



THE UNIVERSITY OF  
**SYDNEY**

**Out of the shadows: unveiling circumstellar structure with  
optical wavelength spectro-polarimetric interferometry.**

Lucinda Elizabeth Mulhearin Lilley

*A thesis submitted to fulfil the requirements of the degree of:*

Doctor of Philosophy (PhD)

Supervisors:  
Professor Peter Tuthill, Dr Barnaby Norris

The University of Sydney  
School of Physics

2025

## Declaration

I, *Lucinda Lilley*, declare that this thesis submitted to fulfil the requirements for the conferral of the degree *Doctor of Philosophy (PhD)* from the University of Sydney, is wholly my own work unless otherwise referenced or acknowledged. This document has not been submitted for qualifications at any other academic institution.

---

*Lucinda Lilley, October 14 2025*

---

*Peter Tuthill, October 14 2025*

## Abstract

Interferometry is the ultimate exercise in information theory. In performing interferometry a trade is made – fewer observables are obtained than from a direct image but each may have intrinsically higher resolution and signal to noise. However, this trade is not free – from a restricted number of observables, information recovery techniques are required to fit models and reconstruct images of the underlying scene for astrophysical inference.

A powerful observational technique is found in the combination of interferometry with spectro-polarimetry – optical wavelength *spectro-polarised interferometric* measurement provides a high resolution spatio-spectral probe of the geometry, chemistry and grain size of the dust within inner-most circumstellar environments. Observing the circumstellar environments of evolved stars remains a core pursuit of modern stellar astrophysics, as these regions encode information on the physical mechanisms driving the intense stellar mass loss that enriches the interstellar medium. Whilst information recovery techniques for traditional optical interferometric observables have a long history, techniques applicable to optical spectro-polarised interferometry have so far been restricted to simple scattering regimes and symmetric geometries, leaving significant room for technical development.

This thesis aims to advance the use of optical-wavelength spectro-polarimetric interferometry as a method for probing the inner circumstellar environments of evolved stars. The present understanding of how stellar winds are driven is unable to explain how the winds of O-rich evolved stars are launched. One theorised solution is that starlight scattered by large and transparent dust grains may impart sufficient momentum to drive the grains outwards (Höfner 2008). Support of this theory relies on the direct detection and characterisation of circumstellar dust grains – a task for which optical-wavelength spectro-polarimetric interferometry is ideally configured. The primary scientific focus of this thesis is O-rich Red Supergiant star  $\mu$  Cephei, which we will use as a test case to investigate how optical spectro-polarised interferometry can constrain circumstellar geometry, chemistry and dust grain size, and whether this informs the plausibility of a scattering driven stellar wind.

This thesis pursues three core goals: (1) the development of novel astronomical instrumentation – the spectro-polarimetric interferometry mode of the VAMPIRES instrument, (2) the creation of a suite of modelling and image reconstruction tools for spectro-polarised interferometric observables which significantly advance upon existing techniques and (3) the application of these methods to  $\mu$  Cephei as a test case, to characterise its circumstellar environment. Together, these efforts aim to establish VAMPIRES as a leading and fully-supported instrument for performing optical wavelength spectro-polarimetric interferometry.

## Acknowledgements

First and foremost I owe enormous thanks to my supervisors, Peter and Barnaby – for their support, guidance, patience and the provision of all the opportunities that made this PhD not just possible but a memorable and rewarding experience. Most seriously, Peter – thank you for teaching me how to use a comma and what the appropriate size of a hot beverage receptacle is. Barnaby – thank you for the life changing experience of my first 7-Eleven Spam Musubi and for that time you took me to Ken’s House of Pancakes in Hilo when I was really jet-lagged. I think I finally understood polarisation in that diner after our 14th refill of coffee. I continue to learn so much from both of you.

Secondly to my unofficial supervisors, Ben, Eckhart, Chris and Callum. Ben – thank you for teaching me that “I want differentiable everything” is not an outrageous statement but is in fact good taste, and for your consistent 20 second response time on Slack to offer help whenever I was stuck. Eckhart – thank you for your methodical checking of my work and for the time you calmly made me a list of tenured academics who also hadn’t published anything until the final year of their PhD. Chris – thank you for always offering to be a rubber duck, for helping me fix everything from my Python environments to the fans in my computer and for all of your encouragement – particularly at the end. Callum – it was your mentorship in 2017 that inspired me to begin this PhD years later – and it was your generosity, guidance and constant faith in me that made it possible for me to complete it. None of you had any formal responsibility over my candidature yet somehow always made time to help me and I am so grateful.

Thank you to my research group who did everything from helping fix my computer to stocking their desks with snacks which at busy times constituted the only food I ate during the work day. Thank you for your support and friendship throughout many nights of observing and engineering and for always bringing me my favourite snacks back from fieldwork in the States – even at the expense of carrying telescope instrumentation in your hands (!!) on the plane home.

To international colleagues – in particular the SCExAO / VAMPIRES instrument and polarimetry team – thank you for all the time and effort you invested into this project, for your encouragement and technical support, and for letting me put so many NRM tasks into the engineering schedule. Working with you all has been one of the greatest highlights of my PhD.

Thank you to the School of Physics IT team for their tireless work and in particular for their support of my “use of every HPC within a 10km radius of the 2006 postcode” in the last few months of my candidature.

I have been so lucky to have had several mentors who have supported me behind the scenes to develop as a scientist and person. You know who you are, thank you.

---

To my wonderful friends, particularly those with whom I have completed my entire University education – thank you for your endless support and encouragement, even from the other side of the world at what I’m sure were very inconvenient hours for you. I hope someday soon we can all share dinner at the same table again.

Thank you to my extended family for asking me politely each Christmas if I’d finished yet – and when the answer was always no, for making me dinner and saying lovely things to me like “well we are really proud of you, whatever it is that you do” (Aunty Judy, 2024).

Thank you to my Grandparents for always making time to talk to me over the phone, for having me to stay so often and for always being able to give me perspective when I needed it. In particular to Grandad, for everything, especially for your final piece of advice that “the answer is in understanding these Stokes vectors” (you were right) and that most problems can be solved by sitting quietly with a pencil and going through things from first principles.

To my little brother – for being the only reason I still have a sense of humour and for reminding me how much I had to look forward to (in particular, seals).

Without question I owe the most thanks of all to my parents, for being the most fierce supporters of my education. I couldn’t have done any of this without you – I hope that the submission of this PhD brings you the same peace it brings me.



---

## Generative AI Declaration

Generative AI did not exist in its current form for most of this candidature and so was largely un-used. Towards the end of this candidature it was used for the following tasks:

- Assistance with debugging Python code – eg. implementation of Bayesian Optimisation package `optuna` in Chapter 7,
- Assistance with Latex code in the generation of `tikz` images – eg. Figure 5.1 and Figure 7.9,
- Assistance with writing complex Python plotting codes – eg. marginalised posterior plots in Appendix A.4.2,
- Generating the logo for the PIPPI Package (Chapter 4).

## Funding Acknowledgements

This research was conducted with the support of an Australian Government Research Training Program Scholarship.

## *Acknowledgements of Country*

This research took place in two countries.

We wish to recognise and acknowledge the Gadigal people of the Eora nation as the Traditional Owners of the land on which the University of Sydney's Camperdown campus stands.

We wish to recognise and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community, and are most fortunate to have the opportunity to conduct observations from this mountain.

