs, PL201) at 520 nm as the source, and a crystal polariser was used whose extinction ratio was assumed to be much higher than that of the PBS. The PBS was placed after the crystal polariser to analyse the field. The resultant intensity is plotted against the incident angle and shown in Figure 9.3. From the figure, it is evident that the intensity varies significantly with respect to the incident angle. Particularly, around normal incidence, the gradient

of intensity with respect to the incident angle is large, explaining the large intensity gradient across the FoV.

To understand how this intensity gradient is turned into an anisotropic polarisation contrast after background correction, consider the measured intensity in general being a function of the three Stokes parameters, i.e.

$$I_p = f_p(S_0, S_1, S_2). \quad (9.16)$$

In cases where the illumination is non-uniform, an extra factor g can be introduced to account for the illumination non-uniformity across the superpixel. When the intensity non-uniformity happens on the illumination side, all Stokes parameters arriving at the PolCam are adjusted by the factor g_p , and the intensity therefore becomes

$$I'_p = f_p(g_p S_0, g_p S_1, g_p S_2). \quad (9.17)$$

For polarisation channels with $g_p > 1$, their contribution to the relevant Stokes parameter S_1 or S_2 is therefore overestimated, and vice versa. This leads to the IFoV error depending on the distribution of g_p and the particular layout of the micro-polariser array (MPA) at the PolCam. A perfect algorithm would eliminate the IFoV error by estimating I_p from I'_p .

Nevertheless, for a diattenuative system, the non-uniformity arises from the detection side. In addition, nothing prevents the S and D waves from experiencing different modulation with respect to the incident angle. In particular, the intensity gradient observed from the PolCam is evidence that the amplitude of the S wave was angularly dependent, but when the background intensity increases due to the S wave, there is no promise that the D wave increases by the same amount. For example, suppose the intensity increase purely comes from an increased amplitude of the S wave by \sqrt{g} . The measured intensity would become

$$I'_p = f_p(g_p S_0, \sqrt{g_p} S_1, \sqrt{g_p} S_2), \quad (9.18)$$

assuming the S wave is much stronger than the D wave. As there is a fundamental ambiguity between intensity variation and the variation due to polarisation state within a superpixel, the

background correction algorithm still estimates a global correction of g for all Stokes parameters. This results in a corrected intensity of

$$I_p = f_p \left(S_0, \frac{S_1}{\sqrt{g_p}}, \frac{S_2}{\sqrt{g_p}} \right), \quad (9.19)$$

leading to an underestimation of their contribution to S_1 and S_2 . It should be noted that this operates in opposition to the IFoV error. The polarisation channels whose contribution to Stokes parameters are overestimated in the presence of IFoV error have their contribution underestimated after correction, and vice versa. This leads to a reduction in polarisation contrast for states orthogonal to the apparent polarisation state produced by the IFoV error.

Finally, it should be noted that this explanation only illustrates how an anisotropic polarisation contrast may arise from the over-compensation of background correction and IFoV-error removal. The exact mechanism depends not only on the particular combination of PFC, IFoV correction, and the Gaussian filtering, but also depends heavily on the actual angular dependence of the intensity and the extinction ratio of the PBS. Although this anisotropic contrast can in principle be calibrated through accurate measurements of the PBS, the underlying compromised noise performance cannot, as an intensity gradient naturally leads to many pixels whose full-well capacities (FWCs) are underutilised. For this image, the mutual information (MI) with the ground truth was estimated by the VIPP to be 2.517 bpv, a significant drop from the 2.964 bpv obtained from a sector in the same configuration using PPC. As the polarisation contrast is no longer isotropic, the assumption of symmetric channel breaks, so the MI could be improved if the source were channel-coded accordingly.

9.3.3 Improved voxel imaging with a custom polariser

A better option would be to build the diattenuative module with a different partial polariser that has a more uniform angular performance, and a design is presented in Figure 9.3(b). A film polariser (Edmund Optics, #86-189) was cut into a 35 mm \times 25 mm rectangle whose transmission axis aligned with the shorter side. The film polariser was then sandwiched between

a pair of right-angled prisms (Comar, 25 RQ 01) with UV cement, whose right-angled sides are coated with anti-reflective coating while the hypotenuses are not. As the film polariser is tilted around its transmission axis, the incident field sees a reduced projection of the absorbing axis, and therefore resulting in a much lower extinction ratio.

Sandwiching the film polariser between two prisms also helps, by providing refractive-index matching between the polariser and the surrounding media, reducing refraction as light enters and exits the film. This ensures that light also propagates at an angle of 45° within the polariser, which would otherwise be reduced by the Snell's law, thereby making the tilt more effective. Although the geometric interpretation is not a rigorous picture and the actual extinction ratio depends on how exactly the polariser interacts with a longitudinal component of the electric field, Korger et al. [228] showed the picture works well for realistic polarisers. In fact, the extinction ratio was measured to be 89, comparable with that of the PBS in the reflection mode.

Placing the film polariser between prisms is also critical to ensuring a minimal wavefront error. Using a Zygo interferometer, the RMS phase error was measured as 0.075, much lower than the quarter-wave criterion. In addition, when refocusing is allowed to compensate for the quadratic term in the wavefront, the RMS error can be further reduced to 0.044.

Compared to the PBS, the custom polariser only slightly compromises in wavefront quality, but the cross-polarised intensity is measured to be much less dependent on the incident angle, as shown in Figure 9.3(a). The much weaker angular dependence also translates to a much smaller intensity gradient across the FoV in voxel imaging, as shown in Figure 9.4(a). The resultant constellation diagram in Figure 9.4(b) is also much more symmetric, although slight asymmetry can still be observed. The MI was estimated as 2.859 bpv. Although it is still not as good as the 2.964 bpv obtained with the PPC, the custom-polariser-based diattenuative module has great potential for further improvements. Firstly, the tilt angle of the film polariser within the cube can be further optimised for a higher extinction ratio that matches the particular characteristics of birefringent voxels, which is much more practical in comparison to manufacturing a phase-contrast objective with a different level of amplitude modulation. Secondly, the disposal of the

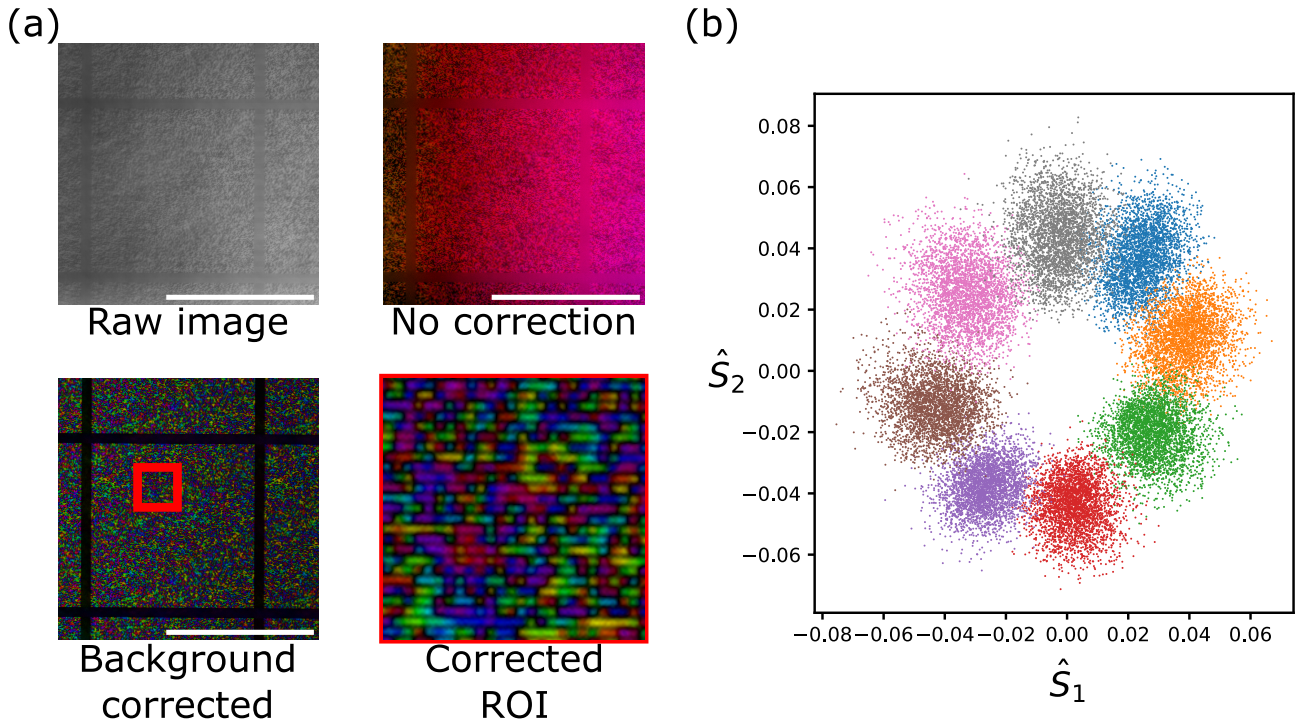


Figure 9.4: Imaging of a sector of birefringent voxels with the diattenuative module implemented with the custom polariser. (a) An intensity gradient can be observed on both the raw image and the false-colour visualisation with no correction, although significantly weaker than the PBS-based system. With correction, the IFoV error was also effectively removed, and the voxels in false colour are shown in a variety of colours, in contrast to the PBS-based system. Hue, AoLP; value, DoLP. Bar=100 μm . (b) Constellation diagram corresponding to the background-corrected image, showing a roughly isotropic polarisation contrast in all directions in the NLSP.

annular illumination scheme allows for the phase-contrast objective to be replaced with others with a higher NA, improving the lateral resolution beyond what is currently achievable. The axial resolution is also expected to be better for clear-aperture illumination, according to which the interlayer spacing may be optimised for.

9.4 Results on samples beyond birefringent voxels

As the diattenuative module is designed for imaging weak polarisation signals, it can be used not only in ODS, but also for imaging samples beyond birefringent voxels. In particular, it can be advantageous compared to PPC when imaging a sample with strong phase features but weak polarisation features, as in PPC the phase signal can dominate over the polarisation signal.

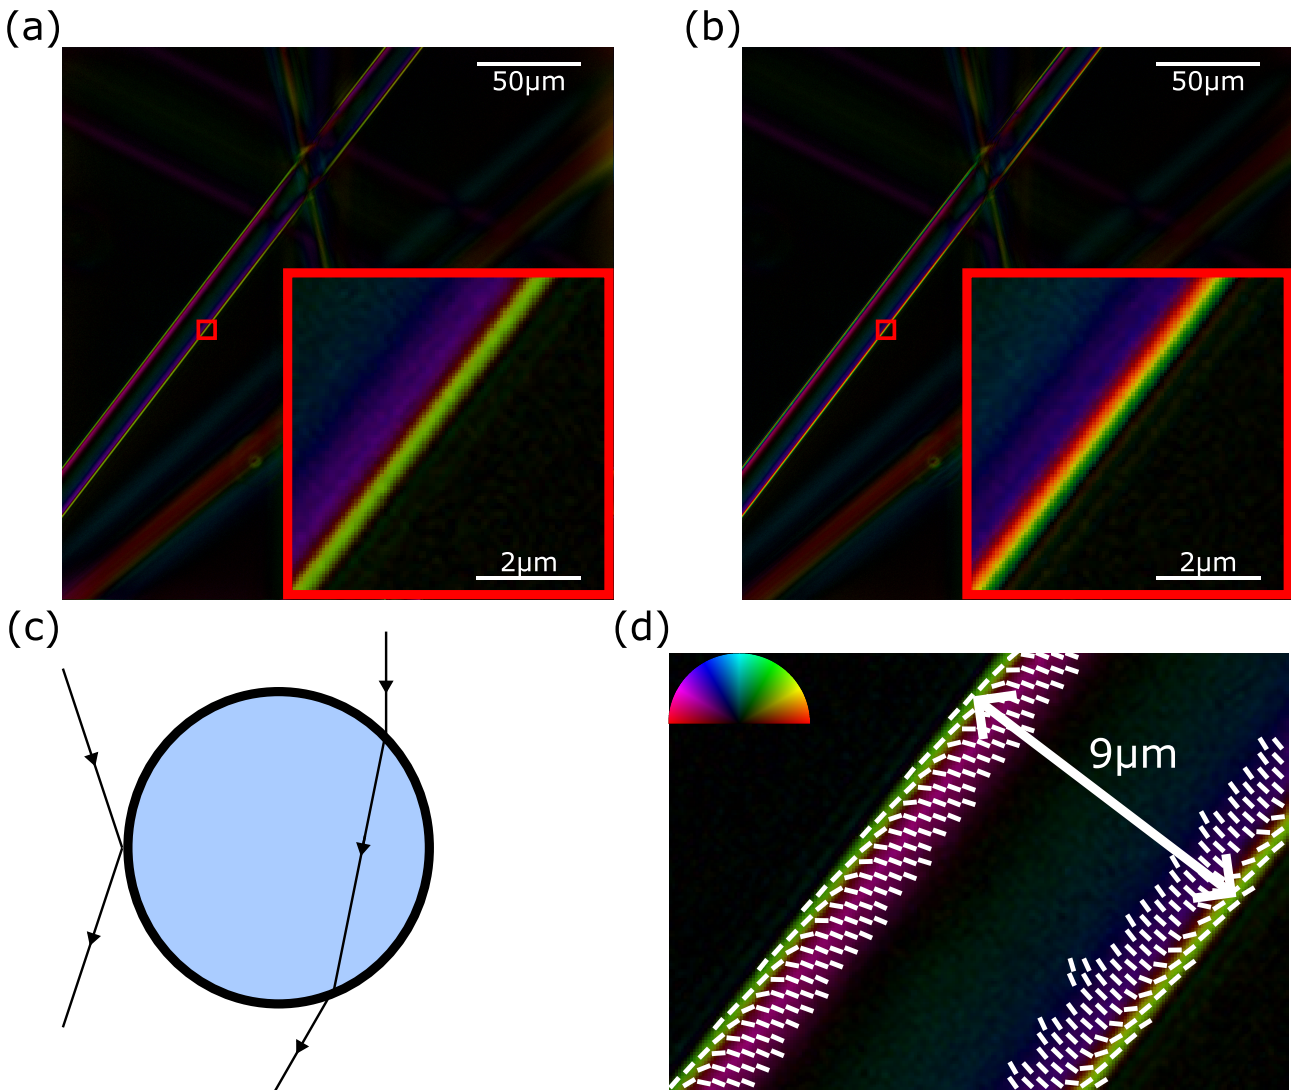


Figure 9.5: False-colour images of glass fibres acquired with the diattenuative module, (a) with and (b) without IFOV-error correction. (c) Schematic of reflected and refracted rays at the glass-media interfaces. (d) Region of interest after IFOV-error correction, with local AoLP visualised as white rods. Hue, AoLP; value, DoLP.

9.4.1 Glass fibres

Glass fibres were imaged to demonstrate this comparative advantage of the diattenuative module. They were mounted in a polymer medium (BIO-133) with a refractive index of 1.33 [229]. Due to the refractive-index mismatch, the sample exhibits strong phase features that can introduce an optical path difference of multiple wavelengths, leading to PPC detecting little polarisation contrast.

Using the diattenuative module, clear polarisation contrast was observed, due to the differing Fresnel transmission coefficients for light passing through the curved surfaces of the glass fibre,

as shown in Figure 9.5. Figure 9.5(a) shows a representative false-colour image after IFoV-error correction. Without IFoV-error correction, the signal is accompanied by significant IFoV error, as shown in Figure 9.5(b).

The polarisation signal arises not from structural anisotropy of the fibre, but from surface reflection and refraction at the glass-media boundary. Given the size of the glass fibre, whose diameter was measured as $9\ \mu\text{m}$, the effect can be explained in the regime of geometric optics. A ray diagram is shown for illustrative purposes in Figure 9.5(c), indicating the plane of incidence for both a reflected and a refracted ray. When light is reflected at the boundary, the S polarisation, which is perpendicular to the plane of incidence, is favoured according to the Fresnel equations [116]. As a result, the local AoLP aligns the fibre direction. In contrast, when light is transmitted through the glass fibre, refraction favours the P polarisation in the plane of incidence, leading to a local AoLP perpendicular to the fibre direction. This effect can be clearly observed in Figure 9.5(d), where local AoLP is visualised as white rods.

Overall, this example shows that the diattenuative module can reveal weak polarisation features that are difficult to observe with PPC in the presence of strong phase contrast. It also demonstrates that the system can capture weak polarisation signals arising from refractive-index mismatch.

9.4.2 Mycobacteria

To demonstrate the potential of such a system in a biological context, *Mycobacterium smegmatis* (*M. smegmatis*) was imaged with the diattenuative module. There is ongoing interest in *M. smegmatis* within our research group [117, 230], due to its similar cell structure and physiology to *Mycobacterium tuberculosis* (*Mtb*) [231], the primary cause of tuberculosis (TB).

Figure 9.6 shows an example image of *M. smegmatis*. The outlines of the cells can clearly be seen in polarisation contrast, and appear to depend on the cell orientation. In many cells, an additional line within the cell can also be observed, with a different measured AoLP.

Regarding the origin of the polarisation signal, two possible mechanisms are proposed. The first




Figure 9.6: False-colour images of *Mycobacterium smegmatis* observed with a diattenuative module. Each mycobacterium is outlined by a polarisation signal that depends on its orientation, and in some cases polarisation signals within the cells can also be observed. Hue, AoLP; value, DoLP. Bar=50 μm .

is structural anisotropy within the mycobacteria. *M. smegmatis* is known to have a complex and unusual cell envelope structure [231, 232], which is critical to its physiology. The polarisation contrast could plausibly arise from anisotropic organisation at the cell wall or envelope, such as mycomembrane lipids, which can be birefringent. In addition, birefringence may also arise from internal organisation within the cell, such as cytosolic proteins.

Besides structural anisotropy, polarisation signals can also be produced by scattering of light due to refractive-index mismatch between the cells and the mounting medium, as observed for glass fibres. For this sample, the mycobacteria were mounted in polyvinyl alcohol (PVA), which has a refractive index of 1.5 [233]. In contrast, mycobacteria have a refractive index much closer to water, which is reported to be 1.34–1.35 for *Mtb* [234].

Nevertheless, since mycobacteria are much smaller, with a typically sub-micron width, geometric optics and scalar diffraction theory are no longer applicable for modelling polarisation-




Figure 9.7: Simulated false-colour image of a Mie scatterer of a diameter of $2\ \mu\text{m}$ and a refractive index of 1.5. Polarisation signal arises despite the scatterer being a homogeneous sphere, with a radially varying AoLP and DoLP profile. Hue, AoLP; value, DoLP.

dependent scattering [235]. Instead, a Mie scatterer was considered using the simulation pipeline described in Chapter 5. For simplicity, the scatterer was modelled as a homogeneous sphere with a refractive index different from its surroundings. Using the analytic form of the scattering amplitudes [236], a scattering kernel can be constructed as a function of the input and output wave vectors in the angular spectrum, following the description in Section 5.1.2. The Mie scattering amplitudes were obtained numerically using the `miepython` package [237], and then transformed to the lab frame for the scattering kernel required by Equation 5.7 for simulations¹.

Figure 9.7 shows the simulated false-colour image for a homogeneous sphere with a refractive index of 1.5. As the simulation was performed without polarisation-channel mosaicking, all polarisation channels are perfectly registered, and no IFoV error is present. Polarisation contrast is most prominent near the outline of the sphere, indicating that such signals can arise from refractive-index mismatch even in the absence of an underlying anisotropic structure. In addition, the AoLP varies radially, which can be understood qualitatively as an additional relative phase accumulated between the S and D components due to the varying thickness of the sphere. Although this model is not intended to represent mycobacteria realistically, it provides supporting evidence that scattering is another plausible mechanism for the polarisation signals observed in mycobacteria.

Further studies on imaging mycobacteria with the diattenuative system could help determine whether the observed polarisation contrast arises from structural anisotropy, scattering due to refractive-index mismatch, or a combination of both. If the signal originates from the

¹Full algebraic working can be found in Appendix A.

anisotropic cell envelope, comparative experiments could be performed in which selected envelope components are chemically removed using detergents such as Triton-X100 [238]. If scattering is a dominant contributor, the refractive-index mismatch could be significantly reduced by mounting the mycobacteria in a medium with a refractive index closer to that of the cells, such as phosphate-buffered saline. In addition, objectives with higher magnification and NA could be used to resolve finer spatial features. Overall, these results demonstrate that the diattenuative system can capture weak polarisation signals in biological specimens, and provide a practical basis for further investigation of their physical origin.

9.5 Conclusion and outlook

This chapter introduced a clear-aperture polarisation-imaging approach based on a diattenuative module, motivated by practical limitations of PPC inherited from the underlying ZPC configuration. By attenuating the left-handed circular component selectively, the module reduces the polarisation background from the S wave while retaining the right-handed D-wave component, and therefore enhancing polarisation contrast without requiring an annular aperture or a phase ring. The theoretical description shows that the resulting imaging configuration has the same STF as the bright-field configuration, while avoiding the AoLP offset introduced by the phase modulation in PPC.

A practical implementation was realised by placing a partial linear polariser between two QWPs at $\pm 45^\circ$, and the choice of the partial polariser was shown to be constrained by the simultaneous requirements of low extinction ratio, good wavefront quality, and low angular dependence. Using a PBS in reflection as the partial polariser enabled voxel imaging, but introduced a strong angular dependence in the infinity space, leading to a large intensity gradient across the FoV and an anisotropic polarisation contrast after background correction and IFoV-error removal. These limitations motivated the design of a custom polariser that improves angular uniformity while maintaining a sufficiently low extinction ratio and acceptable wavefront performance, by placing a tilted polariser between two right-angled prisms. With this custom polariser,

the voxel constellation became substantially more symmetric and the estimated MI improved, demonstrating that the diattenuative approach is practically viable for imaging of birefringent voxels.

