

Illuminating the photonic lantern: coherent spectro-polarimetric characterisation for high angular resolution astronomy

ADAM KRZYSZTOF TARAS

B. Eng. (Hons.), B Sci.



THE UNIVERSITY OF
SYDNEY

Supervisor: Dr. Barnaby R. M. Norris
Associate Supervisor: Prof. Peter G. Tuthill
Associate Supervisor: Prof. Sergio Leon-Saval

A thesis submitted in fulfilment of
the requirements for the degree of
Masters of Philosophy

School of Physics
Faculty of Science
The University of Sydney
Australia

November 2025

Statement of Contributions

I, Adam Taras, hereby declare that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged. Explicitly, the contributions involving collaboration include

- Designing the research question, conducting background research and writing of the literature review was completed with the guidance of my supervisors.
- The white light fringe search algorithm was jointly developed with M. Ireland [1].
- B. Norris provided the simulated results.

Generative artificial intelligence was used only in writing software (GitHub Copilot plugin). All other contributions are entirely my own, including independently collecting data, conducting experiments, generating figures and analysis in this thesis. This thesis has not been submitted for any other degree or purpose.

As part of my engineering contract with Astralis before commencing this degree, I designed and implemented a predecessor to the characterisation system which collected data at a few discrete wavelengths incoherently, and used an existing package (DigHolo [2]) for post-processing. All other contributions, including designing and implementing the current version of the setup with polarisation measurements and fast, continuous wavelength sweeping with custom post-processing code were completed during this degree.

A majority of this thesis has been submitted in a manuscript entitled “Illuminating the lantern: coherent, spectro-polarimetric characterisation of a multimode converter”. This is mentioned at the start of relevant chapters.

Student

Supervisor

Adam K. Taras Dr. Barnaby R. M. Norris

Date

Signature

Abstract

Photonic technologies offer an enticing means to reach exquisite detail with high angular resolution measurements of astrophysical scenes through precision measurement and control. Over the past decade, photonic lanterns have become key components for applications requiring mode sorting ranging from astronomical imaging to free-space communications. In these applications, knowledge of the electric field needs to be inferred from intensity measurements alone or beam-shape control is needed. While these devices can efficiently and uniquely map a set of input modes to single-mode outputs (or vice versa), the optical mode transfer matrix of any particular fabricated device cannot be fully specified at the design stage due to manufacturing limitations. This thesis presents a characterisation system to directly measure the electric field from a photonic lantern using digital off-axis holography, following its evolution over a 73 nm range near 1550 nm and in two orthogonal, linear polarisations. Performance of the testbed is validated on a single-mode fibre and applied to a 19-port, multicore fibre based photonic lantern. The first broadband multi-wavelength, polarisation decomposed characterisation of the mode transfer matrix of a photonic lantern is presented. As part of this characterisation, the system measures the position of zero path difference between different ports and polarisations, reveals the typical wavelength scale over which the modal mapping evolves and quantifies the mode dispersion within the device itself. In addition to detailing the system, empirical mode transfer matrices, raw data and post-processing code are shared, enabling future work in astrophotonics to understand where photonic lanterns fit in the wider picture of high angular resolution astronomical instrumentation.

Glossary

BFGS: Broyden–Fletcher–Goldfarb–Shanno. 55

LP: linearly polarised. 5, 7, 19, 23, 38, 44, 57

MCF: multicore fibre. 7, 8, 12, 21, 34, 38, 51, 52

MMF: multimode fibre. 7

SLD: superluminescent diode. 19, 23

SLM: spatial light modulator. 13

SMF: single-mode fibre. 7, 8, 12, 18, 21, 29, 34, 36, 47, 48

SNR: signal to noise ratio. 24, 31, 32

Acknowledgements

Firstly, I cannot begin to express my thanks to my supervisors: Barnaby Norris, Sergio Leon-Saval and Peter Tuthill. Your invaluable feedback, guidance and support has made this thesis enjoyable, engaging and valuable. I'm particularly grateful for our regular conversations on a wide range of topics that have shaped my view of our field.

I would also like to extend my thanks to the wider research groups including FOCI and the photonic lantern sensing group. Our regular discussions gave me a breadth of knowledge that I would have otherwise missed, and I look forward to collaborating with the friends made along the way.

I'm also deeply indebted to my family, in particular my parents, for their support and love. Your caring nature raised me to understand the value of education and the importance of balance in life. I owe so much to you.

Parts of this work were supported by an Australian Government Research Training Program (RTP) Scholarship in the Research Training Program fee offset, the Faculty of Science through the Research Stipend Scholarship and Astralis through the Postgraduate Research Supplementary Scholarship.

Lastly, I am grateful to my friends from university, school and the wider community for their interest in my work. Finally, I'm especially grateful to Samantha for her consistent support which inspires me to pursue positive, impactful science.

Contents

Statement of Contributions	ii
Abstract	iii
Glossary	iv
Acknowledgements	v
Contents	vi
List of Figures	viii
List of Tables	x
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Overview	2
1.3 Contributions	4
Chapter 2 Background and Literature review	5
2.1 Modes in optical fibres	5
2.2 Photonic Lanterns	7
2.2.1 Operating principle	7
2.2.2 Definition of the mode transfer matrix	8
2.2.3 Applications	9
2.3 Holography	11
2.4 Characterisation of photonic lanterns	12
2.4.1 Direct characterisation	12
2.4.2 Data-driven characterisation methods	15

Chapter 3	Methods	16	
3.1	Overview	16	
3.2	Wavelength calibration	20	
3.3	Exciting a single principal mode	21	
3.4	White light fringe finding	22	
3.5	Digital off-axis holography reconstruction	23	
3.6	Projection onto LP modes	27	
3.7	Simulation	27	
Chapter 4	Results	29	
4.1	Validation on a single-mode fibre	29	
4.2	Transfer matrix of a 19-port lantern	31	
4.3	Orthogonality of photonic lantern principal modes	38	
4.4	Fitting transfer matrices in wavelength	39	
4.5	Differential signals for principal mode dispersion	50	
4.6	Towards relative principal mode phase measurements	52	
Chapter 5	Discussion	57	
5.1	Exploring the characterisation testbed	57	
5.2	Implications for photonic lantern applications	59	
Chapter 6	Conclusion	62	
6.1	Significance	62	
6.2	Directions for future work	63	
	Bibliography	65	
	Appendix A	Component List	73

List of Figures

1.1	Conventional and proposed approaches to astronomical, high-contrast imaging	3
2.1	Simulated LP modes and their linear combinations	6
2.2	Photonic lanterns encode both the phase and amplitude at the multimode end, enabling wavefront sensing	10
2.3	Previous intensity-only characterisation of a photonic lantern	13
2.4	Example of a data-driven reconstructor for wavefront sensing	14
3.1	Visual overview of photonic lantern characterisation	17
3.2	System block diagram for digital off-axis holography	18
3.3	Data capture and analysis block flowchart	19
3.4	Wavelength to frame calibration	20
3.5	Optimisation of injection of light from a single-mode fibre into a single port of a multicore fibre	21
3.6	Parameter fitting for digital off-axis holography	25
4.1	Reconstruction validation on a single-mode fibre, with monochromatic performance and wavelength-coherent reconstruction	30
4.2	Data and reconstruction for the first 11 ports probed (out of 19) in a single wavelength and polarisation	32
4.3	Characterisation of the broadband coherence properties of the photonic lantern	33
4.4	Wavelength evolution of modes in a photonic lantern	35
4.5	Relative throughput and LP mode explainability of fields	36
4.6	Slices through the transfer matrices of the 19 port lantern	37
4.7	Orthogonality of photonic lantern principal modes	39
4.8	Simulated effect of different coherence packets on wavelength evolution	40
4.9	Multiple coherence packets improves the fit of the overlap integral with wavelength	43

4.10	Fitting the wavelength evolution of a single port, polarisation and mode	47
4.11	Limitations of the fitted model	49
4.12	Relative modal dispersion measurements of the photonic lantern	51
4.13	Simulated superpositions of port outputs reveal relative phase dependent behaviour	53
4.14	Contour plot of loss landscape for relative phase solving	54
4.15	Simulated recovery of relative phase between principal modes	55

List of Tables

4.1 Optimisation parameters for global modal transfer matrix fitting	46
A.1 Components in the characterisation system, listed in order of propagation.	73

Introduction

1.1 Motivation

Understanding the cosmos has been an aspirational goal since the dawn of human civilisation. Questions such as “How did the Earth form?” and “Is there life on other planets beyond our solar system?” have inspired debate and philosophical inquiry down the ages. Quite apart from their intellectual appeal, these goals also “serve to organize and measure the best of our energies and skills” (John F. Kennedy, 1962), so that working towards answers has been incredibly fruitful for wider society, contributing to the invention of new technologies in areas such as telecommunications (Wi-Fi) [3], cameras (CCD detectors) [4] and spaceflight. This thesis takes steps towards one of the most critical challenges in contemporary science: the ability to deliver direct images of exoplanets orbiting distant stars. The technical demands of this challenge are analogous to imaging a firefly only a few millimetres from the glare of a lighthouse beam from a distance of 100 km whilst fighting the distortions of the atmosphere [5].

Currently, high contrast astronomical imaging is most commonly achieved using a coronagraph [6], which blocks the starlight with the intent of revealing faint exoplanets and dust in the immediate vicinity. Although this technology has delivered success over a limited parameter space – that of understanding giant planets at large separations – its impact on exoplanetary science has been muted. The richest domain for scientific inquiry lies within innermost region near the star, and in particular the so-called “Habitable Zone”. In reaching extreme contrasts at larger spatial scales, coronagraphs typically deliberately block this region. Scientific goals apart from understanding exoplanets include probing protoplanetary

disks and the regions around dying stars, piecing together the life cycle of systems similar to ours. The proposed approach uses novel photonics to complement the coronagraph, by using the light from the innermost region (including the star), imaging structure in the Habitable Zone. In addition to sensitivity to structure at finer resolutions than the diffraction limit [7], photonic lanterns could also provide a means of disambiguating aberrated starlight from faint companions by exploiting the spatial coherence properties of the light [8, 9], ultimately delivering targets at more extreme contrasts. Furthermore, imaging at such high resolutions requires correction of the atmospheric seeing, which for high contrast imaging requires wavefront sensing at the same wavelengths as the science signal [10], as delivered by such devices.

1.2 Overview

Figure 1.1 visually depicts the conventional and proposed approaches. Typically, coronagraphic systems measure intensity only images of a scene where the starlight is blocked. As the wavefront sensor is split into a different set of optics, non-common path errors are present, which cause speckles in the final image. This forms an ambiguity when looking at intensity only imagery – a persistent, bright pixel may be a true companion or the result of imperfect correction. This thesis forms part of an alternate architecture, proposing the use of a photonic lantern to encode the complex electric field at the image plane into the intensities at single-mode outputs. A reconstruction algorithm then jointly estimates the astrophysical scene and the state of the wavefront. This estimator can exploit both the amplitude, phase and mutual coherence at the image plane, yielding richer information than intensities only. For such an architecture to function and make convincing detections, knowledge of the mapping between the inputs and outputs of the photonic lantern is critical. Previous work has only just begun to measure this mapping, typically using intensity only measurements at a few wavelengths. This thesis focuses on a more complete characterisation of the photonic lanterns in order to validate and inform the framework as a whole, directly measuring the complex electric field over a dense range of wavelengths and in two orthogonal polarisations.




FIGURE 1.1: **Conventional and proposed approaches to astronomical, high-contrast imaging.** In astronomical imaging, the true scene is distorted by atmospheric turbulence. Conventional direct imaging approaches involve separate wavefront sensing/corrections to the science signal and removing starlight using a coronagraph, which also removes information about the region closest to the star (result image from [11]). The proposed approach instead uses a photonic lantern to encode the complex electric field into intensities. A reconstruction algorithm, validated and informed by a separate characterisation process, is then able to infer images from the innermost angles of the system as well as provide real-time feedback to the adaptive optics system to correct for the atmosphere. The characterisation stage is the focus of this thesis. Coronagraph image credit: [NASA](#).

In addition to the benefits outlined above, photonic technologies generally have a smaller form factor, enabling a wider range of use at lower cost, in particular for space missions.

Many astronomical instrument frameworks, such as interferometric combination of beams from multiple telescopes fed with photonic lanterns [12] or beam combination after injection into a single device [13] also benefit greatly from an understanding of this modal mapping. It is also worth noting that, while well motivated by high angular resolution astronomy alone, a fundamental understanding of the behaviour of photonic lanterns benefits many other applications including (but not limited to): space-division multiplexing in telecommunications

[14, 15], wavefront sensing in astrophotonics [16–19], beam shaping for free-space optical communications [20, 21] and computational imaging for microendoscopy [22].

1.3 Contributions

This thesis demonstrates a means of directly characterising a manufactured photonic lantern. Specifically, the key contributions of this thesis are that:

- I design and implement a laboratory system to characterise a photonic lantern using digital off-axis holography, sharing the key insights, lessons learnt, code, and post-processed data (to be published as git repository with the submitted manuscript);
- I provide the first multi-wavelength, polarisation decomposed characterisation of the principal modes of a photonic lantern, highlighting some insights about the modal mappings; and
- I uncover trends in the coherence of the measured fields, including direct measurement of measuring modal dispersion.

In accomplishing this an unprecedented, quantitative picture emerges detailing how fabricated devices deviate from idealised simulations. This reinforces the need for characterisation of the fabricated photonic lantern before downstream use in a wide range of applications. Future instruments that employ photonic lanterns will benefit from empirical measurement providing a prior on the device in addition to higher fidelity simulations that use the performance of simulated devices to develop new applications.

Background and Literature review

2.1 Modes in optical fibres

When light is confined in a waveguide (such as an optical fibre), its propagation and behaviour is described by Maxwell's equations. Each discrete solution to the equations is a *mode* of the waveguide, and has a propagation constant that defines the evolution of the phase of that mode.

In a typical fibre the refractive index is a circularly symmetric step function, where the size of the step is referred to as the index contrast. The arising modes are conventionally approximated by linearly polarised (LP) modes, and naturally expressed in cylindrical coordinates, indexed by a radial index m and an azimuthal index ℓ , commonly denoted as e.g. LP01 (for $m = 0, \ell = 1$). The number and (to a lesser extent) shape of these modes depends on the relationship between the core diameter, index contrast and wavelength λ . Each mode experiences a different effective index, and hence a different optical path difference through the device.

Figure 2.1 shows the first 12 LP modes in a multimode fibre, and how a (complex) linear combination of them can span a useful space of electric fields. This is a subset of basis set for the modes used in this thesis, where the characterised device will support up to 23 modes in total.

In this work, electric fields are visualised on a two dimensional colormap, with the phase axis selected to be perceptually uniform (the rate of color change with value appears constant) and uniform in lightness (all phases appear equally bright), and the amplitude reflected in

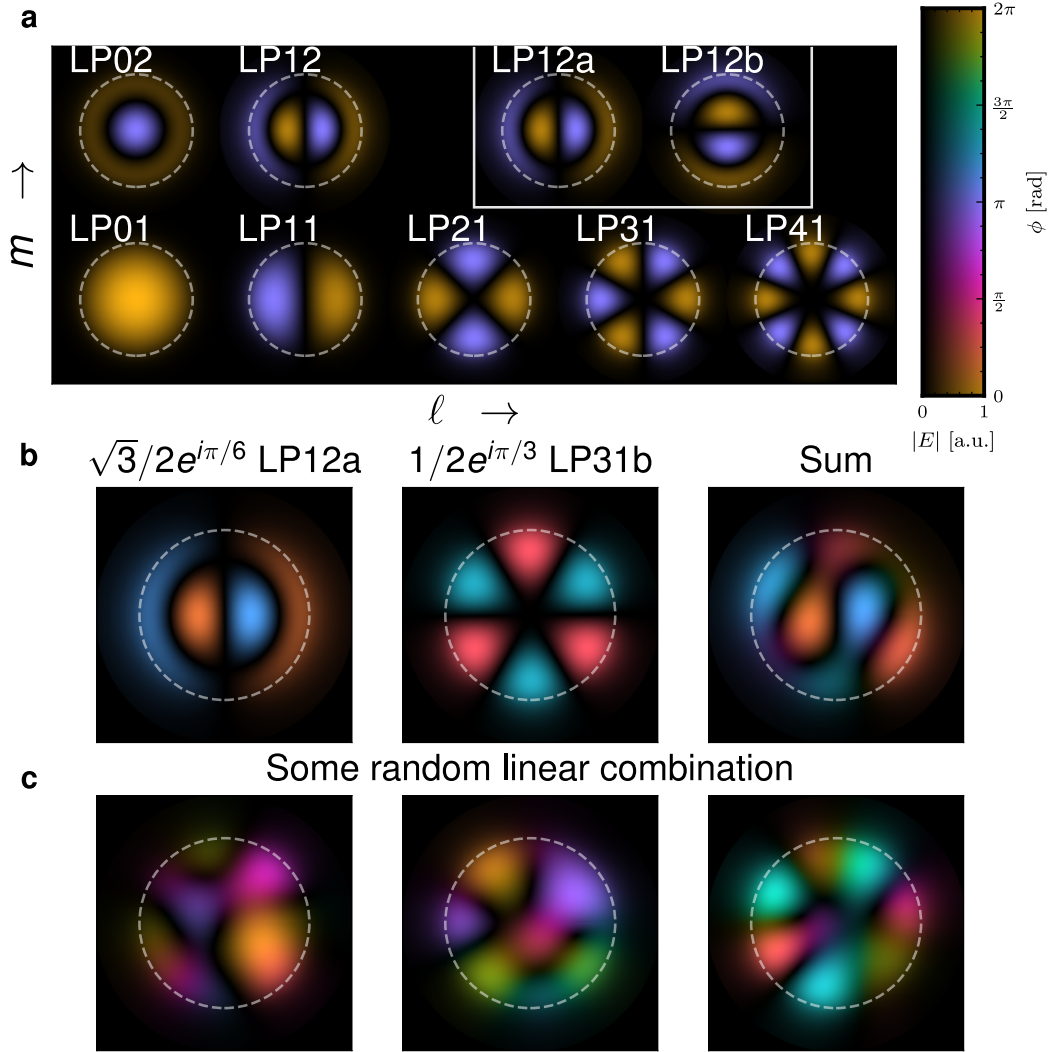


FIGURE 2.1: **Simulated LP modes and their linear combinations.** **a** Electric fields visualised with the amplitude as the brightness of the pixel and the phase as the color. LP modes supported in a $26.7 \mu\text{m}$ core diameter fibre with index contrast 5.5×10^{-3} at $\lambda = 1.55 \mu\text{m}$. The white dashed circle represents the outline of the core. (White box) Modes with $\ell > 1$ actually represent two modes, $\text{LP}^{m\ell a}$ and $\text{LP}^{m\ell b}$, with a rotation between the two. This fibre supports 12 modes in total. **b** A complex linear combination of two modes show that even relatively simple components (left and middle) generate non-trivial superpositions (right). **c** Random linear combinations of all modes also span a variety of electric field profiles.

the brightness of the pixel. This choice ensures all phases appear to vanish when moving towards zero amplitude equally quickly, which is not the case with the HSV colormap that is typically used.

It is useful to equip this vector space with an inner product. In this thesis, I use

$$\langle \mathbf{A}, \mathbf{B} \rangle = \frac{\int \mathbf{A} \mathbf{B}^* dA}{\sqrt{(\int |\mathbf{A}|^2 dA) (\int |\mathbf{B}|^2 dA)}} \quad (2.1)$$

as the normalised inner product between two fields \mathbf{A} and \mathbf{B} , where $*$ denotes the complex conjugate. The integrals are over the interface, i.e. the plane perpendicular to the fibre face. The absolute value of this quantity gives the fraction of total power shared between the two fields. The LP modes form an orthonormal basis set in this space.

2.2 Photonic Lanterns

2.2.1 Operating principle

Photonic lanterns provide a means of converting a superposition of a multi-modal basis (LP in this thesis) into a set of single modes. Such a device was originally conceived as a means of applying Bragg gratings – custom wavelength filters – to light in a multimode fibre (MMF) [23]. This was used to suppress OH lines in astrophotonics applications. Photonic lanterns are low loss [24] and can be made in a variety of ways [25]. While this original work was focused on building high efficiency devices for spectroscopy, the utility of photonic lanterns as mode converters has only recently been widely explored.

In recent applications, an understanding of the optical mode transfer matrix – that is the mapping from modes at one end of the device to modes at the other – is critical. This transfer matrix is linear with complex electric fields (i.e. it satisfies superposition and scalar multiplication), however past approaches typically only characterise devices in intensity. Inferring the mode transfer matrix in this case is not only an under-constrained problem, but a (mathematically) non-linear one.

In this work, the term “port” refers to a waveguide at the single-mode end of the device, whether that be through a single-mode fibre (SMF) – as in a pigtailed photonic lantern – or through a single core in a multicore fibre (MCF).

2.2.2 Definition of the mode transfer matrix

In this subsection, all coefficients and matrices are functions of wavelength and polarisation, however this is omitted for brevity.

A photonic lantern is a mode sorting device that maps from n_{mm} modes supported at the multimode end with complex coefficients $c_{\text{mm}} \in \mathbb{C}^{n_{\text{mm}}}$ (known as the ‘input’ in e.g. wavefront sensing applications, but the ‘output’ in beam shaping applications), to the n_{sm} single modes supported at the MCF or multiple SMF end with coefficients $c_{\text{sm}} \in \mathbb{C}^{n_{\text{sm}}}$. These coefficients form the electric field in each respective basis, with

$$\mathbf{E}_{\text{mm}} = \sum_i c_{\text{mm},i} \mathbf{E}_{\text{mm},i}^b, \quad \mathbf{E}_{\text{sm}} = \sum_i c_{\text{sm},i} \mathbf{E}_{\text{sm},i}^b, \quad (2.2)$$

where the sums are taken over all supported modes, $\mathbf{E}_{\text{mm},i}^b$ is the i -th basis mode supported at the multimode end (assumed to be a LP mode basis in this work), and $\mathbf{E}_{\text{sm},i}^b$ is the i -th basis mode at the MCF end (a simple single mode located at a position on a grid, typically hexagonal). Most photonic lanterns are designed such that $n_{\text{mm}} = n_{\text{sm}}$ at the shorter wavelengths of interest, but this doesn’t need to be the case over the full bandwidth. Note that (in analogy to etendue in classical optics) lossless devices require that the number of modes at the input be no more than the number at the output. Hence the ideal relative value for n_{mm} and n_{sm} depend on the direction of propagation used in the application.