# Publications

The work in this thesis is contained within or contributed to the following publications:

L Lilley et al. (July 2025b). “PIRATES – a machine-learning framework for polarized, interferometric image reconstruction”. In: *JATIS (accepted)* (Chapter 6)

L Lilley et al. (Dec. 2025a). “Beneath the shadows: characterising the circumstellar dust of RSG mu Cephei with VAMPIRES”. in: *A&A (In Prep)* (Chapter 7)

Miles Lucas et al. (Nov. 2024). “Visible-light High-contrast Imaging and Polarimetry with SCEXAO VAMPIRES”. in: *Publications of the Astronomical Society of the Pacific* 136.11, p. 114504. ISSN: 0004-6280 (Chapter 3)

Manxuan Zhang et al. (Oct. 2023). “Characterizing the instrumental polarization of SCEXAO VAMPIRES”. in: *Techniques and Instrumentation for Detection of Exoplanets XI*. ed. by Garreth J. Ruane. SPIE, p. 29. ISBN: 9781510665743 (Chapter 3)

As supervisor for the candidature upon which this thesis is based, I can confirm that the authorship attribution statements above are correct.

---

*Lucinda Lilley, October 14 2025*

---

*Peter Tuthill, October 14 2025*

# Contents

<b>Declaration</b>	<b>ii</b>
<b>Abstract</b>	<b>iii</b>
<b>Acknowledgements</b>	<b>iv</b>
<b>Publications</b>	<b>vii</b>
<b>List of Figures</b>	<b>ix</b>
<b>List of Tables</b>	<b>xii</b>
<b>1 An Introduction to Spectro-Polarimetric Interferometry</b>	<b>1</b>
1.1 Observing through Earth's turbulent atmosphere	1
1.2 Principles of Interferometry	3
1.3 High Contrast Imaging via Polarimetry	5
1.4 Spectro-Polarimetric Interferometry	7
1.5 Scientific Motivations	9
1.5.1 Mass Loss from Evolved Stars	9
1.5.2 $\mu$ Cephei	10
1.6 Thesis Overview	12
<b>2 Constraining the properties of <math>\mu</math> Cephei's circumstellar dust</b>	<b>14</b>
2.1 Methods for constraining expectations	14
2.2 The condensation sequence of O-rich evolved stars	15
2.3 Spectroscopic analysis of ISO-SWS data of $\mu$ Cephei	15
2.4 On the availability of optical constant data	17
2.5 Deriving the Hofner grain size criteria for $\mu$ Cephei	17
2.6 Thermodynamic constraint via radiative transfer simulation	20
2.7 Chapter Summary	25
<b>3 Developing the NRM mode of the VAMPIRES Instrument</b>	<b>26</b>
3.1 Instrumental Mueller Matrix Model of VAMPIRES	27
3.2 A facility data reduction pipeline for VAMPIRES + NRM	33
3.3 Polarimetric calibration limit of upgraded VAMPIRES	37
3.4 A 30% improvement to observing efficiency	38
3.5 Observing Campaign of $\mu$ Cephei	39
3.6 Reduced Data Products of $\mu$ Cephei	39
3.7 Diagnosis and correction of faults in archival VAMPIRES data	44
3.8 Chapter Summary	47

<b>4</b>	<b>PIPPI</b>	<b>48</b>
4.1	A brief introduction to Parametric Modelling . . . . .	48
4.2	PIPPI – Parametric Inference Package for Polarimetric Interferometry . . . . .	50
4.3	Part I – Astronomical Scenes . . . . .	51
4.4	Part II – Instrumental Mueller Matrix Model . . . . .	58
4.5	Part III – Simulating Non Redundant Masking . . . . .	59
4.6	Part IV – Graphical User Interface (GUI) . . . . .	59
4.7	Part V – Inference . . . . .	63
4.8	Chapter Summary . . . . .	64
<b>5</b>	<b>Application of PIPPI to <math>\mu</math> Cephei</b>	<b>65</b>
5.1	Constraining $\mu$ Cephei’s circumstellar geometry with Rayleigh scattering . . . . .	65
5.1.1	Results of Rayleigh scattering inference . . . . .	67
5.1.2	Analysis of Rayleigh scattering inference . . . . .	73
5.2	Constraining $\mu$ Cephei’s dust grain size and chemistry with Mie scattering . . . . .	77
5.2.1	Results of Mie scattering chemical inference . . . . .	79
5.2.2	Analysis of Mie scattering chemical inference . . . . .	80
5.3	A renewed perspective on single scattering models . . . . .	81
5.4	Chapter Summary . . . . .	85
<b>6</b>	<b>The PIRATES Image Reconstruction Algorithm</b>	<b>86</b>
6.1	Motivating the creation of PIRATES . . . . .	86
6.2	A brief introduction to Machine Learning . . . . .	87
6.3	Paper: PIRATES . . . . .	89
<b>7</b>	<b>Application of PIRATES to <math>\mu</math> Cephei</b>	<b>116</b>
7.1	Managing systematic errors in real data . . . . .	117
7.2	Reconstructing images of $\mu$ Cephei with PIRATES . . . . .	120
7.3	Parametrisation of $\mu$ Cephei’s circumstellar geometry . . . . .	127
7.4	Chemical Inference with Radiative Transfer . . . . .	127
7.5	Analysis of Radiative Transfer models . . . . .	133
7.6	Chapter Summary . . . . .	144
<b>8</b>	<b>Conclusions and Future Work</b>	<b>146</b>
8.1	Summary . . . . .	146
8.2	Suggestions for Future Work . . . . .	147
<b>A</b>	<b>Appendix Material</b>	<b>150</b>
A.1	Chapter 2 . . . . .	151
A.1.1	Complex Refractive Indices $m(\lambda)$ . . . . .	151
A.2	Chapter 3 . . . . .	154
A.3	Chapter 5 . . . . .	155
A.3.1	MCMC Hyper-parameters, Priors, Initial Conditions . . . . .	155
A.3.2	Corner Plots . . . . .	159
A.4	Chapter 8 . . . . .	179
A.4.1	EWC Scans . . . . .	179
A.4.2	Bayesian Optimisation Posteriors . . . . .	181
	<b>Bibliography</b>	<b>184</b>