Beyond ODS, the diattenuative module provides a useful basis for imaging weak polarisation features in the presence of strong phase signals, where PPC can be less suitable because phase variations dominate the polarimetric measurements. This was illustrated by the imaging of glass fibres, where polarisation contrast arises from surface reflection and refraction, and by imaging *M. smegmatis*, where the observed polarisation contrast may originate from structural anisotropy, scattering due to a refractive-index mismatch, or a combination of both. A supporting simulation of a homogeneous Mie scatterer further indicates that weak polarisation signals can arise from scattering alone, highlighting the need for future studies.

In the future, the diattenuative system can be improved by optimising the effective diattenuation via the extinction ratio of the partial polariser, for a potentially even stronger polarisation contrast. For ODS, an important next step is to evaluate the clear-aperture configuration on true multi-layer tracks with small interlayer spacings, where simulations predict that axial performance and interlayer crosstalk can improve from PPC. More broadly, the ability to decouple polarisation contrast enhancement from phase-contrast optics, together with the increased flexibility in objective selection, suggests that the diattenuative module could serve as a useful platform for voxel readout and polarisation imaging with improved lateral and axial resolutions, more accurate AoLP measurements, and a broader scope of application.

Chapter 10

Conclusion

10.1 Summary of thesis achievements

In this thesis, polarisation microscopy was developed for readout of birefringent voxels for long-term optical data storage (ODS). In particular, a key challenge was identified, that the voxel-induced polarisation signals are typically very weak. The signal could therefore be compromised by many factors including photon noise, intra- and interlayer crosstalk, and aberrations in the system. This work showed how division of focal plane (DoFP) polarimetry could be developed for imaging weak birefringence. The thesis spanned system design and polarisation control, modelling and numerical simulation, practical implementation, and image calibration and processing methods. Building on both simulation and experimental results, voxel readout was evaluated in an information-theoretic framework. This enabled quantitative analysis of readout quality via mutual information (MI) and of how the performance changed under different conditions. The thesis further demonstrated explicit symbol recovery with a voxel-wise machine learning (ML) decoder, which completed an end-to-end ODS readout pipeline.

After establishing the application context, the thesis first described polarisation camera (Pol-Cam) in Chapter 2, which enabled single-shot DoFP polarimetry. With a micro-polariser array (MPA) at the camera, the incident field was analysed simultaneously in four linear polarisation states. Three of the four Stokes parameters could then be obtained. This gave rise to ambiguity

in the circularly polarised component of the field, which was demonstrated with an example involving mixing of polarisation signals with and without coherence.

We showed that the polarisation measurement was a projection of the Stokes vector in the normalised linear Stokes plane (NLSP), and the degree of linear polarisation (DoLP) and angle of linear polarisation (AoLP) were identified as key obtainable measurements. With circularly polarised illumination, weak linear birefringence produced an orthogonally circular component. Its amplitude and relative phase encoded the retardance and azimuth of the sample birefringence respectively. The capture of three Stokes parameters could also be understood through Fourier analysis of PolCam images. The sampling criterion for PolCam was evaluated to avoid crosstalk between Stokes parameters. A linear noise model was presented, and the noise performance for PolCam was characterised with automated camera characterization via electron noise tool (ACCeNT), concluding that the dominant source was photon shot noise, provided that spatial non-uniformity was effectively corrected with the polarisation-field correction (PFC) routine we proposed.

The analysis of PolCam motivated design considerations for an ideal system for weak birefringence in which the DoLP signal should be maximised. The full-well capacity (FWC) offered by the PolCam should also be utilised as much as possible. These considerations motivated the design of both polarisation-sensitive phase contrast (PPC) and the diattenuative module, which became central to our methodology.

On this foundation, PPC was proposed in Chapter 3 by employing a Zernike's phase contrast (ZPC) setup. Under the weak object assumption, an analytical model could be built by splitting the wave into a surround (S) wave and a diffracted (D) wave. It was shown that selective attenuation of the S wave at the objective phase ring enhanced the DoLP signal and therefore the polarisation contrast. By assuming incoherent illumination at the condenser aperture, a model was derived that described the captured Stokes images with the Stokes transfer functions (STFs). By comparing the STFs in PPC with those in a bright-field setup, imaging characteristics were discussed, such as the compromised spatial resolution and the AoLP bias due to phase modulation.

The analytical model itself, however, was deemed not sufficiently accurate. This was not only due to the weak object approximation and the split of S and D waves, but also due to an inaccurate account of the partially coherent imaging process. Chapter 4 and Chapter 5 addressed this challenge by developing a simulation pipeline that took the second-order statistics of the field into account. Through coherent-mode decomposition of an apertured Gaussian Schell-model (GSM) source, propagation of the partially coherent field could be simulated more accurately and flexibly than the analytical model. An efficient implementation was developed based on multiprocessing for mode generation, Fourier propagation accelerated with a graphics processing unit (GPU), and cluster deployment on high-performance computing (HPC) facilities. Together, this allowed simulations of the partially coherent system at high throughput, which complemented experimental results for quantitative analysis.

The practical implementation of PPC was then introduced in Chapter 6. Particular emphasis was placed on automation and reproducibility. The accompanying software unified acquisition and calibration of 3D sample positioning, automation, and real-time processing and visualisation using the polarisation-camera image processing (PCIP) module. This implementation provided the experimental basis for quantitative evaluation, which also enabled generation of the datasets required for informational analysis and decoding.

With both modelling and implementation in place, the thesis evaluated voxel readout as a communication process in an information-theoretic framework, as presented in Chapter 7. The channel input was identified as the voxel symbol, while the polarisation measurement in the NLSP was regarded as the output. MI between the two quantities then provided a decoder-agnostic metric that quantified how much information was preserved, due to Shannon's noisy channel coding theorem. To make informational analysis operational on polarisation images, the voxel image processing pipeline (VIP) pipeline was developed. It comprised PFC calibration, PCIP-based image processing, correction of instantaneous field-of-view (IFoV) errors, sub-pixel voxel localisation, and MI estimation using a k -nearest neighbour (k -NN)-based estimator. The pipeline was applicable to both simulation and experimental images. An MI of 2.964 bpv was demonstrated experimentally for a 3-bpv sector at a pitch of 0.7 μm , while trade-offs between voxel pitch, crosstalk, and achievable density were identified.

With experimental and simulation results, the framework was used to study the degradation of MI with respect to the photon shot noise. It also established feasible multi-layer readout with resolvable information peaks across multiple layers. By comparing axial profiles obtained from simulations and experiments, it was learnt that a thick-voxel model was necessary to accurately recreate the long- and short-range interlayer crosstalk, which critically limited the interlayer spacing and thus the usable data capacity. In addition, depth-induced aberration was analysed, which was shown to be practically mitigable with a combination of collar correction and sample refocusing. This indicated that the depth-induced aberration did not necessarily impose a fundamental information limit when the system was well calibrated.

The thesis then completed the demonstration of voxel readout by demonstrating explicit symbol recovery through decoding in Chapter 8. A voxel-wise decoding strategy was adopted, taking advantage of the voxel localisation functionality provided by VIPP. By providing a local neighbourhood around a voxel, a shallow multilayer perceptron (MLP) network was able to map these local measurements to voxel symbols. Compared to a sector-based strategy, this method reduced input dimensionality dramatically and enabled efficient training from relatively small experimental datasets, while still achieving near-lossless recovery in a $0.7\ \mu\text{m}$ -3 bpv setup at a symbol MI of 2.995 bpv. For a more challenging configuration where voxel pitch was smaller, the number of voxels in a sector was much less, and the source entropy was higher, a polarisation-based image-augmentation strategy was developed by exploiting the channel symmetry provided by PPC. A data density of $(13.85 \pm 0.06)\ \text{Mbit mm}^{-2}$ per layer was demonstrated, which, assuming Gray coding, corresponded to a bit error rate (BER) of 0.135 ± 0.004 with a MLP network.

Finally, an alternative to PPC was presented, motivated by limitations due to the spatial filtering inherited from the ZPC setup. Chapter 9 introduced the design, implementation, and imaging of a diattenuative module that selectively attenuated the S wave based on polarisation instead, enhancing the DoLP signal without introducing an annular aperture or a phase ring required by the PPC setup. The resulting configuration shared the bright-field STF and avoided the AoLP offsets associated with phase modulation in PPC. A practical implementation was demonstrated using a partial linear polariser placed between two quarter-wave plates (QWPs),

and the work identified key practical constraints on the partial polariser, including its extinction ratio, wavefront quality, and angular uniformity. Limitations of a polarising beam splitter (PBS)-based implementation motivated a custom polariser by placing a tilted film polariser between two right-angled prisms, to simultaneously reduce the extinction ratio and improve wavefront quality, while carrying a much smaller angular dependence than a PBS. The diattenuative system was shown to support voxel imaging and to reveal weak polarisation signals in samples where strong phase features would overwhelm PPC, and provided a platform for further development both for ODS and for biological polarisation imaging. It further validated that the enhancement of polarisation contrast could be realised in other ways besides phase-contrast optics, suggesting a path of improved lateral and axial performance and more faithful AoLP measurements.

Overall, this thesis developed polarisation microscopy for readout of birefringent voxels. It used PolCam-based single-shot DoFP polarimetry and circular illumination to link weak birefringence to measurements of DoLP and AoLP. It then improved readout by controlling the relative amplitudes of the S and D waves, both in PPC and in a clear-aperture diattenuative alternative. Analytical modelling, partially coherent simulation, and a practical implementation enabled an information-theoretic analysis for characterising the read quality under a variety of conditions, and an ML-based decoder showed that symbol information could be recovered.

10.2 Future possibilities

While end-to-end ODS is demonstrated, there are many elements within this process that can be further improved. Firstly, work can be done to enable automated depth correction for multi-layer acquisition. This is applicable to both PPC and the diattenuative configuration, as both benefit from high-throughput multi-layer imaging without depth-induced aberrations. One route is to use a motorised objective collar correction in conjunction with slight refocusing of the sample, as suggested by the wavefront analysis with a well-calibrated collar. Alternatively, remote-refocusing techniques can be useful if the depth-dependent aberration can be corrected

for by matching the remote-refocusing magnification to the refractive index of silica glass. In this case, an aberration-free copy of the track can be constructed in the remote space to be re-imaged. Combined with samples of varying interlayer spacing, automated depth correction enables experimental analysis of the read performance in the presence of substantial interlayer crosstalk. Together with the intra-layer voxel pitch and source entropy, the trade-off between crosstalk and read quality can be comprehensively studied, to optimise the usable data density in a working system.

In parallel, the diattenuative module can be further improved in several practical aspects. This includes optimising its effective extinction ratio by adjusting the tilt angle of the film polariser, and using a light source with a smaller bandwidth to reduce the impact of wave-plate retardance variation. It may also be beneficial to place the diattenuative module outside the infinity space to minimise the impact of angular dependence, for example by introducing a relay system. In ODS applications, the extinction ratio may be optimised for a particular level of glass modification depending on the writing parameters, but the diattenuative system can also find its applications outside ODS such as mycobacteria biology.

With improved depth robustness and a better-optimised readout, the communication channel can then be characterised more rigorously. This characterisation should quantify how the channel depends on interlayer spacing, voxel pitch, source entropy, and any residual non-uniformity across azimuthal directions and the field-of-view (FoV). Together with write-side choices, such channel measurements provide a basis for system-level optimisation that aims to maximise channel capacity.

It is also worth investigating the use of birefringent retardance to encode information in addition to the azimuth, which should lead to a higher data density. In principle, polarisation microscopy and the processing pipelines should be able to measure different levels of DoLP corresponding to the retardance levels of the voxels. Nevertheless, efficient utilisation of this additional degree of freedom requires the polarisation contrast to be as uniform as possible in different azimuthal directions and across the FoV. Challenges are also imposed to the write side and the coding strategy, as the communication channel becomes no longer symmetric.

Finally, the decoding strategy can be optimised beyond the simple MLP network presented in this work. To make decoding more robust against factors such as camera noise, defocus, uncorrected aberrations, and crosstalk, the model should be trained with more diverse datasets beyond a single sector. Acquiring the datasets needed requires not only more glass samples, but also a higher level of automation for functionalities such as sector navigation and focusing. Besides, the simulation pipeline can be used to generate the larger amount of datasets required, which can be used in conjunction with the experimental data for transfer learning. A more robust ML decoder is also likely to require a larger model with a more advanced architecture beyond the MLP.

10.3 Research outputs

- **[Related publication (not submitted as part of this thesis)]** James Allison and Zhonghe Feng *et al.* “Laser writing in glass for dense, fast and efficient archival data storage”. In *Nature* **650**, no. 8102 (2026), pp. 606–612.
- **[Conference presentation]** Zhonghe Feng, James Clegg, Matthew R. Foreman, and Mark A. A. Neil. “Quantitatively Analysing the Read Quality of Polarization Microscopy for Optical Data Storage Using Mutual Information”. Oral presentation at Focus on Microscopy 2024. Mar. 2024.
- **[Conference poster]** Zhonghe Feng, Mark A. A. Neil, Matthew R. Foreman, and James Clegg. “Polarization-sensitive Phase Contrast Microscopy for Optical Data Storage”. Poster presentation at Focus on Microscopy 2023. Apr. 2023.
- **[Invited talk]** Zhonghe Feng. “Polarisation microscopy for optical data storage”. Invited talk at Politecnico di Milano (host: Andrea Bassi, Department of Physics). Apr. 2025.

Appendix A

Simulating a Mie scatterer

This appendix outlines how the scattering kernel $S_{ij}(k_x, k_y, k'_x, k'_y)$ used in Equation 5.7 can be obtained for a Mie scatterer, such that the scatterer can be simulated with the simulation pipeline according to Equation 5.7, which is given as

$$U'_i(k_x, k_y, z) = \sum_j \iint S_{ij}(k_x, k_y, k'_x, k'_y) U_j(k'_x, k'_y, z) dk'_x dk'_y. \quad (\text{A.1})$$

A.1 Wave vectors and field components

Let the incident and scattered plane-wave components be indexed by their wave vectors \mathbf{k}^i and \mathbf{k}^s , with a common wave number $k_0 = 2\pi/\lambda$. The wave vector therefore becomes

$$\mathbf{k} = \begin{bmatrix} k_x \\ k_y \\ k_z \end{bmatrix}, \quad (\text{A.2})$$

where

$$k_z = \sqrt{k_0^2 - k_x^2 - k_y^2}. \quad (\text{A.3})$$

Define the corresponding propagation unit vector $\hat{\mathbf{k}} = \mathbf{k}/k_0$. The spherical angles (θ, ϕ) can then be associated with the normalised wave vector $\hat{\mathbf{k}}$, given as

$$\hat{\mathbf{k}} = \begin{bmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{bmatrix}. \quad (\text{A.4})$$

For each direction $\hat{\mathbf{k}}$, define a local orthonormal basis $(\hat{\mathbf{1}}, \hat{\mathbf{2}}, \hat{\mathbf{k}})$, where $\hat{\mathbf{1}}$ and $\hat{\mathbf{2}}$ are orthogonal to $\hat{\mathbf{k}}$. In particular, $\hat{\mathbf{1}}$ is normal to the plane in which scattering takes place, and $\hat{\mathbf{2}}$ lies within the scattering plane, following the convention in Tsang et al. [236]. Let $U_1(\mathbf{k})$ and $U_2(\mathbf{k})$ denote the complex field amplitudes along $\hat{\mathbf{1}}$ and $\hat{\mathbf{2}}$, respectively. The corresponding field amplitude is therefore

$$\mathbf{U}(\mathbf{k}) = \begin{bmatrix} U_1(\mathbf{k}) \\ U_2(\mathbf{k}) \end{bmatrix}. \quad (\text{A.5})$$

A.2 Scattering in the 1–2 system

For a single incident direction \mathbf{k}^i and observation direction \mathbf{k}^s , the scattered field in the 1–2 system is related to the incident field by a 2×2 scattering amplitude matrix,

$$\mathbf{U}'(\mathbf{k}^s) = \mathbf{S}^{(12)}(\mathbf{k}^s, \mathbf{k}^i) \mathbf{U}(\mathbf{k}^i), \quad (\text{A.6})$$

where

$$\mathbf{S}^{(12)}(\mathbf{k}^s, \mathbf{k}^i) = \begin{bmatrix} S_{11}^{(12)}(\mathbf{k}^s, \mathbf{k}^i) & S_{12}^{(12)}(\mathbf{k}^s, \mathbf{k}^i) \\ S_{21}^{(12)}(\mathbf{k}^s, \mathbf{k}^i) & S_{22}^{(12)}(\mathbf{k}^s, \mathbf{k}^i) \end{bmatrix}. \quad (\text{A.7})$$

For scattering by a homogeneous sphere, $\mathbf{S}^{(12)}$ is diagonal in the conventional 1–2 definition and can be expressed in terms of two angle-dependent amplitude functions $S_1(\Theta)$ and $S_2(\Theta)$,

$$\mathbf{S}^{(12)}(\mathbf{k}^s, \mathbf{k}^i) = \frac{i}{k_0} \begin{bmatrix} S_1(\Theta) & 0 \\ 0 & S_2(\Theta) \end{bmatrix}, \quad (\text{A.8})$$

where Θ is the deflection angle defined by

$$\cos \Theta = \hat{\mathbf{k}}^i \cdot \hat{\mathbf{k}}^s. \quad (\text{A.9})$$

In this work, $S_1(\Theta)$ and $S_2(\Theta)$ can be obtained numerically using the `miepython` package [237].

A.3 Rotation from the lab transverse basis to the 1–2 basis

The scattering relation in Equation A.6 is defined in the 1–2 system $(\hat{\mathbf{1}}, \hat{\mathbf{2}})$. To assemble an angular-spectrum kernel acting on the transverse components in the lab frame, a coordinate transformation is required.

Let the lab basis be $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}})$, and the scattering basis be $(\hat{\mathbf{1}}, \hat{\mathbf{2}}^i, \hat{\mathbf{k}}^i)$ for the incident wave and $(\hat{\mathbf{1}}, \hat{\mathbf{2}}^s, \hat{\mathbf{k}}^s)$ for the scattered wave. Define the 3D matrix $\mathbb{R}_{\text{in}}(\mathbf{k}^i; \mathbf{k}^s)$ that transforms a vector from the lab basis to the scattering basis, requiring

$$\mathbb{R}_{\text{in}}(\hat{\mathbf{k}}^i; \hat{\mathbf{k}}^s) \hat{\mathbf{k}}^i = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad (\text{A.10})$$

and

$$\mathbb{R}_{\text{in}}(\hat{\mathbf{k}}^i; \hat{\mathbf{k}}^s) \hat{\mathbf{k}}^s \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = 0. \quad (\text{A.11})$$

It can be found that the matrix can be broken down into two components

$$\mathbb{R}_{\text{in}}(\hat{\mathbf{k}}^i; \hat{\mathbf{k}}^s) = \mathbb{R}'(\gamma^i) \mathbb{R}^R(\theta_i, \phi_i), \quad (\text{A.12})$$

where $\mathbb{R}^R(\theta_i, \phi_i)$ represents the Rodrigues' rotation [239] given as

$$\mathbb{R}^R(\theta_i, \phi_i) = \begin{bmatrix} \cos \theta_i \cos^2 \phi_i + \sin^2 \phi_i & (\cos \theta_i - 1) \sin \phi_i \cos \phi_i & -\sin \theta_i \cos \phi_i \\ (\cos \theta_i - 1) \sin \phi_i \cos \phi_i & \cos \theta_i \sin^2 \phi_i + \cos^2 \phi_i & -\sin \theta_i \sin \phi_i \\ \sin \theta_i \cos \phi_i & \sin \theta_i \sin \phi_i & \cos \theta_i \end{bmatrix}, \quad (\text{A.13})$$

such that

$$\mathbb{R}^R(\theta_i, \phi_i) \hat{\mathbf{k}}^i = \hat{\mathbf{z}}. \quad (\text{A.14})$$

The other matrix, $\mathbb{R}'(\gamma^i)$, is a simple rotation about $[0, 0, 1]^T$ that then places $\hat{\mathbf{k}}^s$ in the $\hat{\mathbf{z}}^i$ - $\hat{\mathbf{k}}^i$ plane, which can be found as

$$\mathbb{R}'(\gamma^i) = \begin{bmatrix} \cos \gamma^i & -\sin \gamma^i & 0 \\ \sin \gamma^i & \cos \gamma^i & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (\text{A.15})$$

where

$$\gamma^i := \frac{\pi}{2} - (\mathbb{R}^R(\theta_i, \phi_i) \hat{\mathbf{k}}^s)_\phi. \quad (\text{A.16})$$

In general, however, $\mathbb{R}_{\text{in}}(\hat{\mathbf{k}}^i; \hat{\mathbf{k}}^s)$ is a 3×3 matrix, which is incompatible with the simulation pipeline. It can however be noticed that in a non-paraxial picture the condenser lens has a spherical principle surface, leading to a non-zero longitudinal component of the electric field to begin with. The focusing of the condenser lens can be described by the 3D matrix $\mathbb{Q}(\theta, \phi)$, given as [20]

$$\mathbb{Q}(\theta, \phi) = \frac{a(\theta)}{2} \begin{bmatrix} q_1 + q_2 \cos 2\phi & q_2 \sin 2\phi & q_3 \cos \phi \\ q_2 \sin 2\phi & q_1 - q_2 \cos 2\phi & q_3 \sin \phi \\ -q_3 \cos \phi & -q_3 \sin \phi & q_4 \end{bmatrix}, \quad (\text{A.17})$$

where

$$q_1 := \cos \theta + 1, \quad (\text{A.18})$$

$$q_2 := \cos \theta - 1, \quad (\text{A.19})$$

$$q_3 := 2 \sin \theta, \quad (\text{A.20})$$

$$q_4 := 2 \cos \theta, \quad (\text{A.21})$$

$$(\text{A.22})$$

and $a(\theta)$ is the apodisation factor. The full 3D transformation from the condenser aperture plane to the scattering basis can then be written as $\mathbb{R}_{\text{in}}(\hat{\mathbf{k}}^i; \hat{\mathbf{k}}^s)\mathbb{Q}(\theta_i, \phi_i)$.