A transfer matrix $T \in \mathbb{C}^{n_{\text{sm}} \times n_{\text{mm}}}$, then, is the linear transformation applied by the photonic lantern, relating the above sets of coefficients as

$$c_{\text{sm}} = T c_{\text{mm}}. \quad (2.3)$$

Whilst these operations are linear in the modal bases, typical measurement systems capture the intensity $|\sum_i c_i \mathbf{E}_i|^2$, which is mathematically non-linear in the coefficients.

Each row of the transfer matrix determines the response of the multicore fibre port to the electric field at the multimode end. The approach in this thesis measures T^\dagger , the Hermitian transpose of T . The proposed approach injects a single mode at the multicore fibre end and use digital off-axis holography to measure the complex electric field at the multimode end,

known as a principal mode [13]. Between measurements of different ports, the system drifts in phase and hence there is an ambiguous relative phase between different ports. Thus each reconstructed field contains a superposition with

$$c_{\text{mm}} = c_{\text{sm}} T^\dagger = e^{i\phi_e} \mathbf{u}_j T^\dagger = e^{i\phi_e} T_j^\dagger, \quad (2.4)$$

where \mathbf{u}_j is a vector of zeros with a one at index j , and ϕ_e is a (unknown) relative phase between ports. In other words, the proposed approach measures each row of the transfer matrix up to a single absolute phase ambiguity.

In the wavelength dimension, the quantities are measured coherently and rapidly so that there is no information-destroying phase wrapping nor large drift that occurs between samples. Hence the proposed approach is able to capture the relative phase between wavelength samples.

2.2.3 Applications

Arguably the most mature application of photonic lanterns in astronomy is measurement of the aberrations caused by the atmosphere, known as wavefront sensing. This problem is challenging since cameras and photodetectors measure intensity whilst the atmosphere imparts phase aberrations, and the timescale at which the system changes is on the order of milliseconds.

Photonic lanterns offer solutions to these challenges. First, these devices are sensitive to the phase in addition to the amplitude at the multimode end of the device, as illustrated in Figure 2.2. Despite the degeneracy in image intensity with pupil plane aberrations (astigmatism in this case, but there is a degeneracy in general), the intensities of the photonic lantern outputs are not degenerate. In addition, as high throughput devices, integration times can be kept short, enabling adaptive optics systems to keep up with the rapid changes of the atmosphere.

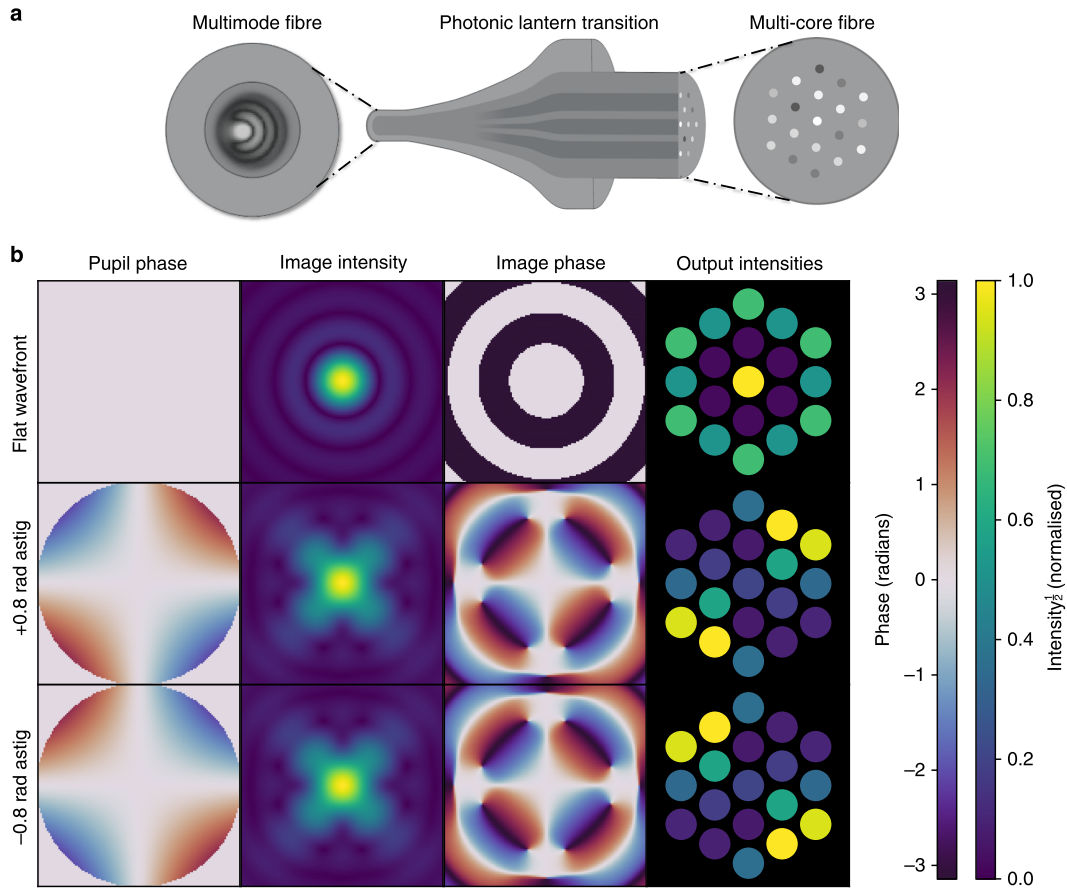


FIGURE 2.2: Photonic lanterns encode both the phase and amplitude at the multimode end, enabling wavefront sensing. Figure from [16]. **a** Aberrated light from a telescope is brought to focus on a photonic lantern and transitions to a grid of single modes, where the intensity is measured. **b** Simulated results for different aberrations. With different signs on pupil phase, there is a degeneracy when observing intensity only at the image plane. The photonic lantern encodes both amplitude and phase, and hence the intensities at the single mode end are not degenerate.

It is then unsurprising, then, that this application in particular has attracted significant interest: from proof-of-concept [16], theoretical developments with simulations [26, 27] to real-time demonstrations at world leading observatories [18].

More recently, progress has been made on the problem of reconstructing astrophysical images with a photonic lantern. In conventional (intensity only) imaging, aberrations of the wavefront irreversibly mix the faint light from the Habitable Zone with light from the star.

The photonic lantern, on the other hand, interferes light in the scene through the propagation in an adiabatic taper. This means that one could (in principle) differentiate between contaminating starlight and other incoherent sources.

In classical imaging, the resolution limit of a telescope with diameter D at a wavelength λ is $1.22\lambda/D$. Recent work [7, 28], however, has shown that measuring the phase (in addition to amplitude) overcomes this limit. Furthermore, systems that sort modes have been shown to be quantum optimal for the problem of exoplanet detection and localisation [7], meaning that they are the most informative way to use the rare and valuable photons from an astrophysical scene, especially when the target is a faint companion. With this in mind, ‘oversampled’ photonic lanterns have been proposed [8], which have more ports than modes supported at the multimode end, and hence have sufficient intensity-only measurements to constrain the amplitude, phase and spatial coherence of the incoming light. This technology could be configured so as to complement typical coronagraphs which block the inner $2 - 3\lambda/D$ of the scene.

2.3 Holography

Holography is a broad term that describes the reconstruction of both the amplitude and phase information of a wavefront. Its importance was highlighted with the 1971 Nobel Prize in Physics, awarded to Gabor for his work in the field [29].

A direct comparison can be drawn between photography (which captures only the amplitude of the wavefront) and holography. By capturing the phase information too, holograms record the complete electric field, in particular including direction information. This means that a hologram of a physical scene appears 3D, with depth from parallax and other view dependent phenomena recreated.

In this thesis, digital off-axis holography is used to reconstruct the electric field of the photonic lantern at the output by interfering the unknown field with a known, tilted reference beam. In post processing, the field is extracted through isolating the relevant area in the

Fourier domain. Digital off-axis holography [30] has been used in a range of applications, including spatial biphoton states [31] and mode sorters/multiplexers [32, 33].

2.4 Characterisation of photonic lanterns

2.4.1 Direct characterisation

The majority of previous work aimed at characterising fabricated photonic lanterns has measured the intensities of modes, typically in monochromatic light [34–36]. Figure 2.3 illustrates one such example in more detail, where essentially two photonic lanterns are fabricated at once with the multimode ends joined. These results provide some limited insight into the wavelength scale at which the modes evolve, and their far-field intensity profiles show a combination of a few basis modes as expected. However, power measurements (Figure 2.3 **b-d**) are blind to common mode mismatch in both photonic lanterns – if both are imperfect from the target design in the same way then the power in one port would still pass through to the same port on the other side despite the target field in the multimode region not matching the target design. The intensity-only images after cleaving (Figure 2.3 **e**) are also blind to the rich phase encodings and mode evolution with wavelength. Furthermore, intensity-only measurements become increasingly ambiguous for devices with more modes, and hence this approach scales poorly. The proposed approach in this thesis instead measures the full electric field using digital off-axis holography, and is able to follow the evolution of the modes in wavelength over a 73 nm bandwidth.

Other work on photonic lanterns [37] has also used digital off-axis holography, though this was at a single wavelength using a photonic lantern supporting much smaller number of modes and where the single-mode end is pigtailed (i.e. multiple SMFs). In addition to using a broad range of wavelengths and two orthogonal polarisations, the characterisation setup in this thesis deals with a MCF-fed photonic lantern, injecting light into an individual port with a fibre alignment stage. Through a fibre bundle as an adapter, the setup can also characterise pigtailed devices.

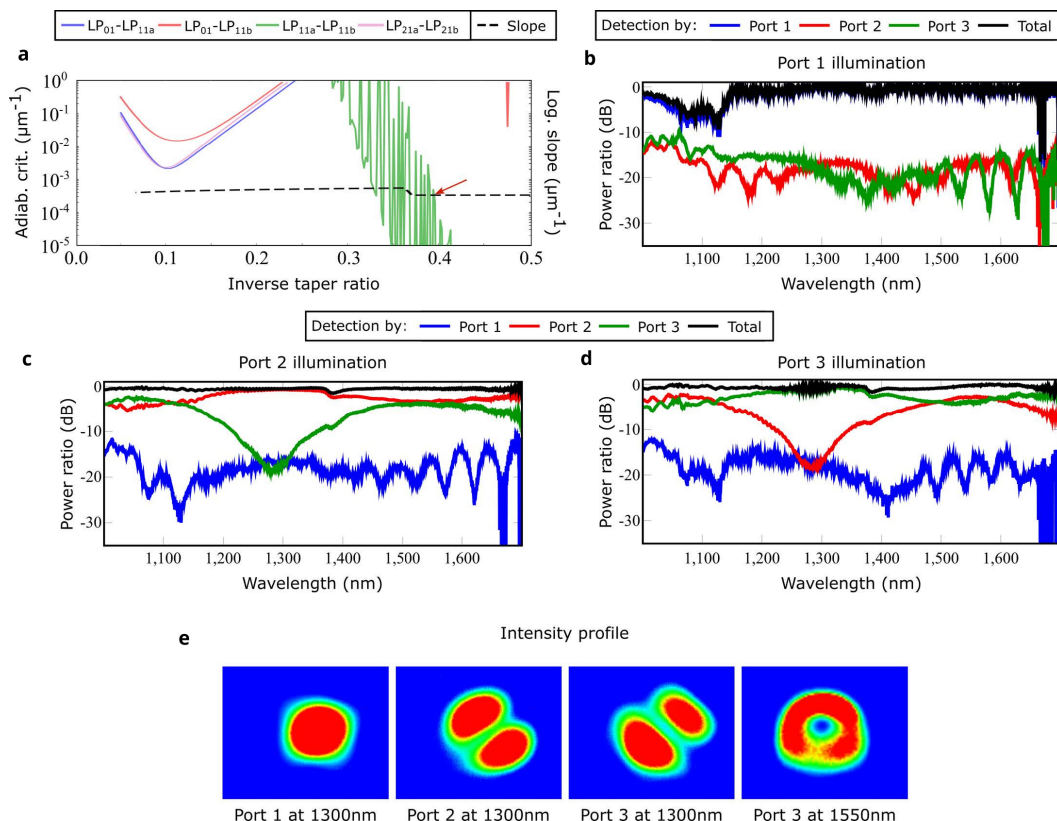


FIGURE 2.3: **Previous intensity-only characterisation of a photonic lantern.** Figure from [36]. **a** Simulated lantern, indicating that the modes do not violate the adiabatic criterion (black dashed line). **b-d**, the lantern is fabricated to be essentially two lanterns with the multimode ends joined. The coupling from each port to each other port measures the mode selectivity, however an ambiguity remains over which lantern has the imperfections. **e** After cleaving the two lanterns apart, intensity-only images are taken at a few wavelengths.

There is also ongoing work to measure the transfer matrix of a photonic lantern [38, 39] using off-axis holography for the initial estimate of the modes at large wavelength steps and then using a carefully calibrated spatial light modulator (SLM) to probe each mode at a denser wavelength spacing. However, it remains to be seen how accurate this method is. Instead, the approach in this thesis moves the complexity from requiring precisely known input electric fields to precisely measuring output electric fields. Ultimately, a combination of these characterisations is needed to assess key properties such as adiabaticity.




FIGURE 2.4: **Example of a data driven reconstructor for wavefront sensing.** Figure from [16]. An estimator is trained to predict (crosses) the Zernike coefficients applied in the pupil plane (black lines) using only the intensities of the 19 ports. Different coloured lines reflect different dataset sizes for training: 48,000 (red), 4,800 (green) and 480 (blue). The aberrations applied are of order π radians, a regime where conventional wavefront sensors fail due to assumptions of linearity. A characterisation of the modal mapping of this device would likely significantly reduce the amount of data needed to train the estimator and better constrain the model.

2.4.2 Data-driven characterisation methods

In some applications, the device has been characterised implicitly as part of the task. Figure 2.4 (taken from [16]) shows one such example, where a neural network estimator is trained to predict the applied aberrations in the pupil plane from the intensities of a 19-port lantern at the focal plane. The predictions continue to provide good estimates even when the aberrations are large, and the ablation study indicates that the (mathematical) non-linearity of the problem is captured by the model. This also has the advantage that any systematics outside of the device are learnt as part of the training process and do not bias performance.

The mapping from Zernike modes at the pupil plane to electric fields at the multimode end of the photonic lantern is related analytically through a Fourier transform, so effectively the above example is trained to learn the modal mapping, and hence is a form of characterisation. This is a fairly inefficient procedure if characterisation is the goal alone: at least 4800 measurements are needed in the final application domain, whereas the proposed setup can capture 19 holograms (one for each port) at the time of device fabrication and infer the mapping. Hence, the characterisation setup should be used to probe devices before task specific applications, and the modal mapping measured should constrain the data driven models used. A more complete, Bayesian approach would use both methods, with a fundamental model built using direct characterisation and a flexible, data-driven model to continuously fine tune in situ during the downstream task to account for changes in e.g. alignment, temperature etc.. Such a framework is part of the broader goal of this thesis, as illustrated in Figure 1.1.

Methods

All sections in this chapter are common with the submitted manuscript.

3.1 Overview

A visual summary of the proposed approach to characterise a photonic lantern using digital off-axis holography is shown in Figure 3.1. Either a broadband or wavelength sweeping source is split into a reference and injection beam. The injected beam is aligned to a single port of the photonic lantern, where the mode is re-imaged onto the detector through a beam displacer, providing two orthogonal polarisations. The reference beam is collimated and tilted, inducing fringes required for digital off-axis holography. The broadband source is used to drive the delay line stage to the path length matched position per port. The system captures holograms using the wavelength sweeping source over all ports, both with and without the reference beam. In post processing (Figure 3.1 c), the coherent component in Fourier space is recovered and the contribution of the reference beam is removed, inferring the complete electric field of the device under test. The pipeline then decomposes this into LP modes, providing the first ever multi-wavelength, polarisation decomposed characterisation of a photonic lantern.

Figure 3.2 illustrates a system level view of the setup. Light from a laser is first split for use with a wavemeter, then split into two beams: one for injection (bottom row) and one for the reference beam (middle row). Switching optics such as the shutter and fibre injection stage are operated remotely for stability. The shutter is used to capture mode field data with and without the reference arm, with the latter used for flux normalisation and validation of the

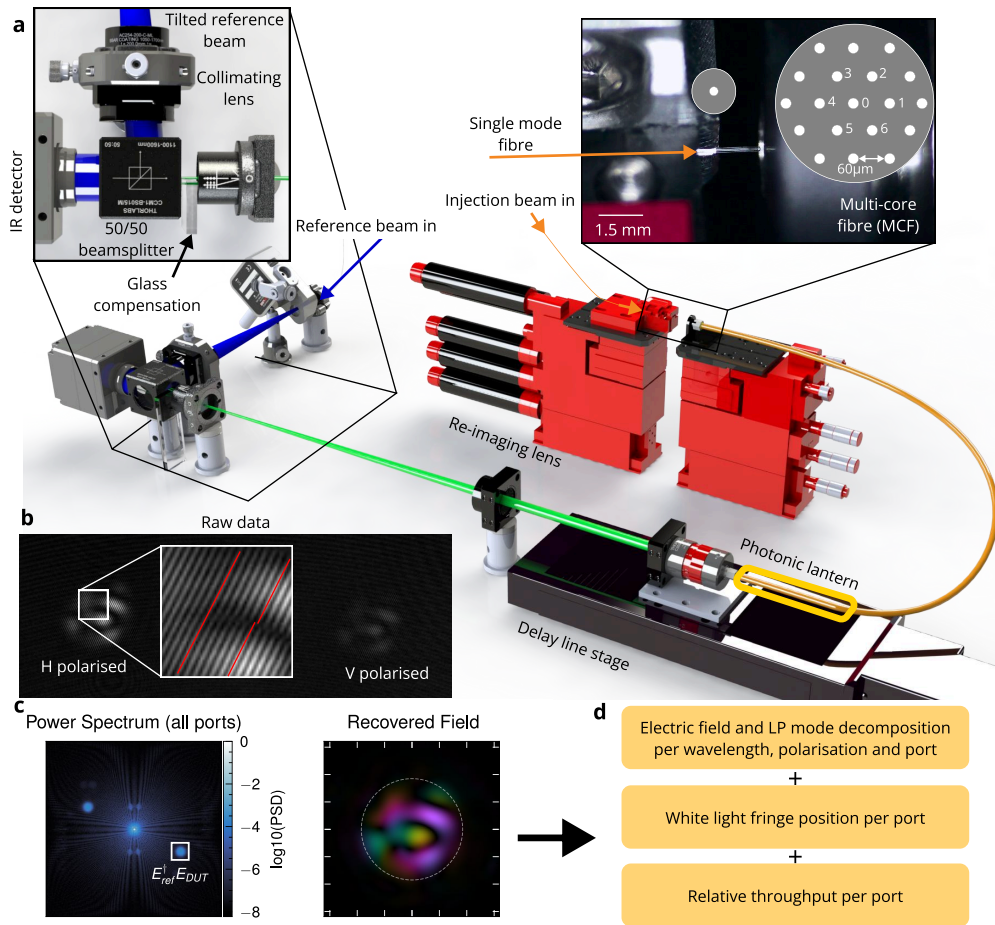


FIGURE 3.1: Visual overview of photonic lantern characterisation. **a** Labeled render of the laboratory setup. Light from a wavelength sweeping source is split into a reference (blue) and injection (orange) fibre. The injection beam excites a single port of the photonic lantern, producing a mode field that is re-imaged on the detector, with interference from a tilted reference beam producing the hologram. One arm of the interferometer has a delay line in order to maintain coherence during the wavelength sweep. **b** An example frame of raw data from the detector. The beam displacer splits the mode field into H and V polarisations. Fringes crests (red lines) reveal how the hologram captures phase information. **c** The holography reconstruction takes the frames, crops to each polarisation (not shown), extracts the coherent component in Fourier space (white box) and then removes the electric field of the reference in image space, resulting in the electric field of the device under test, shown here on a single example. **d** By taking these measurements with both a broadband source and a sweeping laser, the system characterises the port to electric fields and LP mode mapping, white light fringe position and relative throughput of the device under test.

pipeline. When using the laser in sweeping mode, a trigger from the laser drives the start of a burst of frames, captured in rapid succession using the camera’s internal clock. A table of all components is provided in Appendix A.

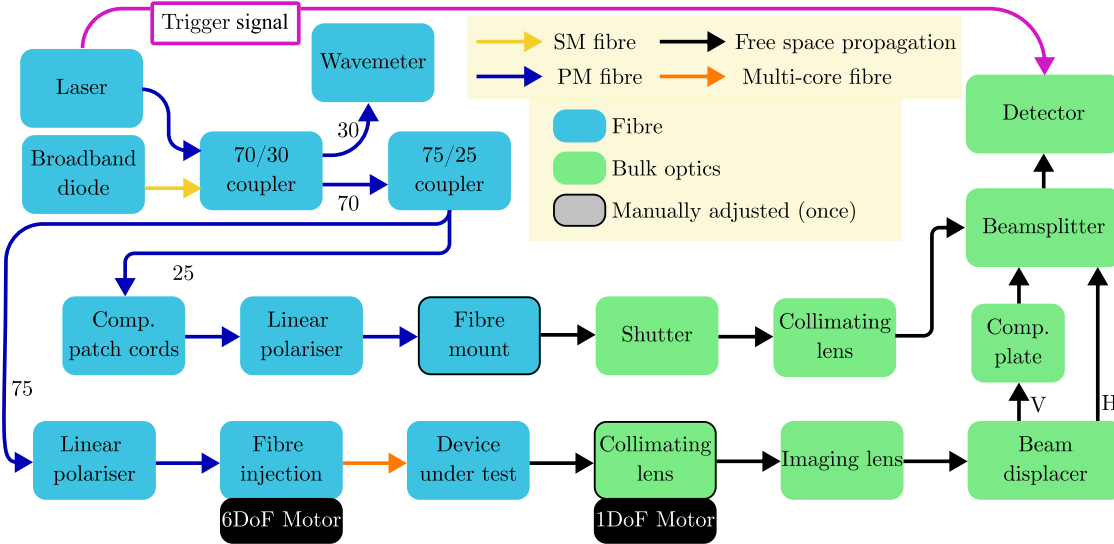


FIGURE 3.2: **System block diagram for digital off-axis holography.** Fibre and free space optical components are shown in blue and green respectively. Motors with degrees of freedom (DoF) shown in black. Path length matching is achieved coarsely with compensation (“Comp.”) patch cords and glass plates, while fine matching is achieved by moving the collimating lens on a linear stage. As the system requires stability for many measurements, all of the switching optics are operated remotely, with only some optics adjusted manually once for focus.