# List of Figures

1.1	The relationship between angular size and coherence . . . . .	3
1.2	The VAMPIRES ‘g18’ aperture mask . . . . .	4
1.3	HR 4796A – PDI reveals disk inclination . . . . .	7
1.4	Simulated VAMPIRES observables for a circumstellar envelope . . . . .	8
1.5	Pre-existing images and models of $\mu$ Cephei . . . . .	11
2.1	ISO-SWS spectra of $\mu$ Cephei, EP Aqr and Betelgeuse . . . . .	16
2.2	Deriving the Höfner grain size criteria for $\mu$ Cephei . . . . .	19
2.3	Radiative transfer for deriving condensation radii – Part 1 . . . . .	21
2.4	Radiative transfer for deriving condensation radii – Part 2 . . . . .	22
2.5	Radiative transfer for deriving condensation radii – Part 3 . . . . .	23
2.6	Condensation radii as a function of grain size . . . . .	24
3.1	Schematic of Subaru telescope and VAMPIRES Instrument . . . . .	27
3.2	Augmenting models using the instrumental Mueller Matrix . . . . .	32
3.3	Flowchart of VAMPIRES NRM data reduction pipeline . . . . .	34
3.4	Alignment of aperture mask to power spectra . . . . .	35
3.5	Measuring the polarimetric calibration limit of upgraded VAMPIRES . . . . .	37
3.6	A 30% increase in observing efficiency . . . . .	38
3.7	Polarised differential visibilities for datasets 1-4 (2017-2020) . . . . .	40
3.8	Polarised differential closure phases for datasets 1-4 (2017-2020) . . . . .	41
3.9	Polarised differential visibilities for dataset 6 (2023) . . . . .	42
3.10	Polarised differential closure phases for dataset 6 (2023) . . . . .	43
3.11	A fault in archival 2018 VAMPIRES NRM data . . . . .	44
3.12	PDI images of AB Aurigae . . . . .	45
3.13	Fitting to polarimetric mis-calibrations in archival VAMPIRES data . . . . .	46
3.14	Applying polarimetric corrections to 2018 data . . . . .	47
4.1	PIPPI – Parametric Inference Package for Polarimetric Interferometry . . . . .	52
4.2	PIPPI coordinate system . . . . .	53
4.3	Dust density distributions generated by PIPPI . . . . .	54
4.4	Deriving vectors for computing scattering grids . . . . .	55
4.5	Deriving the local dust grain coordinate system . . . . .	56
4.6	Generating polarised images from PIPPI . . . . .	57
4.7	Output from PIPPI Part (II) . . . . .	59
4.8	Simulated observation with VAMPIRES and the ‘g18’ mask . . . . .	60
4.9	The impact of a circumstellar dust envelope on Stokes I visibilities . . . . .	61
4.10	How asymmetry distorts polarised visibilities . . . . .	62
4.11	The location of Stokes Q ( $\frac{V_H}{V_V} = 1$ crossing) with rotation . . . . .	63
4.12	Final outputs of PIPPI . . . . .	64
5.1	Geometric parameters fitted with PIPPI . . . . .	67
5.2	Representative fitted models from Table 5.1 . . . . .	74
5.3	Images and models of $\mu$ Cephei’s inner circumstellar environment . . . . .	75

5.4	Polarised scattering cross section as a function of dust grains size	78
5.5	Chemical and dust grain size inference for $\mu$ Cephei (2023) (I)	79
5.6	Chemical and dust grain size inference for $\mu$ Cephei (2023) (II)	80
5.7	Maximum obtainable polarised visibility amplitudes for thin shells (I)	82
5.8	Maximum obtainable polarised visibility amplitudes for shells (II)	83
5.9	Models from Norris et al. 2012	84
6.1	A simple fully connected (dense) neural network	88
6.2	The CNN convolution operation	88
6.3	A simple schematic of a typical CNN	89
7.1	PIRATES reconstructions for a test clumpy circumstellar environment	118
7.2	PIRATES reconstructions for a test inclined disk	119
7.3	PIRATES image reconstructions for VAMPIRES data of $\mu$ Cephei (I)	121
7.4	PIRATES image reconstructions for VAMPIRES data of $\mu$ Cephei (II)	122
7.5	The geometric consistency of PIPPI and PIRATES	123
7.6	$\mu$ Cephei's variable polarisation (V band, 550 nm)	125
7.7	Polarisation strength $P\%$ and angle $\phi_P^\circ$ from PIRATES images	125
7.8	Polarisation fields of PIRATES reconstructed images	126
7.9	Overview of parametrisation fit in radiative transfer code	128
7.10	Density distribution of best fit model of $\mu$ Cephei	129
7.11	Images of the radiative transfer model fit to $\mu$ Cephei 2023 data	131
7.12	Correlation plot of observables – radiative transfer model fit to $\mu$ Cephei	132
7.13	Optical Depth of best fit model – Amorphous Enstatite	133
7.14	$P_{survival}$ as a function of $k$	134
7.15	Simulated thin shell of Amorphous Enstatite around $\mu$ Cephei	136
7.16	The cost of obtaining $\mu$ Cephei's polarised visibility amplitudes	136
7.17	The effect of optical depth on the wavelength dependence of polarisation	138
7.18	Relationship between grain size and polarised scattering anisotropy	140
7.19	Relationship between grain size and scattering cross section	140
7.20	Fitted dust grain distributions scaled by polarisation	141
7.21	Narrow Silicon-carbon grain distributions	142
A.1	Optical constants for dust species used in chemical inference - Part 2	151
A.2	Optical constants for dust species used in chemical inference - Part 3	152
A.3	Optical constants for dust species used in chemical inference - Part 1	153
A.4	Corner plots for Model Class A — thin spherical shell.	159
A.5	Corner plots for Model Class B — thick spherical shell.	160
A.6	Corner plots for Model Class C — power-law spherical shell.	161
A.7	Corner Plots for Model Class D — Two thin spherical shells.	162
A.8	Corner Plots for Model Class E — Thin elliptical shell.	163
A.9	Corner Plots for Model Class F — Thick elliptical shell.	164
A.10	Corner Plots for Model Class G — Power-law elliptical shell.	165
A.11	Corner Plots for Model Class H — Thin offset spherical shell.	166
A.12	Corner Plots for Model Class I — Thin offset elliptical shell.	167
A.13	Corner Plots for Model Class J — Thin spherical shell with blob.	168
A.14	Corner Plots for Model Class K — Thin elliptical shell with blob.	169
A.15	Corner Plots for Model Class L — Thick spherical shell with blob.	170
A.16	Corner Plots for Model Class M — Thick elliptical shell with blob.	171
A.17	Corner Plots for Model Class N — Power-law spherical shell with blob.	172
A.18	Corner Plots for Model Class O — Power-law elliptical shell with blob.	173
A.19	Corner Plots for Model Class P — Offset power-law spherical shell + blob.	174
A.20	Corner Plots for Model Class Q — Offset power-law elliptical shell + blob.	175
A.21	Corner Plots for Model Class R — Enhanced power-law spherical shell + blob.	176
A.22	Corner Plots for Model Class S — Enhanced power-law elliptical shell + blob.	177
A.23	$\mu$ Cephei EWC scans for archival datasets 1-4.	179

---

A.24 $\mu$ Cephei EWC scans for multi-band dataset 6. . . . .	180
A.25 Bayesian Optimisation Posteriors: Amorphous Enstatite – Model 2. . . . .	181
A.26 Bayesian Optimisation Posteriors: Amorphous Enstatite – Model 1. . . . .	181
A.27 Bayesian Optimisation Posteriors: Amorphous Spinel – Model 1. . . . .	181
A.28 Bayesian Optimisation Posteriors: Amorphous Corundum – Model 1. . . . .	182
A.29 Bayesian Optimisation Posteriors: Amorphous Forsterite – Model 1. . . . .	182
A.30 Bayesian Optimisation Posteriors: Pyroxene 95 – Model 1. . . . .	182
A.31 Bayesian Optimisation Posteriors: Amorphous Silica – Model 1. . . . .	183
A.32 Bayesian Optimisation Posteriors: Crystalline Enstatite – Model 1. . . . .	183
A.33 Bayesian Optimisation Posteriors: Crystalline Forsterite – Model 1. . . . .	183