It can then be shown that, in fact,

$$\mathbb{R}^R(\theta_i, \phi_i)\mathbb{Q}(\theta_i, \phi_i) = a(\theta_i)\mathbb{I}. \quad (\text{A.23})$$

Therefore,

$$\mathbb{R}_{\text{in}}(\hat{\mathbf{k}}^i; \hat{\mathbf{k}}^s)\mathbb{Q}(\theta_i, \phi_i) = a(\theta_i) \begin{bmatrix} \cos \gamma^i & -\sin \gamma^i & 0 \\ \sin \gamma^i & \cos \gamma^i & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (\text{A.24})$$

which is trivial in the longitudinal direction. It is therefore possible to relate the field from the condenser aperture plane to the incident scattering basis using a reduced 2×2 transformation, given by

$$a(\theta_i)\mathbb{R}^{2 \times 2}(\gamma^i), \quad (\text{A.25})$$

where $\mathbb{R}^{2 \times 2}$ denotes a 2×2 rotator.

Similarly, the transformation from the scattered basis to the objective pupil plane can be found as a 2×2 matrix as well, given as

$$\frac{1}{a(\theta_i)}\mathbb{R}^{2 \times 2}(\gamma^s), \quad (\text{A.26})$$

where

$$\gamma^s := \frac{\pi}{2} + (\mathbb{R}^R(\theta_s, \phi_s)\hat{\mathbf{k}}^i)_\phi. \quad (\text{A.27})$$

As a result, the full scattering process from the condenser aperture plane to the objective pupil plane can be written as

$$\frac{ia(\theta_i)}{a(\theta_s)k_0} \mathbb{R}^{2 \times 2}(\gamma^s) \begin{bmatrix} S_1(\Theta) & 0 \\ 0 & S_2(\Theta) \end{bmatrix} \mathbb{R}^{2 \times 2}(\gamma^i), \quad (\text{A.28})$$

using only 2D matrix manipulation, which is now compatible with the simulation pipeline. As a result, a scattering kernel of

$$S_{ij}(k_x^s, k_y^s, k_x^i, k_y^i) \quad (\text{A.29})$$

can be computed with Equation A.28 to be provided to the pipeline for simulations.

Bibliography

- [1] Martin Hilbert and Priscila López. “The World’s Technological Capacity to Store, Communicate, and Compute Information”. In: *Science* 332.6025 (Apr. 2011), pp. 60–65. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1200970](https://doi.org/10.1126/science.1200970). URL: <https://www.science.org/doi/10.1126/science.1200970>.
- [2] *Worldwide IDC Global DataSphere Forecast, 2024–2028: AI Everywhere, But Upsurge in Data Will Take Time*. IDC Market Forecast US52076424. Market research report. Framingham, MA, USA: International Data Corporation (IDC), May 7, 2024.
- [3] Jackie Jung. *Giving HDD Rare Earth Elements New Life*. Apr. 17, 2025. URL: <https://blog.westerndigital.com/giving-hdd-rare-earth-elements-new-life/>.
- [4] David Reinsel, John Gantz, and John Rydning. *Data Age 2025: The Evolution of Data to Life-Critical. Don’t Focus on Big Data; Focus on the Data That’s Big*. IDC White Paper. Sponsored by Seagate. Framingham, MA, USA: International Data Corporation, Apr. 2017. URL: <https://www.seagate.com/files/www-content/our-story/trends/files/Seagate-WP-DataAge2025-March-2017.pdf>.
- [5] George Amvrosiadis, Alina Oprea, and Bianca Schroeder. “Practical scrubbing: Getting to the bad sector at the right time”. In: *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2012)*. 2012 42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). Boston, MA, USA: IEEE, June 2012, pp. 1–12. ISBN: 978-1-4673-1625-5 978-1-4673-1624-8 978-1-4673-1623-1. DOI: [10.1109/DSN.2012.6263919](https://doi.org/10.1109/DSN.2012.6263919). URL: <http://ieeexplore.ieee.org/document/6263919/>.

- [6] Patrick Anderson et al. “Project Silica: Towards Sustainable Cloud Archival Storage in Glass”. In: *ACM Transactions on Storage* 21.1 (Feb. 28, 2025), pp. 1–31. ISSN: 1553-3077, 1553-3093. DOI: [10.1145/3708996](https://doi.org/10.1145/3708996). URL: <https://dl.acm.org/doi/10.1145/3708996>.
- [7] Min Gu, Xiangping Li, and Yaoyu Cao. “Optical storage arrays: a perspective for future big data storage”. In: *Light: Science & Applications* 3.5 (May 23, 2014), e177–e177. ISSN: 2047-7538. DOI: [10.1038/lsa.2014.58](https://doi.org/10.1038/lsa.2014.58). URL: <https://www.nature.com/articles/lsa201458>.
- [8] Qiming Zhang, Zhilin Xia, Yi-Bing Cheng, and Min Gu. “High-capacity optical long data memory based on enhanced Young’s modulus in nanoplasmonic hybrid glass composites”. In: *Nature Communications* 9.1 (Mar. 22, 2018), p. 1183. ISSN: 2041-1723. DOI: [10.1038/s41467-018-03589-y](https://doi.org/10.1038/s41467-018-03589-y). URL: <https://www.nature.com/articles/s41467-018-03589-y>.
- [9] Eduardo Pinheiro, Wolf-Dietrich Weber, and Luiz André Barroso. “Failure Trends in a Large Disk Drive Population”. In: *Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST ’07)*. USENIX Association, 2007. URL: <https://www.usenix.org/conference/fast-07/failure-trends-large-disk-drive-population>.
- [10] T.J.E. Schwarz, Qin Xin, E.L. Miller, D.D.E. Long, A. Hospodor, and Spencer Ng. “Disk scrubbing in large archival storage systems”. In: *The IEEE Computer Society’s 12th Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems, 2004. (MASCOTS 2004). Proceedings*. The IEEE Computer Society’s 12th Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems, 2004. (MASCOTS 2004). Proceedings. Volendam, The Netherlands, EU: IEEE, 2004, pp. 409–418. ISBN: 978-0-7695-2251-7. DOI: [10.1109/MASCOT.2004.1348296](https://doi.org/10.1109/MASCOT.2004.1348296). URL: <http://ieeexplore.ieee.org/document/1348296/>.
- [11] Bianca Schroeder, Sotirios Damouras, and Phillipa Gill. “Understanding latent sector errors and how to protect against them”. In: *ACM Transactions on Storage* 6.3 (Sept.

- 2010), pp. 1–23. ISSN: 1553-3077, 1553-3093. DOI: [10.1145/1837915.1837917](https://doi.org/10.1145/1837915.1837917). URL: <https://dl.acm.org/doi/10.1145/1837915.1837917>.
- [12] Patrick Anderson et al. “Glass: A new media for a new era?” In: *10th USENIX Workshop on Hot Topics in Storage and File Systems, HotStorage 2018, co-located with USENIX ATC 2018* (2018).
- [13] Matthias Grawinkel, Lars Nagel, Markus Mäsker, Federico Padua, André Brinkmann, and Lennart Sorth. “Analysis of the ECMWF Storage Landscape”. In: *Proceedings of the 13th USENIX Conference on File and Storage Technologies (FAST ’15)*. USENIX Association, 2015, pp. 15–27. ISBN: 978-1-931971-20-1. URL: <https://www.usenix.org/conference/fast15/technical-sessions/presentation/grawinkel>.
- [14] Babar Khan and Andreas Koch. “Reflecting on the Past 17 Years of Shingled Magnetic Recording for Insights Into Future Disk Transitions: A Survey”. In: *ACM Transactions on Storage* 21.3 (Aug. 31, 2025), pp. 1–50. ISSN: 1553-3077, 1553-3093. DOI: [10.1145/3731453](https://doi.org/10.1145/3731453). URL: <https://dl.acm.org/doi/10.1145/3731453>.
- [15] Mark A. Lantz, Simeon Furrer, Martin Petermann, Hugo Rothuizen, Stella Brach, Luzius Kronig, Ilias Iliadis, Beat Weiss, Ed R. Childers, and David Pease. “Magnetic Tape Storage Technology”. In: *ACM Transactions on Storage* 21.1 (Feb. 28, 2025), pp. 1–70. ISSN: 1553-3077, 1553-3093. DOI: [10.1145/3708997](https://doi.org/10.1145/3708997). URL: <https://dl.acm.org/doi/10.1145/3708997>.
- [16] F. Träger, ed. *Springer handbook of lasers and optics*. New York: Springer, 2007. 1332 pp. ISBN: 978-0-387-95579-7 978-0-387-30420-5 978-0-387-34557-4.
- [17] G. Bowhuis and G. Bouwhuis, eds. *Principles of optical disc systems*. Reprint. Bristol: Hilger, 1986. 277 pp. ISBN: 978-0-85274-785-8.
- [18] M. Mansuripur and G. Sincerbox. “Principles and techniques of optical data storage”. In: *Proceedings of the IEEE* 85.11 (Nov. 1997), pp. 1780–1796. ISSN: 00189219. DOI: [10.1109/5.649657](https://doi.org/10.1109/5.649657). URL: <http://ieeexplore.ieee.org/document/649657/>.

- [19] Peter Zijlstra, James W. M. Chon, and Min Gu. “Five-dimensional optical recording mediated by surface plasmons in gold nanorods”. In: *Nature* 459.7245 (May 2009), pp. 410–413. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/nature08053](https://doi.org/10.1038/nature08053). URL: <https://www.nature.com/articles/nature08053>.
- [20] Matthew R. Foreman and Peter Török. “Information and resolution in electromagnetic optical systems”. In: *Physical Review A* 82.4 (Oct. 25, 2010), p. 043835. ISSN: 1050-2947, 1094-1622. DOI: [10.1103/PhysRevA.82.043835](https://doi.org/10.1103/PhysRevA.82.043835). URL: <https://link.aps.org/doi/10.1103/PhysRevA.82.043835>.
- [21] C. Macias-Romero, P.R.T. Munro, and P. Török. “Polarization–multiplexed encoding at nanometer scales”. In: *Optics Express* 22.21 (Oct. 20, 2014), p. 26240. ISSN: 1094-4087. DOI: [10.1364/OE.22.026240](https://doi.org/10.1364/OE.22.026240). URL: <https://opg.optica.org/oe/abstract.cfm?uri=oe-22-21-26240>.
- [22] C. Macias-Romero, M. R. Foreman, P. R. T. Munro, and P. Török. “Confocal polarization imaging in high-numerical-aperture space”. In: *Optics Letters* 39.8 (Apr. 15, 2014), p. 2322. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.39.002322](https://doi.org/10.1364/OL.39.002322). URL: <https://opg.optica.org/abstract.cfm?URI=ol-39-8-2322>.
- [23] Jingyu Zhang, Mindaugas Gecevičius, Martynas Beresna, and Peter G. Kazansky. “Seemingly Unlimited Lifetime Data Storage in Nanostructured Glass”. In: *Physical Review Letters* 112.3 (Jan. 23, 2014), p. 033901. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.112.033901](https://doi.org/10.1103/PhysRevLett.112.033901). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.033901>.
- [24] Huijun Wang, Yuhao Lei, Lei Wang, Masaaki Sakakura, Yanhao Yu, Gholamreza Shayegannad, and Peter G. Kazansky. “100-Layer Error-Free 5D Optical Data Storage by Ultrafast Laser Nanostructuring in Glass”. In: *Laser & Photonics Reviews* 16.4 (Apr. 2022), p. 2100563. ISSN: 1863-8880, 1863-8899. DOI: [10.1002/lpor.202100563](https://doi.org/10.1002/lpor.202100563). URL: <https://onlinelibrary.wiley.com/doi/10.1002/lpor.202100563>.

- [25] James Allison et al. “Laser writing in glass for dense, fast and efficient archival data storage”. In: *Nature* 650.8102 (Feb. 1, 2026), pp. 606–612. ISSN: 1476-4687. DOI: [10.1038/s41586-025-10042-w](https://doi.org/10.1038/s41586-025-10042-w). URL: <https://doi.org/10.1038/s41586-025-10042-w>.
- [26] K. Miura, Jianrong Qiu, H. Inouye, T. Mitsuyu, and K. Hirao. “Photowritten optical waveguides in various glasses with ultrashort pulse laser”. In: *Applied Physics Letters* 71.23 (Dec. 8, 1997), pp. 3329–3331. ISSN: 0003-6951, 1077-3118. DOI: [10.1063/1.120327](https://pubs.aip.org/apl/article/71/23/3329/67953/Photowritten-optical-waveguides-in-various-glasses). URL: <https://pubs.aip.org/apl/article/71/23/3329/67953/Photowritten-optical-waveguides-in-various-glasses>.
- [27] D. Homoelle, S. Wielandy, Alexander L. Gaeta, N. F. Borrelli, and Charlene Smith. “Infrared photosensitivity in silica glasses exposed to femtosecond laser pulses”. In: *Optics Letters* 24.18 (Sept. 15, 1999), p. 1311. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.24.001311](https://opg.optica.org/abstract.cfm?URI=ol-24-18-1311). URL: <https://opg.optica.org/abstract.cfm?URI=ol-24-18-1311>.
- [28] Yasuhiko Shimotsuma, Peter G. Kazansky, Jiarong Qiu, and Kazuoki Hirao. “Self-Organized Nanogratings in Glass Irradiated by Ultrashort Light Pulses”. In: *Physical Review Letters* 91.24 (Dec. 11, 2003), p. 247405. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.91.247405](https://link.aps.org/doi/10.1103/PhysRevLett.91.247405). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.91.247405>.
- [29] Erica Bricchi, Bruce G. Klappauf, and Peter G. Kazansky. “Form birefringence and negative index change created by femtosecond direct writing in transparent materials”. In: *Optics Letters* 29.1 (Jan. 1, 2004), p. 119. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.29.000119](https://opg.optica.org/abstract.cfm?URI=ol-29-1-119). URL: <https://opg.optica.org/abstract.cfm?URI=ol-29-1-119>.
- [30] Yasuhiko Shimotsuma, Kazuyuki Hirao, Peter G. Kazansky, and Jiarong Qiu. “Three-Dimensional Micro- and Nano-Fabrication in Transparent Materials by Femtosecond Laser”. In: *Japanese Journal of Applied Physics* 44.7 (July 1, 2005), p. 4735. ISSN: 0021-4922, 1347-4065. DOI: [10.1143/JJAP.44.4735](https://iopscience.iop.org/article/10.1143/JJAP.44.4735). URL: <https://iopscience.iop.org/article/10.1143/JJAP.44.4735>.

- [31] Masaaki Sakakura, Yuhao Lei, Lei Wang, Yan-Hao Yu, and Peter G. Kazansky. “Ultralow-loss geometric phase and polarization shaping by ultrafast laser writing in silica glass”. en. In: *Light: Science & Applications* 9.1 (Feb. 2020), p. 15. ISSN: 2047-7538. DOI: [10.1038/s41377-020-0250-y](https://doi.org/10.1038/s41377-020-0250-y). URL: <https://www.nature.com/articles/s41377-020-0250-y>.
- [32] Yuhao Lei, Masaaki Sakakura, Lei Wang, Yanhao Yu, Huijun Wang, Gholamreza Shayegannad, and Peter G. Kazansky. “High speed ultrafast laser anisotropic nanostructuring by energy deposition control via near-field enhancement”. In: *Optica* 8.11 (Nov. 20, 2021), p. 1365. ISSN: 2334-2536. DOI: [10.1364/OPTICA.433765](https://doi.org/10.1364/OPTICA.433765). URL: <https://opg.optica.org/abstract.cfm?URI=optica-8-11-1365>.
- [33] K. M. Davis, K. Miura, N. Sugimoto, and K. Hirao. “Writing waveguides in glass with a femtosecond laser”. In: *Optics Letters* 21.21 (Nov. 1, 1996), p. 1729. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.21.001729](https://doi.org/10.1364/OL.21.001729). URL: <https://opg.optica.org/abstract.cfm?URI=ol-21-21-1729>.
- [34] Thomas Torsten Winkler et al. “Phase-modulated Optical Data Storage”. U.S. pat. 20250087240A1. Mar. 13, 2025. URL: <https://patents.google.com/patent/US20250087240A1/en>.
- [35] Yasuhiko Shimotsuma, Masaaki Sakakura, Peter G. Kazansky, Martynas Beresna, Jiarong Qiu, Kiyotaka Miura, and Kazuyuki Hirao. “Ultrafast Manipulation of Self-Assembled Form Birefringence in Glass”. In: *Advanced Materials* 22.36 (Sept. 22, 2010), pp. 4039–4043. ISSN: 0935-9648, 1521-4095. DOI: [10.1002/adma.201000921](https://doi.org/10.1002/adma.201000921). URL: <https://advanced.onlinelibrary.wiley.com/doi/10.1002/adma.201000921>.
- [36] Martynas Beresna, Mindaugas Gecevičius, Peter G. Kazansky, Thomas Taylor, and Alexey V. Kavokin. “Exciton mediated self-organization in glass driven by ultrashort light pulses”. In: *Applied Physics Letters* 101.5 (July 30, 2012), p. 053120. ISSN: 0003-6951, 1077-3118. DOI: [10.1063/1.4742899](https://doi.org/10.1063/1.4742899). URL: <https://pubs.aip.org/aip/apl/article/1111408>.

- [37] Yuhao Lei, Huijun Wang, Gholamreza Shayeganrad, and Peter G. Kazansky. “Ultrafast laser nanostructuring in transparent materials for beam shaping and data storage [Invited]”. In: *Optical Materials Express* 12.9 (Sept. 1, 2022), p. 3327. ISSN: 2159-3930. DOI: [10.1364/OME.463151](https://doi.org/10.1364/OME.463151). URL: <https://opg.optica.org/abstract.cfm?URI=ome-12-9-3327>.
- [38] Yuhao Lei, Gholamreza Shayeganrad, Huijun Wang, Masaaki Sakakura, Yanhao Yu, Lei Wang, Dmitrii Kliukin, Linards Skuja, Yuri Svirko, and Peter G. Kazansky. “Efficient ultrafast laser writing with elliptical polarization”. en. In: *Light: Science & Applications* 12.1 (Mar. 2023), p. 74. ISSN: 2047-7538. DOI: [10.1038/s41377-023-01098-2](https://doi.org/10.1038/s41377-023-01098-2). URL: <https://www.nature.com/articles/s41377-023-01098-2>.
- [39] Yang Liao, Yinglong Shen, Lingling Qiao, Danping Chen, Ya Cheng, Koji Sugioka, and Katsumi Midorikawa. “Femtosecond laser nanostructuring in porous glass with sub-50 nm feature sizes”. In: *Optics Letters* 38.2 (Jan. 15, 2013), p. 187. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.38.000187](https://doi.org/10.1364/OL.38.000187). URL: <https://opg.optica.org/abstract.cfm?URI=ol-38-2-187>.
- [40] Robert Buschlinger, Stefan Nolte, and Ulf Peschel. “Self-organized pattern formation in laser-induced multiphoton ionization”. In: *Physical Review B* 89.18 (May 29, 2014), p. 184306. ISSN: 1098-0121, 1550-235X. DOI: [10.1103/PhysRevB.89.184306](https://doi.org/10.1103/PhysRevB.89.184306). URL: <https://link.aps.org/doi/10.1103/PhysRevB.89.184306>.
- [41] Zhonghe Feng. “Microscopy of Phase and Polarisation in Birefringent Structures”. en. MRes Thesis. London, UK: Imperial College London, Sept. 2021.
- [42] Alexandre Rothen. “The Ellipsometer, an Apparatus to Measure Thicknesses of Thin Surface Films”. en. In: *Review of Scientific Instruments* 16.2 (Feb. 1945), pp. 26–30. ISSN: 0034-6748, 1089-7623. DOI: [10.1063/1.1770315](https://doi.org/10.1063/1.1770315). URL: <https://pubs.aip.org/rsi/article/16/2/26/296199/The-Ellipsometer-an-Apparatus-to-Measure>.
- [43] Hiroyuki Fujiwara. *Spectroscopic ellipsometry: principles and applications*. Printed with corrections. Chichester: Wiley, 2009. 369 pp. ISBN: 978-0-470-01608-4.