The system captures frames during a continuous wavelength sweep for several reasons. Firstly, the IMX990 InGaAs detector used in the setup showed signs of internal fringing (either from the cover window or within the photosensitive element itself) that was dependent on the wavelength, pixel position and angle of incidence on the light. This is not possible to calibrate, and hence the imaging is moved to a shorter coherence length regime, where such behaviour is smoothed out. An equivalent setup could use a monochromator. In sweeping over wavelengths continuously, the system is also able to sample a wide bandwidth over a very short (< 1 s) window, reducing the effects of phase drifts from between the two arms of the interferometer. When using the testbed over the scale of minutes with a SMF, the phase drifts over tens of seconds. This is likely due to a combination of factors: slow variations

in temperature causing different effective indices in each arm, vibrations from outside the laboratory slowly changing the path length of fibres and turbulence caused by air conditioning altering the differential air path of each arm. Herein all of these factors are referred to as “testbed instabilities”. Finally, the wavelength sweep captures data at the white light fringe, hence probing the dispersion of the device under test.

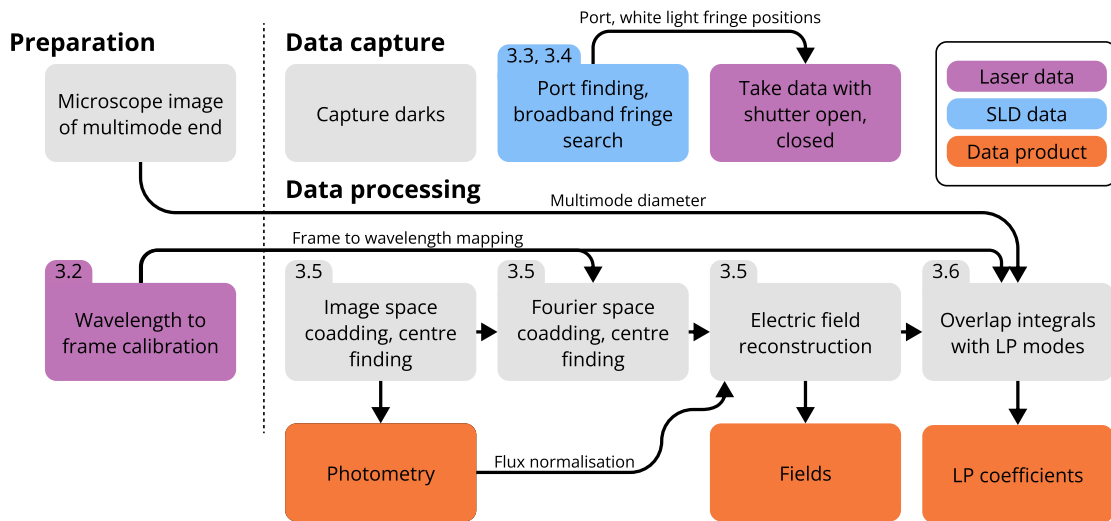


FIGURE 3.3: **Data capture and analysis block flowchart.** Numbered tabs show the relevant section. In the preparatory stage, the multimode diameter is measured directly and the laser is calibrated in time against the camera. Next, each port is excited using the broadband superluminescent diode (SLD) source, and a fringe search is conducted. The motor configurations are then used to take data with the wavelength swept source, and this is the raw data product. In the processing pipeline, centres of each polarisation are found in image space, and the off-axis component is fit in Fourier space (see Figure 3.6). In each case, a smooth function is fit against wavelength. The data captured without the reference arm becomes the photometry. The electric field is then reconstructed using the centres found previously. Finally, the projection onto LP modes is computed, forming the transfer matrix that describes the modal mapping of the device.

In this chapter, I cover the key steps that are required to collect and process this data: wavelength calibration, exciting one (and only one) port of the photonic lantern, white light fringe finding, digital off-axis holography reconstruction and projection into the LP mode space. Figure 3.3 summarises the interplay between all of these steps, as well as showing the data products produced.

3.2 Wavelength calibration

The frame to wavelength calibration is done in two steps, illustrated in Figure 3.4. First, the functional form of $\lambda(t)$ is fit using data from a wavemeter, polling wavelength measurements during the sweep. The sweep is weakly non-linear – with a quadratic term present that causes the sweep speed to decrease from an initial value of 105 nm/s to a final value of 90 nm/s. Next, the start and end times of the sweep are carefully measured by monitoring metrics like flux per polarisation or phase, with the camera providing timestamps. The total duration of the sweep is less than 1 s. This mapping is then used for all sweeps later, and the laser calibration routine is run at the start of every set of measurements.

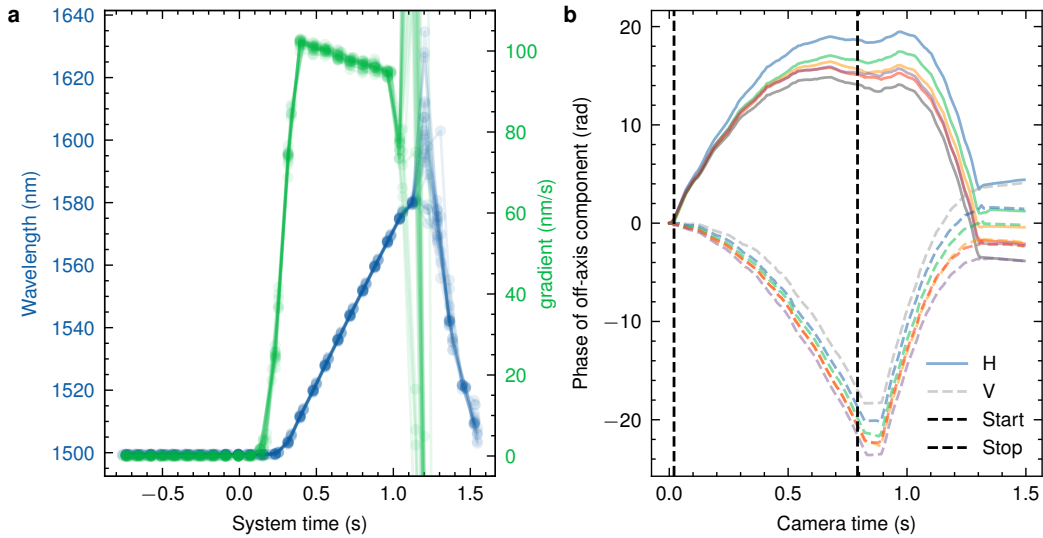


FIGURE 3.4: **Wavelength to frame calibration.** **a** Wavelength (blue) and its gradient (green) captured using a wavemeter as a function of time relative to issuing the start of the sweep command, sampled on the machine that issues commands. The sweep itself is weakly non-linear, with the gradient changing from just over 100 nm/s to 90 nm/s. **b** Phase of the off-axis holographic component (central pixels of the white box in Figure 3.1 **c**) depends on the wavelength, and hence serves as a good metric to find where the sweep starts and ends. Black dashed lines mark the first and last frame, and different colors show different sweeps (with testbed instability causing differences in phase) .

3.3 Exciting a single principal mode

The characterisation system excites a principal mode of the photonic lantern from the multicore end via butt coupling a SMF to one core of a MCF on a motorized fibre alignment stage. Using a MCF creates the additional challenge – that of aligning fibres rather than simply using a fibre switcher. This involves aligning the injection fibre to within a few micrometers while the spacing between different ports is $60\ \mu\text{m}$. It is known that the flux coupled to a waveguide is Gaussian function of position [40], so the total flux on the camera is measured in a given motor position and a Gaussian function is fit to the resulting distribution.

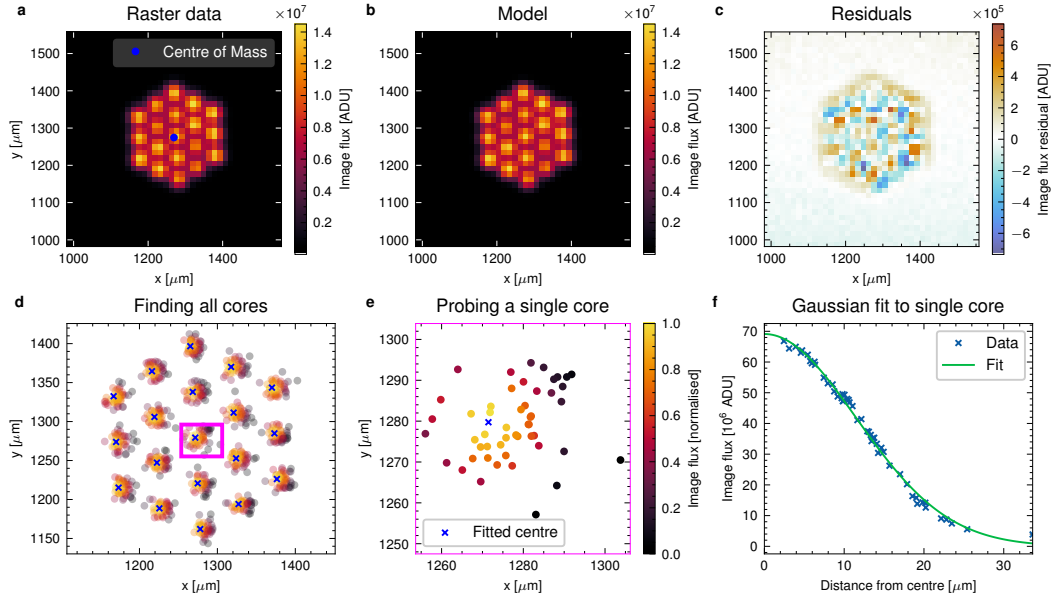


FIGURE 3.5: Optimisation of injection of light from a single-mode fibre into a single port of a multicore fibre. **a** Measuring total flux through a raster of (x, y) positions, taken after pulling away from the fibre facet. **b** Fit from a sum of Gaussians model, where the centres lie on a finite hexagonal grid. **c** Map of residuals between the data and model, note the difference in the order of magnitude on the colorbar. **d** After approach to the final distance, the system samples points near ports. **e** Zoom of sampling near a single port. **f** Normalised fluxes for a single port as a function of distance from the fitted centre. Each port is well fit by a Gaussian model.

To excite a single principal mode with minimal crosstalk, the system must operate the fibre alignment stage at a distance that has an effective mode field diameter of less than $10\ \mu\text{m}$. The

separation between cores is $60\mu\text{m}$, and hence the set of motor states that excites a single port clearly is sparse. The alignment system finds this set in two stages, summarised in each row of Figure 3.5. First, $130\mu\text{m}$ of distance is added between the fibres to widen the regions in (x, y) where injection occurs, enabling a coarser sampling to take place, decreasing the time taken. Here a raster scan is conducted, measuring the total flux on the camera at each point, shown in Figure 3.5 **a**. The flux is then fit as the sum of Gaussian functions, with centres lying on a finite hexagonal grid. The origin, separation and rotation of the grid, as well as the nuisance parameters of the mode field diameter and the individual amplitude are all fit using this data. There is no need to fit an offset term as all images taken are dark subtracted; images with no coupling have an expected total flux value of zero. The fibre stage then takes a step closer to the final distance and visits each port, sampling points around it and fitting the overall grid again. This is repeated until the fibres are close enough such that the width of the Gaussian around each peak is around $1/5$ the port separation, sufficiently rejecting the light in other modes. The system then validates the alignment by visiting each port, refining the positions and estimate of the grid, as well as adjusting the laser power used on the whole set to avoid saturation. The system then fine tunes the alignment using a local search. This involves sampling points drawn from a truncated Gaussian distribution centred on an initial guess, then fitting a Gaussian to the flux measured at each position (shown in Figure 3.5 **d-f**). This is done twice – first with a wider Gaussian distribution to get a coarse estimate, and then refined using both previous samples and more samples from a narrower distribution. The number of samples was tuned empirically, with robust behaviour at 20 coarse and 20 fine samples. The fit of a hexagonal grid is refined further using these measurements.

3.4 White light fringe finding

In order to capture the evolution of the mode transfer matrix with wavelength completely, the phase of the measured field must change by less than $\pm\pi$ between frames in order to unwrap this evolution correctly. The *white light fringe* refers to the case where this phase change is virtually zero over many wavelengths, such that interference fringes are visible even at large bandwidth.

The position of the white light fringe in the interferometer is found by changing the length of the injection arm using the delay line stage shown in Figure 3.1 **a** and a broadband (50 nm) SLD as the source. The coherence length – that is, the range of delay line positions that will have significant interference – is $45\ \mu\text{m}$. The setup captures frames at motor positions that offset are a fraction of this length and monitors the maximum off-axis power in the power spectrum after applying a smoothing Gaussian filter. This maximum value is reported as the “fringe power”. The values along the range of the motor are saved and the system selects the midpoint between the peaks of the fringe power in H and V polarisations as the location for the sweep, thus capturing high contrast fringes in both polarisations simultaneously. The optical path difference between the polarisations is smaller than 0.5 mm due to the compensating glass plate.

3.5 Digital off-axis holography reconstruction

After capturing data, the following step is to extract the coherent, off-axis component to reconstruct the complex electric field from the holographic data. An existing library (`DigHolo` [2]) implements this in an extremely computationally efficient way, however requires a selection of the modal basis (such as LP modes) which is then used in the optimisation of parameters. Also, the package ultimately represents all fields as superpositions of Hermite-Gaussian modes – modes arising from propagation in free space rather than from a circularly symmetric waveguide. Instead, the pipeline developed in this thesis is agnostic to the basis and seeks to reconstruct a field, only assuming that all fields in the set are supported on the same circle. This avoids truncation errors or large numbers of intermediate modes.

In order to reconstruct the field, the algorithm must first centre in the image and Fourier planes. The process for this is shown in Figure 3.6. For each polarisation and at each wavelength, the photometry data is co-added, the outline of the core found through Canny edge detection [41] and the centre estimated using a least squares circle estimator [42]. This uses geometric features and hence is more robust to asymmetries in the sum of the mode fields.

The chromaticity of the optics in the beam train must be considered when fitting these parameters. From refractive index tables, the calcite beam displacer alone causes a shift of around 2 pixels ($10\ \mu\text{m}$) over a 50 nm change. This is accounted for by fitting a degree 2 polynomial in λ to each of the x and y coordinates at each output polarisation, shown in panel **d** of Figure 3.6.

The above process also trivially gives the radius of the circular feature, which is then used to determine the size of the window before taking a Fourier transform.

The centre and radius must now be fit in Fourier space. The structure follows similarly to the image space, where the centre ($\theta(\lambda)$) and radius of the off-axis component of the power spectrum are found, and then fit this as a smooth function of wavelength.

To find the geometric centre with sub-pixel precision, the optimisation problem

$$\arg \min_{\mathbf{x}} (r(\mathbf{x})), \quad (3.1)$$

$$\text{s.t. } \sum \left(\mathbf{C}_{\mathbf{x}}(r(\mathbf{x})) \hat{\mathbf{P}} \right) = \rho, \quad (3.2)$$

is solved, where $\mathbf{C}_{\mathbf{x}}(r(\mathbf{x}))$ is a soft edged circle with centre $\mathbf{x} = (x_c, y_c)$ and radius $r(\mathbf{x})$, which is defined implicitly as the radius at which the circle encloses a fraction $\rho \in (0, 1)$ of the normalised power spectrum $\hat{\mathbf{P}}$. The power spectrum is also smoothed with a Gaussian filter beforehand to improve convergence. Intuitively, the above finds the smallest possible circle that encloses a significant fraction of the energy in the image, identifying the boundary the off-axis feature and the background. I use the ∂Lux [43] implementation of a soft edged circle and the solver in `Optimstix` [44] and `Jax` [45], solving a root finding problem within a minimisation problem. The fraction ρ is selected empirically to be 0.9 which gives a good balance of robustness to asymmetry in the circle and noise outside the circle.

For each frame taken with the shutter open (over all polarisations), the incoherent sum of the individual power spectra builds a single high signal to noise ratio (SNR) power spectrum, plotted on a logarithmic stretch in panel **e** of Figure 3.6. Next, noting that the off-axis component is half the radius of the central component [2], an off-axis centre is initialised. The scale of the central peak is given by the first point of inflection in the encircled energy

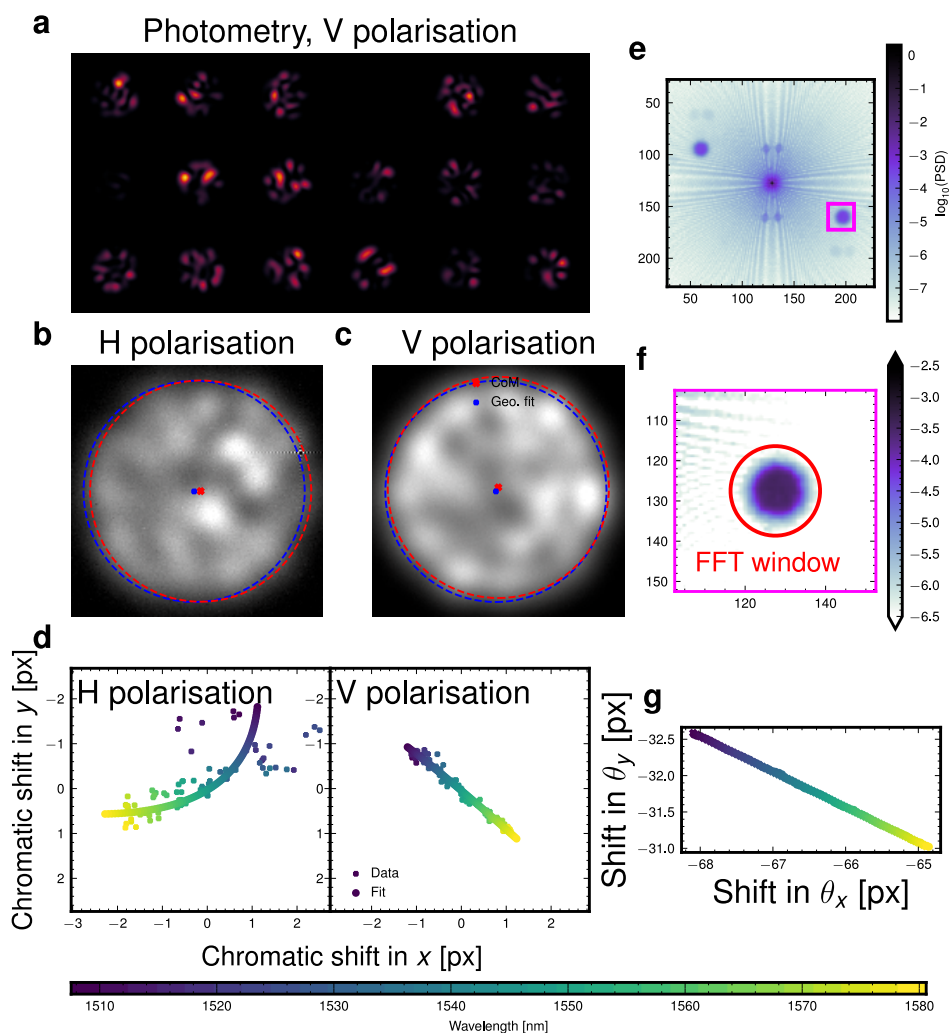


FIGURE 3.6: **Parameter fitting for digital off-axis holography.** All data in the first two rows is at a single wavelength. **a** The photometry for a single wavelength and polarisation for 18/19 ports (for compactness) shows a wide range of patterns. **b, c** When summed, the outline of the core as a soft edged circle becomes apparent. The geometric fit employed is more robust to bright regions of the co-added image than the standard centre of mass (CoM). **d** A low order polynomial is fit to the centres at each wavelength. **f** Co-added power spectrum, with off-axis component highlighted in magenta. **g** The centre is found at each wavelength. The size of the FFT window used in the reconstruction is shown for scale. **h** The linear fit to the centres in frequency space shows excellent agreement.

as a function of radius. This region is then masked and the resulting power spectrum is convolved with a Gaussian with the expected scale. A simple maximum then yields the centre of the off-axis component to within a few pixels, enabling us to apply the above

optimisation procedure in Equation 3.1 to fine tune the estimate. Again, this centre must shift with wavelength, however this is much more constrained than in the image plane: the motion should be in a line, with spacing in wavelength obeying

$$\boldsymbol{\theta}(\lambda) = \frac{\boldsymbol{\theta}_0 \lambda_0}{\lambda}, \quad (3.3)$$

where the zero subscript refers to some arbitrarily chosen zero point. The resulting line fit doesn't precisely go through the origin, and instead another centre is fit as an additional two parameters. This could be eliminated through the use of reflective instead of refractive optics to collimate the reference arm. The centres of these off-axis components are illustrated in panel **g** of Figure 3.6.

The reconstruction of the electric field \mathbf{E}_i from an image at a known wavelength is given by

$$\mathbf{E}_i = \mathbf{E}'_{\text{ref}} \mathcal{F}^{-1}(\mathbf{T}_{\boldsymbol{\theta}}(\mathbf{W}_{\text{Fourier}})(\mathcal{F}(\mathbf{W}_{\text{image}} \mathbf{I}_i))), \quad (3.4)$$

where \mathbf{E}'_{ref} is the estimated electric field of the reference beam, ' denotes the conjugate, \mathcal{F} denotes the Fourier transform (with inverse \mathcal{F}^{-1}), $\mathbf{T}_{\boldsymbol{\theta}}$ is a translation operator with parameter $\boldsymbol{\theta} = (\theta_u, \theta_v)$ (implemented using a phase ramp in Fourier space), $\mathbf{W}_{\text{image}}$ and $\mathbf{W}_{\text{Fourier}}$ are generalised Gaussian windows with width and shape parameters, and \mathbf{I}_i is the image from the holography setup of port i . The electric field of the reference \mathbf{E}'_{ref} is estimated with the phase determined by a Fourier transform of a delta function located at the peak of the power spectrum, and a uniform amplitude.