# List of Tables

2.1	Summary of candidate dust species properties . . . . .	25
3.1	Empirically fitted Mueller matrix coefficients for original VAMPIRES . . . . .	31
3.2	Non-ideal coefficients for Mueller matrix model of upgraded VAMPIRES . . . . .	31
3.3	Outputs of the <code>vampires_dpp</code> NRM pipeline for the g18 mask . . . . .	36
3.4	Summary of our observing campaign of $\mu$ Cephei . . . . .	39
4.1	Outline of PIPPI's core components . . . . .	51
5.1	Overview of parametric models fitted with PIPPI . . . . .	66
5.2	Summary of Rayleigh scattering fits (I) . . . . .	68
5.3	Summary of Rayleigh scattering fits (II) . . . . .	69
5.4	Summary of Rayleigh scattering fits (III) . . . . .	70
5.5	Summary of Rayleigh scattering fits (IV) . . . . .	71
5.6	Summary of Rayleigh scattering fits (V) . . . . .	72
7.1	The improvement yielded by PIRATES over PIPPI . . . . .	120
7.2	Bayesian Optimisation Parameters . . . . .	129
7.3	Bayesian Optimisation Results . . . . .	130
A.1	MCMC steps, walkers and burn in discard . . . . .	155
A.2	Model Priors and Initial Conditions – Rayleigh Scattering with PIPPI (I) . . . . .	156
A.3	Model Priors and Initial Conditions – Rayleigh Scattering with PIPPI (II) . . . . .	157

# Chapter 1

## An Introduction to Spectro-Polarimetric Interferometry

*“There is no such thing as a diffraction limit, only noise.”*

---

Peter G. Tuthill

### 1.1 Observing through Earth’s turbulent atmosphere

Study of the cosmos has always fascinated human kind – the practice of astronomy is woven into historical records across cultures and continents (Ruggles 2015; Selin 2000). Our knowledge of ancient astronomical practices is rapidly increasing as greater attention is given to astronomical observations embedded in ancient texts, monuments and the dominant practice of oral history records (Norris et al. 2015b). Our existing understanding indicates that study of the sky was born out of both necessity and curiosity (Kelley et al. 2005); as a means to understand weather patterns, agriculture and to navigate, but also from a deep desire to understand the vast depth surrounding us – which existed long before humanity and will persist long after life on Earth has ceased.

The observation that stars ‘twinkle’ (defined in modern astronomy as *scintillation*) is widely recorded as a method for understanding weather variability (Hamacher et al. 2019; Kerrison et al. 2024). Scintillation is caused by Earth’s turbulent atmosphere, which spatio-temporally randomises the phase of incident starlight at scales between  $\sim 1$  mm - 10 m (Kolmogorov 1941). At the optical wavelengths sensed by the human eye, these spatial and temporal scales are  $r_0 \sim 10$  cm (the Fried parameter) and  $\tau_0 \sim 10$  ms (the atmospheric coherence time) (Labeyrie et al. 2006a). The interaction of a wavefront with the atmosphere results in modulation to the observed amplitude (or intensity) of the object – this causes unresolved stars to appear to ‘jitter’ about their true position in the sky. For sources such as planets that occupy a larger angular size, the effects of this scintillation can be dramatically reduced (leading to the well-known tip for distinguishing stars from planets), though the quality of imagery recovered remains badly affected.

The same atmospheric variability that causes scintillation (amplitude error) also causes an effect called *seeing* (phase error) that blurs the wavefront containing information on the astronomical scene, impacting the performance of ground based telescopes with severity proportional to observing wavelength. The commonly understood resolution of a ground based telescope is usually quantified by Rayleigh’s criteria, where the diffraction limit  $\theta$  of a telescope of diameter  $D$  and observing wavelength  $\lambda$  defines the smallest resolvable angular separation of two point sources (Equation 1.1).

$$\theta \approx 1.22 \frac{\lambda}{D} \tag{1.1}$$

In principle, finer detail can be resolved by decreasing the observing wavelength or increasing the telescope diameter. However, in reality the diffraction limited resolution of large optical telescopes (with diameter above 10 cm) is not attainable without intervention as the atmosphere lowers the inherent resolution of

ground based telescopes to what is known as *seeing limited* performance. At optical wavelengths the effects of seeing typically limit a telescope’s resolution to  $\sim \mathcal{O}(1'')$  (Labeyrie et al. 2006a). Given this is our starting point, performing high resolution optical astronomy is a remarkably challenging pursuit.

One way to overcome atmospheric seeing is to put telescopes in space, a solution with significant interest and increasing investment (National Academies of Sciences 2023; Kouveliotou et al. 2014). Existing observatories like the Hubble Space Telescope (HST) (Lauer et al. 1991) and the James Web Space Telescope (JWST) (Gardner et al. 2023) capture remarkably high resolution images of astrophysical scenes. Planned space-based observatories like the Habitable Worlds Observatory (HWO) and the PLAnetary Transits and Oscillations of stars (PLATO) mission will transform space-domain detection and characterisation of exoplanets (Stark et al. 2024; Rauer et al. 2014). These missions carry high risk though promise great reward, entailing high expense and unique challenges in troubleshooting and upgrade. Complementary to space missions are ground based techniques designed to mitigate seeing, which include specialised instrument design, active wavefront control and data post-processing techniques. Within the present thesis, two of these techniques are used – adaptive optics and single telescope interferometry.

Adaptive optics (AO) systems entail control loops which use wavefront sensing to measure perturbations in the phase of an incident wave-front, correcting for phase errors by varying instrumental optical path lengths via a deformable mirror (DM) (Guyon 2018; Hampson et al. 2021). A well tuned AO system can drive a seeing limited telescope towards a diffraction limited performance (Ahn et al. 2021), although the diffraction limit still represents a practical ceiling for the resolution of classical imaging systems (Max Born 1975). Within this work we use adaptive optics as an end-user, benefiting from its improvement to the resolution of interferometric measurements without delving into its complexities.

The focus of this thesis is *interferometry* – a technique responsible for some of the highest resolution measurements made in modern astronomy (Eisenhauer et al. 2023). Instead of recording images of an astrophysical scene, light from the scene is interfered to produce an interferogram. Most famously, long-baseline interferometry interferes recombined light from multiple separated telescopes to obtain higher resolution by virtue of creating a greater effective telescope diameter – this separation between telescope diameters is called a ‘baseline length’ (Hofmann 2005; Ten Brummelaar 1999). Notable examples of long-baseline interferometers include the Very Large Telescope Interferometer (VLTI) (Glindemann et al. 2000), the Center for High Angular Resolution Astronomy (CHARA) (Ten Brummelaar 1999), the Atacama Large Millimeter/submillimeter Array (ALMA) (Wootten et al. 2009) and the Event Horizon Telescope (EHT) (Akiyama et al. 2019a). These facilities perform interferometry across the electromagnetic spectrum, permitting investigation of astrophysical phenomena ranging from the surfaces of evolved stars (Witkowski et al. 2016) to the environments around black holes (Akiyama et al. 2019a).

However, it is also possible to perform *single telescope* interferometry by placing an opaque mask called a Sparse Aperture Mask (SAM) or Non Redundant aperture Mask (NRM) into the pupil plane of a single telescope. The advantage of doing so is counter-intuitive at first glance, as adding an aperture mask neither changes the observing wavelength nor increases the telescope’s diameter (Equation 1.1). As we will subsequently explore, NRM’s advantage is that it removes spatial *redundancy noise* from the measured signal – which allows us to drive the resolution of a single telescope below the diffraction limit. NRM is so successful at removing redundancy noise that even space based observatories like the James Webb Space Telescope (JWST) carry an aperture mask, as it can provide an observational advantage for a subset of targets of appropriate size and sufficient brightness to tolerate the significant loss incurred by the low throughput of an aperture mask (Sivaramakrishnan et al. 2014). In contrast to the expense of long baseline interferometry, NRM requires only the addition of an opaque mask to the pupil plane of an existing instrument (Lacour et al. 2011) and does not suffer from signal loss along the long beam paths required for recombining signal in long baseline interferometers (Kawaguchi 1983).