- [44] N. H. HARTSHORNE. “MODERN APPLICATIONS OF POLARISATION MICROSCOPY. Part I”. In: *Science Progress (1933-)* 49.196 (1961), pp. 601–618. ISSN: 00368504, 20477163. URL: <http://www.jstor.org/stable/43418181>.
- [45] Cornelis Klein and Anthony Philpotts. *Earth Materials 2nd Edition: Introduction to Mineralogy and Petrology*. 2nd ed. Cambridge University Press, Dec. 15, 2016. ISBN: 978-1-107-15540-4 978-1-316-65290-9 978-1-316-60885-2. DOI: [10.1017/9781316652909](https://doi.org/10.1017/9781316652909). URL: <https://www.cambridge.org/highereducation/books/earth-materials-2nd-edition/02B13B6EA8B5F61D07DB410E0FE3A701#contents>.
- [46] Michael Shribak, James LaFountain, David Biggs, and Shinya Inouè. “Orientation-independent differential interference contrast microscopy and its combination with an orientation-independent polarization system”. In: *Journal of Biomedical Optics* 13.1 (2008). ISSN: 10833668. DOI: [10.1117/1.2837406](https://doi.org/10.1117/1.2837406).
- [47] Yang Dong, Ji Qi, Honghui He, Chao He, Shaoxiong Liu, Jian Wu, Daniel S. Elson, and Hui Ma. “Quantitatively characterizing the microstructural features of breast ductal carcinoma tissues in different progression stages by Mueller matrix microscope”. en. In: *Biomedical Optics Express* 8.8 (Aug. 2017), p. 3643. ISSN: 2156-7085, 2156-7085. DOI: [10.1364/BOE.8.003643](https://doi.org/10.1364/BOE.8.003643). URL: <https://opg.optica.org/abstract.cfm?URI=boe-8-8-3643>.
- [48] Matthew R. Foreman, Carlos Macías Romero, and Peter Török. “Determination of the three-dimensional orientation of single molecules”. In: *Optics Letters* 33.9 (May 1, 2008), p. 1020. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.33.001020](https://doi.org/10.1364/OL.33.001020). URL: <https://opg.optica.org/abstract.cfm?URI=ol-33-9-1020>.
- [49] Shalin B. Mehta, Molly McQuilken, Patrick J. La Riviere, Patricia Occhipinti, Amitabh Verma, Rudolf Oldenbourg, Amy S. Gladfelter, and Tomomi Tani. “Dissection of molecular assembly dynamics by tracking orientation and position of single molecules in live cells”. In: *Proceedings of the National Academy of Sciences* 113.42 (Oct. 18, 2016). ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1607674113](https://doi.org/10.1073/pnas.1607674113). URL: <https://pnas.org/doi/full/10.1073/pnas.1607674113>.

- [50] Ezra Bruggeman et al. “POLCAM: instant molecular orientation microscopy for the life sciences”. en. In: *Nature Methods* 21.10 (Oct. 2024), pp. 1873–1883. ISSN: 1548-7091, 1548-7105. DOI: [10.1038/s41592-024-02382-8](https://doi.org/10.1038/s41592-024-02382-8). URL: <https://www.nature.com/articles/s41592-024-02382-8>.
- [51] Michael Shribak and Rudolf Oldenbourg. “Techniques for fast and sensitive measurements of two-dimensional birefringence distributions”. In: *Applied Optics* 42.16 (June 2003), pp. 3009–3017. DOI: [10.1364/AO.42.003009](https://doi.org/10.1364/AO.42.003009). URL: <http://ao.osa.org/abstract.cfm?URI=ao-42-16-3009>.
- [52] Shalin B. Mehta, Michael Shribak, and Rudolf Oldenbourg. “Polarized light imaging of birefringence and diattenuation at high resolution and high sensitivity”. In: *Journal of Optics (United Kingdom)* 15.9 (2013). ISSN: 20408986. DOI: [10.1088/2040-8978/15/9/094007](https://doi.org/10.1088/2040-8978/15/9/094007).
- [53] Sunwoong Hur, Seungri Song, Soocheol Kim, and Chulmin Joo. “Polarization-sensitive differential phase-contrast microscopy”. In: *Optics Letters* 46.2 (2021), pp. 392–395.
- [54] Patrick Ferrand, Marc Allain, and Virginie Chamard. “Ptychography in anisotropic media”. en. In: *Optics Letters* 40.22 (Nov. 2015). Publisher: Optica Publishing Group, p. 5144. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/ol.40.005144](https://doi.org/10.1364/ol.40.005144). URL: <https://opg.optica.org/abstract.cfm?URI=ol-40-22-5144>.
- [55] Arthur Baroni, Marc Allain, Peng Li, Virginie Chamard, and Patrick Ferrand. “Joint estimation of object and probes in vectorial ptychography”. en. In: *Optics Express* 27.6 (Mar. 2019). Publisher: Optica Publishing Group, p. 8143. ISSN: 1094-4087. DOI: [10.1364/oe.27.008143](https://doi.org/10.1364/oe.27.008143). URL: <https://opg.optica.org/abstract.cfm?URI=oe-27-6-8143>.
- [56] Matthew H. Smith, James D. Howe, Jacob B. Woodruff, Miranda A. Miller, George R. Ax Jr., Thomas E. Petty, and Elizabeth A. Sornsin. “Multispectral infrared Stokes imaging polarimeter”. en. In: *SPIE Proceedings*. Ed. by Dennis H. Goldstein and David B. Chenault. Vol. 3754. ISSN: 0277-786X. Denver, CO: SPIE, Oct. 1999, pp. 137–143.

- DOI: 10.1117/12.366324. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=994913>.
- [57] Matthew H. Smith, Jacob B. Woodruff, and James D. Howe. “Beam wander considerations in imaging polarimetry”. en. In: *SPIE Proceedings*. Ed. by Dennis H. Goldstein and David B. Chenault. Vol. 3754. ISSN: 0277-786X. Denver, CO: SPIE, Oct. 1999, pp. 50–54. DOI: 10.1117/12.366359. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=994889>.
- [58] Craig A. Farlow, David B. Chenault, J. L. Pezzaniti, Kevin D. Spradley, and Michael G. Gulley. “Imaging polarimeter development and applications”. In: *SPIE Proceedings*. Ed. by Dennis H. Goldstein, David B. Chenault, Walter G. Egan, and Michael J. Duggin. Vol. 4481. ISSN: 0277-786X. San Diego, CA, USA: SPIE, Jan. 2002, p. 118. DOI: 10.1117/12.452880. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.452880>.
- [59] R.M.A. Azzam. “Division-of-amplitude Photopolarimeter (DOAP) for the Simultaneous Measurement of All Four Stokes Parameters of Light”. In: *Optica Acta: International Journal of Optics* 29.5 (May 1982), pp. 685–689. ISSN: 0030-3909. DOI: 10.1080/713820903. URL: <https://www.tandfonline.com/doi/full/10.1080/713820903>.
- [60] J. Scott Tyo. “Hybrid division of aperture/division of a focal-plane polarimeter for real-time polarization imagery without an instantaneous field-of-view error”. en. In: *Optics Letters* 31.20 (Oct. 2006). Publisher: Optica Publishing Group, p. 2984. ISSN: 0146-9592, 1539-4794. DOI: 10.1364/ol.31.002984. URL: <https://opg.optica.org/abstract.cfm?URI=ol-31-20-2984>.
- [61] J. Larry Pezzaniti and David B. Chenault. “A division of aperture MWIR imaging polarimeter”. en. In: *SPIE Proceedings*. Ed. by Joseph A. Shaw and J. Scott Tyo. ISSN: 0277-786X. San Diego, California, USA: SPIE, Aug. 2005, 58880K. DOI: 10.1117/12.623543. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.623543>.

- [62] Tingkui Mu, Chunmin Zhang, Qiwei Li, and Rongguang Liang. “Error analysis of single-snapshot full-Stokes division-of-aperture imaging polarimeters”. en. In: *Optics Express* 23.8 (Apr. 2015). Publisher: Optica Publishing Group, p. 10822. ISSN: 1094-4087. DOI: [10.1364/oe.23.010822](https://doi.org/10.1364/oe.23.010822). URL: <https://opg.optica.org/abstract.cfm?URI=oe-23-8-10822>.
- [63] Tomohiro Yamazaki, Yasushi Maruyama, Yusuke Uesaka, Motoaki Nakamura, and Yoshihisa Matoba. “Four-Directional Pixel-Wise Polarization CMOS Image Sensor Using Air-Gap Wire Grid”. In: *IEEE International Electron Devices Meeting*. San Francisco, CA, USA: IEEE, 2016, pp. 220–223. ISBN: 978-1-5090-3902-9. DOI: [10.1109/IEDM.2016.7838378](https://doi.org/10.1109/IEDM.2016.7838378).
- [64] Yasushi Maruyama et al. “3.2-MP Back-Illuminated Polarization Image Sensor With Four-Directional Air-Gap Wire Grid and 2.5- μm Pixels”. en. In: *IEEE Transactions on Electron Devices* 65.6 (June 2018). Publisher: Institute of Electrical and Electronics Engineers (IEEE), pp. 2544–2551. ISSN: 0018-9383, 1557-9646. DOI: [10.1109/ted.2018.2829190](https://doi.org/10.1109/ted.2018.2829190). URL: <https://ieeexplore.ieee.org/document/8354948/>.
- [65] Xingzhou Tu, Oliver J. Spires, Xiaobo Tian, Neal Brock, Rongguang Liang, and Stanley Pau. “Division of amplitude RGB full-Stokes camera using micro-polarizer arrays”. en. In: *Optics Express* 25.26 (Dec. 2017). Publisher: Optica Publishing Group, p. 33160. ISSN: 1094-4087. DOI: [10.1364/oe.25.033160](https://doi.org/10.1364/oe.25.033160). URL: <https://opg.optica.org/abstract.cfm?URI=oe-25-26-33160>.
- [66] Sony Semiconductor Solutions Corporation. *PolarSensTM — Polarization Image Sensor*. 2025. URL: <https://www.sony-semicon.com/en/technology/industry/polarsens.html>.
- [67] Bijie Bai, Hongda Wang, Tairan Liu, Yair Rivenson, John FitzGerald, and Aydogan Ozcan. “Pathological crystal imaging with single-shot computational polarized light microscopy”. In: *Journal of Biophotonics* 13.1 (2020), pp. 1–7. ISSN: 18640648. DOI: [10.1002/jbio.201960036](https://doi.org/10.1002/jbio.201960036).

- [68] F. Zernike. “How I discovered phase contrast”. In: *Science* 121.3141 (1955), pp. 345–349. ISSN: 00368075. DOI: [10.1126/science.121.3141.345](https://doi.org/10.1126/science.121.3141.345).
- [69] Douglas B Murphy and Michael W Davidson. *Fundamentals of Light Microscopy and Electronic Imaging*. Somerset, United States: John Wiley & Sons, Incorporated, 2012. ISBN: 978-1-118-38291-2. URL: <http://ebookcentral.proquest.com/lib/imperial/detail.action?docID=918267>.
- [70] Haoran Liu, Zhenyang Li, Xuefeng Lei, Shuming Shi, Zhenhai Liu, Maoxin Song, Zhiyuan Zhou, Zhenwei Qiu, Qiang Cong, and Jin Hong. “Image interpolation methods for division of focal plane polarimeters: a review”. en. In: *Second Conference on Space, Atmosphere, Marine, and Environmental Optics (SAME 2024)*. Ed. by Weibiao Chen, Yongchao Zheng, and Dengxin Hua. Hangzhou, China: SPIE, July 2024, p. 22. DOI: [10.1117/12.3032319](https://doi.org/10.1117/12.3032319). URL: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/13189/3032319/Image-interpolation-methods-for-division-of-focal-plane-polarimeters/10.1117/12.3032319.full>.
- [71] P. R. T. Munro and P. Török. “Properties of high-numerical-aperture Mueller-matrix polarimeters”. In: *Optics Letters* 33.21 (Nov. 1, 2008), p. 2428. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.33.002428](https://doi.org/10.1364/OL.33.002428). URL: <https://opg.optica.org/abstract.cfm?URI=ol-33-21-2428>.
- [72] Russell A. Chipman, Wai-Sze Tiffany Lam, and Garam Young. *Polarized light and optical systems*. eng. Optical sciences and applications of light. Boca Raton London New York: CRC Press, 2019. ISBN: 978-1-4987-0056-6 978-1-4987-0057-3 978-1-351-12907-7.
- [73] Dennis Goldstein. *Polarized Light*. New York: CRC Press, 2003. ISBN: 0-8247-4053-X.
- [74] Jun J. Sakurai and Jim Napolitano. *Modern quantum mechanics*. eng. 2. ed., international ed. Boston, Mass.: Addison-Wesley, Pearson, 2011. ISBN: 978-0-8053-8291-4.
- [75] YongKeun Park, Christian Depeursinge, and Gabriel Popescu. “Quantitative phase imaging in biomedicine”. en. In: *Nature Photonics* 12.10 (Oct. 2018), pp. 578–589. ISSN: 1749-4885, 1749-4893. DOI: [10.1038/s41566-018-0253-x](https://doi.org/10.1038/s41566-018-0253-x). URL: <https://www.nature.com/articles/s41566-018-0253-x>.

- [76] Zhuo Wang, Larry Millet, Mustafa Mir, Huafeng Ding, Sakulsuk Unarunotai, John Rogers, Martha U Gillette, and Gabriel Popescu. “Spatial light interference microscopy (SLIM)”. In: *Optics Express* 19.2 (2011), pp. 1016–1026.
- [77] Ulrich Kubitscheck. *Fluorescence microscopy: from principles to biological applications*. eng. 2nd ed. Weinheim: Wiley-VCH, 2017. ISBN: 978-3-527-68773-2 978-3-527-68772-5.
- [78] Richard W. Taylor and Vahid Sandoghdar. “Interferometric Scattering (iSCAT) Microscopy and Related Techniques”. en. In: *Label-Free Super-Resolution Microscopy*. Ed. by Vasily Astratov. Series Title: Biological and Medical Physics, Biomedical Engineering. Cham: Springer International Publishing, 2019, pp. 25–65. ISBN: 978-3-030-21721-1. DOI: [10.1007/978-3-030-21722-8_2](https://doi.org/10.1007/978-3-030-21722-8_2). URL: http://link.springer.com/10.1007/978-3-030-21722-8_2.
- [79] Cheng Liu, Shouyu Wang, and Suhas P Veetil. *Computational Optical Phase Imaging*. Singapore: Springer, 2022. ISBN: 978-981-19-1640-3. DOI: [10.1007/978-981-19-1641-0](https://doi.org/10.1007/978-981-19-1641-0).
- [80] Joseph W Goodman. *Introduction to Fourier Optics*. New York: McGraw-Hill, 1996. ISBN: 0-07-024254-2.
- [81] C.E. Shannon. “Communication in the Presence of Noise”. In: *Proceedings of the IRE* 37.1 (Jan. 1949), pp. 10–21. ISSN: 0096-8390. DOI: [10.1109/JRPROC.1949.232969](https://doi.org/10.1109/JRPROC.1949.232969). URL: <http://ieeexplore.ieee.org/document/1697831/>.
- [82] Stephen K. Park, Robert Schowengerdt, and Mary-Anne Kaczynski. “Modulation-transfer-function analysis for sampled image systems”. In: *Applied Optics* 23.15 (Aug. 1, 1984), p. 2572. ISSN: 0003-6935, 1539-4522. DOI: [10.1364/AO.23.002572](https://doi.org/10.1364/AO.23.002572). URL: <https://opg.optica.org/abstract.cfm?URI=ao-23-15-2572>.
- [83] Mike Marchywka and Dennis G. Socker. “Modulation transfer function measurement technique for small-pixel detectors”. In: *Applied Optics* 31.34 (Dec. 1, 1992), p. 7198. ISSN: 0003-6935, 1539-4522. DOI: [10.1364/AO.31.007198](https://doi.org/10.1364/AO.31.007198). URL: <https://opg.optica.org/abstract.cfm?URI=ao-31-34-7198>.

- [84] J. Scott Tyo, Charles F. LaCasse, and Bradley M. Ratliff. “Total elimination of sampling errors in polarization imagery obtained with integrated microgrid polarimeters”. en. In: *Optics Letters* 34.20 (Oct. 2009), p. 3187. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.34.003187](https://doi.org/10.1364/OL.34.003187). URL: <https://opg.optica.org/abstract.cfm?URI=ol-34-20-3187>.
- [85] Alan V. Oppenheim and Roland W. Schaffer. *Discrete-time signal processing*. Third edition, Pearson New international edition. Always learning. Harlow: Pearson, 2014. 1 p. ISBN: 978-1-292-03815-5.
- [86] Sebastiano Battiato. *Image Processing for Embedded Devices*. eng. 1st ed. Sharjah: Bentham Science Publishers, 2010. ISBN: 978-1-60805-170-0.
- [87] European Machine Vision Association. *EMVA 1288 – Standard for Characterization of Image Sensors and Cameras, Release 3.1a*. Tech. rep. 1288, Release 3.1a. Version 3.1a. Barcelona, Spain: European Machine Vision Association, Dec. 2016. DOI: [10.5281/zenodo.235942](https://doi.org/10.5281/zenodo.235942). URL: <https://doi.org/10.5281/zenodo.235942>.
- [88] Joseph W. Goodman. *Statistical optics*. Second edition. Wiley series in pure and applied optics. Hoboken, New Jersey: John Wiley & Sons Inc, 2015. 1 p. ISBN: 978-1-119-00945-0 978-1-119-00946-7.
- [89] James R. Janesick, Tom Elliott, Stewart Collins, Morley M. Blouke, and Jack Freeman. “Scientific Charge-Coupled Devices”. en. In: *Optical Engineering* 26.8 (Aug. 1987). ISSN: 0091-3286. DOI: [10.1117/12.7974139](https://doi.org/10.1117/12.7974139). URL: <http://opticalengineering.spiedigitallibrary.org/article.aspx?doi=10.1117/12.7974139>.
- [90] Christopher W Dunsby. “Wide-field coherence-gated imaging techniques including photorefractive holography”. PhD Thesis. London, UK: University of London, 2003.
- [91] Robin Diekmann, Joran Deschamps, Yiming Li, Aline Tschanz, Maurice Kahnwald, Ulf Matti, and Jonas Ries. “Photon-free (s)CMOS camera characterization for artifact reduction in high- and super-resolution microscopy”. In: *Nature Communications* 13.3362 (2022). Publisher: Springer Nature, pp. 1–9. DOI: [10.1038/s41467-022-30907-2](https://doi.org/10.1038/s41467-022-30907-2).

- [92] Arthur D. Edelstein, Mark A. Tsuchida, Nenad Amodaj, Henry Pinkard, Ronald D. Vale, and Nico Stuurman. “Advanced methods of microscope control using μ Manager software”. In: *Journal of Biological Methods* 1.2 (Nov. 2014), p. 1. ISSN: 2326-9901. DOI: [10.14440/jbm.2014.36](https://doi.org/10.14440/jbm.2014.36). URL: <https://polscientific.com/journal/JBM/1/2/10.14440/jbm.2014.36>.
- [93] Johannes Schindelin et al. “Fiji: an open-source platform for biological-image analysis”. en. In: *Nature Methods* 9.7 (July 2012), pp. 676–682. ISSN: 1548-7091, 1548-7105. DOI: [10.1038/nmeth.2019](https://doi.org/10.1038/nmeth.2019). URL: <https://www.nature.com/articles/nmeth.2019>.
- [94] Kleckner Lab. *simple_pyspin: A Simplified Interface for FLIR Spinnaker SDK*. 2020. URL: https://github.com/klecknerlab/simple_pyspin.
- [95] FLIR Integrated Imaging Solutions Inc. *EMVA 1288 Imaging Performance: Blackfly S BFS-U3-51S5*. Tech. rep. Specification sheet. FLIR Integrated Imaging Solutions Inc., July 2023.
- [96] Lingfei Song and Hua Huang. “Fixed Pattern Noise Removal Based on a Semi-Calibration Method”. en. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45.10 (Oct. 2023), pp. 11842–11855. ISSN: 0162-8828, 2160-9292, 1939-3539. DOI: [10.1109/TPAMI.2023.3274826](https://doi.org/10.1109/TPAMI.2023.3274826). URL: <https://ieeexplore.ieee.org/document/10122709/>.
- [97] Qiang Wen, Siqu Zhu, Shichang Liu, Houwei Guo, Jingwen Jin, and Yaoxin Zhu. “Image sensor correction algorithm for photon transfer curve based on neural network”. en. In: *Optical Engineering* 59.06 (June 2020), p. 1. ISSN: 0091-3286. DOI: [10.1117/1.OE.59.6.067102](https://doi.org/10.1117/1.OE.59.6.067102). URL: <https://www.spiedigitallibrary.org/journals/optical-engineering/volume-59/issue-06/067102/Image-sensor-correction-algorithm-for-photon-transfer-curve-based-on/10.1117/1.OE.59.6.067102.full>.
- [98] James R. Janesick. *Photon transfer: DN [λ]*. Bellingham, Wash: SPIE, 2007. ISBN: 978-0-8194-6722-5.
- [99] James A. Seibert, John M. Boone, and Karen K. Lindfors. “Flat-field correction technique for digital detectors”. In: *Medical Imaging 1998: Physics of Medical Imaging* 3336.May (1998), p. 348. ISSN: 0277786X. DOI: [10.1117/12.317034](https://doi.org/10.1117/12.317034).