The filtering of only off-axis components means that low order aberrations and background effects manifesting near zero spatial frequency are removed, resulting strong robustness to these effects. Intuitively the reconstruction of off-axis holography data can be thought of as using the intensity to recover the amplitude and the curvature of the fringes to infer the derivative of the phase, allowing a complex quantity to be recovered.

3.6 Projection onto LP modes

After reconstructing the fields, the overlap integral η is calculated, given by

$$\eta_{i,j} = \langle \mathbf{E}_i, \mathbf{E}_{\text{mm},j}^b \rangle, \quad (3.5)$$

recalling that the inner product is defined in Equation 2.1, i is the index of the lantern port used and again $\mathbf{E}_{\text{mm},j}^b$ is the j -th basis mode supported at the multimode end. The integrals are approximated as the sum of the relevant quantities over a grid, with mode profiles found using the manufactured parameters and the `ofiber` package [46]. The physical fibre diameter used to generate the LP modes is measured using a microscope beforehand, and the refractive index contrast is known from manufacturing specifications as 5.5×10^{-3} . The plate scale of the imaging setup is fit such that the maximum (power weighted) overlap value is achieved.

An important metric I monitor is the “explainability” of the field in the modal basis projected into, defined as $\sqrt{\sum_j |\eta_{i,j}|}$. This represents the fraction (in an amplitude sense) of the field that is consistent with basis modes alone, with values close to unity indicating a good match. The square of this value is the total power explained by the underlying basis set.

This thesis only employs LP modes, however it is worth noting that this is only an approximation. Due to structures from the multicore fibre remaining in the multimode end, the refractive index is not a simple top hat function. Future work could directly apply basis modes that are solved numerically to better reflect the structure.

3.7 Simulation

This paragraph does not cover my work during this thesis, however it is included for completeness.

The results for the simulated device in the results chapter are generated using `Rsoft Beamprop`. A model of the 19-core PL was created using a linear taper, with core size

and separation, taper ratio and length, and refractive indexes based on known values for the physical device. Light was launched into single-mode ports one at a time and propagated to the multimode end. The resulting mode field was decomposed into LP modes via an overlap integral and their complex coefficients recorded. This was repeated for 10 wavelengths spread evenly between 1507.5 nm and 1580.5 nm.

Results

Sections 4.1, 4.2, and 4.5 are common with the submitted manuscript.

4.1 Validation on a single-mode fibre

The performance of the system is first evaluated on a device with a known mode field: a SMF. This is shown in Figure 4.1. The electric fields reconstructed appear near uniform in phase and Gaussian in amplitude, as expected. The dominant term in the phase residual is a tilt with a small amount of defocus. For the SMF this could be optimised further, however this is a realistic impression on how finely the focus and angular alignment of a photonic lantern can be adjusted (where the field is not known to be flat in phase). The phase as a function of wavelength is consistent with a common differential air path of -165.3 mm and a glass path of 141.4 mm or 144.8 mm for H and V respectively, which is expected due to the compensating plate.

There are, however, oscillations in both the phase and amplitude as a function of wavelength. This is consistent with the fringe power in delay space, where a secondary peak of relative height $A_r = 0.25$ at a distance of $d = 192 \mu\text{m}$ (as is seen in V, using the two peaks around 29.5 mm in Figure 4.1 **d**). This secondary peak contributes another phasor to the measured signal, which rotates faster in the complex plane than the phasor from the main peak due to the different optical path difference. Hence, the coherent field from the secondary peak induces oscillations in amplitude, with a factor $A_r = 0.25$ and in phase, with size $\arctan(A_r) \approx 0.25 \text{ rad}$. Both have a period given by $\lambda^2/d = 11.25 \text{ nm}$, corresponding to between 6 and 7 oscillations over the 75 nm bandpass. This is all very consistent with the data

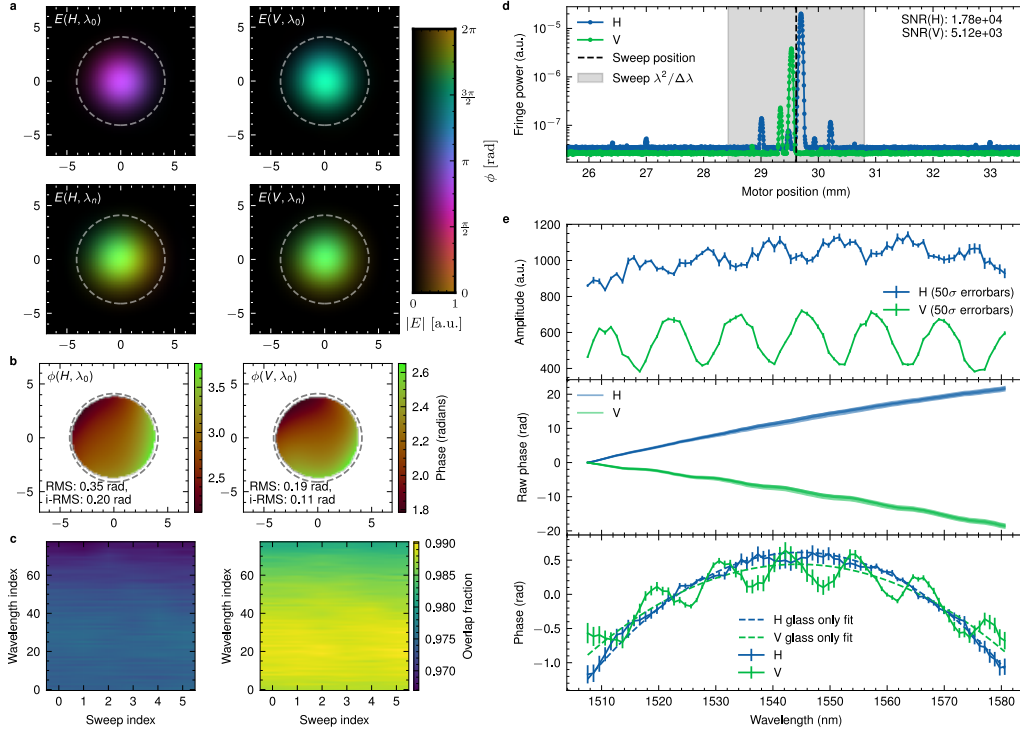


FIGURE 4.1: Reconstruction validation on a single-mode fibre, with monochromatic performance (left) and wavelength-coherent reconstruction (right). **a** Reconstructed electric field for two polarisations at two different wavelengths show phase uniformity and a Gaussian structure. All image plane axes units are μm . **b** The phase over the face of the fibre is flat to within 0.20 rad rms, weighted by intensity. **c** The overlap integral explainability shows that the reconstructed field is at worst 97% explained by a single LP01 mode, as expected. **d** The setup also probes the position of the white light fringe. The grey shaded region shows the effective coherence length during the laser sweep, and the dashed black line shows the position of the motor during the sweep. The fringe visibility using a broadband source shows clear peaks and validates the phase unwrapping used in the following step. **e** Complex LP01 modal coefficients as a function of wavelength. The average amplitude over six sweeps is very consistent – error bars show 50 \times the standard deviation. Raw phases (middle row) of the reconstructed mode are dominated by linear trends induced from path difference between polarisations. Once removed, the dispersion properties are consistent with a common differential air path of -165.3 mm and a glass path of 141.4 mm or 144.8 mm for H and V respectively, plus oscillations consistent with birefringence as evident in **d** (see main text for details).

shown in the top and bottom panels of Figure 4.1 **e**. These oscillations are similarly present in H polarisation, however are much smaller in amplitude as the main peak has higher power.

The peaks arise from birefringence, likely from polarisation maintaining fibres that are misaligned with the electric field polarisation direction. Further analysis of the oscillations and the impact on photonic lantern characterisation is conducted in section 4.4.

4.2 Transfer matrix of a 19-port lantern

The 19-port photonic lantern characterised in this thesis has a multimode diameter of $34.7 \mu\text{m}$, a port refractive index of 1.44 and a cladding index with a contrast of 5×10^{-3} . Over the wavelengths swept by the characterisation system, the output supports 23 LP modes below $\lambda_c = 1562 \text{ nm}$ and 21 LP modes above λ_c . Hence, this is a marginally undersampled lantern, manufactured in order to guarantee that all the modes of interest are strongly bound over the whole operational wavelength range.

The characterisation setup takes data for this device in just over one hour, with time primarily spent on finding good injection and the position of the white light fringe. The data capture with the swept wavelength source takes less than 10 minutes. This feature is important in enabling a tight feedback loop with the manufacture of devices.

Figure 4.2 shows the steps of the reconstruction. The first two rows are dark subtracted data, taken without and with the reference beam respectively. Despite each port having the same input beam (including polarisation), after propagation through the device the polarisation changes in a port dependent way. In the cases where a port appears dim in H (such as port 3), there is good signal in V. The following rows illustrate the reconstruction. The intensity of the reconstructed field $|E|^2$ is in good agreement with the photometry (mode fields with no reference beam). The reconstructed fields show smoothness, phase diversity and are well contained by the core outline, as expected. The fields are also well explained by LP modes, with the decomposition showing visual similarity to the field itself. This decomposition is reasonable but imperfect, suggesting future work in using higher fidelity modes. Finally, I also visualise the fields on an Argand diagram, illustrating a unique structure for each port. Importantly, the reconstruction is of reasonable quality even in low SNR – with the resulting field still well explained by LP modes. This is due to the robustness of digital off-axis

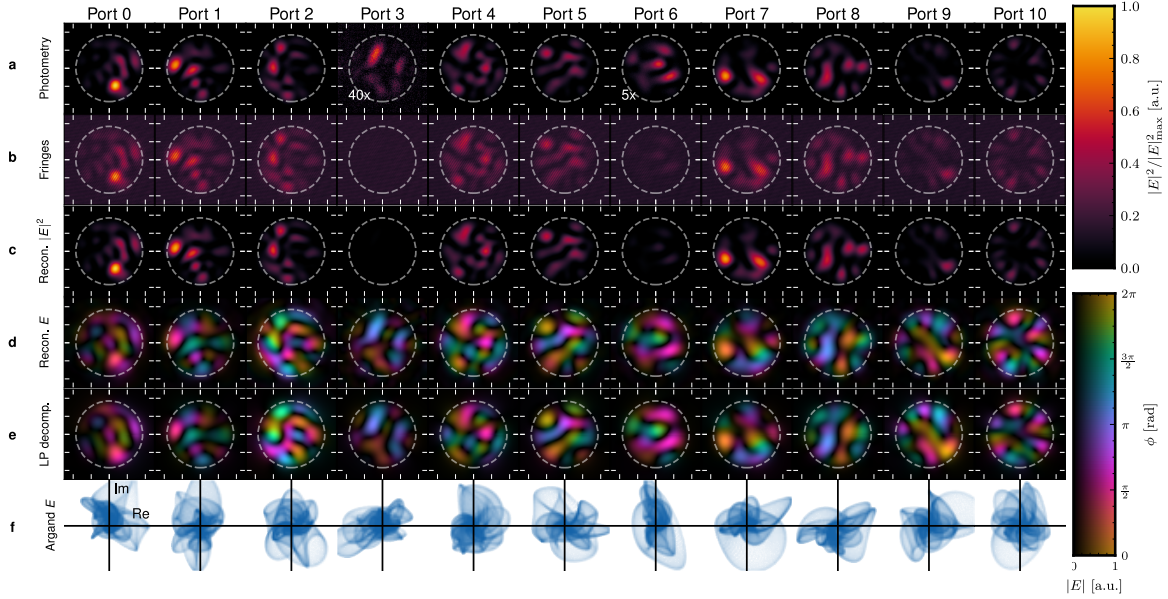


FIGURE 4.2: Data and reconstruction for the first 11 ports probed (out of 19) in a single wavelength and polarisation. Two sets of data are captured: **a** photometry and **b** fringes. Plots are all shown on the same stretch such that the maximum of the row is common and the floor is zero counts. Data shown is at $\lambda = 1513.9$ nm and in H polarisation, so some ports appear dim as power is concentrated in V. Distance between white ticks corresponds to $10 \mu\text{m}$ and the white circle has the diameter of the multimode end as measured by a microscope. **c** The digital off-axis holography reconstruction has an intensity that is similar to the photometry data. **d** The reconstructed electric field is well recovered, even at low counts such as port 3. **e** The decomposition into LP modes matches the reconstruction well and are a good first order approximation. **f** Visualisation of the modes in an Argand diagram, where each pixel is counted in a 2D histogram. While some electric fields might appear qualitatively similar in image space, this alternative visualisation further highlights the diversity of the modes supported.

holography where only the coherent information is extracted, and hence contamination such as background and read noise are less noticeable.

Next, the coherence properties of the lantern are illustrated in Figure 4.3. Each row is formed from taking a single port from the outer edges of the hexagonal grid, plus the central core. This is visualised in the bottom row, where the circle edge colour matches the axis colour. The fringe power metric finds the white light fringe, located at positions ii and iii for V and H respectively. All white light fringes are found with a SNR greater than 10^3 using a median and median absolute deviation. This metric is selected to avoid bias from secondary

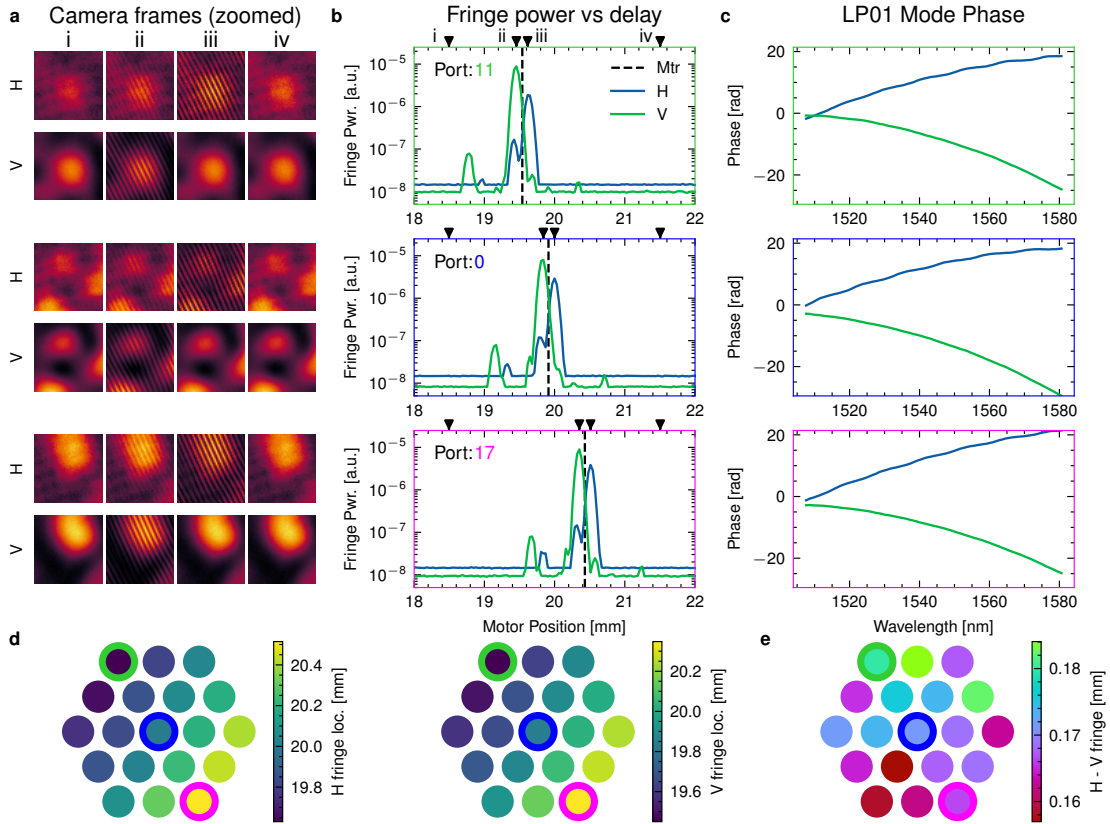


FIGURE 4.3: Characterisation of the broadband coherence properties of the photonic lantern. **a** Images of both polarisations in broadband light for different motor positions for 3 different ports in both polarisations. **b** The fringe power as a function of motor position identifies the white light fringe. Arrows above plots indicate the positions of the frames shown on the left, with the maximum fringe power position shown in ii for V polarisation and iii for H polarisation. The final motor position for data capture is indicated with a black dashed line. **c** The reconstructed phase of the LP01 mode as a function of wavelength. This is consistent with the broadband results, with oscillations more apparent in the case of stronger secondary peaks in fringe power. **d** Measuring the peak position over all ports, a trend emerges over the fringe location consistent with a bend dependent change in the refractive index of the multicore fibre that feeds the photonic lantern (see text for details). Coloured edges around circles correspond to the ports highlighted in other rows.

peaks. The picture in this delay space demonstrates a clear shift of about $750\ \mu\text{m}$ in the position of the white light fringe along a direction on the multicore fibre. This is much larger than any tip/tilt induced effects at the fibre stage, where even a 2° misalignment would only induce a difference of $10\ \mu\text{m}$. Such a large effect is likely due to a bend effect in the 1 m multicore fibre rather than the 7 cm photonic lantern itself. Again, secondary peaks are

present in delay space, however, the peaks are at different relative heights between ports. As the peaks depend on the port, their effect cannot simply be calibrated out relative to a single-mode fibre. Qualitatively the same behaviour of oscillations is observed, where the polarisation with a relatively larger secondary peak (H in this case, was V previously with the SMF) has more visible oscillations in the phase of the modes when using the swept wavelength source. The difference in the peak positions for each polarisation is also shown, using a cubic spline interpolation to sub-sample the motor steps ($32 \mu\text{m}$) and estimate the peak location. Differences in white light position between polarisations are weakly present, with a maximum difference of $27 \mu\text{m}$ and no dominant direction.

Figure 4.4 shows the breadth of wavelength dependent data available from the system. First, I show the electric field for one polarisation for different ports. All ports change at approximately the same rate, with neighbouring panels appearing identical up to an overall phase, while distant panels are clearly different. The overall phase evolution is generated from being offset from the white light fringe, recalling that the system took the midpoint between the two polarisations and hence each are $\sim 85 \mu\text{m}$ from either white light fringe.

Next, I quantify the wavelength scale of field evolution in the device by measuring the overlap integral of the field at a particular wavelength (black dashed line in Figure 4.4 **b**) compared to all other wavelengths. The characteristic scale for the mode field to change noticeably is of order 20 nm , broadly consistent with previous results in wavefront sensing [16]. The falloff is near linear around 20 nm from the fixed wavelength. Both the simulation and measurement agree in the functional form (smooth turning point with linear decay) but the rate of decay can be slower or faster in the actual device, depending on the port. Finally, the rate of decay in this linear region is visualised on the port map of the device. Generally, ports towards the centre of the MCF decay faster, however the device clearly does not have the symmetry of the idealised simulation. The difference in wavelength needed for principal modes to become orthogonal could be as little as $100/1.8 = 56 \text{ nm}$ or as much as $100/0.6 = 167 \text{ nm}$.

Figure 4.5 presents the throughput and LP mode explainability per port and wavelength. There is no significant trend in the total throughput, but it is included for completeness. The

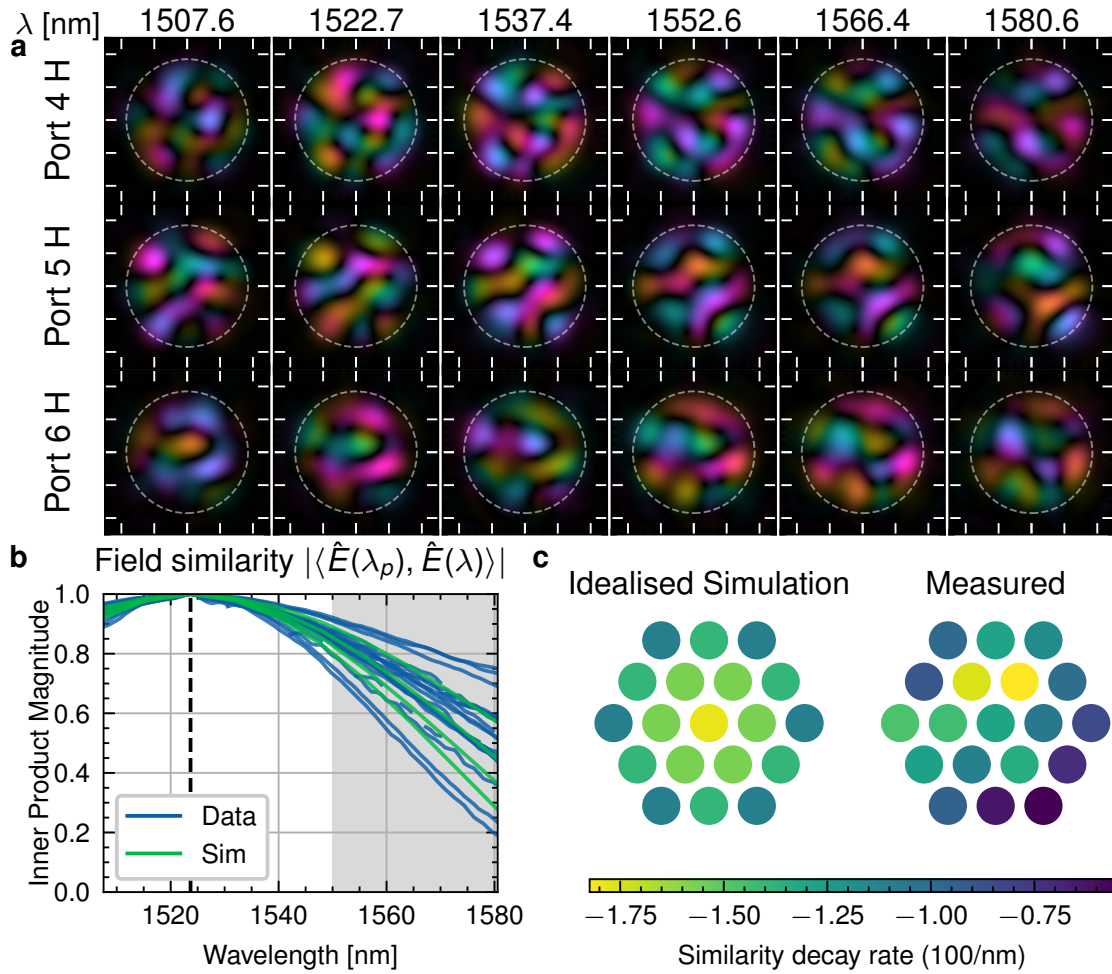


FIGURE 4.4: **Wavelength evolution of modes in a photonic lantern.** **a** Electric field plots over ports 4-6 over the full bandwidth. **b** Field similarity between a fixed wavelength λ_p (black dashed line) and all other wavelengths within the same port, for data and simulation. Each curve is one port. The falloff becomes near linear after 20 nm (grey box). **c** Gradient of similarity decay for simulated and measured fields over all ports.