NRM was first implemented in 1851 in Fizeau’s earliest demonstrations of interferometry and was then used by Michelson in the 1920s to measure stellar diameters (Glindemann 2012). Subsequent interferometric experimentation quickly sought to combine light from multiple telescopes to obtain higher resolutions (Glindemann 2012). NRM’s modern resurgence began with a series of experiments performed at the Keck Observatory in the late 1990s (Tuthill et al. 2000b; Woodruff et al. 2008; Monnier et al. 2007a), permitting study of the dusty circumstellar environments of evolved stars. With maturation of

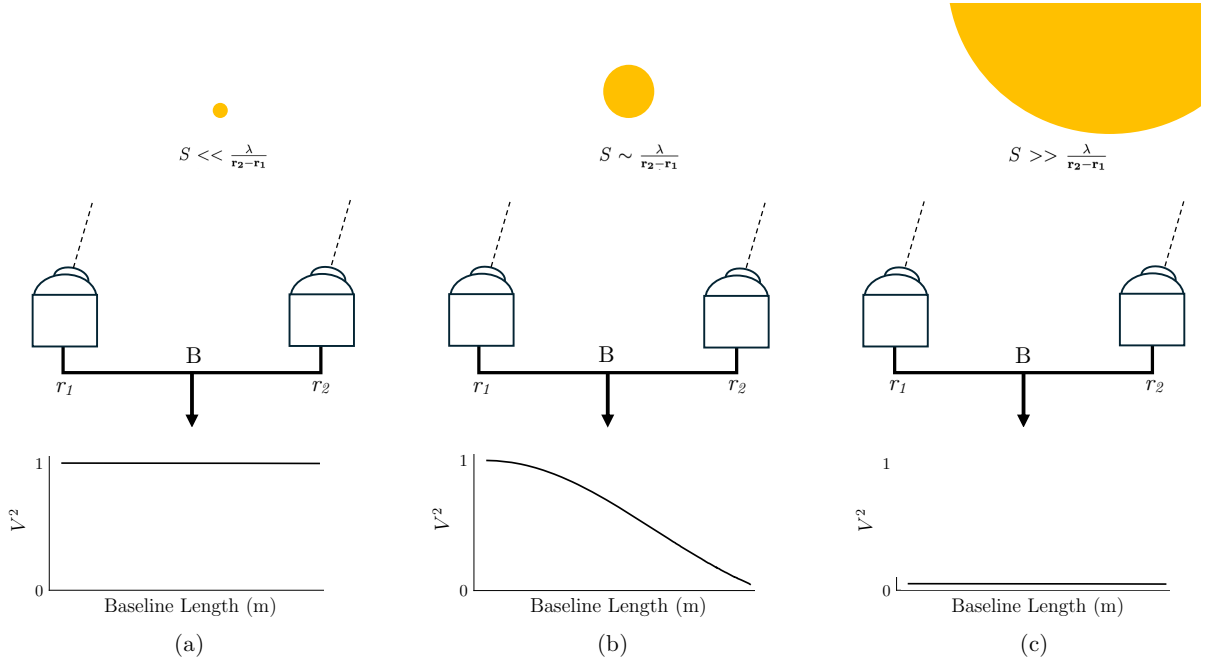
the technique it spread to other ground-based facilities including the VLT where it has been used on NACO (Tuthill et al. 2010b; Lacour, Sylvestre, Tuthill, Peter 2023), SPHERE (Cheetham et al. 2016), and at the SUBARU Telescope on the VAMPIRES Instrument (Norris et al. 2015a; Lucas et al. 2024). The development and use of the NRM mode of VAMPIRES (VAMPIRES + NRM) is the subject of this thesis.

## 1.2 Principles of Interferometry

At the core of interferometry is the concept of *complex coherence*  $\gamma$  – the *time averaged* correlation of two electric fields  $E_1(\mathbf{r}_1, t, \lambda \pm \Delta\lambda)$  and  $E_2(\mathbf{r}_2, t, \lambda \pm \Delta\lambda)$  emitted from spatial locations  $\mathbf{r}_1$  and  $\mathbf{r}_2$  at time  $t$  and wavelength  $\lambda \pm \Delta\lambda$  (Equation 1.2). For a baseline vector  $\mathbf{r}_2 - \mathbf{r}_1$ , an unresolved object with size  $S \ll \frac{\lambda}{\mathbf{r}_2 - \mathbf{r}_1}$  has full coherence ( $\gamma = 1$ ) – all light from the object is degenerate and has perfect time averaged correlation (Figure 1.1, (a)). Conversely, an object which is over-resolved on  $\mathbf{r}_2 - \mathbf{r}_1$  with  $S \gg \frac{\lambda}{\mathbf{r}_2 - \mathbf{r}_1}$  is incoherent ( $\gamma = 0$ ) (Figure 1.1, (c)) – there is no time averaged correlation between the two signals and thus no measurable interference fringes. Between these two extremes is an object that is partially resolved on baseline  $\mathbf{r}_2 - \mathbf{r}_1$  with  $S \sim \frac{\lambda}{\mathbf{r}_2 - \mathbf{r}_1}$ , for which the measured electric fields will have a coherence of between  $0 < \gamma < 1$  (Figure 1.1, (b)).

$$\gamma(\mathbf{r}_1, \mathbf{r}_2, \lambda \pm \Delta\lambda) = \frac{\langle E_1(\mathbf{r}_1, t, \lambda \pm \Delta\lambda) E_2^*(\mathbf{r}_2, t, \lambda \pm \Delta\lambda) \rangle}{[\langle |E_1(\mathbf{r}_1, t, \lambda \pm \Delta\lambda)|^2 \rangle \langle |E_2(\mathbf{r}_2, t, \lambda \pm \Delta\lambda)|^2 \rangle]^{\frac{1}{2}}} \quad (1.2)$$

As the baseline length and angle defined by  $\mathbf{r}_2 - \mathbf{r}_1$  vary, the complex coherence  $\gamma$  varies between 0 and 1 as a function of the size and orientation of the underlying astronomical scene (Figure 1.1 (b)). As there is no sharp transition from ‘resolved’ to ‘unresolved’, partially resolved objects with features smaller than the diffraction limit may be characterised. The fidelity of these measurements depends predominantly on the limitations imposed by noise and calibration (Ireland 2013). As a result of these mathematical properties, for an otherwise identical apparatus and a partially resolved astronomical target of sufficient brightness, interferometry can deliver higher angular resolution measurements than a direct image.



**Figure 1.1:** Interferometric measurement of objects of angular size  $S$  at wavelength  $\lambda \pm \Delta\lambda$ . (a) Unresolved objects (where  $S \ll \lambda/B$ ) exhibit full coherence and visibility measurements are  $\sim 1$ . (b) Partially resolved objects (where  $S \sim \lambda/B$ ) exhibit partial coherence and have visibilities between 0 and 1, which are a function of the size and orientation of the object. (c) Completely resolved objects (where  $S \gg \lambda/B$ ) will not exhibit time averaged coherence and have visibilities of zero.