- [100] Matthew R Foreman and Peter Török. “Fundamental limits in single-molecule orientation measurements”. In: *New Journal of Physics* 13.9 (Sept. 7, 2011), p. 093013. ISSN: 1367-2630. DOI: [10.1088/1367-2630/13/9/093013](https://doi.org/10.1088/1367-2630/13/9/093013). URL: <https://iopscience.iop.org/article/10.1088/1367-2630/13/9/093013>.
- [101] Rasheed M. Azzam, Rasheed Mohammed Abdel-Goward Azzam, Nicholas Mitchell Bashara, and Nicholas M. Bashara. *Ellipsometry and polarized light*. eng. Repr. 1989. North-Holland personal library. Amsterdam: North-Holland, 1989. ISBN: 978-0-7204-0694-8.
- [102] Xavier Theillier, Sylvain Rivet, Matthieu Dubreuil, and Yann Le Grand. “Swept-wavelength null polarimeter for high-speed weak anisotropy measurements”. en. In: *Optics Express* 30.11 (May 2022), p. 18889. ISSN: 1094-4087. DOI: [10.1364/OE.454193](https://doi.org/10.1364/OE.454193). URL: <https://opg.optica.org/abstract.cfm?URI=oe-30-11-18889>.
- [103] Victor Twersky. “Form and Intrinsic Birefringence.” In: *Journal of the Optical Society of America* 65.3 (1975), pp. 239–245. ISSN: 0030-3941. DOI: [10.1364/JOSA.65.000239](https://doi.org/10.1364/JOSA.65.000239).
- [104] C.J.R. Sheppard. “Electromagnetic field in the focal region of wide-angular annular lens and mirror systems”. en. In: *IEE Journal on Microwaves, Optics and Acoustics* 2.5 (1978), p. 163. ISSN: 03086976. DOI: [10.1049/ij-moa.1978.0035](https://doi.org/10.1049/ij-moa.1978.0035). URL: <https://digital-library.theiet.org/content/journals/10.1049/ij-moa.1978.0035>.
- [105] Emil Wolf. “Electromagnetic diffraction in optical systems - I. An integral representation of the image field”. en. In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 253.1274 (Dec. 1959), pp. 349–357. ISSN: 0080-4630, 2053-9169. DOI: [10.1098/rspa.1959.0199](https://doi.org/10.1098/rspa.1959.0199). URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1959.0199>.
- [106] C. J. R. Sheppard and H. J. Matthews. “Imaging in high-aperture optical systems”. en. In: *Journal of the Optical Society of America A* 4.8 (Aug. 1987), p. 1354. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/JOSAA.4.001354](https://doi.org/10.1364/JOSAA.4.001354). URL: <https://opg.optica.org/abstract.cfm?URI=josaa-4-8-1354>.
- [107] John B. DeVelis and George B. Parrent. “Transfer Function for Cascaded Optical Systems”. en. In: *Journal of the Optical Society of America* 57.12 (Dec. 1967), p. 1486.

- ISSN: 0030-3941. DOI: [10.1364/JOSA.57.001486](https://doi.org/10.1364/JOSA.57.001486). URL: <https://opg.optica.org/abstract.cfm?URI=josa-57-12-1486>.
- [108] P.H Van Cittert. “Die Wahrscheinliche Schwingungsverteilung in Einer von Einer Lichtquelle Direkt Oder Mittels Einer Linse Beleuchteten Ebene”. de. In: *Physica* 1.1-6 (Jan. 1934), pp. 201–210. ISSN: 00318914. DOI: [10.1016/S0031-8914\(34\)90026-4](https://doi.org/10.1016/S0031-8914(34)90026-4). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0031891434900264>.
- [109] Andrey S. Ostrovsky, Gabriel Martínez-Niconoff, Patricia Martínez-Vara, and Miguel A. Olvera-Santamaría. “The van Cittert-Zernike theorem for electromagnetic fields”. In: *Optics Express* 17.3 (2009), p. 1746. ISSN: 10944087. DOI: [10.1364/oe.17.001746](https://doi.org/10.1364/oe.17.001746).
- [110] H. Gamo. “Matrix Treatment of Partial Coherence”. In: *Progress in Optics, Volume III*. Ed. by Emil Wolf. Amsterdam: North-Holland Publishing Co., 1964. Chap. III, pp. 187–332.
- [111] Andrey S. Ostrovsky. *Coherent-mode representations in optics*. Bellingham, Wash: SPIE Press, 2006. 86 pp. ISBN: 978-0-8194-6350-0 978-0-8194-7835-1.
- [112] Brian J. Thompson. “Image Formation with Partially Coherent Light”. In: *Progress in Optics, Volume VII*. Ed. by Emil Wolf. Section: IV. Amsterdam and London: North-Holland Publishing Company, 1969, pp. 169–230.
- [113] F. Zernike. “The concept of degree of coherence and its application to optical problems”. In: *Physica* 5.8 (Aug. 1938), pp. 785–795. ISSN: 00318914. DOI: [10.1016/S0031-8914\(38\)80203-2](https://doi.org/10.1016/S0031-8914(38)80203-2). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0031891438802032>.
- [114] Chao Zuo, Jiaji Li, Jiasong Sun, Yao Fan, Jialin Zhang, Linpeng Lu, Runnan Zhang, Bowen Wang, Lei Huang, and Qian Chen. “Transport of intensity equation: a tutorial”. In: *Optics and Lasers in Engineering* 135 (June 2020). Publisher: Elsevier Ltd. ISSN: 01438166. DOI: [10.1016/j.optlaseng.2020.106187](https://doi.org/10.1016/j.optlaseng.2020.106187).
- [115] Heinrich Kirchauer. “Photolithography Simulation”. PhD dissertation. Vienna, Austria: Technische Universität Wien, Fakultät für Elektrotechnik, Apr. 17, 1998. URL: <https://www.iue.tuwien.ac.at/phd/kirchauer/diss.html>.

- [116] Max Born and Emil Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. eng. 6. ed., reprinted (with corrections). Oxford: Pergamon Press, 1993. ISBN: 978-0-08-026481-3 978-0-08-026482-0.
- [117] Huihui Liu, Sunil Kumar, Edwin Garcia, William Flanagan, Jonathan Lightley, Christopher Dunsby, and Paul M.W. French. “Open-source implementation of polarisation-resolved single-shot differential phase contrast microscopy (pDPC) on a modular openFrame-based microscope”. en. In: *HardwareX* 21 (Mar. 2025), e00622. ISSN: 24680672. DOI: [10.1016/j.ohx.2024.e00622](https://doi.org/10.1016/j.ohx.2024.e00622). URL: <https://linkinghub.elsevier.com/retrieve/pii/S2468067224001160>.
- [118] Shalin B. Mehta and Colin J.R. Sheppard. “Using the phase-space imager to analyze partially coherent imaging systems: bright-field, phase contrast, differential interference contrast, differential phase contrast, and spiral phase contrast”. en. In: *Journal of Modern Optics* 57.9 (May 2010), pp. 718–739. ISSN: 0950-0340, 1362-3044. DOI: [10.1080/09500340.2010.481729](https://doi.org/10.1080/09500340.2010.481729). URL: <http://www.tandfonline.com/doi/abs/10.1080/09500340.2010.481729>.
- [119] Lei Tian and Laura Waller. “Quantitative differential phase contrast imaging in an LED array microscope”. en. In: *Optics Express* 23.9 (May 2015), p. 11394. ISSN: 1094-4087. DOI: [10.1364/OE.23.011394](https://doi.org/10.1364/OE.23.011394). URL: <https://opg.optica.org/abstract.cfm?URI=oe-23-9-11394>.
- [120] Ranjan Kalita, William Flanagan, Jonathan Lightley, Sunil Kumar, Yuriy Alexandrov, Edwin Garcia, Mark Hintze, Michalis Barkoulas, Chris Dunsby, and Paul M. W. French. “Single-shot phase contrast microscopy using polarisation-resolved differential phase contrast”. en. In: *Journal of Biophotonics* 14.12 (Dec. 2021), e202100144. ISSN: 1864-063X, 1864-0648. DOI: [10.1002/jbio.202100144](https://doi.org/10.1002/jbio.202100144). URL: <https://onlinelibrary.wiley.com/doi/10.1002/jbio.202100144>.
- [121] Guoan Zheng, Roarke Horstmeyer, and Changhuei Yang. “Wide-field, high-resolution Fourier ptychographic microscopy”. en. In: *Nature Photonics* 7.9 (Sept. 2013), pp. 739–

745. ISSN: 1749-4885, 1749-4893. DOI: [10.1038/nphoton.2013.187](https://doi.org/10.1038/nphoton.2013.187). URL: <https://www.nature.com/articles/nphoton.2013.187>.
- [122] H H Hopkins and P M Barham. “The Influence of the Condenser on Microscopic Resolution”. In: *Proceedings of the Physical Society. Section B* 63.10 (Oct. 1, 1950), pp. 737–744. ISSN: 0370-1301. DOI: [10.1088/0370-1301/63/10/301](https://doi.org/10.1088/0370-1301/63/10/301). URL: <https://iopscience.iop.org/article/10.1088/0370-1301/63/10/301>.
- [123] H. H. Hopkins. “On the diffraction theory of optical images”. In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 217.1130 (May 7, 1953), pp. 408–432. ISSN: 0080-4630, 2053-9169. DOI: [10.1098/rspa.1953.0071](https://doi.org/10.1098/rspa.1953.0071). URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1953.0071>.
- [124] Tatsuro Otaki. “Artifact halo reduction in phase contrast microscopy using apodization”. In: *Optical Review* 7.2 (2000), pp. 119–122. ISSN: 13406000. DOI: [10.1007/s10043-000-0119-5](https://doi.org/10.1007/s10043-000-0119-5).
- [125] Christian Maurer, Alexander Jesacher, Stefan Bernet, and Monika Ritsch-Marte. “Phase contrast microscopy with full numerical aperture illumination”. In: *Optics Express* 16.24 (2008), p. 19821. ISSN: 1094-4087. DOI: [10.1364/oe.16.019821](https://doi.org/10.1364/oe.16.019821).
- [126] Tan H. Nguyen, Mikhail Kandel, Haadi M. Shakir, Catherine Best-Popescu, Jyothi Arikath, Minh N. Do, and Gabriel Popescu. “Halo-free Phase Contrast Microscopy”. In: *Scientific Reports* 7 (2017), pp. 1–10. ISSN: 20452322. DOI: [10.1038/srep44034](https://doi.org/10.1038/srep44034).
- [127] Ki-Sik Kim and Dae-Yoon Park. “Variation of Global Coherence on Propagation in Coherent Mode Representation”. In: *J. Opt. Soc. Korea* 10.4 (Dec. 2006). Publisher: Optica Publishing Group, pp. 162–168. URL: <https://opg.optica.org/josk/abstract.cfm?URI=josk-10-4-162>.
- [128] David Voelz. *Computational Fourier Optics*. Bellingham, Washington: SPIE Press, 2011. 232 pp.
- [129] F.J. Harris. “On the use of windows for harmonic analysis with the discrete Fourier transform”. In: *Proceedings of the IEEE* 66.1 (1978), pp. 51–83. ISSN: 0018-9219. DOI: [10.1109/PROC.1978.10837](https://doi.org/10.1109/PROC.1978.10837). URL: <http://ieeexplore.ieee.org/document/1455106/>.

- [130] E. Wolf. “A new description of second-order coherence phenomena in the space-frequency domain”. In: *Optics in Four Dimensions-1980*. Ensenada, Mexico, 1980, pp. 42–48. DOI: [10.1063/1.32286](https://doi.org/10.1063/1.32286). URL: <https://pubs.aip.org/aip/acp/article/65/1/42-48/583901>.
- [131] E. Wolf. “New spectral representation of random sources and of the partially coherent fields that they generate”. In: *Optics Communications* 38.1 (July 1981), pp. 3–6. ISSN: 00304018. DOI: [10.1016/0030-4018\(81\)90295-9](https://doi.org/10.1016/0030-4018(81)90295-9). URL: <https://linkinghub.elsevier.com/retrieve/pii/0030401881902959>.
- [132] Jack D. Gaskill. *Linear systems, Fourier transforms and optics*. Wiley series in pure and applied optics. New York, NY: Wiley, 1978. 554 pp. ISBN: 978-0-471-29288-3.
- [133] A Starikov and E Wolf. “Coherent-mode representation of Gaussian Schell-model Sources and of their radiation fields”. In: *Journal of the Optical Society of America* 72.7 (1982), pp. 923–928.
- [134] Bahaa E.A. Saleh. “Optical Bilinear Transformations: General Properties”. In: *Optica Acta: International Journal of Optics* 26.6 (June 1979), pp. 777–799. ISSN: 0030-3909. DOI: [10.1080/713820055](https://doi.org/10.1080/713820055). URL: <https://www.tandfonline.com/doi/full/10.1080/713820055>.
- [135] R. Castañeda and F.F. Medina. “Partially coherent imaging with Schell-model beams”. In: *Optics & Laser Technology* 29.4 (June 1997), pp. 165–170. ISSN: 00303992. DOI: [10.1016/S0030-3992\(97\)00006-6](https://doi.org/10.1016/S0030-3992(97)00006-6). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0030399297000066>.
- [136] Tiago E.C. Magalhães and José M. Rebordão. “PyWolf: A PyOpenCL implementation for simulating the propagation of partially coherent light”. In: *Computer Physics Communications* 276 (July 2022), p. 108336. ISSN: 00104655. DOI: [10.1016/j.cpc.2022.108336](https://doi.org/10.1016/j.cpc.2022.108336). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0010465522000546>.
- [137] Tiago E.C. Magalhães and Daniela O. Santos. “A new version of PyWolf for the propagation of partially coherent light in media other than free space”. In: *Computer Physics*

- Communications* 294 (Jan. 2024), p. 108899. ISSN: 00104655. DOI: [10.1016/j.cpc.2023.108899](https://doi.org/10.1016/j.cpc.2023.108899). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0010465523002448>.
- [138] Kenji Yamazoe. “Two matrix approaches for aerial image formation obtained by extending and modifying the transmission cross coefficients”. In: *Journal of the Optical Society of America A* 27.6 (June 1, 2010), p. 1311. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/JOSAA.27.001311](https://doi.org/10.1364/JOSAA.27.001311). URL: <https://opg.optica.org/abstract.cfm?URI=josaa-27-6-1311>.
- [139] Emil Wolf. *Introduction to the theory of coherence and polarization of light*. 1. publ. Cambridge: Cambridge Univ. Press, 2007. 222 pp. ISBN: 978-0-521-82211-4.
- [140] James Mercer. “XVI. Functions of positive and negative type, and their connection the theory of integral equations”. In: *Philosophical Transactions of the Royal Society A* 209.441 (Jan. 1909), pp. 415–446. ISSN: 0264-3952, 2053-9258. DOI: [10.1098/rsta.1909.0016](https://doi.org/10.1098/rsta.1909.0016). URL: <https://royalsocietypublishing.org/doi/10.1098/rsta.1909.0016>.
- [141] Norbert Wiener. “Coherency Matrices and Quantum”. In: *Journal of Mathematics and Physics* 7.1 (Apr. 1928), pp. 109–125. ISSN: 0097-1421. DOI: [10.1002/sapm192871109](https://doi.org/10.1002/sapm192871109). URL: <https://onlinelibrary.wiley.com/doi/10.1002/sapm192871109>.
- [142] E. Wolf. “Coherence properties of partially polarized electromagnetic radiation”. In: *Il Nuovo Cimento* 13.6 (Sept. 1959), pp. 1165–1181. ISSN: 0029-6341, 1827-6121. DOI: [10.1007/BF02725127](https://doi.org/10.1007/BF02725127). URL: <http://link.springer.com/10.1007/BF02725127>.
- [143] Franco Gori, Massimo Santarsiero, Raja Simon, Gemma Piquero, Riccardo Borghi, and Giorgio Guattari. “Coherent-mode decomposition of partially polarized, partially coherent sources”. In: *Journal of the Optical Society of America A* 20.1 (Jan. 1, 2003), p. 78. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/JOSAA.20.000078](https://doi.org/10.1364/JOSAA.20.000078). URL: <https://opg.optica.org/abstract.cfm?URI=josaa-20-1-78>.
- [144] Matthew R. Foreman and Peter Török. “Focusing of spatially inhomogeneous partially coherent, partially polarized electromagnetic fields”. In: *Journal of the Optical Society of America A* 26.11 (Nov. 1, 2009), p. 2470. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/](https://doi.org/10.1364/)

- JOSAA.26.002470. URL: <https://opg.optica.org/abstract.cfm?URI=josaa-26-11-2470>.
- [145] A. Schell. “A technique for the determination of the radiation pattern of a partially coherent aperture”. In: *IEEE Transactions on Antennas and Propagation* 15.1 (Jan. 1967), pp. 187–188. ISSN: 0018-926X, 1558-2221. DOI: [10.1109/TAP.1967.1138864](https://doi.org/10.1109/TAP.1967.1138864). URL: <http://ieeexplore.ieee.org/document/1138864/>.
- [146] K.A. Nugent. “A generalization of Schell’s theorem”. In: *Optics Communications* 79.5 (Nov. 1990), pp. 267–269. ISSN: 00304018. DOI: [10.1016/0030-4018\(90\)90066-3](https://doi.org/10.1016/0030-4018(90)90066-3). URL: <https://linkinghub.elsevier.com/retrieve/pii/0030401890900663>.
- [147] F. Gori, G. Guattari, and C. Padovani. “Modal expansion for J_0 -correlated Schell-model sources”. In: *Optics Communications* 64.4 (1987), pp. 311–316. ISSN: 00304018. DOI: [10.1016/0030-4018\(87\)90242-2](https://doi.org/10.1016/0030-4018(87)90242-2).
- [148] F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds. *NIST Digital Library of Mathematical Functions*. Version 1.2.4. Release 1.2.4. Mar. 15, 2025. URL: <https://dlmf.nist.gov/>.
- [149] P. De Santis, F. Gori, G. Guattari, and C. Palma. “Synthesis of partially coherent fields”. In: *Journal of the Optical Society of America A* 3.8 (1986), p. 1258. ISSN: 1084-7529. DOI: [10.1364/josaa.3.001258](https://doi.org/10.1364/josaa.3.001258).
- [150] E Wolf and E Collett. “Partially coherent sources which produce the same far-field intensity distribution as a laser”. In: *Optics Communications* 25.3 (1978), pp. 293–296.
- [151] F Gori. “Mode Propagation of the Field Generated by Collett-Wolf Schell-model Sources”. In: *Optics Communications* 46.3 (1983), pp. 149–154. DOI: [10.1016/0030-4018\(83\)90266-3](https://doi.org/10.1016/0030-4018(83)90266-3).
- [152] Ari T. Friberg and Ronald J. Sudol. “The spatial coherence properties of gaussian schell-model beams”. In: *Optica Acta* 30.8 (1983), pp. 1075–1097. ISSN: 00303909. DOI: [10.1080/713821334](https://doi.org/10.1080/713821334).

- [153] Ari T. Friberg and Jari Turunen. “Imaging of Gaussian Schell-model sources”. In: *Journal of the Optical Society of America A* 5.5 (May 1, 1988), p. 713. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/JOSAA.5.000713](https://doi.org/10.1364/JOSAA.5.000713). URL: <https://opg.optica.org/abstract.cfm?URI=josaa-5-5-713>.
- [154] Eero Tervonen, Jari Turunen, and Ari T. Friberg. “Gaussian Schell-model beams generated with synthetic acousto-optic holograms”. In: *Journal of the Optical Society of America A* 9.5 (May 1, 1992), p. 796. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/JOSAA.9.000796](https://doi.org/10.1364/JOSAA.9.000796). URL: <https://opg.optica.org/abstract.cfm?URI=josaa-9-5-796>.
- [155] Hanna Lajunen and Toni Saastamoinen. “Propagation characteristics of partially coherent beams with spatially varying correlations”. In: *Optics Letters* 36.20 (Oct. 15, 2011), p. 4104. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.36.004104](https://doi.org/10.1364/OL.36.004104). URL: <https://opg.optica.org/abstract.cfm?URI=ol-36-20-4104>.
- [156] W. H. Carter and E. Wolf. “Coherence and radiometry with quasihomogeneous planar sources”. In: *Journal of the Optical Society of America* 67.6 (1977), p. 785. ISSN: 0030-3941. DOI: [10.1364/josa.67.000785](https://doi.org/10.1364/josa.67.000785).
- [157] H. Kogelnik and T. Li. “Laser Beams and Resonators”. In: *Appl. Opt.* 5.10 (Oct. 1966). Publisher: Optica Publishing Group, pp. 1550–1567. DOI: [10.1364/AO.5.001550](https://doi.org/10.1364/AO.5.001550). URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-5-10-1550>.
- [158] Dieter Meschede. *Optics, Light and Lasers: The Practical Approach to Modern Aspects of Photonics and Laser Physics: Second Edition*. Publication Title: Optics, Light and Lasers: The Practical Approach to Modern Aspects of Photonics and Laser Physics: Second Edition. 2008. DOI: [10.1002/9783527618873](https://doi.org/10.1002/9783527618873).
- [159] M. Plancherel and W. Rotach. “Sur les valeurs asymptotiques des polynômes d’Hermite $H_n(x) = (-i)^n e^{x^2/2} \frac{d^n}{dx^n} (e^{-x^2/2})$ ”. In: *Commentarii Mathematici Helvetici* 1.1 (Dec. 1929), pp. 227–254. ISSN: 0010-2571, 1420-8946. DOI: [10.1007/BF01208365](https://doi.org/10.1007/BF01208365). URL: <https://link.springer.com/10.1007/BF01208365>.
- [160] Walter Rotach. “Reihenentwicklungen einer willkürlichen Funktion nach Hermite’schen und Laguerre’schen Polynomen”. Artwork Size: 33 S. Medium: application/pdf Pages:

- 33 S. PhD thesis. ETH Zurich, 1925. DOI: [10.3929/ETHZ-A-000092029](https://doi.org/10.3929/ETHZ-A-000092029). URL: <http://hdl.handle.net/20.500.11850/133495>.
- [161] Gabor Szegő. *Orthogonal polynomials*. 4th ed. American mathematical society 23. Providence (R.I.): American mathematical society, 1939. ISBN: 978-0-8218-1023-1.
- [162] Bin Zhang and Baida Lü. “Transformation of Gaussian Schell-model beams and their coherent-mode representation”. In: *Journal of Optics - Nouvelle Revue d’Optique* 27.3 (1996), pp. 99–103. ISSN: 0150536X. DOI: [10.1088/0150-536X/27/3/001](https://doi.org/10.1088/0150-536X/27/3/001).
- [163] Egon Moecklin. “Asymptotische Entwicklungen der Laguerreschen Polynome”. Artwork Size: 23 S. Medium: application/pdf Pages: 23 S. PhD thesis. ETH Zurich, 1934. DOI: [10.3929/ETHZ-A-000092417](https://doi.org/10.3929/ETHZ-A-000092417). URL: <http://hdl.handle.net/20.500.11850/133650>.
- [164] Gilbert Strang. *Linear algebra and its applications*. 4. ed., internat. student ed. Belmont, Calif.: Thomson, Brooks/Cole, 2006. 487 pp. ISBN: 978-0-534-42200-4 978-0-03-010567-8 978-81-315-0172-6.
- [165] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (Sept. 17, 2020), pp. 357–362. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2). URL: <https://www.nature.com/articles/s41586-020-2649-2>.
- [166] Hai Gong, Wenjun Guo, and Mark A. A. Neil. “GPU-accelerated real-time reconstruction in Python of three-dimensional datasets from structured illumination microscopy with hexagonal patterns”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2199 (June 14, 2021), p. 20200162. ISSN: 1364-503X, 1471-2962. DOI: [10.1098/rsta.2020.0162](https://doi.org/10.1098/rsta.2020.0162). URL: <https://royalsocietypublishing.org/doi/10.1098/rsta.2020.0162>.
- [167] B.R. Boruah and M.A.A. Neil. “Focal field computation of an arbitrarily polarized beam using fast Fourier transforms”. In: *Optics Communications* 282.24 (Dec. 2009), pp. 4660–4667. ISSN: 00304018. DOI: [10.1016/j.optcom.2009.09.019](https://doi.org/10.1016/j.optcom.2009.09.019). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0030401809008773>.