LP mode explainability shows a high agreement, with median value of 0.937 and 0.949 for H and V over all ports and wavelengths. Lower values seen in oscillations for port 7 in H are indicative of low flux. Visualising this in the hexagonal grid of the multicore fibre (bottom row) reveals a correlation between ports, and a weak gradient where the ports that are physically lower have higher LP mode overlap. The drop in explainability at cutoff wavelength (where the model only supports 21 modes) indicates that the physical device still has some power in these modes.

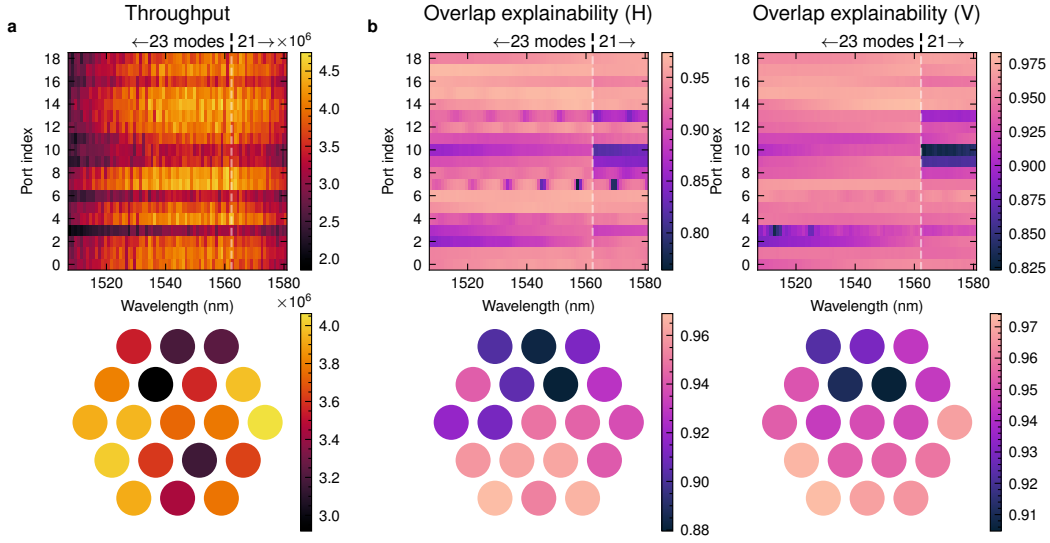


FIGURE 4.5: **Relative throughput and LP mode explainability of fields.** **a** Throughput for each port and wavelength, measured by photometry. The median value is illustrated on a hexagonal grid. There is no obvious structure in either case. **b** The overlap integral $\eta_{i,j}$ over all ports and wavelengths, and plotted on a hexagonal grid. The white dashed line indicates the wavelength at which a multimode fibre with a diameter equal to that of the device begins to guide 21 instead of 23 modes. For some fields there is a noticeable drop in the explainability of the field as a superposition of LP modes, however for others this change has little to no effect.

Figure 4.6 summarises the axes available in the post processed transfer matrices. Slices through wavelength (Figure 4.6 **b**) show the evolution of modes in amplitude and phase, which qualitatively are consistent with polynomials with common mode oscillations. Relative measurements between the phase are very precise and explored further in Figure 4.12. Also, the quadratic term is stronger in the photonic lantern than the SMF as about 116 mm more compensating fibre in the injection arm was required to find the white light fringe, despite the length of the photonic lantern with multicore fibre being close to that of the SMF used in validation. Slices at a particular wavelength (Figure 4.6 **c**) yield monochromatic transfer matrices in a given polarisation. Previous work has typically used simulations to generate these, and hence shows strong symmetry between ports on opposite sides of the multicore fibre grid. This is not observed, highlighting the significance of manufacturing variations from a target design. The proposed characterisation system would be capable of

discerning if this variation is a fixed or variable offset from the design target given multiple instances of devices, and this is left as future work.

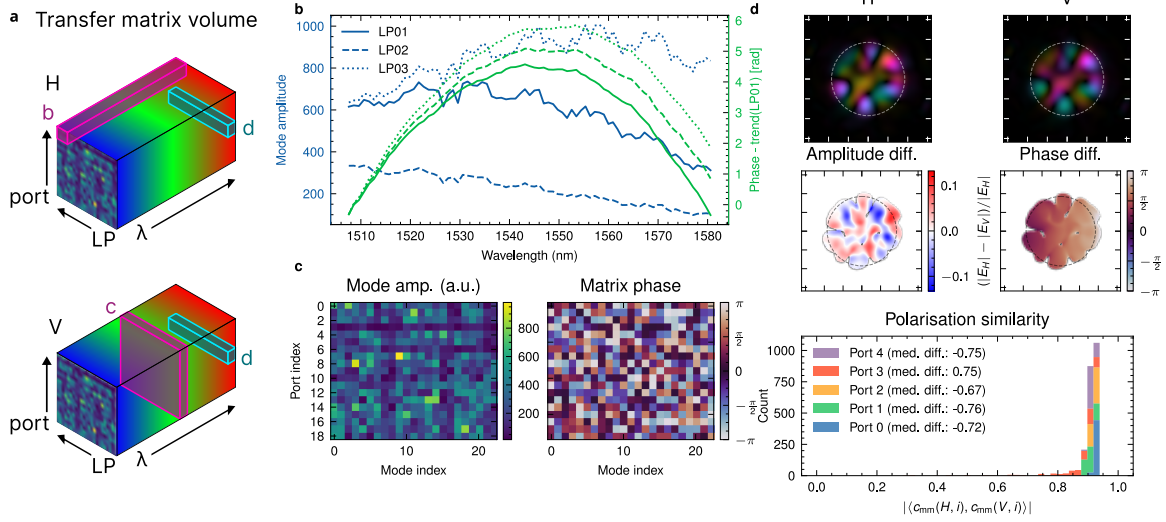


FIGURE 4.6: **Slices through the transfer matrices of the 19 port lantern.** **a** The characterisation setup produces transfer matrix cubes, with coherent properties across wavelength. This figure explores three slices through the cube. **b** Mode evolution along wavelength. **c** Monochromatic transfer matrices. **d** Slices of the same port and wavelength in different polarisations. The fields appear visually similar (top row), however subtle differences exist (middle row). Over ports and wavelengths the fields of different polarisations (stacked histogram, bottom row) exhibit a dissimilarity that is not explained by a difference in signal alone (see text for details).

Furthermore, slices in polarisation can be compared, visualised in Figure 4.6 **d**. The electric fields from each polarisation appear identical up to a spatially constant phase offset. Taking differences in amplitude, however, changes by up to 10% are present in some regions, with a structure that is not consistent with sub-pixel centring residuals alone. In phase there is a near uniform difference, with a weak tilt likely caused by a deviation in angle from the beam displacer. In the bottom row, a similarity statistic is presented over a few ports and all wavelengths in a stacked histogram. I normalise the LP coefficients for a given port and wavelength, and then compute the magnitude of the inner product between the two polarisations. For identical fields with a spatially constant phase offset, this metric is unity. Since the quality of the LP overlap is related to the flux, I also report the median flux difference between H and V, i.e. the median of $(f_H - f_V)/(f_H + f_V)$ over all wavelengths. Values closer

of ± 1 indicate less flux, and hence lower value inner product. A port independent bound of around 0.9 is observed, consistent with the phase tilt and amplitude differences shown above. There is, however, port dependence that indicates a weak difference in the electric fields of different polarisations. This is most evident between ports 0 and 2, from which a LP quality argument would favour higher similarity in port 2 but the data suggests port 0 has higher similarity.

4.3 Orthogonality of photonic lantern principal modes

One important property of mode sorting devices is the orthogonality of the supported modes. In applications where the modal mapping is inverted to make an inference of the electric field at the multimode end, mode sorters with orthogonal modes will provide better conditioned solvers. Direct measurement of the orthogonality is shown in Figure 4.7. For each combination of normalised electric fields from two ports, the overlap integral $\left| \langle \hat{E}_i, \hat{E}_j \rangle \right|$ is computed, which corresponds to the cosine of the angle between the two complex vectors that represent each field. The majority of cases have values close to zero, indicating orthogonality. Figure 4.7 c illustrates that the less orthogonal principal modes tend to be those towards the outside of the MCF (higher port index) and tend to be from neighbouring ports. Statistics are also taken over all wavelengths. To demonstrate that this is not an effect of the dimension of the space (noting that random vectors in large directions are nearly orthogonal) I compute the distribution for randomly draw normalised vectors in 23 dimensions (equal to the number of LP modes that are supported), which has a much larger value. An adiabatic photonic lantern must have orthogonal principal modes, and hence the result is consistent with a near-adiabatic device, corroborating the high throughput measurements made after manufacture. A relatively simple summary of a device's orthogonality is proposed: the median and 5-95 percentile range of the overlap integral over all different port combinations and all wavelengths. For this device, these values are 3.3% and 10% respectively.

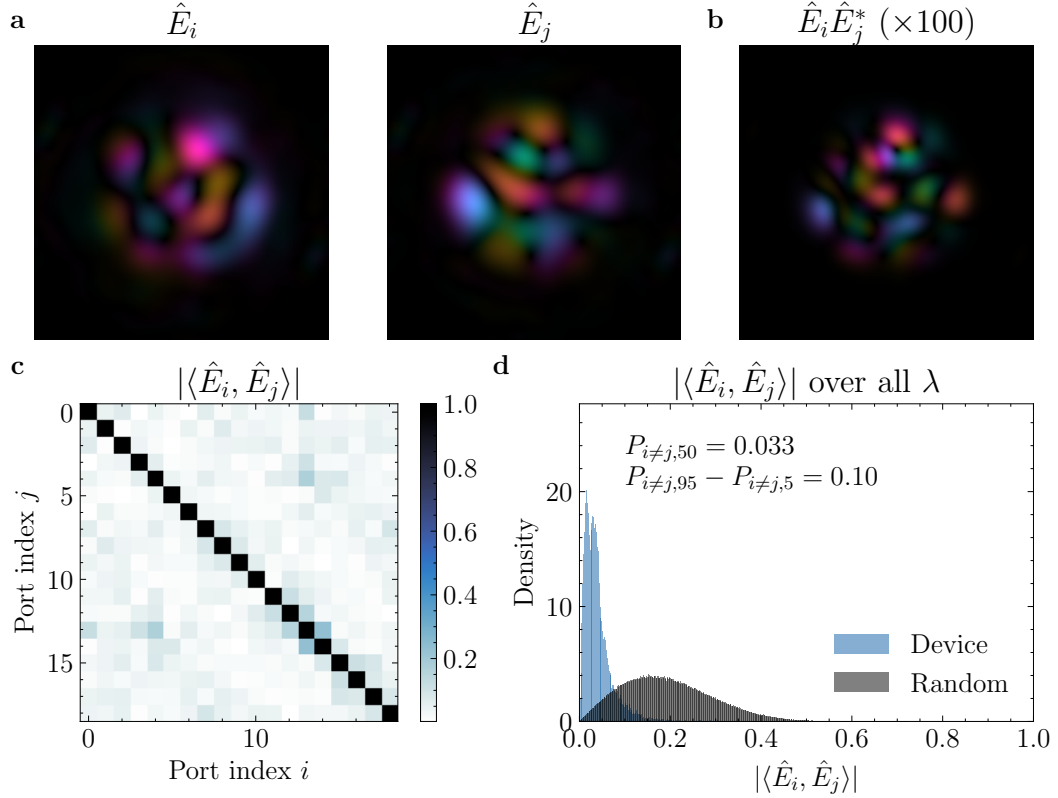


FIGURE 4.7: **Orthogonality of photonic lantern principal modes.** **a** Normalised electric fields for two different cores at the same wavelength and polarisation. **b** The complex product produces a result that has a mix of values and phases. The inner product is the complex sum of this image. **c** The inner product for all combinations of ports measures the orthogonality of the modes. **d** The distribution from repeating this over all wavelengths measures the orthogonality of the photonic lantern principal modes. The distribution is summarised with a median P_{50} and 5-95 percentile range $P_{95} - P_5$ of overlap integrals that are taken over all combinations of different ports ($i \neq j$). For comparison, a distribution from random normalised vectors (black) is also shown.

4.4 Fitting transfer matrices in wavelength

In the interpretation of Figure 4.1, the oscillations of the amplitude and phase of the LP01 mode fit was found to be qualitatively consistent with the sum of a discrete number of phasors. These could arise from ghost reflections or birefringence in the system outside of the device under test. In this section, I quantify this further, extending the interpretation to the

photonic lantern case, discussing different ways of fitting this behaviour and ultimately finding limitations in the system. Having understood these limitations, a robust observable – the differential mode phase – is identified and explored in the following section.

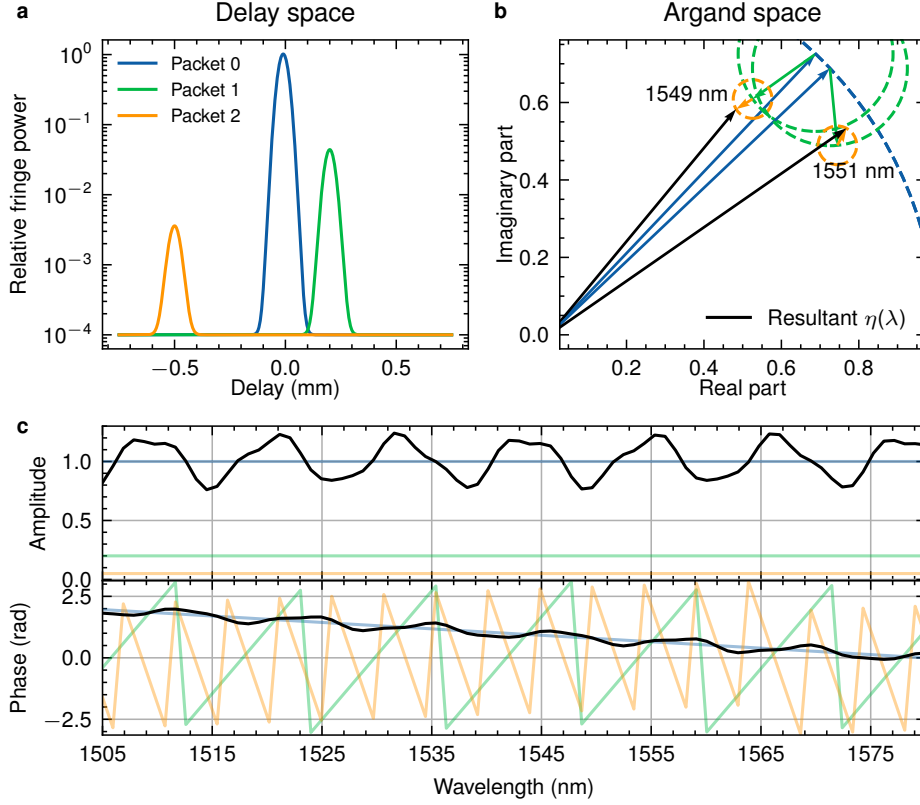


FIGURE 4.8: Simulated effect of different coherence packets on wavelength evolution. **a** Toy model of three coherence packets in delay space (plot analogous to Figure 4.3 b). Each packet of coherence has a different height and distance compared to the (blue) fundamental packet. **b** Sums of phasors in Argand space. The resultant vector is the only value observed by the characterisation system and represents the overlap integral at that wavelength $\eta(\lambda)$. The size of each phasor is the square root of the relative amplitude in fringe power, and the rotation rate depends on the distance from zero delay. **c** Simulated results of $\eta(\lambda)$ from the toy model over the wavelengths covered in this work.

Figure 4.8 outlines how a toy model of a few coherence packets at different values in delay space can lead to non-linear, oscillatory behaviour of the overlap integral as a function of wavelength $\eta(\lambda)$. The digital off-axis holography measurement is sensitive to all coherent components, and hence retrieves a complex sum of phasors from all nearby delays. Each phasor rotates at an angular rate d/λ^2 , where d is the distance from zero delay. The amplitude

of the phasor is given by the square root of the value of the peak in delay space, hence just a 1% reflection in intensity becomes a 10% variation in the resulting value. This is observed in Figure 4.8 **b**, where the phasor from packet 2 rotates around π radians over 2 nm of wavelength change, but the phasor from packet 1 (which is also longer) rotates less. Figure 4.8 **c** illustrates this behaviour over a larger bandwidth. Even though each amplitude is constant and phase is linear, the addition in complex space yields non-linear, oscillating functions of amplitude and phase as a function of wavelength.

Observe, however, that the somewhat complex evolution of $\eta(\lambda)$ becomes predictable given only a few additional parameters: the relative amplitude of the packets (dictating the size of the oscillation), their distance from zero delay (dictating the period of oscillation) and the phase of the phasor at a single wavelength. In this section I discuss algorithms to fit this behaviour. I start with the simple case of fitting this behaviour for a single port, mode and polarisation of the photonic lantern before generalising to a joint solution over all dimensions of the transfer volume measured in this work.

I separate the single packet evolution of the mode $\eta_{\text{fund}}(\lambda)$ from the effective coefficient that the coherence packets provide $C_{\text{pack}}(\lambda)$, resulting in the expansion

$$\eta(\lambda) = \eta_{\text{fund}}(\lambda)C_{\text{pack}}(\lambda), \quad (4.1)$$

$$\eta_{\text{fund}}(\lambda) = A(\lambda) \exp(i(a_2(\lambda - \lambda_c)^2 + \dots + a_{q_\phi}(\lambda - \lambda_c)^{q_\phi})), \quad (4.2)$$

$$C_{\text{coh}}(\lambda) = \sum_{k=0}^{n_{\text{coh}}} A_k \exp(i(m_k(\lambda - \lambda_c) + \phi_k)), \quad (4.3)$$

where $A(\lambda)$ is the single packet amplitude evolution and a_i are the $q_\phi = 3$ degree polynomial coefficients for the single packet evolution in phase taken as a Taylor series around a central wavelength λ_c , truncated to not be degenerate with the linear terms in the fundamental ($k = 0$) packet. The amplitude values are free to vary arbitrarily at each wavelength, though this will be restricted at a later step. The coherence packet contribution holds the three parameters mentioned previously: the relative amplitude A_k , the rate of phase change with wavelength m_k and the phase offset at λ_c as ϕ_k . I define the fundamental packet by convention with $A_0 = 1$. To determine the significance of the additional packets, it is also useful to model a

minimal single packet case, where $n_{\text{coh}} = 0$ (and hence the sum is only of a single term so $C_{\text{coh}} \equiv 1$ for all wavelengths).

The next step is to test how well the above model fits the data. The goal is to minimise the discrepancy between the data η_j and the model $\eta(\lambda_j)$ over all wavelength indices j , and hence the mean squared error is used as a loss function. There is a computational trick significantly speeds up the convergence for this problem. Notice that, when all other parameters are fixed, the problem of finding the $A(\lambda)$ values is a classical least squares problem. Explicitly, this is

$$\min_x \|\tilde{A}x - b\|^2, \quad (4.4)$$

with

$$\tilde{A} = \begin{pmatrix} A_T \cos(\Phi_T) \\ A_T \sin(\Phi_T) \end{pmatrix}, \quad (4.5)$$

$$x = \begin{pmatrix} A(\lambda) \end{pmatrix}, \quad (4.6)$$

$$b = \begin{pmatrix} \mathcal{R}(\eta(\lambda)) \\ \mathcal{I}(\eta(\lambda)) \end{pmatrix}, \quad (4.7)$$

where A_T , Φ_T are the amplitude and phase from all factors other than $A(\lambda)$ in Equation 4.1. This trick is leveraged by using a hybrid optimisation algorithm, where at each step other parameters are adjusted by gradient descent and then all amplitude coefficients are obtained using a simple least squares solve. The optimiser used for the gradient descent step is adam [47].