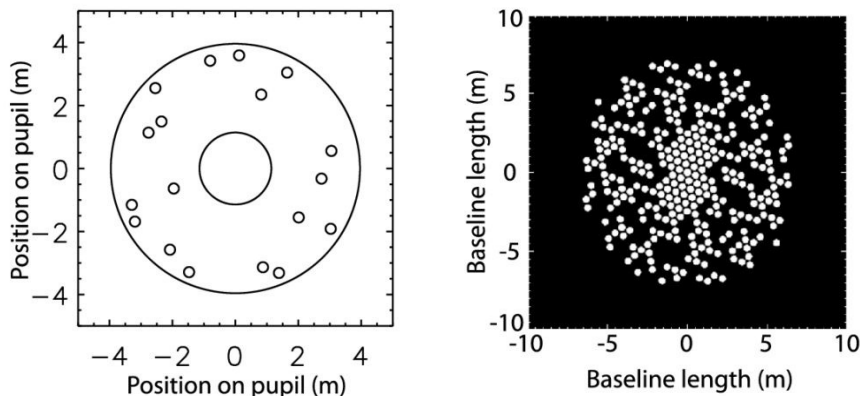
In practice, the complex coherence function is measured via a proxy for the contrast of the interference fringes, the *squared visibilities*  $V^2$ , as a function of  $\mathbf{r}_2 - \mathbf{r}_1$  (Figure 1.1 (b)). The relationship between the squared visibilities and the complex coherence function is given by Equation 1.3, where  $A_1$  and  $A_2$  are the amplitudes of the electric fields  $E_1(\mathbf{r}_1, t, \lambda \pm \Delta\lambda)$  and  $E_2(\mathbf{r}_2, t, \lambda \pm \Delta\lambda)$ , and  $I_{\max}$  and  $I_{\min}$  are the maximum and minimum fringe visibilities generated by the two electric fields (Greenbaum et al. 2014).

$$V^2(\mathbf{r}_1, \mathbf{r}_2, \lambda \pm \Delta\lambda) = \left( \frac{2A_1A_2}{A_1^2 + A_2^2} |\gamma(\mathbf{r}_1, \mathbf{r}_2, \lambda \pm \Delta\lambda)| \right)^2 = \left( \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}} \right)^2 \quad (1.3)$$

The relationship between the complex coherence function  $\gamma$  and the underlying source structure is given formally by the Van Cittert Zernike Theorem (Equation 1.4) – “for a monochromatic, incoherent source, the Fourier transform  $\mathcal{F}$  of the complex spatial coherence function is the angular intensity distribution of the source  $I(\mathbf{r})$ ” (Labeyrie et al. 2006b). Measuring the squared visibilities as a function of baseline length and angle is thus equivalent to probing the two dimensional Fourier spectrum of the imaged astronomical scene, where each baseline vector measures a single Fourier component.

$$I(\mathbf{r}_1, \mathbf{r}_2, \lambda \pm \Delta\lambda) = \mathcal{F}(\gamma(\mathbf{r}_1, \mathbf{r}_2, \lambda \pm \Delta\lambda)) = \frac{1}{\lambda} \int \int \gamma(\mathbf{r}_1, \mathbf{r}_2, \lambda \pm \Delta\lambda) e^{i2\pi(\mathbf{r}_1 - \mathbf{r}_2) \cdot \mathbf{u}} d^2\mathbf{r}, \quad (1.4)$$

The design of aperture masks is driven by observational intention – different numbers, sizes and orientations of sub-apertures have different throughput and Fourier coverage which must be matched to the astronomical target. The aperture mask used within the present work is the ‘g18’ – an 18 hole mask with 153 independent baseline vectors and 3% throughput (Golay 1971; Norris et al. 2015a). The ‘g18’ mask design is displayed in Figure 1.2 (left) alongside its corresponding power spectrum (right) – the Fourier coverage of the g18 mask is reasonably complete.



**Figure 1.2:** The VAMPIRES ‘g18’ 18-hole aperture mask design (left), and the associated power spectrum (right) – from which interferometric observables (visibilities and closure phases) are measured. Figure is reproduced from (Norris et al. 2015a).

The phase-randomisation induced by atmospheric turbulence is an ever-present contribution. It can be modelled to first order as a piston term above each sub-aperture, which manifests as varying spatial displacement of each pair of interference fringes on the detector (Tuthill et al. 2000a). Although the fringe amplitudes are in principle mathematically immune to atmospheric phase noise, the rapid and random displacement of all pairs of fringes leads to a noisy incoherent sum on the detector, limiting our ability to capitalise on this theoretical immunity. This fringe displacement can be reduced with a well tuned AO system which will stabilise the fringes on the detector thereby improving the signal to noise of the measured interferometric quantities (Tuthill et al. 2006).

The sub-diffraction limited resolution promised by interferometry is most powerful when interferometry is performed *non redundantly*. When two identical Fourier components are measured twice by the same interferometric array, they sum incoherently on the detector and at the same location in the  $(u,v)$  plane

(power spectrum). Earth’s atmosphere has both temporal ( $\tau_0$ ) and spatial ( $r_0$ ) scales to its phase corruption. If a Fourier component is measured twice by an interferometric array then each pair of apertures will experience a different phase error on account of differing location in the pupil plane. When their fringes are summed incoherently, each pair of sub-apertures contributes a unique random error term. The resulting noise is called *redundancy noise* – by measuring the same baseline twice (ie, a redundant measurement), the amount of noise in our measurement is also increased. Redundancy noise grows as  $\sqrt{N}$  for N redundant measurements of the same Fourier component. As such, signal to noise is maximised with an array design that probes each Fourier component only a single time in a ‘non-redundant’ fashion (Tuthill 1994).

The second significant benefit of interferometry is that it allows for some degree of phase retrieval. Whilst the observed phase of incident light is corrupted by the atmosphere, the conservative nature of the phase field can be used to remove phase error by constructing phase quantities known as *closure phases* (Readhead et al. 1988). Each triplet of telescopes (or sub apertures) each with measured phase  $\phi_n$  and error  $\epsilon_n$  (Equations 1.5-1.7) can be used to construct a closure phase  $\phi_{CP}$ , a function of the three measured phases which is immune to the individual error terms (Equation 1.8).

$$\begin{aligned}\Phi_{1-2} &= \phi_1 + \epsilon_1 - \phi_2 - \epsilon_2 \\ &= \phi_{1-2} + \epsilon_1 - \epsilon_2\end{aligned}\tag{1.5}$$

$$\begin{aligned}\Phi_{2-3} &= \phi_2 + \epsilon_2 - \phi_3 - \epsilon_3 \\ &= \phi_{2-3} + \epsilon_2 - \epsilon_3\end{aligned}\tag{1.6}$$

$$\begin{aligned}\Phi_{3-1} &= \phi_3 + \epsilon_3 - \phi_1 - \epsilon_1 \\ &= \phi_{3-1} + \epsilon_3 - \epsilon_1\end{aligned}\tag{1.7}$$

$$\phi_{CP} = \Phi_{1-2} + \Phi_{2-3} + \Phi_{3-1} = \phi_{1-2} + \phi_{2-3} + \phi_{3-1}\tag{1.8}$$

Several assumptions are made in the construction of closure phases, including the perfect conservation of the phase field and the uniformity of phase across each sub-aperture. However even with the approximate satisfaction of these conditions, closure phases provide a very effective method for phase recovery. For N non-redundant sub-apertures there are  $\binom{N}{2}$  independent Fourier phases and  $\binom{N-1}{2}$  independent closure phases. The proportion of recovered phase information is the ratio between independent closure phases and Fourier phases – simple combinatorics can show that by the time  $N = 50$ , 96 % of the phase information has been recovered (Monnier 2002). For the g18 mask used in this thesis we may recover as much as 89 %. Closure phases provide critical information on the edges, orientations, and asymmetry of the underlying source structure, breaking the symmetric degeneracy present within visibility measurements (Monnier 2003). Closure phases are also immune to image translations and to varying phase centre, making them a robust source of information on asymmetry (Monnier 2003; Readhead et al. 1988).