- [168] Charles N. Christensen, Edward N. Ward, Meng Lu, Pietro Lio, and Clemens F. Kaminski. “ML-SIM: universal reconstruction of structured illumination microscopy images using transfer learning”. In: *Biomedical Optics Express* 12.5 (May 1, 2021), p. 2720. ISSN: 2156-7085, 2156-7085. DOI: [10.1364/BOE.414680](https://doi.org/10.1364/BOE.414680). URL: <https://opg.optica.org/abstract.cfm?URI=boe-12-5-2720>.
- [169] Chang Qiao et al. “Zero-shot learning enables instant denoising and super-resolution in optical fluorescence microscopy”. In: *Nature Communications* 15.1 (May 16, 2024), p. 4180. ISSN: 2041-1723. DOI: [10.1038/s41467-024-48575-9](https://doi.org/10.1038/s41467-024-48575-9). URL: <https://www.nature.com/articles/s41467-024-48575-9>.
- [170] *Imperial College Research Computing Service*. Imperial College London. 2021. URL: <http://doi.org/10.14469/hpc/2232>.
- [171] Shalin B. Mehta and Rudolf Oldenbourg. “Image simulation for biological microscopy: microlith”. In: *Biomedical Optics Express* 5.6 (June 1, 2014), p. 1822. ISSN: 2156-7085, 2156-7085. DOI: [10.1364/BOE.5.001822](https://doi.org/10.1364/BOE.5.001822). URL: <https://opg.optica.org/abstract.cfm?URI=boe-5-6-1822>.
- [172] Craig F. Bohren and Donald R. Huffman. *Absorption and scattering of light by small particles*. A Wiley-Interscience publication. New York: Wiley-CH Verlag GmbH & Co KGaA, 2004. 1 p. ISBN: 978-0-471-29340-8 978-3-527-61815-6. DOI: [10.1002/9783527618156](https://doi.org/10.1002/9783527618156).
- [173] M. J. Booth, M. A.A. Neil, and T. Wilson. “Aberration correction for confocal imaging in refractive-index-mismatched media”. In: *Journal of Microscopy* 192.2 (1998), pp. 90–98. ISSN: 00222720. DOI: [10.1111/j.1365-2818.1998.99999.x](https://doi.org/10.1111/j.1365-2818.1998.99999.x).
- [174] E. Abbe. “VII.— On the Estimation of Aperture in the Microscope.” In: *Journal of the Royal Microscopical Society* 1.3 (June 1881), pp. 388–423. ISSN: 0368-3974. DOI: [10.1111/j.1365-2818.1881.tb05909.x](https://doi.org/10.1111/j.1365-2818.1881.tb05909.x). URL: <https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2818.1881.tb05909.x>.
- [175] *CUDA Toolkit Documentation*. NVIDIA Corporation. 2025. URL: <https://docs.nvidia.com/cuda/>.

- [176] Ryosuke Okuta, Yuya Unno, Daisuke Nishino, Shohei Hido, and Crissman Loomis. “CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations”. In: *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*. 2017. URL: http://learningsys.org/nips17/assets/papers/paper_16.pdf.
- [177] Daniel Smith, Johnnie Gray, et al. *opt_einsum: Optimizing Einsum Expressions in Python*. Documentation. URL: <https://optimized-einsum.readthedocs.io/en/stable/>.
- [178] *cuBLAS Library*. CUDA Basic Linear Algebra Subprograms library. NVIDIA Corporation. 2025. URL: <https://docs.nvidia.com/cuda/cublas/>.
- [179] The mpmath development team. *mpmath: A Python Library for Arbitrary-Precision Floating-Point Arithmetic*. 2023. URL: <https://mpmath.org/>.
- [180] Edward S. Barnard, Benedikt Ursprung, Nick Borys, Jonas Zipfel, Frank Ogletree, Clarice Aiello, Hao Wu, Alan Buckley, Peter Ercius, and Mark Hager. *ScopeFoundry: A Python Platform for Controlling Custom Laboratory Experiments and Visualizing Scientific Data*. 2025. URL: <https://scopefoundry.org/>.
- [181] C. E. Shannon. “A Mathematical Theory of Communication”. In: *Bell System Technical Journal* 27.3 (July 1948), pp. 379–423. ISSN: 00058580. DOI: 10.1002/j.1538-7305.1948.tb01338.x. URL: <https://ieeexplore.ieee.org/document/6773024>.
- [182] Frank Gray. “Pulse Code Communication”. U.S. pat. 2,632,058. Mar. 17, 1953. URL: <https://patents.google.com/patent/US2632058A>.
- [183] David J. C. MacKay. *Information theory, inference, and learning algorithms*. Cambridge, UK ; New York: Cambridge University Press, 2003. 628 pp. ISBN: 978-0-521-64298-9.
- [184] T. M. Cover and Joy A. Thomas. *Elements of information theory*. 2nd ed. Hoboken, N.J: Wiley-Interscience, 2006. 748 pp. ISBN: 978-0-471-24195-9.
- [185] Bradley M. Ratliff, Charles F. LaCasse, and J. Scott Tyo. “Interpolation strategies for reducing IFOV artifacts in microgrid polarimeter imagery”. In: *Optics Express* 17.11 (May 25, 2009), p. 9112. ISSN: 1094-4087. DOI: 10.1364/OE.17.009112. URL: <https://opg.optica.org/abstract.cfm?URI=oe-17-11-9112>.

- [186] Shengkui Gao and Viktor Gruev. “Bilinear and bicubic interpolation methods for division of focal plane polarimeters”. In: *Optics Express* 19.27 (Dec. 19, 2011), p. 26161. ISSN: 1094-4087. DOI: [10.1364/OE.19.026161](https://doi.org/10.1364/OE.19.026161). URL: <https://opg.optica.org/oe/abstract.cfm?uri=oe-19-27-26161>.
- [187] Rafael C. Gonzalez and Richard E. Woods. *Digital image processing*. New York, NY: Pearson, 2018. 1168 pp. ISBN: 978-0-13-335672-4.
- [188] Jean Paul Serra. *Image analysis and mathematical morphology*. London ; New York: Academic Press, 1982. 1 p. ISBN: 978-0-12-637240-3.
- [189] Richard O. Duda and Peter E. Hart. “Use of the Hough transformation to detect lines and curves in pictures”. In: *Communications of the ACM* 15.1 (Jan. 1972), pp. 11–15. ISSN: 0001-0782, 1557-7317. DOI: [10.1145/361237.361242](https://doi.org/10.1145/361237.361242). URL: <https://dl.acm.org/doi/10.1145/361237.361242>.
- [190] Charles Kittel. *Introduction to solid state physics*. 8. ed., [repr.] Hoboken, NJ: Wiley. 680 pp. ISBN: 978-0-471-41526-8.
- [191] Lawrence R. Rabiner and Bernard Gold. *Theory and application of digital signal processing*. 18. [print.] Englewood Cliffs, N.J: Prentice-Hall, 1975. 762 pp. ISBN: 978-0-13-914101-0.
- [192] The MathWorks, Inc. *czt (Chirp Z-Transform)*. *Signal Processing Toolbox Function Reference*. Original implementation by C. Denham, J. McClellan, and T. Krauss. URL: <https://uk.mathworks.com/help/signal/ref/czt.html>.
- [193] Alexander Kraskov, Harald Stoegbauer, and Peter Grassberger. “Estimating Mutual Information”. In: *Physical Review E* 69.6 (June 23, 2004), p. 066138. ISSN: 1539-3755, 1550-2376. DOI: [10.1103/PhysRevE.69.066138](https://doi.org/10.1103/PhysRevE.69.066138). arXiv: [cond-mat/0305641](https://arxiv.org/abs/cond-mat/0305641). URL: <http://arxiv.org/abs/cond-mat/0305641>.
- [194] Brian C. Ross. “Mutual Information between Discrete and Continuous Data Sets”. In: *PLoS ONE* 9.2 (Feb. 19, 2014). Ed. by Daniele Marinazzo, e87357. ISSN: 1932-6203. DOI: [10.1371/journal.pone.0087357](https://doi.org/10.1371/journal.pone.0087357). URL: <https://dx.plos.org/10.1371/journal.pone.0087357>.

- [195] L. F. Kozachenko and N. N. Leonenko. “Sample Estimate of the Entropy of a Random Vector”. In: *Problemy Peredachi Informatsii* 23.2 (1987), pp. 9–16.
- [196] John M. Senior and M. Yousif Jamro. *Optical fiber communications: principles and practice*. 3rd ed. Harlow, England ; New York: Financial Times/Prentice Hall, 2009. 1075 pp. ISBN: 978-0-13-032681-2.
- [197] E. N. Glezer, M. Milosavljevic, L. Huang, R. J. Finlay, T.-H. Her, J. P. Callan, and E. Mazur. “Three-dimensional optical storage inside transparent materials”. In: *Optics Letters* 21.24 (Dec. 15, 1996), p. 2023. ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.21.002023](https://doi.org/10.1364/OL.21.002023). URL: <https://opg.optica.org/abstract.cfm?URI=ol-21-24-2023>.
- [198] Brian Everitt. *The Cambridge dictionary of statistics*. 3rd ed. Cambridge, UK New York: Cambridge University Press, 2006. 1 p. ISBN: 978-0-521-86039-0 978-0-511-24688-3 978-0-511-24619-7 978-0-511-24323-3.
- [199] C. W. McCutchen. “Generalized Aperture and the Three-Dimensional Diffraction Image”. In: *Journal of the Optical Society of America* 54.2 (Feb. 1, 1964), p. 240. ISSN: 0030-3941. DOI: [10.1364/JOSA.54.000240](https://doi.org/10.1364/JOSA.54.000240). URL: <https://opg.optica.org/abstract.cfm?URI=josa-54-2-240>.
- [200] Peter R. T. Munro and Peter Török. “Vectorial, high-numerical-aperture study of phase-contrast microscopes”. In: *Journal of the Optical Society of America A* 21.9 (Sept. 1, 2004), p. 1714. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/JOSAA.21.001714](https://doi.org/10.1364/JOSAA.21.001714). URL: <https://opg.optica.org/abstract.cfm?URI=josaa-21-9-1714>.
- [201] I. H. Malitson. “Interspecimen Comparison of the Refractive Index of Fused Silica”. In: *Journal of the Optical Society of America* 55.10 (Oct. 1, 1965), p. 1205. ISSN: 0030-3941. DOI: [10.1364/JOSA.55.001205](https://doi.org/10.1364/JOSA.55.001205). URL: <https://opg.optica.org/abstract.cfm?URI=josa-55-10-1205>.
- [202] Mikhail N. Polyanskiy. “Refractiveindex.info database of optical constants”. In: *Scientific Data* 11.1 (Jan. 18, 2024), p. 94. ISSN: 2052-4463. DOI: [10.1038/s41597-023-02898-2](https://doi.org/10.1038/s41597-023-02898-2). URL: <https://www.nature.com/articles/s41597-023-02898-2>.

- [203] Matthew J. Parry-Hill and Michael W. Davidson. *Adjustment of Objective Correction Collars*. Nikon Instruments Inc. 2025. URL: <https://www.microscopyu.com/tutorials/adjustment-of-objective-correction-collars>.
- [204] Karl Strehl. “Aplanatische und fehlerhafte Abbildung im Fernrohr”. German. In: *Zeitschrift für Instrumentenkunde* 15 (Oct. 1895), pp. 362–370.
- [205] Alfred Millett-Sikking. *Any immersion remote refocus (AIRR) microscopy*. Version Number: v1.1.0. Dec. 11, 2022. DOI: [10.5281/ZENODO.7425649](https://doi.org/10.5281/ZENODO.7425649). URL: <https://zenodo.org/doi/10.5281/zenodo.7425649>.
- [206] Christopher M. Bishop. *Neural networks for pattern recognition*. Oxford : New York: Clarendon Press ; Oxford University Press, 1995. 482 pp. ISBN: 978-0-19-853849-3 978-0-19-853864-6.
- [207] Vinod Nair and Geoffrey E Hinton. “Rectified Linear Units Improve Restricted Boltzmann Machines”. In: ().
- [208] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. “Multilayer feedforward networks are universal approximators”. In: *Neural Networks* 2.5 (Jan. 1989), pp. 359–366. ISSN: 08936080. DOI: [10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8). URL: <https://linkinghub.elsevier.com/retrieve/pii/0893608089900208>.
- [209] Stinchcombe and White. “Universal approximation using feedforward networks with non-sigmoid hidden layer activation functions”. In: *International Joint Conference on Neural Networks*. International Joint Conference on Neural Networks. Washington, DC, USA: IEEE, 1989, 613–617 vol.1. DOI: [10.1109/IJCNN.1989.118640](https://doi.org/10.1109/IJCNN.1989.118640). URL: <http://ieeexplore.ieee.org/document/118640/>.
- [210] Christopher M. Bishop. *Pattern recognition and machine learning*. Information science and statistics. New York: Springer, 2006. 738 pp. ISBN: 978-0-387-31073-2.
- [211] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *Advances in Neural Information Processing Systems* 32. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc., 2019, pp. 8024–8035.

- [212] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [213] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. “On the importance of initialization and momentum in deep learning”. In: *Proceedings of the 30th International Conference on Machine Learning*. Ed. by Sanjoy Dasgupta and David McAllester. Vol. 28. Proceedings of Machine Learning Research 3. Atlanta, Georgia, USA: PMLR, June 17, 2013, pp. 1139–1147. URL: <https://proceedings.mlr.press/v28/sutskever13.html>.
- [214] P. A. Franaszek. “Sequence-state Methods for Run-length-limited Coding”. In: *IBM Journal of Research and Development* 14.4 (July 1970), pp. 376–383. ISSN: 0018-8646, 0018-8646. DOI: [10.1147/rd.144.0376](https://doi.org/10.1147/rd.144.0376). URL: <https://ieeexplore.ieee.org/document/5391624/>.
- [215] M. Siala and G.K. Kaleh. “Joint multilevel RLL and error correction coding”. In: *Proceedings of 1995 IEEE International Symposium on Information Theory*. 1995 IEEE International Symposium on Information Theory. Whistler, BC, Canada: IEEE, 1995, p. 201. ISBN: 978-0-7803-2453-4. DOI: [10.1109/ISIT.1995.531875](https://doi.org/10.1109/ISIT.1995.531875). URL: <http://ieeexplore.ieee.org/document/531875/>.
- [216] John G. Proakis and Masoud Salehi. *Digital communications*. 5th ed. Boston: McGraw-Hill, 2008. 1150 pp. ISBN: 978-0-07-295716-7.
- [217] R. W. Hamming. “Error Detecting and Error Correcting Codes”. In: *Bell System Technical Journal* 29.2 (Apr. 1950), pp. 147–160. ISSN: 00058580. DOI: [10.1002/j.1538-7305.1950.tb00463.x](https://doi.org/10.1002/j.1538-7305.1950.tb00463.x). URL: <https://ieeexplore.ieee.org/document/6772729>.
- [218] Michael Shribak. “Polychromatic polarization microscope: Bringing colors to a colorless world”. In: *Scientific Reports* 5 (2015), pp. 1–10. ISSN: 20452322. DOI: [10.1038/srep17340](https://doi.org/10.1038/srep17340).
- [219] Nirmalya Ghosh. “Tissue polarimetry: concepts, challenges, applications, and outlook”. In: *Journal of Biomedical Optics* 16.11 (Nov. 1, 2011), p. 110801. ISSN: 1083-3668. DOI: [10.1117/1.3652896](https://doi.org/10.1117/1.3652896). URL: <http://biomedicaloptics.spiedigitallibrary.org/article.aspx?doi=10.1117/1.3652896>.

- [220] Richard Kasprowicz, Rakesh Suman, and Peter O'Toole. "Characterising live cell behaviour: Traditional label-free and quantitative phase imaging approaches". In: *The International Journal of Biochemistry & Cell Biology* 84 (Mar. 2017), pp. 89–95. ISSN: 13572725. DOI: [10.1016/j.biocel.2017.01.004](https://doi.org/10.1016/j.biocel.2017.01.004). URL: <https://linkinghub.elsevier.com/retrieve/pii/S1357272517300055>.
- [221] Garam Yun, Karlton Crabtree, and Russell A. Chipman. "Three-dimensional polarization ray-tracing calculus I: definition and diattenuation". In: *Applied Optics* 50.18 (June 20, 2011), p. 2855. ISSN: 0003-6935, 1539-4522. DOI: [10.1364/AO.50.002855](https://doi.org/10.1364/AO.50.002855). URL: <https://opg.optica.org/abstract.cfm?URI=ao-50-18-2855>.
- [222] Arthur Francis Hallimond. *The Polarizing Microscope*. 3rd ed. York: Vickers Ltd., Vickers Instruments, 1970. 302 pp.
- [223] N. H. Hartshorne and Alan Stuart. *Crystals and the polarising microscope*. 4th ed. London: Edward Arnold, 1970. 614 pp. ISBN: 978-0-7131-2256-5.
- [224] *Economy Film Polarizers*. Product overview and specifications page. Thorlabs. URL: <https://www.thorlabs.com/economy-film-polarizers>.
- [225] *Glan-Taylor Calcite Polarizers*. Product specifications page. Thorlabs. URL: <https://www.thorlabs.com/glan-taylor-calcite-polarizers>.
- [226] *Wire Grid Polarizers on Glass Substrates*. Product overview page. Thorlabs. URL: <https://www.thorlabs.com/wire-grid-polarizers-on-glass-substrates>.
- [227] *Broadband Polarizing Beamsplitter Cubes*. Product specifications page. Thorlabs. URL: <https://www.thorlabs.com/broadband-polarizing-beamsplitter-cubes>.
- [228] Jan Korger, Tobias Kolb, Peter Banzer, Andrea Aiello, Christoffer Wittmann, Christoph Marquardt, and Gerd Leuchs. "The polarization properties of a tilted polarizer". In: *Optics Express* 21.22 (Nov. 4, 2013), p. 27032. ISSN: 1094-4087. DOI: [10.1364/OE.21.027032](https://doi.org/10.1364/OE.21.027032). URL: <https://opg.optica.org/oe/abstract.cfm?uri=oe-21-22-27032>.