Figure 4.9 summarises the results of the fit for a single port, LP mode and polarisation. As this simple model is very expressive (in particular with no restrictions on the amplitude at every wavelength), it is compared to an almost equally expressive model that has only the fundamental coherence packet but nonetheless a comparable number of parameters. Without looking at the residuals, both appear to fit the data very well. Upon inspection of the residuals, however, the benefit to the fit quality from the few additional parameters is significant, with the mean square error differing by almost an order of magnitude. Figure 4.9 c illustrates the model parameters at convergence. The model that uses 3 coherence packets converges

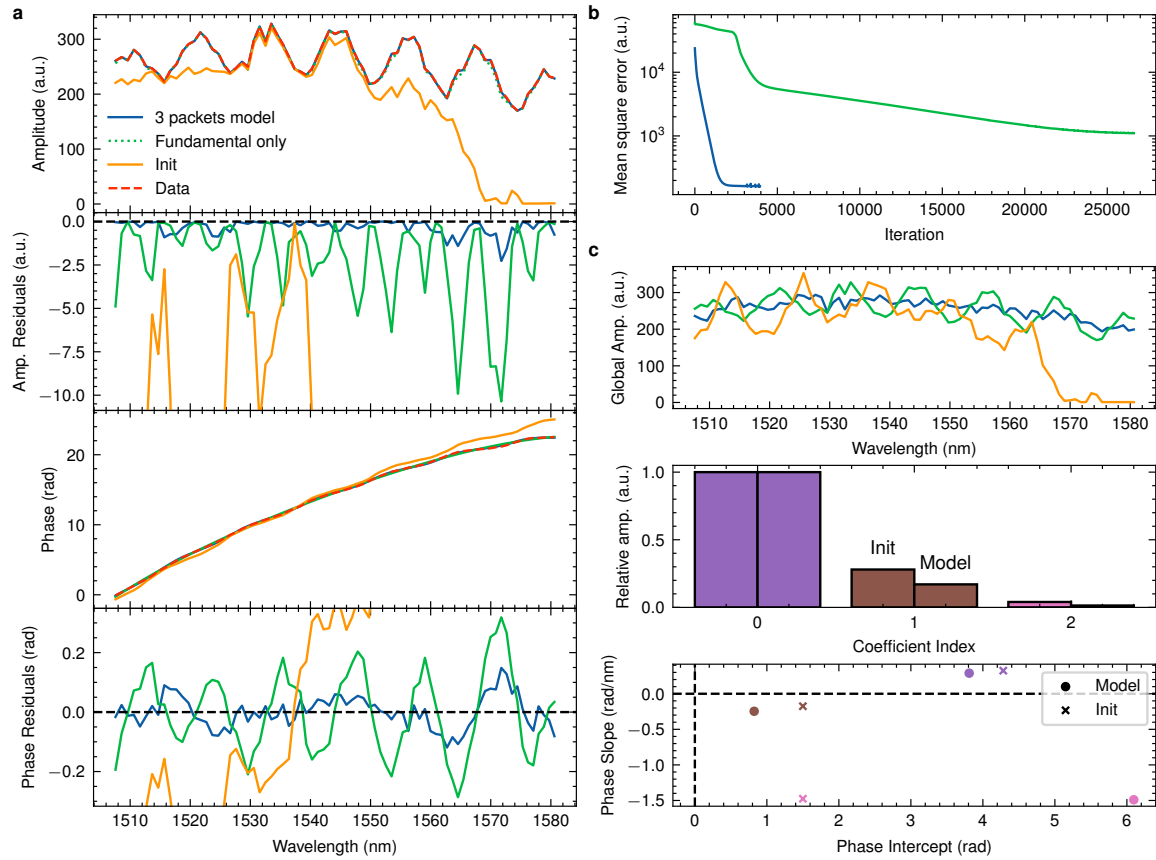


FIGURE 4.9: Multiple coherence packets improves the fit of the overlap integral with wavelength. **a** Converged model with three packets (blue), converged model with only the fundamental packet (green), initialisation (orange) and data (red) in amplitude and phase, with residuals. The initialisation is done by eye using the broadband fringe locations to estimate the amplitude and phase slope of each coherence packet. The converged model shows good agreement with the data, and the residuals appear more stochastic than the oscillations seen for the no packet model. The amplitude is always underestimated as this is the solution to the least squares problem Equation 4.4 when there is a mismatch in phase. **b** Mean square error loss during optimisation. The model first fits the easy parameters like the phase offset of all the data, and then slowly converges on the other parameters. The model with multiple coherence packets is almost an order of magnitude better than the model with only the fundamental packet. **c** Plots of the model parameters. The global amplitude (top) converges to a smooth function, indicating that the coherence packet model explains the oscillations seen in the data well.

on a smoother function of $A(\lambda)$ than the model with only the fundamental (which absorbs some of the oscillation with changes in amplitude). The residual jitter is not stochastic – multiple sweeps show the same effect at each wavelength. This could be caused by residual

internal fringing in the camera and/or additional coherence packets that are aliased at the wavelengths sampled and hence difficult to predict. Nonetheless, this result indicates that the fundamental evolution of overlap integrals with LP modes in photonic lanterns is well approximated by slowly evolving polynomials.

I now turn to the harder problem – fitting the device transfer function $\eta(c, \lambda, P, s, m) \in \mathbb{C}$, where c is the port number, λ is the wavelength, P is the polarisation measured on the camera (H or V), s is the sweep index (a clearly nuisance parameter, but the testbed instabilities causes drifts between sweeps), and m the mode index e.g. LP01.

Again, the natural decomposition of the problem is into one of amplitude and phase functions, where both are slowly varying. As above, multiple coherent packets (n_{coh} in total) are modelled that add in \mathbb{C} to produce the observed value. Hence the decomposition must be of the form

$$\eta(c, \lambda, P, s, m) = A(c, \lambda, P, m) \exp(i\phi_{\text{ti}}(c, t(\lambda), s)) \quad (4.8)$$

$$\left(\sum_{k=0}^{n_{\text{coh}}} A_r(c, k) \exp(i\phi(k, c, \lambda, P, m)) \right), \quad (4.9)$$

where $A_r(c, k) \in [0, 1]$ is a wavelength independent constant that signifies the relative strength of the coherence packet, again with $A_0 = 1$ for all c, k . ϕ_{ti} is a term that represents the drift due to testbed instabilities, and hence depends on the time (which is almost, but not quite, linear with respect to λ due to the sweeping function of the laser).

Polynomials are a natural choice for each term, which are expected to be smoothly varying. The goodness of fit in the results will validate this assumption. For amplitude, this is a simple Taylor expansion around a central wavelength λ_c

$$A(c, \lambda, P, m) = a_{q_A}(c, P, m)(\lambda - \lambda_c)^{q_A} + \dots + a_0(c, P, m), \quad (4.10)$$

for some degree q_A .

Each coherent packet has a different gradient and offset term in phase, but higher order terms between coherence packets should be the same (as the dominant dispersion is from the

common glass/air difference and also passing through the same device). Hence the natural phase decomposition is into two terms, with

$$\phi(k, c, \lambda, P, m) = \phi_l(k, c, \lambda, P, m) + \phi_h(c, \lambda, P, m), \quad (4.11)$$

$$\phi_l(k, c, \lambda, P, m) = b_1(k, c, P, m)(\lambda - \lambda_c) + b_0(k, c, P, m) \quad (4.12)$$

$$\phi_h(c, \lambda, P, m) = b_{q_\phi}(c, P, m)(\lambda - \lambda_c)^{q_\phi} + \dots + b_2(c, P, m)(\lambda - \lambda_c)^2, \quad (4.13)$$

where the subscripts l and h denote the lower (affine) and higher order contributions respectively, and q_ϕ is the degree of the polynomial. The higher order phase ϕ_h can be taken out of the sum over the coherence packets in Equation 4.8 as it doesn't depend on the packet index k .

Finally, the testbed instabilities are common to all modes as they are dominant in the arms of the interferometer, so they cannot depend on the m (mode index) term. To first order, this term is assumed to be linear in time, with

$$\phi_{ti}(c, t(\lambda), s) = d_1(c, s)t(\lambda) + d_0(c, s), \quad (4.14)$$

where the average value of d_1 is expected to be close to zero and $d_0(c, 0) = 0$ is set to fix the absolute phase degree of freedom.

This solver reduces the data from complex numbers with dimensions $(n_{\text{ports}}, n_\lambda, n_P, n_S, w, h)$ – total size on the order of 10^9 – into just a handful of real valued parameters:

- $n_{\text{ports}}(n_{\text{coh}} - 1)$ relative amplitudes A_r ,
- $q_A n_{\text{ports}} n_P n_m$ global (in k) amplitude polynomial coefficients a_q ,
- $2n_{\text{coh}} n_{\text{ports}} n_P n_m$ low order phase coefficients b_1 or b_0 ,
- $(q_\phi - 2)n_{\text{ports}} n_P n_m$ high order phase coefficients b_p, \dots, b_2 ,
- $2n_{\text{ports}} n_s$ sweep testbed instabilities coefficients d_1 and d_0 .

This gives of order 10^4 parameters for a total compression factor of about 10^{-5} . The majority of this reduction in dimensionality is delivered by representing wh pixels with n_m modal coefficients, and then compression over λ to a few polynomial coefficients. This fitting also

better constrains the absolute evolution, as testbed instabilities are removed in a better way than averaging alone.

The model is initialised using different measurements. First, the fundamental packet amplitude and phase is fit directly as a polynomial. Next, coherence packets have A_r and b_1 set using peaks found in the delay space data, following the same calculation as discussed with Figure 4.1. Finally, the phase of each packet is adjusted individually to minimise the mean squared error.

Table 4.1 summarises the `adam` [47] optimiser configuration used. Different learning rates are selected for different parameters to guide convergence, as each have different scales.

TABLE 4.1: Optimisation parameters for global modal transfer matrix fitting.

Term	Description	Learning rate	Units
A_r	Relative amplitude	1×10^{-4}	-
b_0	Phase offset	5×10^{-1}	rad
b_1	Phase linear term	2×10^{-3}	rad/nm
b_2, \dots, b_i	Phase higher order terms	1×10^{-5}	rad/nm ^{i}
a_i	Amplitude coefficients	1×10^{-3}	a.u./nm ^{i}
d_0	Testbed instabilities offset	1×10^{-3}	rad
d_1	Testbed instabilities linear term	1×10^{-3}	rad/s

Figure 4.10 shows the results of the global fit to all overlap values from the 19-port photonic lantern. In the first row, I report the loss progression throughout the optimisation and the estimates of testbed instabilities. The phase offset is consistent with a uniform distribution (noting the low number of counts in total) and the linear drift terms appear with no structure or bias, as expected. Next, the loss term for each port and mode is presented. Two successful fits are highlighted in red, investigated in further detail in subsequent rows. Three fits that will show limitations in the following figure are highlighted in white. In high signal modes shown in Figure 4.10 **c** and **d**, the model correctly matches the data, revealing that the underlying evolution is mostly explained by a simple polynomial in both amplitude and phase, with no oscillations in the transfer function of the device itself.

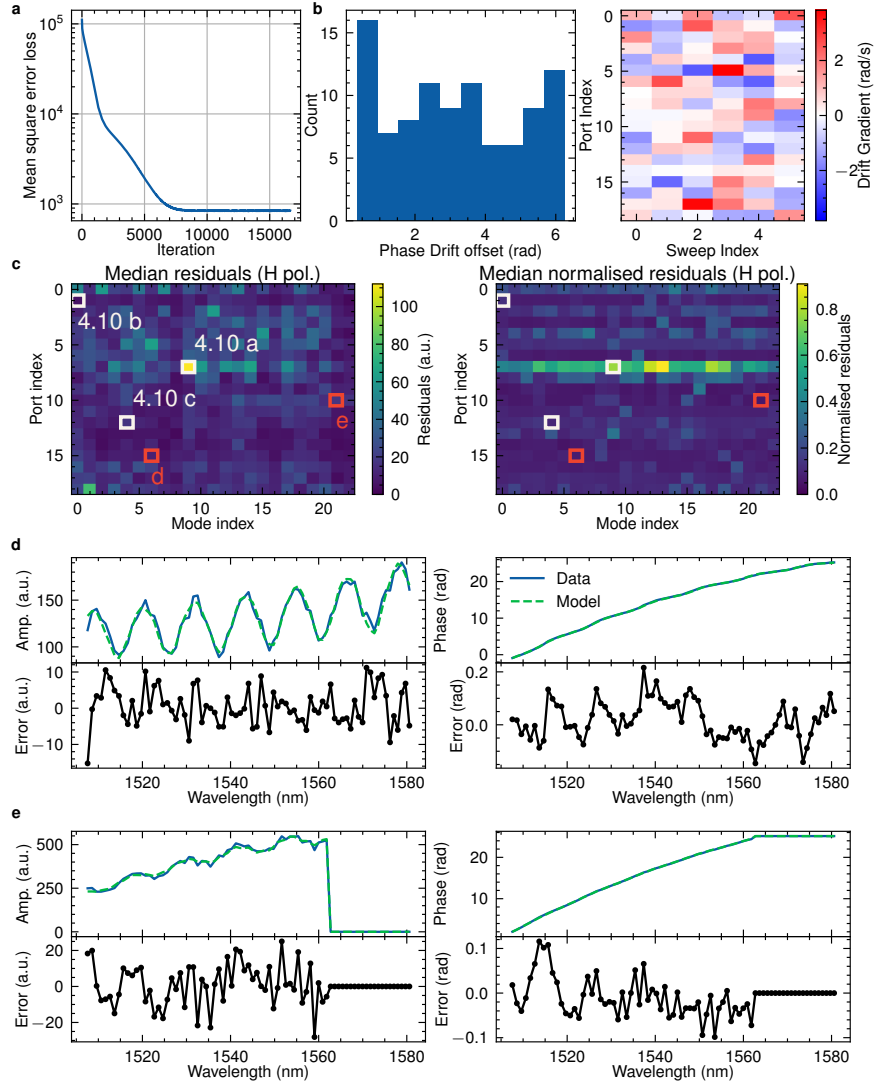


FIGURE 4.10: Fitting the wavelength evolution of all ports, polarisations and modes. **a** Loss evolution during fitting. **b** Testbed instabilities fitting is consistent with priors, with offsets (left) near uniform and trends (right) near zero average. **c** Median residuals over all wavelengths and sweeps for each port and mode (left), with normalisation to the median amplitude (right). Most mode evolution is well explained by the model. Example fits are highlighted in red. **d** An example of a successful fit, with both amplitude and phase residuals (bottom row) showing little structure. **e** A successful fit for a mode that is cut-off. It is clear that, despite the fit matching the data, the LP mode behaviour around the cut-off is unphysical. The model is unbiased by this region and fits the remaining wavelengths well.

The key question is: are the coherence packet oscillations intrinsic to the lantern device? It has been established that the oscillations are formed by different discrete packets of coherence, but there is a need to further establish whether they could be caused by the characterisation setup (in which case they should be discarded). The SMF results in Figure 4.1

are conclusive – these arise from the rest of the testbed, since a single bend in a SMF could contribute to (at most) a single peak, while Figure 4.1 **d** shows multiple peaks. The source of these peaks is likely from imperfect alignment between the axes of the polarisation maintaining fibres used (which are birefringent by design) and the polarisation axis of the source. When the laser polarisation axis is slightly misaligned with the fibre axes, the packet sees two different refractive indices, splitting into two with one delayed with respect to the other. Whether this appears ahead or behind the main peak when sweeping the delay line depends on the arm it originates in. A 3° misalignment between the polarisation direction and the axes in a polarisation maintaining fibre would cause a coherence packet with $A_r = 0.05$, and the distance between peaks depends on the length of birefringent fibre, with one meter producing a packet delay of $340 \mu\text{m}$ (using a birefringence of 3.4×10^{-4} [48]). This is consistent with the results seen and is also supported by the movement/elimination of peaks as different patch cords are employed to make up each arm.

I now turn to understanding some of the limitations of these fits. Figure 4.11 illustrates an obvious failure case, and two more subtle subtle cases. These are all part of the same joint model fit above. Ultimately this points to incongruence between the model expressiveness and the data.

First, the model used to fit the data (including initialisation) assumes that the sum of all non-fundamental coherence packets will not overcome the fundamental signal. This is not the case in the first row of Figure 4.11, where the overall signal oscillates around the origin in complex space, wrapping the phase in a way that is not possible to unwrap naïvely. This causes the model to initialise incorrectly, and find a local minimum at an incorrect packet phase evolution (with the wrong oscillation frequency). A more subtle version of the same effect is shown in the second row, where the model diverges at low amplitude. This points to a slowly evolving background term for each mode, and this is missed by the model. In the case of Figure 4.11 **a**, the background term is in a direction that pushes the fundamental packet enough to come very close to the origin. The misfit for low amplitude modes doesn't significantly degrade the others, since the loss function is the mean square error, and hence modes with low amplitudes have small contributions.

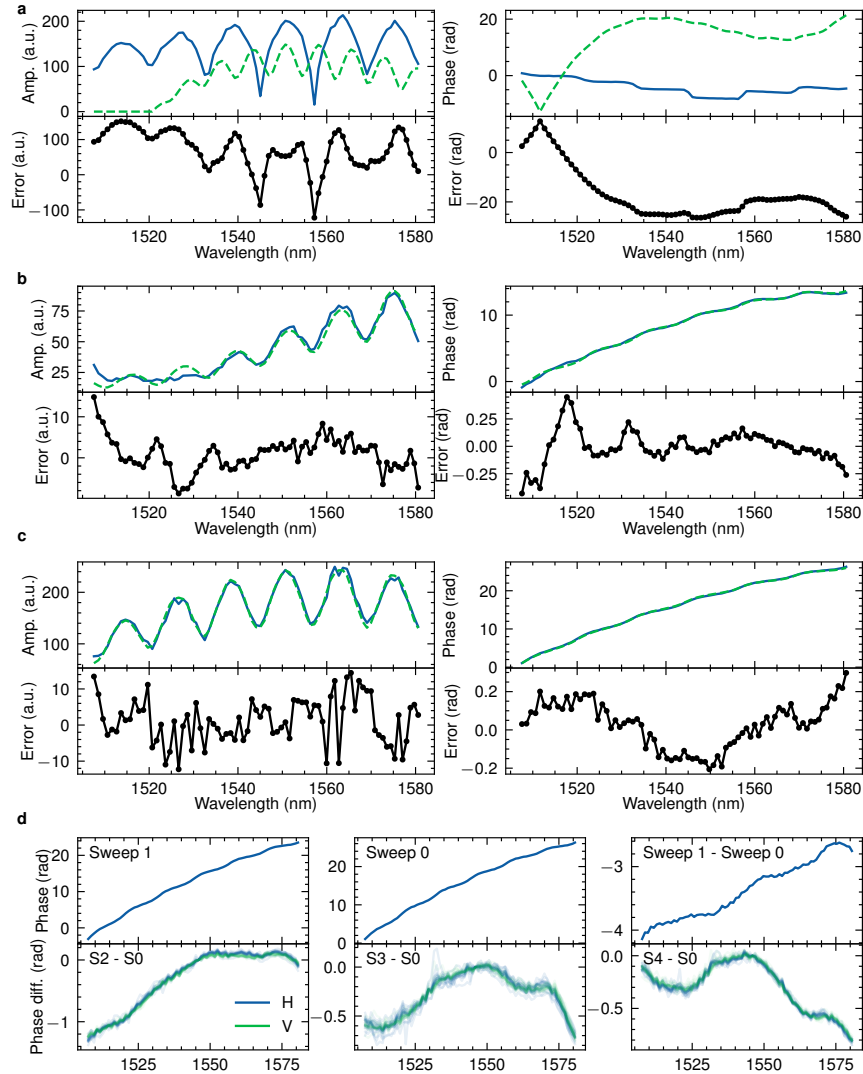


FIGURE 4.11: **Limitations of the fitted model.** **a** Failure at low amplitude, whilst also converging to a local minimum in the linear term of the packet phase (resulting in the wrong oscillation frequency). **b** A more subtle manifestation of the low amplitude contamination, with the fit degrading significantly in phase below data amplitudes of 25 a.u. **c** Misfit in phase only. The amplitude residuals are near stochastic, but the phase shows a clear, oscillatory structure. **d** Investigating the phase between sweeps over a single mode (top row) and many modes (bottom row) highlights non-linear phase evolutions due to testbed instabilities. If the drift was linear, than differences between sweeps would also be linear. This challenges the assumptions in the model and results in the misfit in **c**.

The other limitation seen in the global fits is investigated in the last two rows of Figure 4.11.

In Figure 4.11 **c** there is a good amplitude fit, however there is a clear structure in the residual

of the phase. While it is possible to have a good fit in amplitude but not in phase with a combination of positive and negative frequencies (i.e. coherence packets equally spaced on either side of the fundamental), this is not the most likely cause. Instead, the drift due to testbed instabilities throughout the sweep is inferred by taking differential signals between different sweeps. In time, these are less than a few seconds apart. If, as assumed in the model above, the testbed instabilities drift linearly with time (which is weakly quadratic but close to linear in wavelength), then differences between sweeps should still be linear. Instead, in Figure 4.11 **d**, non-linear variations are present that are common to both polarisations and all modes. The difference seems to have the same timescale as the phase residual seen in Figure 4.11 **c**.

Future work will need to better characterise these or upgrade testbed stability to remove the effects. A starting point for each problem would be to use a different basis set of functions to fit the phase drift from testbed instabilities, while low amplitude background might be able to be characterised by measurements of the reference beam only (and the way this signal propagates through the pipeline to the final modes). Background signal could also be inferred by rotating the lantern and only fitting field components that are consistent with a rotation, analogous to angular differential imaging in exoplanet imaging [49].

Given major upgrades in instrumental stability required are beyond the scope of the present work, I instead pause to take stock of the insights that can be gained from the system as is. By filtering out low amplitude modes and taking differential measurements between modes, the effects of both above limitations would be mitigated, gaining insight into the behaviour of a “typical” mode in phase. This approach is pursued in the following section.

4.5 Differential signals for principal mode dispersion

Since the characterisation setup probes the device response to coherent light, the system is able to directly measure differential observables of the modal dispersion of the photonic lanterns. Oscillations in phase are common between modes (as seen in Figure 4.6**b**), reinforcing the view that these are present from weak birefringence outside of the photonic lantern

itself, as modelled in the previous section. Hence, the approach is extremely sensitive to differential measurements between modes. Only modes relative to high signal LP01 modes are considered for brevity, however this decision is arbitrary and further data relative to other modes (with perhaps higher signal) could piece together a more complete picture. Note that this result is insensitive to any birefringence in the system and the MCF part of the photonic lantern.