### 1.3 High Contrast Imaging via Polarimetry

A central challenge in observational astronomy is the ability to make measurements within high-contrast environments, to detect and characterise faint structures whose signals are dwarfed by a bright counterpart. Two of the most notable examples of faint structures in high contrast environments are: circumstellar dust shells and exoplanetary signatures (Bjorkman 2000; Hough 2006; Kolokolova et al. 2015; Claudi et al. 2025; Zurlo 2024). A powerful strategy for observing these faint components is *polarimetry*. The polarisation state of transverse waves describes the orientation of the oscillating electric field component with respect to the direction of propagation (Max Born 1975). The polarisation state of light is modified by physical processes where symmetry is broken, such as: reflection, transmission, dichroism, double refraction and scattering (Jenkins 1981). Each process imparts a unique signature into the polarisation state, providing information on the physical process that is unavailable within an intensity measurement.

A convenient mathematical framework for describing partially polarised light is the Stokes formalism, where the polarisation state is given by a four vector  $\mathbf{S} = (S_0, S_1, S_2, S_3) = (I, Q, U, V)$  (Stokes 1851)

(Equation 1.9). Each component of the Stokes vector is derived from measurement of orthogonal electric field components  $E_x$  and  $E_y$  which are separated by an angle  $\delta$  in the plane perpendicular to the wavefront's propagation (Snik et al. 2013). Stokes I defines the total intensity of incident light, whilst Stokes Q and U define the orthogonal linear polarisation states. Stokes V describes circularly polarised light which is unmeasured in the present work. Equation 1.10 is an inequality when light is partially polarised (most astronomical contexts, including the present). From a Stokes vector, the polarisation vector  $\vec{P}$  may be constructed – with length  $P$  (Equation 1.11) and angle  $\phi_P$  (Equation 1.12) (Trippe 2014).

$$\mathbf{S} = \begin{bmatrix} E_x^2 + E_y^2 \\ E_x^2 - E_y^2 \\ 2E_x E_y \cos \delta \\ 2E_x E_y \sin \delta \end{bmatrix} = \begin{bmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{bmatrix} = \begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} \quad (1.9)$$

$$S_0^2 \geq S_1^2 + S_2^2 + S_3^2 \quad (1.10)$$


$$P = \sqrt{S_1^2 + S_2^2} = \sqrt{Q^2 + U^2} \quad (1.11)$$

$$\phi_P = \frac{1}{2} \arctan \frac{U}{Q} \quad (1.12)$$

Polarimetry is a particularly effective high contrast imaging technique because the polarisation properties of high contrast structures are usually very different (Snik et al. 2018). At optical wavelengths, light scattered off circumstellar dust grains is highly polarised compared to light originating from the stellar surface (Snik et al. 2013). This permits polarimetry to reveal dust structures close to the stellar surface that would otherwise be dominated by the stellar signal in an intensity measurement (Gledhill 1991; Canovas et al. 2011; Holstein 2021). Polarimetry has been used to study a wide variety of dusty astronomical objects, including: protoplanetary disks (Hashimoto et al. 2012), debris discs (Vandepoort et al. 2019), planetary nebulae (Gledhill et al. 2001), circumbinary disks (Andrych et al. 2023; Andrych et al. 2024) and the circumstellar environments of evolved stars (González Delgado et al. 2003; Haubois et al. 2019) which will be the focus of the present work.

In addition to providing reach into high contrast environments, polarimetric images can reveal three-dimensional dust structure if the dust exhibits *anisotropic* scattering. In these cases, the magnitude and orientation of the emergent polarisation vector depends on the scattering angle and physical properties of the dust grains, enabling inference of the relative distances of different dusty components along the line of sight. An example of this is the use of polarimetry to probe the inclination angles of circumstellar disks (Perrin et al. 2009a; Avenhaus et al. 2018a; Wahhaj et al. 2024; Quanz et al. 2011; Perrin et al. 2015) – Figure 1.3 displays a striking increase in information from polarimetry for disk HR 4796A.

When imaged in Stokes I (ADI, Figure 1.3 left), HR 4796A's dusty ring is partially occluded by speckle noise around a coronagraph and the ring's 3D orientation is unclear. When imaged in polarised light (Figure 1.3 right), the polarised image suffers less from speckle noise and the image reveals the disk's inclination – ie, which face of the disk is closest to us (the brighter limb). This brightened forwards sector of the ring is the result of *forward scattering* – light scattered off large ( $\sim \lambda$ ) dust grains at large scattering angles (closer to us at  $\sim 180^\circ$ ) exhibit stronger polarisation than dust grains at small scattering angles (furthest away from us at  $\sim 0^\circ$ ) which exhibit *back scattering*. Polarimetry is made even more powerful when performed over multiple wavelengths – a technique known as *spectro-polarimetry*. Information on the dust, such as its chemical composition and grain size, are degenerate with the total amount of scattered light at a single wavelength. Spectro-polarimetry enables these properties to be disentangled as chemistry and grain size have unique wavelength dependent signatures (Honda et al. 2022; Hartquist et al. 2006; Hunziker et al. 2021; Dykes et al. 2024; Haubois et al. 2019; Norris et al. 2012).




**Figure 1.3:** HR 4796A in ADI and PDI – obtained via the Gemini Planet Imager (GPI) in the K1 band ( $1.90\text{--}2.19\ \mu\text{m}$ ). Figure is reproduced from Perrin et al. 2015. Whilst the disk’s shape is evident in ADI (left), in PDI (right) the inclination of the disk is revealed via strong anisotropic forward scattering induced by large dust grains (large meaning of size comparable to the observing wavelength).

## 1.4 Spectro-Polarimetric Interferometry

A innovative observational technique is forged by combining interferometry and spectro-polarimetry – the latter yielding multi-wavelength characterisation of high contrast regions while the former enables us to probe them with sub-diffraction limited resolution. When performed at optical wavelengths where polarisation signatures induced by dust scattered starlight are strong, *spectro-polarimetric interferometry* provides a sub-diffraction limited spatio-spectral probe of innermost circumstellar environments (Norris et al. 2015a; Norris 2015). Within this thesis we will perform spectro-polarimetric interferometry using the VAMPIRES + NRM instrument (Norris et al. 2015a; Lucas et al. 2024), described in detail in Chapter 3. Here we introduce the spectro-polarimetric interferometry data product (depicted in Figure 1.4) and summarise its history and scientific potential. Figure 1.4 depicts a simulated spherical circumstellar dust shell as observed with each mode of VAMPIRES. VAMPIRES produces both a polarised differential imaging (PDI) data product and a polarimetric differential interferometry data product (PDI + NRM), both are displayed for comparison. The polarimetric data product is described as *differential* because its computation requires the subtraction (for PDI) or division (for PDI + NRM) of orthogonal linear polarisation states H/V (Stokes Q) and H45/V45 (Stokes U). The polarised differential visibilities of a spherical shell (like Figure 1.4) have a characteristic sinusoidal and cosinusoidal appearance. In conventional measurement of high contrast circumstellar environments, intensity (Stokes I) is dominated by the stellar signal (column 1), whilst in polarimetric measurements the dust dominates (columns 2-3).