- [229] Xiaofei Han et al. “A polymer index-matched to water enables diverse applications in fluorescence microscopy”. In: *Lab on a Chip* 21.8 (2021), pp. 1549–1562. ISSN: 1473-0197, 1473-0189. DOI: [10.1039/D0LC01233E](https://doi.org/10.1039/D0LC01233E). URL: <https://xlink.rsc.org/?DOI=D0LC01233E>.
- [230] Huihui Liu. “Polarisation-based Approaches to Quantitative Phase and Hyperspectral Imaging”. PhD thesis. London, UK: Imperial College London, 2025.
- [231] Ian L. Sparks, Keith M. Derbyshire, William R. Jacobs, and Yasu S. Morita. “Mycobacterium smegmatis: The Vanguard of Mycobacterial Research”. In: *Journal of Bacteriology* 205.1 (Jan. 26, 2023). Ed. by Patricia A. Champion, e00337–22. ISSN: 0021-9193, 1098-5530. DOI: [10.1128/jb.00337-22](https://doi.org/10.1128/jb.00337-22). URL: <https://journals.asm.org/doi/10.1128/jb.00337-22>.
- [232] Stephanie R. Smelyansky, Chi-Wang Ma, Victoria M. Marando, Gregory H. Babunovic, So Young Lee, Bryan D. Bryson, and Laura L. Kiessling. “Exploiting thioether reactivity to label mycobacterial glycans”. In: *Proceedings of the National Academy of Sciences* 122.19 (May 13, 2025), e2422185122. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2422185122](https://doi.org/10.1073/pnas.2422185122). URL: <https://pnas.org/doi/10.1073/pnas.2422185122>.
- [233] Dmitri B. Papkovsky, ed. *Live Cell Imaging*. Vol. 591. Methods in Molecular Biology. Totowa, NJ: Humana Press, 2010. ISBN: 978-1-60761-403-6 978-1-60761-404-3. DOI: [10.1007/978-1-60761-404-3](https://doi.org/10.1007/978-1-60761-404-3). URL: <http://link.springer.com/10.1007/978-1-60761-404-3>.
- [234] N.R. Ramanujam, Shobhit K. Patel, N. Manohar Reddy, Sofyan A. Taya, D. Vigneswaran, and M.S. Mani Rajan. “One-dimensional ring mirror-defect photonic crystal for detection of mycobacterium tuberculosis bacteria”. In: *Optik* 219 (Oct. 2020), p. 165097. ISSN: 00304026. DOI: [10.1016/j.ijleo.2020.165097](https://doi.org/10.1016/j.ijleo.2020.165097). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0030402620309335>.
- [235] Drew A. Pommet, M. G. Moharam, and Eric B. Grann. “Limits of scalar diffraction theory for diffractive phase elements”. In: *Journal of the Optical Society of America A*

- 11.6 (June 1, 1994), p. 1827. ISSN: 1084-7529, 1520-8532. DOI: [10.1364/JOSAA.11.001827](https://doi.org/10.1364/JOSAA.11.001827). URL: <https://opg.optica.org/abstract.cfm?URI=josaa-11-6-1827>.
- [236] Leung Tsang, Jin Au Kong, and Kung-Hau Ding. *Scattering of Electromagnetic Waves: Theories and Applications*. 1st ed. Wiley, July 14, 2000. ISBN: 978-0-471-38799-2 978-0-471-22428-0. DOI: [10.1002/0471224286](https://doi.org/10.1002/0471224286). URL: <https://onlinelibrary.wiley.com/doi/book/10.1002/0471224286>.
- [237] Scott Prahl. *miepython: Pure python calculation of Mie scattering*. Version 3.1.0. Feb. 8, 2026. DOI: [10.5281/ZENODO.7949263](https://doi.org/10.5281/ZENODO.7949263). URL: <https://zenodo.org/doi/10.5281/zenodo.7949263>.
- [238] Hwa-Jung Kim, Eun-Kyeong Jo, Jeong-Kyu Park, Jae-Hyun Lim, Dullei Min, and Tae-Hyun Paik. “Isolation and partial characterisation of the Triton X-100 solubilised protein antigen from *Mycobacterium tuberculosis*”. In: *Journal of Medical Microbiology* 48.6 (June 1, 1999), pp. 585–591. ISSN: 0022-2615, 1473-5644. DOI: [10.1099/00222615-48-6-585](https://doi.org/10.1099/00222615-48-6-585). URL: <https://www.microbiologyresearch.org/content/journal/jmm/10.1099/00222615-48-6-585>.
- [239] Luc Fraiture. “A History of the Description of the Three-Dimensional Finite Rotation”. In: *The Journal of the Astronautical Sciences* 57.1 (Jan. 2009), pp. 207–232. ISSN: 0021-9142, 2195-0571. DOI: [10.1007/BF03321502](https://doi.org/10.1007/BF03321502). URL: <http://link.springer.com/10.1007/BF03321502>.

Permission document for Figure 1.1(a).

AIP PUBLISHING LICENSE TERMS AND CONDITIONS

Feb 14, 2026

This Agreement between Zhonghe Feng ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	6207620024608
License date	Feb 14, 2026
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Applied Physics Letters
Licensed Content Title	Photowritten optical waveguides in various glasses with ultrashort pulse laser
Licensed Content Author	Miura, K.; Qiu, Jianrong
Licensed Content Date	Dec 8, 1997
Licensed Content Volume	71
Licensed Content Issue	23
Type of Use	Thesis/Dissertation
Requestor type	Student
Format	Print and electronic
Portion	Figure/Table

Permission document for Figure 1.1(a).

Number of figures/tables	1
Will you be translating?	No
Title of new work	Polarisation Microscopy and Its Application to Optical Data Storage: Design, Informational Analysis, and ML-Based Decoding
Institution name	Imperial College London
Expected presentation date	Feb 2026
Portions	Figure 3
The Requesting Person / Organization to Appear on the License	Zhonghe Feng
Requestor Location	Blackett Laboratory Prince Consort Road London, SW7 2BW United Kingdom
Payment Type	Invoice
Email Address	zf1115@ic.ac.uk
Billing Address	Mr. Zhonghe Feng Blackett Laboratory Prince Consort Road London, United Kingdom SW7 2BW
Total	0.00 USD

Terms and Conditions

AIP Publishing -- Terms and Conditions: Permissions Uses

Permission document for Figure 1.1(a).

AIP Publishing hereby grants to you the non-exclusive right and license to use and/or distribute the Material according to the use specified in your order, on a one-time basis, for the specified term, with a maximum distribution equal to the number that you have ordered. Any links or other content accompanying the Material are not the subject of this license.

1. You agree to include the following copyright and permission notice with the reproduction of the Material: "Reprinted from [FULL CITATION], with the permission of AIP Publishing." For an article, the credit line and permission notice must be printed on the first page of the article or book chapter. For photographs, covers, or tables, the notice may appear with the Material, in a footnote, or in the reference list.
2. If you have licensed reuse of a figure, photograph, cover, or table, it is your responsibility to ensure that the material is original to AIP Publishing and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover, or table does not indicate that it was reprinted by AIP Publishing, with permission, from another source. Under no circumstances does AIP Publishing purport or intend to grant permission to reuse material to which it does not hold appropriate rights.
You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may not use the Material for promotional purposes.
3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.
4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.
5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.
6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.
7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

Other Terms and Conditions:

V1.2

Permission document for Figure 1.1(a).

Questions? customercare@copyright.com.



Permission document for Figure 1.1(b).



This is a License Agreement between Zhonghe Feng ("User") and Copyright Clearance Center, Inc. ("CCC") on behalf of the Rightsholder identified in the order details below. The license consists of the order details, the Marketplace Permissions General Terms and Conditions below, and any Rightsholder Terms and Conditions which are included below.

All payments must be made in full to CCC in accordance with the Marketplace Permissions General Terms and Conditions below.

Order Date	19-Feb-2026	Type of Use	Republish in a thesis/dissertation
Order License ID	1701693-1	Publisher	IOP Publishing on behalf of the Japan Society of Applied Physics (JSAP)
ISSN	1347-4065	Portion	Image/photo/illustration

LICENSED CONTENT

Publication Title	Japanese journal of applied physics : JJAP online	Country	Japan
Author / Editor	Institute of Pure and Applied Physics.	Rightsholder	IOP Publishing Limited
Date	01/01/1962	Publication Type	e-Journal
Language	English	URL	http://www.ipap.jp/jjap/index.htm

REQUEST DETAILS

Portion Type	Image/photo/illustration	Distribution	Worldwide
Number of Images / Photos / Illustrations	1	Translation	Original language of publication
Format (select all that apply)	Print, Electronic	Copies for the Disabled?	No
Who Will Republish the Content?	Academic institution	Minor Editing Privileges?	Yes
Duration of Use	Life of current edition	Incidental Promotional Use?	No
Lifetime Unit Quantity	Up to 499	Currency	GBP
Rights Requested	Main product		

NEW WORK DETAILS

Title	Polarisation Microscopy and Its Application to Optical Data Storage: Design, Informational Analysis, and ML-Based Decoding	Institution Name	Imperial College London
Instructor Name	Zhonghe Feng	Expected Presentation Date	2026-02-20

ADDITIONAL DETAILS

Order Reference Number	N/A
-------------------------------	-----

Permission document for Figure 1.1(b).

The Requesting Person /
Organization to Appear
on the License Zhonghe Feng

REQUESTED CONTENT DETAILS

Title, Description or Numeric Reference of the Portion(s)	Fig. 16. Backscattering electron images of the polished silicate glass sample after being polished close to the depth of the focal spot location. The magnifications of (a) and (b) are 10000 and 30000 , respectively	Title of the Article / Chapter the Portion Is From	Three-Dimensional Micro- and Nano-Fabrication in Transparent Materials by Femtosecond Laser
Editor of Portion(s)	N/A	Author of Portion(s)	Yasuhiko Shimotsuma, Kazuyuki Hirao, Peter G. Kazansky and Jiarong Qiu
Volume / Edition	44	Issue, if Republishing an Article From a Serial	7R
Page or Page Range of Portion	4743	Publication Date of Portion	2005-07-08

RIGHTSHOLDER TERMS AND CONDITIONS

These special terms and conditions are in addition to the standard terms and conditions for CCC's Republication Service and, together with those standard terms and conditions, govern the use of the Works. As the User you will make all reasonable efforts to contact the author(s) of the article which the Work is to be reused from, to seek consent for your intended use. Contacting one author who is acting expressly as authorised agent for their co-author(s) is acceptable. User will reproduce the following wording prominently alongside the Work: the source of the Work, including author, article title, title of journal, volume number, issue number (if relevant), page range (or first page if this is the only information available) and date of first publication; and a link back to the article (via DOI); and if practicable, and IN ALL CASES for new works published under any of the Creative Commons licences, the words "© The Japan Society of Applied Physics. Reproduced by permission of IOP Publishing Ltd. All rights reserved" Without the express permission of the author(s) and the Rightsholder of the article from which the Work is to be reused, User shall not use it in any way which, in the opinion of IOP Publishing Ltd, could: (i) distort or alter the author(s)' original intention(s) and meaning; (ii) be prejudicial to the honour or reputation of the author(s); and/or (iii) imply endorsement by the author(s) and/or the Rightsholder and/or IOP Publishing Ltd. This licence does not apply to any article which is credited to another source and which does not have the copyright line '© The Japan Society of Applied Physics'. User must check the copyright line of the article from which the Work is to be reused to check that the Rightsholder and IOP Publishing Ltd have all the necessary rights to be able to grant permission. User is solely responsible for identifying and obtaining separate licences and permissions from the copyright owner for reuse of any such third party material/figures which the Rightsholder is not the copyright owner of. The Rightsholder shall not reimburse any fees which User pays for a republication license for such third party content. This licence does not apply to any material/figure which is credited to another source in the Rightsholder's publication or has been obtained from a third party. User must check the Version of Record of the article from which the Work is to be reused, to check whether any of the material in the Work is third party material. Third party citations and/or copyright notices and/or permissions statements may not be included in any other version of the article from which the Work is to be reused and so cannot be relied upon by the User. User is solely responsible for identifying and obtaining separate licences and permissions from the copyright owner for reuse of any such third party material/figures where the Rightsholder is not the copyright owner. The Rightsholder shall not reimburse any fees which User pays for a republication license for such third party content. User and CCC acknowledge that IOP Publishing Ltd and/or the Rightsholder may, from time to time, make changes or additions to these special terms and conditions without express notification, provided that these shall not apply to permissions already secured and paid for by User prior to such change or addition. User acknowledges that the Rightsholder and IOP Publishing Ltd (which includes companies within its group and third parties for whom it publishes its titles) may make use of personal data collected through the service in the course of their business. If User is the author of the Work, User may automatically have the right to reuse it under the rights granted back when User transferred the copyright in the article to the Rightsholder. User should check the copyright form and the relevant author rights policy to check whether permission is required. If User is the author of the Work and does require permission for proposed reuse of the Work, User should select 'Author of requested content' as the Requestor Type. The Rightsholder shall not reimburse any fees which User pays for a republication license. As well as CCC, the Rightsholder and IOP Publishing Ltd shall have the right to bring any legal action that they deem necessary to enforce their rights should they consider that the Work infringes those rights in any way. For content reuse requests that qualify for permission under the STM Permissions Guidelines, which may be updated from time to time, the STM

Permission document for Figure 1.1(b).

Permissions Guidelines supplement the terms and conditions contained in this license. CCC Republication Terms and Conditions 1. Description of Service; Defined Terms. This Republication License enables the User to obtain licenses for republication of one or more copyrighted works as described in detail on the relevant Order Confirmation (the "Work(s)"). Copyright Clearance Center, Inc. ("CCC") grants licenses through the Service on behalf of the rightsholder identified on the Order Confirmation (the "Rightsholder"). "Republication", as used herein, generally means the inclusion of a Work, in whole or in part, in a new work or works, also as described on the Order Confirmation. "User", as used herein, means the person or entity making such republication. 2. The terms set forth in the relevant Order Confirmation, and any terms set by the Rightsholder with respect to a particular Work, govern the terms of use of Works in connection with the Service. By using the Service, the person transacting for a republication license on behalf of the User represents and warrants that he/she/it (a) has been duly authorized by the User to accept, and hereby does accept, all such terms and conditions on behalf of User, and (b) shall inform User of all such terms and conditions. In the event such person is a "freelancer" or other third party independent of User and CCC, such party shall be deemed jointly a "User" for purposes of these terms and conditions. In any event, User shall be deemed to have accepted and agreed to all such terms and conditions if User republishes the Work in any fashion. 3. Scope of License; Limitations and Obligations. 1. All Works and all rights therein, including copyright rights, remain the sole and exclusive property of the Rightsholder. The license created by the exchange of an Order Confirmation (and/or any invoice) and payment by User of the full amount set forth on that document includes only those rights expressly set forth in the Order Confirmation and in these terms and conditions, and conveys no other rights in the Work(s) to User. All rights not expressly granted are hereby reserved. 2. General Payment Terms: You may pay by credit card or through an account with us payable at the end of the month. If you and we agree that you may establish a standing account with CCC, then the following terms apply: Remit Payment to: Copyright Clearance Center, 29118 Network Place, Chicago, IL 60673-1291. Payments Due: Invoices are payable upon their delivery to you (or upon our notice to you that they are available to you for downloading). After 30 days, outstanding amounts will be subject to a service charge of 1-1/2% per month or, if less, the maximum rate allowed by applicable law. Unless otherwise specifically set forth in the Order Confirmation or in a separate written agreement signed by CCC, invoices are due and payable on "net 30" terms. While User may exercise the rights licensed immediately upon issuance of the Order Confirmation, the license is automatically revoked and is null and void, as if it had never been issued, if complete payment for the license is not received on a timely basis either from User directly or through a payment agent, such as a credit card company. 3. Unless otherwise provided in the Order Confirmation, any grant of rights to User (i) is "one-time" (including the editions and product family specified in the license), (ii) is non-exclusive and non-transferable and (iii) is subject to any and all limitations and restrictions (such as, but not limited to, limitations on duration of use or circulation) included in the Order Confirmation or invoice and/or in these terms and conditions. Upon completion of the licensed use, User shall either secure a new permission for further use of the Work(s) or immediately cease any new use of the Work(s) and shall render inaccessible (such as by deleting or by removing or severing links or other locators) any further copies of the Work (except for copies printed on paper in accordance with this license and still in User's stock at the end of such period). 4. In the event that the material for which a republication license is sought includes third party materials (such as photographs, illustrations, graphs, inserts and similar materials) which are identified in such material as having been used by permission, User is responsible for identifying, and seeking separate licenses (under this Service or otherwise) for, any of such third party materials; without a separate license, such third party materials may not be used. 5. Use of proper copyright notice for a Work is required as a condition of any license granted under the Service. Unless otherwise provided in the Order Confirmation, a proper copyright notice will read substantially as follows: "Republished with permission of [Rightsholder's name], from [Work's title, author, volume, edition number and year of copyright]; permission conveyed through Copyright Clearance Center, Inc. " Such notice must be provided in a reasonably legible font size and must be placed either immediately adjacent to the Work as used (for example, as part of a by-line or footnote but not as a separate electronic link) or in the place where substantially all other credits or notices for the new work containing the republished Work are located. Failure to include the required notice results in loss to the Rightsholder and CCC, and the User shall be liable to pay liquidated damages for each such failure equal to twice the use fee specified in the Order Confirmation, in addition to the use fee itself and any other fees and charges specified. 6. User may only make alterations to the Work if and as expressly set forth in the Order Confirmation. No Work may be used in any way that is defamatory, violates the rights of third parties (including such third parties' rights of copyright, privacy, publicity, or other tangible or intangible property), or is otherwise illegal, sexually explicit or obscene. In addition, User may not conjoin a Work with any other material that may result in damage to the reputation of the Rightsholder. User agrees to inform CCC if it becomes aware of any infringement of any rights in a Work and to cooperate with any reasonable request of CCC or the Rightsholder in connection therewith. 4. Indemnity. User hereby indemnifies and agrees to defend the Rightsholder and CCC, and their respective employees and directors, against all claims, liability, damages, costs and expenses, including legal fees and expenses, arising out of any use of a Work beyond the scope of the rights granted herein, or any use of a Work which has been altered in any unauthorized way by User, including claims of defamation or infringement of rights of copyright, publicity, privacy or other tangible or intangible property. 5. Limitation of Liability. UNDER NO CIRCUMSTANCES WILL CCC OR THE RIGHTSHOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL OR INCIDENTAL DAMAGES (INCLUDING WITHOUT LIMITATION DAMAGES FOR LOSS OF BUSINESS PROFITS OR INFORMATION, OR FOR BUSINESS INTERRUPTION) ARISING OUT OF THE USE OR INABILITY TO USE A WORK, EVEN IF ONE OF THEM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In any event, the total liability of the Rightsholder and CCC (including their respective employees and directors) shall not exceed the total amount actually paid by User for this license. User assumes full liability for the actions and omissions of its principals, employees, agents, affiliates, successors and assigns. 6. Limited

Permission document for Figure 1.1(b).

Warranties. THE WORK(S) AND RIGHT(S) ARE PROVIDED "AS IS". CCC HAS THE RIGHT TO GRANT TO USER THE RIGHTS GRANTED IN THE ORDER CONFIRMATION DOCUMENT. CCC AND THE RIGHTSHOLDER DISCLAIM ALL OTHER WARRANTIES RELATING TO THE WORK(S) AND RIGHT(S), EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. ADDITIONAL RIGHTS MAY BE REQUIRED TO USE ILLUSTRATIONS, GRAPHS, PHOTOGRAPHS, ABSTRACTS, INSERTS OR OTHER PORTIONS OF THE WORK (AS OPPOSED TO THE ENTIRE WORK) IN A MANNER CONTEMPLATED BY USER; USER UNDERSTANDS AND AGREES THAT NEITHER CCC NOR THE RIGHTSHOLDER MAY HAVE SUCH ADDITIONAL RIGHTS TO GRANT. 7. Effect of Breach. Any failure by User to pay any amount when due, or any use by User of a Work beyond the scope of the license set forth in the Order Confirmation and/or these terms and conditions, shall be a material breach of the license created by the Order Confirmation and these terms and conditions. Any breach not cured within 30 days of written notice thereof shall result in immediate termination of such license without further notice. Any unauthorized (but licensable) use of a Work that is terminated immediately upon notice thereof may be liquidated by payment of the Rightsholder's ordinary license price therefor; any unauthorized (and unlicensable) use that is not terminated immediately for any reason (including, for example, because materials containing the Work cannot reasonably be recalled) will be subject to all remedies available at law or in equity, but in no event to a payment of less than three times the Rightsholder's ordinary license price for the most closely analogous licensable use plus Rightsholder's and/or CCC's costs and expenses incurred in collecting such payment. 8. Miscellaneous. 1. User acknowledges that CCC may, from time to time, make changes or additions to the Service or to these terms and conditions, and CCC reserves the right to send notice to the User by electronic mail or otherwise for the purposes of notifying User of such changes or additions; provided that any such changes or additions shall not apply to permissions already secured and paid for. 2. Use of User-related information collected through the Service is governed by CCC's privacy policy, available online here: <https://marketplace.copyright.com/rs-ui-web/mp/privacy-policy> 3. The licensing transaction described in the Order Confirmation is personal to User. Therefore, User may not assign or transfer to any other person (whether a natural person or an organization of any kind) the license created by the Order Confirmation and these terms and conditions or any rights granted hereunder; provided, however, that User may assign such license in its entirety on written notice to CCC in the event of a transfer of all or substantially all of User's rights in the new material which includes the Work(s) licensed under this Service. 4. No amendment or waiver of any terms is binding unless set forth in writing and signed by the parties. The Rightsholder and CCC hereby object to any terms contained in any writing prepared by the User or its principals, employees, agents or affiliates and purporting to govern or otherwise relate to the licensing transaction described in the Order Confirmation, which terms are in any way inconsistent with any terms set forth in the Order Confirmation and/or in these terms and conditions or CCC's standard operating procedures, whether such writing is prepared prior to, simultaneously with or subsequent to the Order Confirmation, and whether such writing appears on a copy of the Order Confirmation or in a separate instrument. 5. The licensing transaction described in the Order Confirmation document shall be governed by and construed under the law of the State of New York, USA, without regard to the principles thereof of conflicts of law. Any case, controversy, suit, action, or proceeding arising out of, in connection with, or related to such licensing transaction shall be brought, at CCC's sole discretion, in any federal or state court located in the County of New York, State of New York, USA, or in any federal or state court whose geographical jurisdiction covers the location of the Rightsholder set forth in the Order Confirmation. The parties expressly submit to the personal jurisdiction and venue of each such federal or state court. If you have any comments or questions about the Service or Copyright Clearance Center, please contact us at 978-750-8400 or send an e-mail to support@copyright.com.

If you are producing a new book or article to be published by another STM Signatory publisher, please choose "Reuse in a book under STM Guidelines" or "Reuse in a journal under STM Guidelines." You can consult the STM Guidelines website to see whether your new publisher is a signatory to the STM Permissions Guidelines.