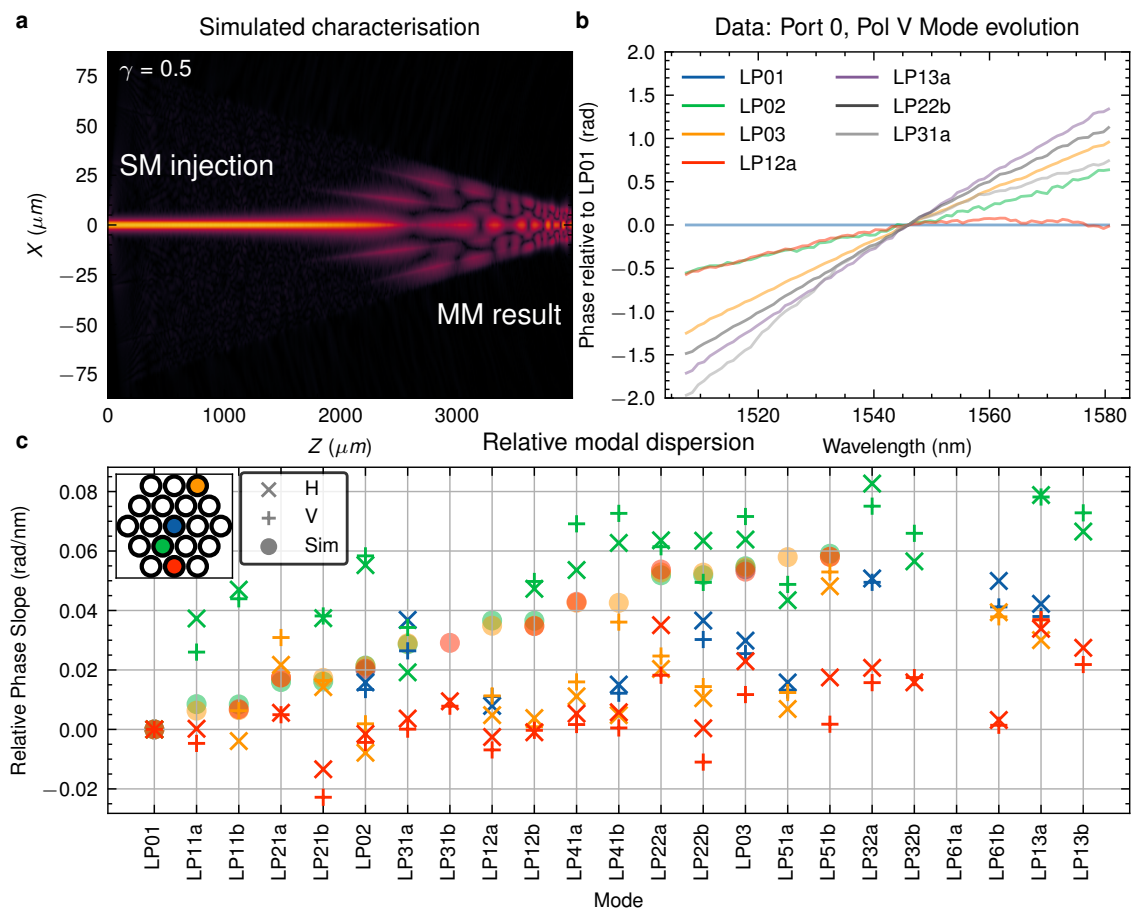


FIGURE 4.12: **Relative modal dispersion measurements of the photonic lantern.** **a** Cut through of a photonic lantern simulation, showing amplitude of the electric field when exciting the central port. **b** Experimental data of differential phase of a few modes in a single port. Evolution is nearly linear, as expected. **c** Gradient of the phase slope in **b** for every mode, in both polarisations for the ports indicated. A cut is taken to only include modes with all amplitudes above 80 data units to avoid low signal contamination. Standard deviations over six sweeps are shown as errorbars, which are virtually invisible.

Figure 4.12 shows this analysis with a comparison to an idealised, yet typically used, scalar simulation. For each mode, a line is fit to the phase relative to LP01, the gradient of which is the relative modal dispersion. This is repeated for different polarisations and with the simulated data and summarised in Figure 4.12 c. Each measurement is extremely repeatable, with the median scatter (standard deviation from 6 sweeps) over all modes above the cut as 10^{-4} rad/nm. An upper bound for the systematic error introduced by background contamination at low signal can be estimated: a 5% background signal can make at most a $2 \times \arctan(0.05)$ rad contribution to the phase over the whole bandwidth, or a change $0.1/73 = 1.5 \times 10^{-3}$ rad/nm in the signal plotted, an order of magnitude smaller than the differences observed between modes. Also, note that this background contamination upper bound would not apply to differences between ports, as all have the same background. Hence, a majority of the scatter between ports and modes is from the device itself, and that the simulation is too idealised. There are also subtle differences between the polarisations that again warrant further investigation with different devices.

4.6 Towards relative principal mode phase measurements

The proposed approach is limited in that the phase between different principal modes cannot be constrained. This is for two reasons. First, the phase is not constant due to the drifts from testbed instabilities, and measurements at different times confound the effect of these drifts with genuine differences between ports. Secondly, to excite different ports (with neighbours spaced at $60 \mu\text{m}$ in the MCF) with no more than $\pi/10$ radians error requires tolerances of no more than $\sim 0.1^\circ$ in relative tip/tilt between the two fibre alignment stages.

This limitation has no consequences when the device is used with intensity measurements of the single-mode end (such as in wavefront sensing), but matters when recombining different ports [13] or using the device for beam shaping. In this section, I use simulation to explore a remedy to this limitation by collecting data when multiple ports are excited simultaneously. The simulation used in this section is not the same as the `RsOfT` simulation above; instead

a random, asymmetric, unitary transfer matrix is assumed which better reflects the diversity of principal modes observed in fabricated devices.

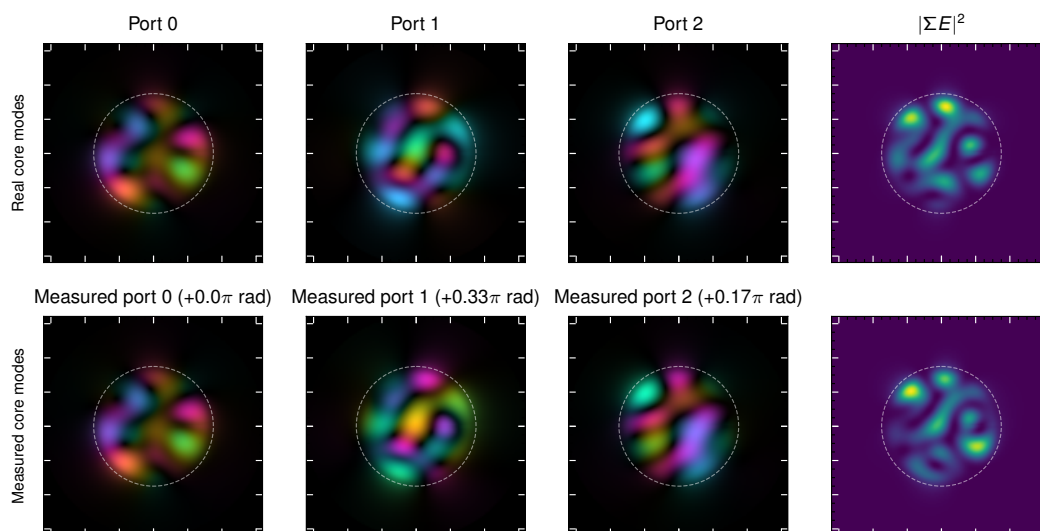


FIGURE 4.13: **Simulated superpositions of port outputs reveal relative phase dependent behaviour.** **a** Simulating three ports of a 19 port photonic lantern and apply a global phase offset to all ports relative to port 0 (bottom row) from true, relative port mapping (top row). **b** The intensity of the superposition of the three fields is different in each case, providing a means to infer relative phase at a sub-radian level.

The key insight leveraged is that it is possible to inject multiple ports at once in a phase deterministic way – by pulling the fibre injection stage back and turning up the laser power, a superposition of the known principal modes for multiple ports are excited. As the number of pixels is larger than the number of unknown, spatially constant phases between ports, there could be enough information to disentangle the phases.

To demonstrate this in a minimal case, I start with an example of 3 ports, where I will recover the relative phases (i.e. 2 phases total). Figure 4.13 shows the results of a simulation of 3 ports of a photonic lantern, with real modes and relative phases (top row) compared to the same modes but with a phase offset relative to port 0 (bottom row). The intensity of the sum of the fields $|\sum E|^2$ is visually different and hence a single measurement of the mode intensity when illuminating 3 ports provides information on the relative phase between ports.

The next questions are: given this additional intensity measurement, is there enough information to constrain the relative phases? If so, is such a solution unique? Figure 4.14 illustrates the loss landscape to answer these questions. Each measured electric field from the port E_i is multiplied by a global phase ϕ_i , and the resulting photometric error between the measured intensity and the predicted intensity is shown. Even for complicated mode fields, the loss space is smooth and has a unique minimum. There is also a non-zero covariance between solutions, with the contour profile not aligned with either axis.

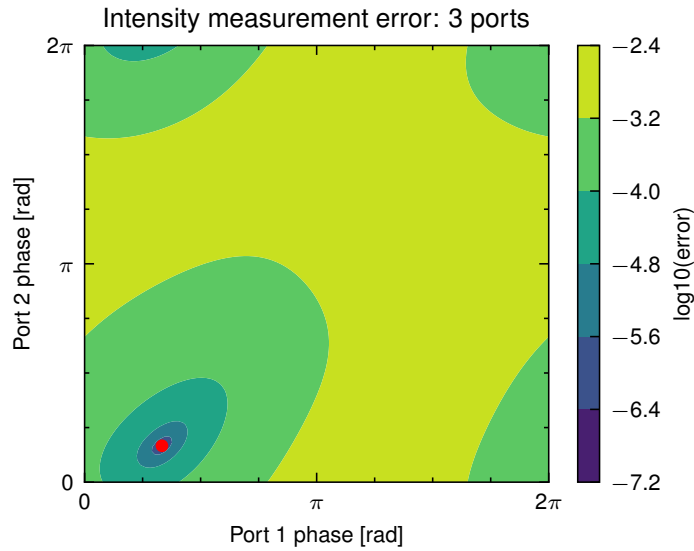


FIGURE 4.14: **Contour plot of loss landscape for relative phase solving.** The sum of squared errors $(I_{\text{meas}} - |\sum_i E_i \exp(i\phi_i)|^2)^2$ for different port phases ϕ_1 and ϕ_2 . The loss landscape is smooth and shows a unique minimum (red dot) when taking into account phase wrapping (and solving on a torus).

Finally, the simulation is generalised to solve for all relative phases, with many measurements. For the hexagonal grid, the optimal injection points are equidistant to the ports. In this part of the simulation, I include the Gaussian beam profile when illuminating ports rather than the idealised, pure 3 port problem shown above. Future work could determine how well this applies to real data.

Figure 4.15 reports the success of the recovery. Additional simulated measurements are taken from the locations shown by taking the superposition of the true modes formed by each port. The measurement matrix is known to the solver, and can be used to infer the relative

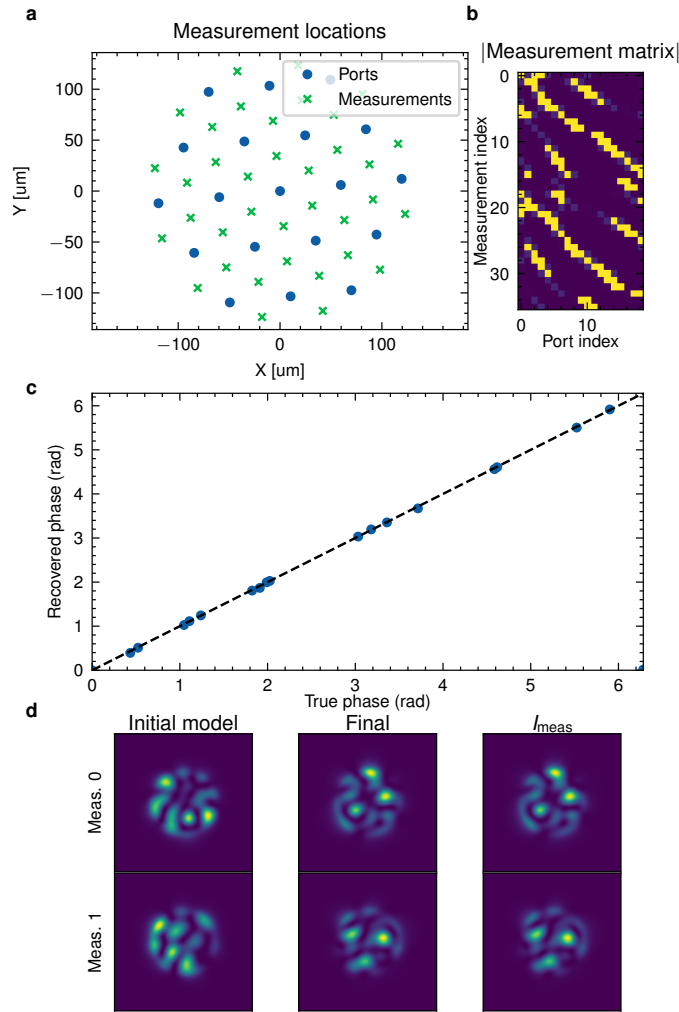


FIGURE 4.15: **Simulated recovery of relative phase between principal modes.** **a** Capturing additional measurements for a 19 port lantern involves sampling all possible combinations of 2 or 3 ports. **b** Absolute value of injected electric field into each port for each measurement. The injection is such that superpositions of different ports are excited at each measurement, with a known (zero) differential phase. The model of a Gaussian beam has small leakage in other ports. **c** The recovered phase vs the true phase, showing there is enough information to recover the relative phase between all ports. **d** Two measurement examples (rows), showing the initial model (left), final model (centre) and the measured intensity (right), showing good agreement as well as the diversity of combined mode fields possible which enables precise phase measurements.

phases. This optimisation problem is solved using the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm (implemented in `scipy` [50]) and initialise from all phases as zero. The

recovered phase is consistent with the true phase despite the distant initialisation in some places, indicating that the insights from Figure 4.14 generalise to higher dimensions. The predictions of the model fit the additional measurements well.

The results presented in this section show that the problem of relative phases between principal modes is not insurmountable with a testbed similar to ours, assuming practical matters like temporal stability and tip/tilt tolerances can be met. While the principle modes do contain enough information to constrain this information, additional work is needed to verify the precision that this can be achieved in the presence of noise and possible cladding modes.

Discussion

5.1 Exploring the characterisation testbed

The characterisation of a photonic lantern through direct measurements of the electric field represents a significant step beyond previous intensity-only studies. Employing digital off-axis holography yields deeper insight and more thorough device characterisation. In combination with coherent spectral characterisation (at the white light fringe) and polarisation splitting at the output of the device, the data products delivered present many more dimensions than previously available, best summarised in Figure 4.6. Ultimately, this illuminates (literally and figuratively) how a photonic lantern maps modes from input to output bases, with a focus on understanding wavelength evolution and polarisation dependent behaviour.

The digital off-axis holography method is well suited for this characterisation problem for several reasons. By attenuating everything except the off-axis terms in Fourier space, low spatial frequency contamination and incoherent components such as varying background are attenuated. As the number of frequencies filtered through is small, the effects of read noise in the image space is also reduced. Both of these advantages provide excellent sensitivity, highlighted by reconstruction in the low signal regime in Figure 4.2. In addition, the particular approach explored in this thesis is basis mode agnostic, assuming only that electric fields are supported on a circle (with the assumption used in the centre finding and filtering), enabling a validation that the modes supported in the device are mostly LP modes.

The system is also able to characterise a device relatively quickly, with all six sweeps in wavelength over 19 ports taken within 6.5 minutes. The port and white light fringe finding

takes longer at about 40 minutes. For higher port count devices, these times would increase approximately linearly, and hence would remain much faster than the production cycle of a photonic lantern (a few days). Hence, this characterisation can be performed easily as part of the manufacturing and design loop.

Through these strengths it is possible to make several key findings that would otherwise be out of reach, including that: photonic lanterns are well (but not completely) described by LP modes, mode evolution is governed by a low order polynomial (hence a device can be well summarised by just a few parameters), and the modes from each port are almost orthogonal. The implications of these and other conclusions are far ranging and summarised in section 5.2.

The characterisation system nonetheless has some limitations. Firstly, the oscillations present as a function of wavelength have the potential to confound sensitive, absolute measurements. Characterisation systems that do not measure different wavelengths coherently (as done in this thesis) would struggle to quantify this effect to the same extent. Other systems which capture data at very narrow line widths would also be more sensitive to distant coherence packets, as the coherence length becomes very large. There are several strategies to mitigate the effects of these oscillations: fitting these to remove them (as shown in section 4.4), or simply applying a bandpass filter at the known frequency, which is valid as the results in Figure 4.4 show the underlying evolution is slow and has a wavelength scale of tens of nanometres.

As highlighted in section 4.4, there are low amplitude components to the signals that are not consistent with the proposed model, and there is also a phase drift due to testbed instabilities throughout the sweep. These are detrimental to measurements of the absolute amplitude and phase. This is mitigated by deriving robust observables from differential measurements with an amplitude threshold.

Finally, the setup is limited in that it can not directly measure the relative phase between principal modes since the testbed instabilities cause the two arms to drift in phase between measurements of different ports, and that the required tolerances on relative tip/tilt at fibre

injection are too stringent. I explore how one could, in principle, still make this measurement if these practical conditions are met, illustrating the process in simulation in section 4.6. However, such a measurement is very sensitive to the exact setup of the multicore fibre, as evidenced in Figure 4.3, where differences of order 10^3 radians between ports are induced by a gentle bend with radius 50 cm. If the application after characterisation requires coherent combination (where the phase between single-mode ports matters), then the system is sensitive to the layout and environment near the multicore fibre and should have this aspect characterisable in-situ. The success of the recovery in section 4.6 also illustrates how tractable this inference problem becomes – relative phase of principal modes can be inferred from intensity only combinations when the underlying electric field of principal modes has been previously characterised.

5.2 Implications for photonic lantern applications

As mode sorting devices find widespread use where electric fields must be measured or controlled, the findings presented have universal implications spanning virtually all downstream applications of photonic lanterns in addition to photonics theory.

The characterisation setup produces an empirical measurement illustrating how a photonic lantern works, directly measuring the mapping between single-mode ports and the multimode face. An analysis of the evolution of the fields as a function of wavelength finds that the behaviour of each mode is well explained as a low dimensional polynomial, evolving with a typical λ scale of a few tens of nanometers. The relative success of the low dimensional representation in section 4.4 suggest that perhaps a fundamental, low dimensional representation of a given device could encapsulate all of the effects. This opens new questions such as “what are the subsets of this space that are manufacturable, and which are useful for a given application?”

The results in this thesis also demonstrate that photonic lantern principal modes are, for the most part, orthogonal. This provides a means of benchmarking devices during fabrication

with some simple metrics where the goal is inversion in sensing applications or expressivity in beam shaping applications.

The existence and extent of birefringence and polarisation-dependent loss within a photonic lantern are still open questions; weak differences between fields with different polarisations are present, and dispersions are different, warranting further investigation with a wider range of devices. The characterisation approach in this thesis offers sensitivity to these quantities, however upgrades to reduce system birefringence are required to convincingly yield robust results. Another possible upgrade would be to vary the input polarisation at device injection, in addition to the existing measurement of output polarisation. Such an experiment would verify to what extent orthogonal polarisations remain orthogonal, with consequences for polarisation multiplexing in telecommunications.

The results presented also deliver a timely warning for the use of simulation alone in modelling these devices. Previous results [17] have shown high symmetry in the transfer matrix between ports on opposite sides of the hexagonal grid. This symmetry is not observed in reality, highlighting the need for physical device characterisation over simulation alone. The characterisation system would be capable of discerning if this variation is a fixed or variable offset from the design target given multiple instances of devices, and this is left as future work.

In enabling the measurement of devices within the manufacturing and design loop, target designs will be better achieved and rapidly verifiable.

Instrument concepts that involve coherent combination of photonic lantern outputs, such as coherent imaging [13] require careful consideration of the port dependent optical path difference induced by the multicore fibre (measured at up to $750 \mu\text{m}$ per metre of fibre in this case) as well as the modal dispersion. The modal dispersion findings (Figure 4.12) suggest that such an instrument must be spectrally dispersed with a sampling finer than 16 nm (for 1 rad difference) to maintain coherence between different modes in different ports. More broadly, the findings in this thesis should be applied to provide higher fidelity simulations

for instrument concepts, and both the electric fields and transfer matrices for the 19-port lantern are shared.

Next, data driven approaches, especially where data is expensive such as astronomy, should leverage characterisation results as a prior to fitting. For example, an image reconstruction algorithm for high angular resolution astrophotonics should use the measured transfer matrix as the heart of a model that can be fine tuned in-situ using conventional intensity measurements. More generally, estimators using the characterised photonic lantern should be regularised to be consistent with the characterisation. The transfer matrix can also serve in optimal experiment design, where observations could be planned around maximising sensitivity to scientific outcomes depending on the location of objects in the device's field of view.

Previous work for computational imaging with photonic lanterns for areas such as microendoscopy has shown promise, e.g. [22]. More advanced and powerful modalities leverage coherent combination in an active illumination system, or many ports in a passive imaging system. Both cases require knowledge of the relative phases and amplitudes of the fields at the multimode end. The characterisation setup provides almost all of this information directly, with the only exception of relative phases between principal modes. For microendoscopy specifically, the mode/port counts must be considerably higher than the 19 port device in this work, however the approach can scale whilst still remaining tractable. The time needed for recording swept data is around 10s per port, which is equivalent to characterising a 1000 port device in under 3 hours. In this case, time taken for white light fringe finding can be decreased significantly by using a bandpass filter for a longer coherence length (hence less sampling of the delay line motor) and a more innovative sampler that leverages correlation between ports (rather than a fixed step size).

Conclusion

6.1 Significance

This thesis presents a means of characterising the mapping between input and output modes of a photonic lantern. The presented laboratory system employs digital off-axis holography to capture the electric fields directly (as opposed to intensity only) from an individual port. A full description of this system is accompanied by the sharing of key insights, lessons learned, code, and post-processed data. To the best of my knowledge, this thesis provides the first ever multi-wavelength, polarisation decomposed characterisation of the principal modes of a photonic lantern. The analysis of the resulting data enables observation of the evolution of the modal mapping with wavelength, with a smooth change on the scale of tens of nanometres. The system generates robust differential measurements of relative dispersion of components of principle modes, showing significant variation. Finally, there are subtle differences in polarisation in the transfer function, warranting further investigation.

The characterisation setup and sharing of extremely accurate, realistic mode transfer matrices enables future work in a broad range of applications including astrophotonic design, computational imaging, device fabrication feedback loops and beam shaping to an unprecedented level of precision.