Marketplace Permissions General Terms and Conditions

The following terms and conditions ("General Terms"), together with any applicable Publisher Terms and Conditions, govern User's use of Works pursuant to the Licenses granted by Copyright Clearance Center, Inc. ("CCC") on behalf of the applicable Rightsholders of such Works through CCC's applicable Marketplace transactional licensing services (each, a "Service").

1) **Definitions.** For purposes of these General Terms, the following definitions apply:

"License" is the licensed use the User obtains via the Marketplace platform in a particular licensing transaction, as set forth in the Order Confirmation.

"Order Confirmation" is the confirmation CCC provides to the User at the conclusion of each Marketplace transaction. "Order Confirmation Terms" are additional terms set forth on specific Order Confirmations not set forth in the General Terms that can include terms applicable to a particular CCC transactional licensing service and/or any Rightsholder-specific terms.

"Rightsholder(s)" are the holders of copyright rights in the Works for which a User obtains licenses via the Marketplace platform, which are displayed on specific Order Confirmations.

Permission document for Figure 1.1(b).

"Terms" means the terms and conditions set forth in these General Terms and any additional Order Confirmation Terms collectively.

"User" or "you" is the person or entity making the use granted under the relevant License. Where the person accepting the Terms on behalf of a User is a freelancer or other third party who the User authorized to accept the General Terms on the User's behalf, such person shall be deemed jointly a User for purposes of such Terms.

"Work(s)" are the copyright protected works described in relevant Order Confirmations.

2) **Description of Service.** CCC's Marketplace enables Users to obtain Licenses to use one or more Works in accordance with all relevant Terms. CCC grants Licenses as an agent on behalf of the copyright rightsholder identified in the relevant Order Confirmation.

3) **Applicability of Terms.** The Terms govern User's use of Works in connection with the relevant License. In the event of any conflict between General Terms and Order Confirmation Terms, the latter shall govern. User acknowledges that Rightsholders have complete discretion whether to grant any permission, and whether to place any limitations on any grant, and that CCC has no right to supersede or to modify any such discretionary act by a Rightsholder.

4) **Representations; Acceptance.** By using the Service, User represents and warrants that User has been duly authorized by the User to accept, and hereby does accept, all Terms.

5) **Scope of License; Limitations and Obligations.** All Works and all rights therein, including copyright rights, remain the sole and exclusive property of the Rightsholder. The License provides only those rights expressly set forth in the terms and conveys no other rights in any Works

6) **General Payment Terms.** User may pay at time of checkout by credit card or choose to be invoiced. If the User chooses to be invoiced, the User shall: (i) remit payments in the manner identified on specific invoices, (ii) unless otherwise specifically stated in an Order Confirmation or separate written agreement, Users shall remit payments upon receipt of the relevant invoice from CCC, either by delivery or notification of availability of the invoice via the Marketplace platform, and (iii) if the User does not pay the invoice within 30 days of receipt, the User may incur a service charge of 1.5% per month or the maximum rate allowed by applicable law, whichever is less. While User may exercise the rights in the License immediately upon receiving the Order Confirmation, the License is automatically revoked and is null and void, as if it had never been issued, if CCC does not receive complete payment on a timely basis.

7) **General Limits on Use.** Unless otherwise provided in the Order Confirmation, any grant of rights to User (i) involves only the rights set forth in the Terms and does not include subsequent or additional uses, (ii) is non-exclusive and non-transferable, and (iii) is subject to any and all limitations and restrictions (such as, but not limited to, limitations on duration of use or circulation) included in the Terms. Upon completion of the licensed use as set forth in the Order Confirmation, User shall either secure a new permission for further use of the Work(s) or immediately cease any new use of the Work(s) and shall render inaccessible (such as by deleting or by removing or severing links or other locators) any further copies of the Work. User may only make alterations to the Work if and as expressly set forth in the Order Confirmation. No Work may be used in any way that is unlawful, including without limitation if such use would violate applicable sanctions laws or regulations, would be defamatory, violate the rights of third parties (including such third parties' rights of copyright, privacy, publicity, or other tangible or intangible property), or is otherwise illegal, sexually explicit, or obscene. In addition, User may not conjoin a Work with any other material that may result in damage to the reputation of the Rightsholder. Any unlawful use will render any licenses hereunder null and void. User agrees to inform CCC if it becomes aware of any infringement of any rights in a Work and to cooperate with any reasonable request of CCC or the Rightsholder in connection therewith.

8) **Third Party Materials.** In the event that the material for which a License is sought includes third party materials (such as photographs, illustrations, graphs, inserts and similar materials) that are identified in such material as having been used by permission (or a similar indicator), User is responsible for identifying, and seeking separate licenses (under this Service, if available, or otherwise) for any of such third party materials; without a separate license, User may not use such third party materials via the License.

9) **Copyright Notice.** Use of proper copyright notice for a Work is required as a condition of any License granted under the Service. Unless otherwise provided in the Order Confirmation, a proper copyright notice will read substantially as follows: "Used with permission of [Rightsholder's name], from [Work's title, author, volume, edition number and year of copyright]; permission conveyed through Copyright Clearance Center, Inc." Such notice must be provided in a reasonably legible font size and must be placed either on a cover page or in another location that any person, upon gaining access to the material which is the subject of a permission, shall see, or in the case of republication Licenses, immediately adjacent to the Work as used (for example, as part of a by-line or footnote) or in the place where substantially all other credits or notices for the new work containing the republished Work are located. Failure to include the required notice results in loss to the Rightsholder and CCC, and the User shall be liable to pay liquidated damages for each such failure equal to

Permission document for Figure 1.1(b).

twice the use fee specified in the Order Confirmation, in addition to the use fee itself and any other fees and charges specified.

10) **Indemnity.** User hereby indemnifies and agrees to defend the Rightsholder and CCC, and their respective employees and directors, against all claims, liability, damages, costs, and expenses, including legal fees and expenses, arising out of any use of a Work beyond the scope of the rights granted herein and in the Order Confirmation, or any use of a Work which has been altered in any unauthorized way by User, including claims of defamation or infringement of rights of copyright, publicity, privacy, or other tangible or intangible property.

11) **Limitation of Liability.** UNDER NO CIRCUMSTANCES WILL CCC OR THE RIGHTSHOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, OR INCIDENTAL DAMAGES (INCLUDING WITHOUT LIMITATION DAMAGES FOR LOSS OF BUSINESS PROFITS OR INFORMATION, OR FOR BUSINESS INTERRUPTION) ARISING OUT OF THE USE OR INABILITY TO USE A WORK, EVEN IF ONE OR BOTH OF THEM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In any event, the total liability of the Rightsholder and CCC (including their respective employees and directors) shall not exceed the total amount actually paid by User for the relevant License. User assumes full liability for the actions and omissions of its principals, employees, agents, affiliates, successors, and assigns.

12) **Limited Warranties.** THE WORK(S) AND RIGHT(S) ARE PROVIDED "AS IS." CCC HAS THE RIGHT TO GRANT TO USER THE RIGHTS GRANTED IN THE ORDER CONFIRMATION DOCUMENT. CCC AND THE RIGHTSHOLDER DISCLAIM ALL OTHER WARRANTIES RELATING TO THE WORK(S) AND RIGHT(S), EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. ADDITIONAL RIGHTS MAY BE REQUIRED TO USE ILLUSTRATIONS, GRAPHS, PHOTOGRAPHS, ABSTRACTS, INSERTS, OR OTHER PORTIONS OF THE WORK (AS OPPOSED TO THE ENTIRE WORK) IN A MANNER CONTEMPLATED BY USER; USER UNDERSTANDS AND AGREES THAT NEITHER CCC NOR THE RIGHTSHOLDER MAY HAVE SUCH ADDITIONAL RIGHTS TO GRANT.

13) **Effect of Breach.** Any failure by User to pay any amount when due, or any use by User of a Work beyond the scope of the License set forth in the Order Confirmation and/or the Terms, shall be a material breach of such License. Any breach not cured within 10 days of written notice thereof shall result in immediate termination of such License without further notice. Any unauthorized (but licensable) use of a Work that is terminated immediately upon notice thereof may be liquidated by payment of the Rightsholder's ordinary license price therefor; any unauthorized (and unlicensable) use that is not terminated immediately for any reason (including, for example, because materials containing the Work cannot reasonably be recalled) will be subject to all remedies available at law or in equity, but in no event to a payment of less than three times the Rightsholder's ordinary license price for the most closely analogous licensable use plus Rightsholder's and/or CCC's costs and expenses incurred in collecting such payment.

14) **Additional Terms for Specific Products and Services.** If a User is making one of the uses described in this Section 14, the additional terms and conditions apply:

a) **Print Uses of Academic Course Content and Materials (photocopies for academic coursepacks or classroom handouts).** For photocopies for academic coursepacks or classroom handouts the following additional terms apply:

i) The copies and anthologies created under this License may be made and assembled by faculty members individually or at their request by on-campus bookstores or copy centers, or by off-campus copy shops and other similar entities.

ii) No License granted shall in any way: (i) include any right by User to create a substantively non-identical copy of the Work or to edit or in any other way modify the Work (except by means of deleting material immediately preceding or following the entire portion of the Work copied) (ii) permit "publishing ventures" where any particular anthology would be systematically marketed at multiple institutions.

iii) Subject to any Publisher Terms (and notwithstanding any apparent contradiction in the Order Confirmation arising from data provided by User), any use authorized under the academic pay-per-use service is limited as follows:

A) any License granted shall apply to only one class (bearing a unique identifier as assigned by the institution, and thereby including all sections or other subparts of the class) at one institution;

B) use is limited to not more than 25% of the text of a book or of the items in a published collection of essays, poems or articles;

C) use is limited to no more than the greater of (a) 25% of the text of an issue of a journal or other periodical or (b) two articles from such an issue;

D) no User may sell or distribute any particular anthology, whether photocopied or electronic, at more than one institution of learning;

Permission document for Figure 1.1(b).

E) in the case of a photocopy permission, no materials may be entered into electronic memory by User except in order to produce an identical copy of a Work before or during the academic term (or analogous period) as to which any particular permission is granted. In the event that User shall choose to retain materials that are the subject of a photocopy permission in electronic memory for purposes of producing identical copies more than one day after such retention (but still within the scope of any permission granted), User must notify CCC of such fact in the applicable permission request and such retention shall constitute one copy actually sold for purposes of calculating permission fees due; and

F) any permission granted shall expire at the end of the class. No permission granted shall in any way include any right by User to create a substantively non-identical copy of the Work or to edit or in any other way modify the Work (except by means of deleting material immediately preceding or following the entire portion of the Work copied).

iv) Books and Records; Right to Audit. As to each permission granted under the academic pay-per-use Service, User shall maintain for at least four full calendar years books and records sufficient for CCC to determine the numbers of copies made by User under such permission. CCC and any representatives it may designate shall have the right to audit such books and records at any time during User's ordinary business hours, upon two days' prior notice. If any such audit shall determine that User shall have underpaid for, or underreported, any photocopies sold or by three percent (3%) or more, then User shall bear all the costs of any such audit; otherwise, CCC shall bear the costs of any such audit. Any amount determined by such audit to have been underpaid by User shall immediately be paid to CCC by User, together with interest thereon at the rate of 10% per annum from the date such amount was originally due. The provisions of this paragraph shall survive the termination of this License for any reason.

b) *Digital Pay-Per-Uses of Academic Course Content and Materials (e-coursepacks, electronic reserves, learning management systems, academic institution intranets)*. For uses in e-coursepacks, posts in electronic reserves, posts in learning management systems, or posts on academic institution intranets, the following additional terms apply:

i) The pay-per-uses subject to this Section 14(b) include:

A) **Posting e-reserves, course management systems, e-coursepacks for text-based content**, which grants authorizations to import requested material in electronic format, and allows electronic access to this material to members of a designated college or university class, under the direction of an instructor designated by the college or university, accessible only under appropriate electronic controls (e.g., password);

B) **Posting e-reserves, course management systems, e-coursepacks for material consisting of photographs or other still images not embedded in text**, which grants not only the authorizations described in Section 14(b)(i)(A) above, but also the following authorization: to include the requested material in course materials for use consistent with Section 14(b)(i)(A) above, including any necessary resizing, reformatting or modification of the resolution of such requested material (provided that such modification does not alter the underlying editorial content or meaning of the requested material, and provided that the resulting modified content is used solely within the scope of, and in a manner consistent with, the particular authorization described in the Order Confirmation and the Terms), but not including any other form of manipulation, alteration or editing of the requested material;

C) **Posting e-reserves, course management systems, e-coursepacks or other academic distribution for audiovisual content**, which grants not only the authorizations described in Section 14(b)(i)(A) above, but also the following authorizations: (i) to include the requested material in course materials for use consistent with Section 14(b)(i)(A) above; (ii) to display and perform the requested material to such members of such class in the physical classroom or remotely by means of streaming media or other video formats; and (iii) to "clip" or reformat the requested material for purposes of time or content management or ease of delivery, provided that such "clipping" or reformatting does not alter the underlying editorial content or meaning of the requested material and that the resulting material is used solely within the scope of, and in a manner consistent with, the particular authorization described in the Order Confirmation and the Terms. Unless expressly set forth in the relevant Order Confirmation, the License does not authorize any other form of manipulation, alteration or editing of the requested material.

ii) Unless expressly set forth in the relevant Order Confirmation, no License granted shall in any way: (i) include any right by User to create a substantively non-identical copy of the Work or to edit or in any other way modify the Work (except by means of deleting material immediately preceding or following the entire portion of the Work copied or, in the case of Works subject to Sections 14(b)(1)(B) or (C) above, as described in such Sections) (ii) permit "publishing ventures" where any particular course materials would be systematically marketed at multiple institutions.

Permission document for Figure 1.1(b).

iii) Subject to any further limitations determined in the Rightsholder Terms (and notwithstanding any apparent contradiction in the Order Confirmation arising from data provided by User), any use authorized under the electronic course content pay-per-use service is limited as follows:

A) any License granted shall apply to only one class (bearing a unique identifier as assigned by the institution, and thereby including all sections or other subparts of the class) at one institution;

B) use is limited to not more than 25% of the text of a book or of the items in a published collection of essays, poems or articles;

C) use is limited to not more than the greater of (a) 25% of the text of an issue of a journal or other periodical or (b) two articles from such an issue;

D) no User may sell or distribute any particular materials, whether photocopied or electronic, at more than one institution of learning;

E) electronic access to material which is the subject of an electronic-use permission must be limited by means of electronic password, student identification or other control permitting access solely to students and instructors in the class;

F) User must ensure (through use of an electronic cover page or other appropriate means) that any person, upon gaining electronic access to the material, which is the subject of a permission, shall see:

- o a proper copyright notice, identifying the Rightsholder in whose name CCC has granted permission,
- o a statement to the effect that such copy was made pursuant to permission,
- o a statement identifying the class to which the material applies and notifying the reader that the material has been made available electronically solely for use in the class, and
- o a statement to the effect that the material may not be further distributed to any person outside the class, whether by copying or by transmission and whether electronically or in paper form, and User must also ensure that such cover page or other means will print out in the event that the person accessing the material chooses to print out the material or any part thereof.

G) any permission granted shall expire at the end of the class and, absent some other form of authorization, User is thereupon required to delete the applicable material from any electronic storage or to block electronic access to the applicable material.

iv) Uses of separate portions of a Work, even if they are to be included in the same course material or the same university or college class, require separate permissions under the electronic course content pay-per-use Service. Unless otherwise provided in the Order Confirmation, any grant of rights to User is limited to use completed no later than the end of the academic term (or analogous period) as to which any particular permission is granted.

v) Books and Records; Right to Audit. As to each permission granted under the electronic course content Service, User shall maintain for at least four full calendar years books and records sufficient for CCC to determine the numbers of copies made by User under such permission. CCC and any representatives it may designate shall have the right to audit such books and records at any time during User's ordinary business hours, upon two days' prior notice. If any such audit shall determine that User shall have underpaid for, or underreported, any electronic copies used by three percent (3%) or more, then User shall bear all the costs of any such audit; otherwise, CCC shall bear the costs of any such audit. Any amount determined by such audit to have been underpaid by User shall immediately be paid to CCC by User, together with interest thereon at the rate of 10% per annum from the date such amount was originally due. The provisions of this paragraph shall survive the termination of this license for any reason.

c) ***Pay-Per-Use Permissions for Certain Reproductions (Academic photocopies for library reserves and interlibrary loan reporting) (Non-academic internal/external business uses and commercial document delivery).*** The License expressly excludes the uses listed in Section (c)(i)-(v) below (which must be subject to separate license from the applicable Rightsholder) for: academic photocopies for library reserves and interlibrary loan reporting; and non-academic internal/external business uses and commercial document delivery.

i) electronic storage of any reproduction (whether in plain-text, PDF, or any other format) other than on a transitory basis;

ii) the input of Works or reproductions thereof into any computerized database;

Permission document for Figure 1.1(b).

- iii) reproduction of an entire Work (cover-to-cover copying) except where the Work is a single article;
- iv) reproduction for resale to anyone other than a specific customer of User;
- v) republication in any different form. Please obtain authorizations for these uses through other CCC services or directly from the rightsholder.

Any license granted is further limited as set forth in any restrictions included in the Order Confirmation and/or in these Terms.

d) *Electronic Reproductions in Online Environments (Non-Academic-email, intranet, internet and extranet)*. For "electronic reproductions", which generally includes e-mail use (including instant messaging or other electronic transmission to a defined group of recipients) or posting on an intranet, extranet or Intranet site (including any display or performance incidental thereto), the following additional terms apply:

i) Unless otherwise set forth in the Order Confirmation, the License is limited to use completed within 30 days for any use on the Internet, 60 days for any use on an intranet or extranet and one year for any other use, all as measured from the "republication date" as identified in the Order Confirmation, if any, and otherwise from the date of the Order Confirmation.

ii) User may not make or permit any alterations to the Work, unless expressly set forth in the Order Confirmation (after request by User and approval by Rightsholder); provided, however, that a Work consisting of photographs or other still images not embedded in text may, if necessary, be resized, reformatted or have its resolution modified without additional express permission, and a Work consisting of audiovisual content may, if necessary, be "clipped" or reformatted for purposes of time or content management or ease of delivery (provided that any such resizing, reformatting, resolution modification or "clipping" does not alter the underlying editorial content or meaning of the Work used, and that the resulting material is used solely within the scope of, and in a manner consistent with, the particular License described in the Order Confirmation and the Terms.

15) Miscellaneous.

a) User acknowledges that CCC may, from time to time, make changes or additions to the Service or to the Terms, and that Rightsholder may make changes or additions to the Rightsholder Terms. Such updated Terms will replace the prior terms and conditions in the order workflow and shall be effective as to any subsequent Licenses but shall not apply to Licenses already granted and paid for under a prior set of terms.

b) Use of User-related information collected through the Service is governed by CCC's privacy policy, available online at www.copyright.com/about/privacy-policy/.

c) The License is personal to User. Therefore, User may not assign or transfer to any other person (whether a natural person or an organization of any kind) the License or any rights granted thereunder; provided, however, that, where applicable, User may assign such License in its entirety on written notice to CCC in the event of a transfer of all or substantially all of User's rights in any new material which includes the Work(s) licensed under this Service.

d) No amendment or waiver of any Terms is binding unless set forth in writing and signed by the appropriate parties, including, where applicable, the Rightsholder. The Rightsholder and CCC hereby object to any terms contained in any writing prepared by or on behalf of the User or its principals, employees, agents or affiliates and purporting to govern or otherwise relate to the License described in the Order Confirmation, which terms are in any way inconsistent with any Terms set forth in the Order Confirmation, and/or in CCC's standard operating procedures, whether such writing is prepared prior to, simultaneously with or subsequent to the Order Confirmation, and whether such writing appears on a copy of the Order Confirmation or in a separate instrument.

e) The License described in the Order Confirmation shall be governed by and construed under the law of the State of New York, USA, without regard to the principles thereof of conflicts of law. Any case, controversy, suit, action, or proceeding arising out of, in connection with, or related to such License shall be brought, at CCC's sole discretion, in any federal or state court located in the County of New York, State of New York, USA, or in any federal or state court whose geographical jurisdiction covers the location of the Rightsholder set forth in the Order Confirmation. The parties expressly submit to the personal jurisdiction and venue of each such federal or state court.

Last updated October 2022

Permission statement for Figure 2.1(b).

The IEEE does not require individuals working on a thesis to obtain a formal reuse license, however, you may print out this statement to be used as a permission grant:

Requirements to be followed when using any portion (e.g., figure, graph, table, or textual material) of an IEEE copyrighted paper in a thesis:

- 1) In the case of textual material (e.g., using short quotes or referring to the work within these papers) users must give full credit to the original source (author, paper, publication) followed by the IEEE copyright line © 2011 IEEE.
- 2) In the case of illustrations or tabular material, we require that the copyright line © [Year of original publication] IEEE appear prominently with each reprinted figure and/or table.
- 3) If a substantial portion of the original paper is to be used, and if you are not the senior author, also obtain the senior author's approval.

Requirements to be followed when using an entire IEEE copyrighted paper in a thesis:

- 1) The following IEEE copyright/ credit notice should be placed prominently in the references: © [year of original publication] IEEE. Reprinted, with permission, from [author names, paper title, IEEE publication title, and month/year of publication]
- 2) Only the accepted version of an IEEE copyrighted paper can be used when posting the paper or your thesis on-line.
- 3) In placing the thesis on the author's university website, please display the following message in a prominent place on the website: In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of [university/educational entity's name goes here]'s products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribu-

tion, please go to http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a License from RightsLink.

If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.