This thesis is motivated in particular by high angular resolution astronomy. The broad framework presented in Figure 1.1 clearly requires knowledge of the modal mapping of a given, fabricated device. This knowledge is valuable from several perspectives. First, designs of future instruments can use higher fidelity devices in simulation by substituting the measured

mode transfer matrix instead of an ideal, simulated device. Indeed, the results in this thesis demonstrated that idealised simulations inadequately capture the complexities of fabricated devices, most notably in asymmetries in principal mode fields. Furthermore, the proposed approach as a whole is easily incorporated into a manufacturing design loop. When an instrument provides specifications (such as the amount of cross talk between modes or the scale of wavelength change), fabrication can take place in a closed loop on a similar timescale to open loop fabrication without this testbed. Finally, data-driven approaches to wavefront sensing and imaging should use these results to constrain/regularise estimators. This would reduce the data and time needed to develop the pipeline, including on-sky time which is particularly valuable. This would also increase confidence in results from such instruments, as these approaches typically employ machine learning whose functional elements are notoriously opaque to intuitive interpretation.

6.2 Directions for future work

There are exciting avenues of future work with this characterisation setup. It is now possible to study how the transfer matrix varies with different architectures and to compare this to simulations. Furthermore, the manufacturing repeatability of these devices can be measured directly.

Also, further work is needed to understand polarisation dependence in the photonic lantern, and how this changes with parameters such as the number of ports or modes supported at either end.

In addition, there are some hardware upgrades that would improve the system. Using a fibre bundle adapter would expose useable interfaces for pigtailed lanterns, which are desired for some applications. The bundle should have all ports path length matched to within a few millimetres to ensure the white light fringe finding can be completed in a timely fashion. A more involved upgrade could address some of the current limits on the system: drift from testbed instabilities and bandwidth. A broadband source such as a super-continuum laser

combined with an integral field unit (packaged in the camera for convenience) would measure all wavelengths simultaneously, rendering drift from testbed instabilities negligible. To reap the rewards of the increased bandwidth of the source, a bulk optics split between the two arms of the interferometer would overcome the limit by fibre components in the system. The final upgrade is one that could remove the oscillations by providing the degrees of freedom in the polarisation space of each to align the coherence packets with the fundamental.

With these improvements such a characterisation system would illuminate these novel multimode converters, paving the way for future instruments to perhaps directly image the glimmer of exoplanets orbiting distant stars.

Bibliography

- [1] A. K. Taras *et al.*, ‘Heimdallr and solarstein: Alignment, calibration, and correction in the asgard suite at the VLTI,’ in *Ground-based and Airborne Instrumentation for Astronomy X*, SPIE, vol. 13096, 2024, pp. 2178–2182.
- [2] J. Carpenter, *digHolo : High-speed library for off-axis digital holography and Hermite-Gaussian decomposition*, arXiv:2204.02348 [eess], Apr. 2022. [Online]. Available: <http://arxiv.org/abs/2204.02348>.
- [3] J. O’Sullivan, ‘How we made the wireless network,’ en, *Nature Electronics*, vol. 1, no. 2, pp. 147–147, Feb. 2018, ISSN: 2520-1131. DOI: [10.1038/s41928-018-0027-y](https://doi.org/10.1038/s41928-018-0027-y). [Online]. Available: <https://www.nature.com/articles/s41928-018-0027-y>.
- [4] W. S. Boyle and G. E. Smith, *Charge coupled semiconductor devices (Bell System Technical Journal 1970)*. Jan. 1993, Pages: 475 Publication Title: Selected Papers on Instrumentation in Astronomy ADS Bibcode: 1993inas.book..475B. [Online]. Available: <https://ui.adsabs.harvard.edu/abs/1993inas.book..475B>.
- [5] K. B. Follette, ‘An Introduction to High Contrast Differential Imaging of Exoplanets and Disks,’ en, *Publications of the Astronomical Society of the Pacific*, vol. 135, no. 1051, p. 093 001, Sep. 2023, ISSN: 0004-6280, 1538-3873. DOI: [10.1088/1538-3873/aceb31](https://doi.org/10.1088/1538-3873/aceb31). [Online]. Available: <https://iopscience.iop.org/article/10.1088/1538-3873/aceb31>.
- [6] Lyot, M, ‘A study of the Solar Corona And Prominences Without Eclipses,’ *MNRAS*, vol. 99, no. 580, 1939.
- [7] N. Deshler, S. Haffert and A. Ashok, *Quantum Limits of Exoplanet Detection and Localization*, en, arXiv:2403.17988 [quant-ph], Jun. 2025. DOI: [10.48550/arXiv.2403.17988](https://doi.org/10.48550/arXiv.2403.17988). [Online]. Available: <http://arxiv.org/abs/2403.17988>.

- [8] B. R. M. Norris *et al.*, ‘The photonic lantern wavefront sensor and imager: Focal plane wavefront sensing and optimal imaging at the diffraction limit and beyond,’ en, in *Adaptive Optics Systems IX*, D. Schmidt, E. Vernet and K. J. Jackson, Eds., Yokohama, Japan: SPIE, Aug. 2024, p. 54, ISBN: 978-1-5106-7517-9. DOI: [10.1117/12.3019643](https://doi.org/10.1117/12.3019643). [Online]. Available: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/13097/3019643/The-photonic-lantern-wavefront-sensor-and-imager--focal-plane/10.1117/12.3019643.full>.
- [9] M. Bottom, J. K. Wallace, R. D. Bartos, J. C. Shelton and E. Serabyn, ‘Speckle suppression and companion detection using coherent differential imaging,’ *Monthly Notices of the Royal Astronomical Society*, vol. 464, no. 3, pp. 2937–2951, 2017.
- [10] O. Guyon, ‘Limits of Adaptive Optics for High-Contrast Imaging,’ en, *The Astrophysical Journal*, vol. 629, no. 1, pp. 592–614, Aug. 2005, ISSN: 0004-637X, 1538-4357. DOI: [10.1086/431209](https://doi.org/10.1086/431209). [Online]. Available: <https://iopscience.iop.org/article/10.1086/431209>.
- [11] J. Chilcote *et al.*, ‘SCEXAO/CHARIS Direct Imaging of A Low-mass Companion At A Saturn-like Separation from an Accelerating Young A7 Star,’ en, *The Astronomical Journal*, vol. 162, no. 6, p. 251, Dec. 2021, ISSN: 0004-6256, 1538-3881. DOI: [10.3847/1538-3881/ac29ba](https://doi.org/10.3847/1538-3881/ac29ba). [Online]. Available: <https://iopscience.iop.org/article/10.3847/1538-3881/ac29ba>.
- [12] A. K. Taras *et al.*, ‘Kernel nulling at vlti with photonic lanterns for optimal fibre injection,’ in *Optical and Infrared Interferometry and Imaging IX*, SPIE, vol. 13095, 2024, pp. 242–250.
- [13] Y. J. Kim *et al.*, ‘Coherent Imaging with Photonic Lanterns,’ en, *The Astrophysical Journal*, vol. 964, no. 2, p. 113, Apr. 2024, ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/ad245e](https://doi.org/10.3847/1538-4357/ad245e). [Online]. Available: <https://iopscience.iop.org/article/10.3847/1538-4357/ad245e>.
- [14] N. K. Fontaine *et al.*, ‘Photonic lanterns, 3-D waveguides, multiplane light conversion, and other components that enable space-division multiplexing,’ *Proceedings of the IEEE*, vol. 110, no. 11, pp. 1821–1834, 2022.

- [15] A. M. Velázquez-Benítez *et al.*, ‘Scaling photonic lanterns for space-division multiplexing,’ en, *Scientific Reports*, vol. 8, no. 1, Jun. 2018, Publisher: Springer Science and Business Media LLC, ISSN: 2045-2322. DOI: [10.1038/s41598-018-27072-2](https://doi.org/10.1038/s41598-018-27072-2). [Online]. Available: <https://www.nature.com/articles/s41598-018-27072-2>.
- [16] B. R. M. Norris, J. Wei, C. H. Betters, A. Wong and S. G. Leon-Saval, ‘An all-photonic focal-plane wavefront sensor,’ en, *Nature Communications*, vol. 11, no. 1, p. 5335, Oct. 2020, ISSN: 2041-1723. DOI: [10.1038/s41467-020-19117-w](https://doi.org/10.1038/s41467-020-19117-w). [Online]. Available: <https://www.nature.com/articles/s41467-020-19117-w>.
- [17] B. Norris, C. Betters, J. Wei, S. Yerolatsitis, R. Amezcua-Correa and S. Leon-Saval, ‘Optimal broadband starlight injection into a single-mode fibre with integrated photonic wavefront sensing,’ en, *Optics Express*, vol. 30, no. 19, p. 34 908, Sep. 2022, ISSN: 1094-4087. DOI: [10.1364/OE.465639](https://doi.org/10.1364/OE.465639). [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=oe-30-19-34908>.
- [18] J. W. Lin *et al.*, ‘Real-time Experimental Demonstrations of a Photonic Lantern Wavefront Sensor,’ en, *The Astrophysical Journal Letters*, vol. 959, no. 2, p. L34, Dec. 2023, ISSN: 2041-8205, 2041-8213. DOI: [10.3847/2041-8213/ad12a4](https://doi.org/10.3847/2041-8213/ad12a4). [Online]. Available: <https://iopscience.iop.org/article/10.3847/2041-8213/ad12a4>.
- [19] N. Jovanovic *et al.*, ‘2023 Astrophotonics Roadmap: Pathways to realizing multifunctional integrated astrophotonic instruments,’ en, *Journal of Physics: Photonics*, vol. 5, no. 4, p. 042 501, Oct. 2023, ISSN: 2515-7647. DOI: [10.1088/2515-7647/ace869](https://doi.org/10.1088/2515-7647/ace869). [Online]. Available: <https://iopscience.iop.org/article/10.1088/2515-7647/ace869>.
- [20] A. Milne *et al.*, ‘Coherent Beam Shaping with Multicore Fiber Photonic Lanterns,’ en, in *2023 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC)*, Munich, Germany: IEEE, Jun. 2023, pp. 1–1, ISBN: 9798350345995. DOI: [10.1109/CLEO/Europe-EQEC57999](https://doi.org/10.1109/CLEO/Europe-EQEC57999).

- 2023.10232791. [Online]. Available: <https://ieeexplore.ieee.org/document/10232791/>.
- [21] H. K. Chandrasekharan and R. Donaldson, *High-throughput polarization-independent spatial mode shaping with mode-selective photonic lanterns*, en, arXiv:2506.08595 [physics], Jun. 2025. DOI: [10.48550/arXiv.2506.08595](https://doi.org/10.48550/arXiv.2506.08595). [Online]. Available: <http://arxiv.org/abs/2506.08595>.
- [22] D. Choudhury *et al.*, ‘Computational optical imaging with a photonic lantern,’ en, *Nature Communications*, vol. 11, no. 1, p. 5217, Oct. 2020, ISSN: 2041-1723. DOI: [10.1038/s41467-020-18818-6](https://doi.org/10.1038/s41467-020-18818-6). [Online]. Available: <https://www.nature.com/articles/s41467-020-18818-6>.
- [23] S. G. Leon-Saval, T. A. Birks, J. Bland-Hawthorn and M. Englund, ‘Multimode fiber devices with single-mode performance,’ en, *Optics Letters*, vol. 30, no. 19, p. 2545, Oct. 2005, ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.30.002545](https://doi.org/10.1364/OL.30.002545). [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=ol-30-19-2545>.
- [24] S. G. Leon-Saval, A. Argyros and J. Bland-Hawthorn, ‘Photonic lanterns,’ en, *Nanophotonics*, vol. 2, no. 5-6, pp. 429–440, Dec. 2013, ISSN: 2192-8614, 2192-8606. DOI: [10.1515/nanoph-2013-0035](https://doi.org/10.1515/nanoph-2013-0035). [Online]. Available: <https://www.degruyter.com/document/doi/10.1515/nanoph-2013-0035/html>.
- [25] T. A. Birks, I. Gris-Sánchez, S. Yerolatsitis, S. G. Leon-Saval and R. R. Thomson, ‘The photonic lantern,’ en, *Advances in Optics and Photonics*, vol. 7, no. 2, p. 107, Jun. 2015, ISSN: 1943-8206. DOI: [10.1364/AOP.7.000107](https://doi.org/10.1364/AOP.7.000107). [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=aop-7-2-107>.
- [26] J. Lin *et al.*, ‘Focal-plane wavefront sensing with photonic lanterns: Theoretical framework,’ en, *Journal of the Optical Society of America B*, vol. 39, no. 10, p. 2643, Oct. 2022, ISSN: 0740-3224, 1520-8540. DOI: [10.1364/JOSAB.466227](https://doi.org/10.1364/JOSAB.466227). [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=josab-39-10-2643>.

- [27] J. Lin *et al.*, ‘Focal-plane wavefront sensing with photonic lanterns ii: Numerical characterization and optimization,’ *Journal of the Optical Society of America B*, vol. 40, no. 12, pp. 3196–3208, 2023.
- [28] M. Tsang, R. Nair and X.-M. Lu, ‘Quantum theory of superresolution for two incoherent optical point sources,’ *Physical Review X*, vol. 6, no. 3, p. 031 033, 2016.
- [29] D. Gabor, ‘A new microscopic principle.,’ *Nature*, 1948.
- [30] J. Zhang, S. Dai, C. Ma, T. Xi, J. Di and J. Zhao, ‘A review of common-path off-axis digital holography: Towards high stable optical instrument manufacturing,’ *Light: advanced manufacturing*, vol. 2, no. 3, pp. 333–349, 2021.
- [31] D. Zia, N. Dehghan, A. D’Errico, F. Sciarrino and E. Karimi, ‘Interferometric imaging of amplitude and phase of spatial biphoton states,’ en, *Nature Photonics*, vol. 17, no. 11, pp. 1009–1016, Nov. 2023, ISSN: 1749-4885, 1749-4893. DOI: [10.1038/s41566-023-01272-3](https://doi.org/10.1038/s41566-023-01272-3). [Online]. Available: <https://www.nature.com/articles/s41566-023-01272-3>.
- [32] J. Carpenter and T. D. Wilkinson, ‘All Optical Mode-Multiplexing Using Holography and Multimode Fiber Couplers,’ en, *Journal of Lightwave Technology*, vol. 30, no. 12, pp. 1978–1984, Jun. 2012, ISSN: 0733-8724, 1558-2213. DOI: [10.1109/JLT.2012.2191586](https://doi.org/10.1109/JLT.2012.2191586). [Online]. Available: <http://ieeexplore.ieee.org/document/6177203/>.
- [33] N. K. Fontaine *et al.*, ‘Hermite-Gaussian mode multiplexer supporting 1035 modes,’ en, in *Optical Fiber Communication Conference (OFC) 2021*, Washington, DC: Optica Publishing Group, 2021, p. M3D.4, ISBN: 978-1-943580-86-6. DOI: [10.1364/OFC.2021.M3D.4](https://doi.org/10.1364/OFC.2021.M3D.4). [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=OFC-2021-M3D.4>.
- [34] L. Zhao, W. Li, Y. Chen, T. Yu, E. Zhao and J. Tang, ‘Design and characterization of a self-matching photonic lantern for all few-mode fiber laser systems,’ en, *Optics Express*, vol. 32, no. 10, p. 16 799, May 2024, ISSN: 1094-4087. DOI: [10.1364/OE.520588](https://doi.org/10.1364/OE.520588). [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=oe-32-10-16799>.

- [35] D. Yu *et al.*, ‘Mode-dependent characterization of photonic lanterns,’ en, *Optics Letters*, vol. 41, no. 10, p. 2302, May 2016, ISSN: 0146-9592, 1539-4794. DOI: [10.1364/OL.41.002302](https://doi.org/10.1364/OL.41.002302). [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=ol-41-10-2302>.
- [36] R. I. Becerra-Deana *et al.*, *Fabrication and Characterization of Photonic Lanterns Using Coupled-Mode Theory*, en, arXiv:2411.02182 [physics], Nov. 2024. DOI: [10.48550/arXiv.2411.02182](https://doi.org/10.48550/arXiv.2411.02182). [Online]. Available: <http://arxiv.org/abs/2411.02182>.
- [37] Y. Xin *et al.*, ‘Laboratory demonstration of a Photonic Lantern Nuller in monochromatic and broadband light,’ en, *Journal of Astronomical Telescopes, Instruments, and Systems*, vol. 10, no. 02, Apr. 2024, ISSN: 2329-4124. DOI: [10.1117/1.JATIS.10.2.025001](https://doi.org/10.1117/1.JATIS.10.2.025001). [Online]. Available: <https://www.spiedigitallibrary.org/journals/Journal-of-Astronomical-Telescopes-Instruments-and-Systems/volume-10/issue-02/025001/Laboratory-demonstration-of-a-Photonic-Lantern-Nuller-in-monochromatic-and/10.1117/1.JATIS.10.2.025001.full>.
- [38] S. S. Eikenberry *et al.*, ‘Photonic quantum-inspired sub-diffraction imager,’ en, in *Ground-based and Airborne Instrumentation for Astronomy X*, J. R. Vernet, J. J. Bryant and K. Motohara, Eds., Yokohama, Japan: SPIE, Jul. 2024, p. 26, ISBN: 978-1-5106-7515-5. DOI: [10.1117/12.3019152](https://doi.org/10.1117/12.3019152). [Online]. Available: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/13096/3019152/Photonic-quantum-inspired-sub-diffraction-imager/10.1117/12.3019152.full>.
- [39] M. A. Romer *et al.*, ‘Broadband photonic lantern transfer matrix characterization for wavefront sensing,’ en, in *Photonic Instrumentation Engineering XII*, Y. Soskind and L. E. Busse, Eds., San Francisco, United States: SPIE, Mar. 2025, p. 26, ISBN: 978-1-5106-8494-2. DOI: [10.1117/12.3043722](https://doi.org/10.1117/12.3043722). [Online]. Available: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/13373/3043722/Broadband-photonic-lantern-transfer>

- [matrix-characterization-for-wavefront-sensing/10.1117/12.3043722.full](#).
- [40] D. Marcuse, ‘Loss Analysis of Single-Mode Fiber Splices,’ en, *Bell System Technical Journal*, vol. 56, no. 5, pp. 703–718, May 1977, ISSN: 00058580. DOI: [10.1002/j.1538-7305.1977.tb00534.x](#). [Online]. Available: <https://ieeexplore.ieee.org/document/6768444>.
- [41] J. Canny, ‘A computational approach to edge detection,’ *IEEE Transactions on pattern analysis and machine intelligence*, no. 6, pp. 679–698, 2009.
- [42] C. F. Jekel, ‘Obtaining non-linear orthotropic material models for pvc-coated polyester via inverse bubble inflation,’ Ph.D. dissertation, Stellenbosch: Stellenbosch University, 2016.
- [43] L. Desdoigts, B. J. S. Pope, J. Dennis and P. G. Tuthill, ‘Differentiable optics with ∂ Lux: I—deep calibration of flat field and phase retrieval with automatic differentiation,’ en, *Journal of Astronomical Telescopes, Instruments, and Systems*, vol. 9, no. 02, Jun. 2023, ISSN: 2329-4124. DOI: [10.1117/1.JATIS.9.2.028007](#). [Online]. Available: <https://www.spiedigitallibrary.org/journals/Journal-of-Astronomical-Telescopes-Instruments-and-Systems/volume-9/issue-02/028007/Differentiable-optics-with-Lux--Ideep-calibration-of-flat-field/10.1117/1.JATIS.9.2.028007.full>.
- [44] J. Rader, T. Lyons and P. Kidger, *Optimistix: Modular optimisation in JAX and Equinox*, en, arXiv:2402.09983 [math], Feb. 2024. DOI: [10.48550/arXiv.2402.09983](#). [Online]. Available: <http://arxiv.org/abs/2402.09983>.
- [45] J. Bradbury *et al.*, *JAX: Composable transformations of Python+NumPy programs*, version 0.3.13, 2018. [Online]. Available: <http://github.com/google/jax>.
- [46] S. Prahl, *Ofiber: A python module for light propagation in optical fibers*, 2024.
- [47] D. P. Kingma and J. Ba, ‘Adam: A method for stochastic optimization,’ *arXiv preprint arXiv:1412.6980*, 2014.

- [48] L. Polz, A. Jarsen, H. Bartelt and J. Roths, ‘Birefringence properties of a polarization maintaining panda fibre during bragg grating regeneration,’ in *24th International Conference on Optical Fibre Sensors*, SPIE, vol. 9634, 2015, pp. 753–756.
- [49] C. Marois, D. Lafreniere, R. Doyon, B. Macintosh and D. Nadeau, ‘Angular differential imaging: A powerful high-contrast imaging technique,’ *The Astrophysical Journal*, vol. 641, no. 1, p. 556, 2006.
- [50] P. Virtanen *et al.*, ‘Scipy 1.0: Fundamental algorithms for scientific computing in python,’ *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.

APPENDIX A

Component List

Table A.1 lists the major components used in the system.

TABLE A.1: Components in the characterisation system, listed in order of propagation.

Description	Supplier	Part number	Notes
<i>Pre-split components</i>			
Swept wavelength source	Anritsu	MG9637A	
Broadband SLD	Thorlabs	SLD1005S	Driven by CLD1015
Coupler	Thorlabs	PW1550R3F1	
<i>Injection beam</i>			
Inline fibre polariser	AFW	ILP-15-L-1-5-1-SA	
Fibre alignment stage	Luminos	I6000	Custom interface plates
Delay line stage	Newport	UTS150CC	Driven by ESP300 controller
Collimating lens	Thorlabs	C20SMA-C	
ND filter	Thorlabs	NENIR20A-C	
Re-imaging lens	Thorlabs	AC254-400-C-ML	
Beam displacer	Thorlabs	BD27	
Compensating plate	Custom		Microscope slides, NOA61
<i>Reference beam</i>			
Inline fibre polariser	AFW	ILP-15-C-1-1-11-SA	
Shutter	Thorlabs	SH05/M	
Collimating lens	Thorlabs	AC254-200-C-ML	
<i>Beam recombination</i>			
Beamsplitter	Thorlabs	CCM1-BS015/M	
Detector	Lucid	Atlas SWIR 1.3MP	
<i>Electronics</i>			
Trigger controller	Arduino	Nano	