Optical and infrared wavelength spectro-polarimetric interferometry data is relatively novel. The only facilities to currently perform such measurements are SPHERE/ZIMPOL (Cheetham et al. 2016) and VAMPIRES/SCEXAO (Norris et al. 2015a; Lucas et al. 2024). Historically it was performed on NACO/VLT (Tuthill et al. 2010a; Norris et al. 2012; Girard 2010) and the Sydney University Stellar Interferometer (SUSI) (Ireland et al. 2005), both of which are now decommissioned. Interest in developing polarimetric modes of existing interferometers is growing due to the great scientific potential of the data type. Work on instrumental characterisation (the first step to performing polarimetric measurements) is currently underway at the CHARA array (Setterholm et al. 2020). As such, development of methods for extracting, reducing and modelling spectro-polarimetric interferometry data are of wide interest to the broader community as the technique expands and develops maturity. At the time of writing this thesis there is no published VAMPIRES + NRM data beyond instrumental polarisation calibration measurements (Norris

et al. 2015a). There are only three existing studies that perform and model polarimetric interferometry at optical and infrared wavelengths, all of which use decommissioned instruments that pre-date VAMPIRES.



**Figure 1.4:** Simulated VAMPIRES observables for a thick spherically symmetric circumstellar envelope: polarimetric differential images (PDI) of Stokes I, Q and U (top row) and corresponding polarised differential interferometric visibilities (PDINRM) of Stokes Q and U, along with conventional visibilities for Stokes I (bottom row). Stokes Q and U visibilities are polarised and self-differential – formed from division of orthogonal linear polarisation states. Polarimetric visibilities of symmetric objects have a characteristic sinusoidal and co-sinusoidal appearance.

The first successful implementation of polarimetric interferometry was performed with the SUSI instrument (Ireland et al. 2005). Although the technology required for long-baseline interferometry differs greatly from NRM, the astrophysical domain addressed is largely the same. This study characterised the circumstellar material around Mira variables R Car and RR Sco (Ireland et al. 2005). Subsequently, two further studies employed the SAMPoL mode of NACO/VLT – fitting to the geometry, chemical composition and grain size of the circumstellar dust of asymptotic giant branch Mira Variables R Dor, W Hya and R Leo (Norris et al. 2012), and then of red super-giant star Betelgeuse (Haubois et al. 2019).

Several modelling simplifications were employed within these previous studies. Firstly, all three papers assumed the geometry of the circumstellar dust to be symmetric – modelling it as a thin spherical shell. As we will review in Section 1.5.1, there is growing observational evidence that circumstellar environments are asymmetric and inhomogeneous. There is thus great un-tapped potential for spectro-polarised interferometry to probe the inhomogeneity of inner circumstellar environments. The second feature of these existing studies is that they model the circumstellar dust as optically thin and use simple scattering models with either Rayleigh (Ireland et al. 2005) or Mie scattering approximations (Norris et al. 2012; Haubois et al. 2019). The validity of assuming optically thin physics has not been fully tested in this context – no existing studies have systematically injected these simple models into fully physical radiative transfer code to confirm they represent physically valid solutions. Whilst these studies have successfully demonstrated the great potential which spectro-polarimetric interferometry has for characterising the circumstellar environments of evolved stars, their modelling practices leave room for technical development to increase the information content extracted from signals obtained from these critical circumstellar regions.

## 1.5 Scientific Motivations

### 1.5.1 Mass Loss from Evolved Stars

Asymptotic Giant Branch (AGB) and Red Supergiant (RSG) stars are physically extended objects in the advanced evolutionary phases of the stellar life-cycle. Once residing on the main sequence, their histories are dictated primarily by their initial mass which governs how long they fuse Hydrogen before moving into successive phases of Helium and heavier element burning (Carroll et al. 2007; Fraknoi et al. 2016). As they evolve, dredge-up processes enrich their outer envelopes with newly synthesised material (Carroll et al. 2007) and their outer layers expand dramatically, leading to intense mass loss and the eventual end of their stellar lives (Dorfi et al. 1998; Stancliffe et al. 2007). The final fate of these stars is driven strongly by mass loss – the primary mechanism for facilitating stellar mass loss is called the *stellar wind*.

The observation that evolved stars lose mass as dust is a well studied phenomena – the presence of circumstellar dust is indicated spectroscopically by an excess in the infra-red (IR) (Höfner et al. 2018). However, despite wide observation of the end-product of stellar mass loss (circumstellar dust), the mechanisms which drive the stellar wind are not well understood across all stellar sub-classes. Mass loss from carbon-rich AGB stars is reasonably explained by a combination of pulsation, convection and radiation pressure (Mehner 2021a) – pulsation and convection elevate material to sufficient heights for dust nucleation and then radiative pressure drives the material outwards from the star (Bladh et al. 2013; Janet Akyuz Mattei 1997; Soker et al. 2002; Dorfi et al. 1998; Höfner et al. 2018). The role played by radiation pressure is a function of the star’s temperature, variability and the composition of circumstellar dust. Whilst radiation pressure is sufficient to drive the opaque dust of a carbon-rich AGB, it is insufficient at driving the transparent silicate dust of oxygen-rich AGBs (Arroyo-Torres et al. 2015; Mehner 2021b; Adam et al. 2019). This issue extends to oxygen-rich RSG stars which are less luminous and variable, thus having a more extreme deficit in the mechanisms supporting mass loss after pulsation, convection and radiation are modelled (Arroyo-Torres et al. 2015). At the time of writing, the primary mechanism driving mass loss from these O-rich evolved stars is unknown.

General research on RSG stars highlights several remaining questions concerning their circumstellar mass loss. Evolutionary models demonstrate that the stellar winds launched by an RSG are capable of removing the circumstellar envelope (Meynet et al. 2015; Zapartas et al. 2024), but contradict empirical mass loss relations which predict lesser winds (Beasor et al. 2020; Decin et al. 2023). Empirical mass loss relations themselves are conflicted – varying by up to three orders of magnitude (Wit et al. 2024; Walmswell et al. 2012). We comment here that such relations employed to approximate ejected dust typically assume a homogeneous and spherically symmetric outflow, which is in contrast to the growing number of observational studies that indicate extensive irregularity and inhomogeneity in inner circumstellar environments (Ohnaka et al. 2016; Khouri et al. 2015a; Khouri et al. 2024; Adam et al. 2019; Sacuto et al. 2013; Planquart et al. 2024; Karovicova et al. 2013).

The significance of this approximation depends on the research goal – many of these studies aim to explain the total material contribution of dust to the ISM, for which this approximation may suffice. However, given that these relations cannot explain the high rates of mass loss from red supergiants, it is possible that the assumption of a spherically symmetric outflow is a significant limitation to current prescriptions. In addition, mass loss relations which assume symmetry may not be capable of explaining the mechanisms contributing to acute episodic mass loss events, which are typically very asymmetric. There is growing evidence that episodic mass ejections play a dominant role in mass loss from evolved stars (Munoz-Sanchez et al. 2024a; Wit et al. 2024; Humphreys et al. 2022; Bruch et al. 2021; Zapartas et al. 2024; Alharbi 2018), in some cases contributing 17-33% of the total dust mass lost from the star (O’Gorman et al. 2015; Ohnaka 2014).

Understanding RSG mass loss has wider implications as well. The elusive nature of red supergiants extends into their observed frequency. There is an observational lack of evidence for high mass RSG Type II P core-collapse supernova progenitors, a problem described as the ‘Red Supergiant Problem’ (Smartt et al. 2009; Smartt 2015). One possible solution to this problem is that mass loss may be a function of luminosity – suggesting that higher mass stars have higher mass loss rates (Yang et al. 2023; Vink et al. 2023) and that blue-ward evolution via mass loss can prevent these very high mass RSGs from

exploding as Type II P core-collapse supernova (Meynet et al. 2015; Walmswell et al. 2012). As a result, these stars would no longer appear as RSGs when they explode and thus are absent from the observed population of Type II P progenitors.

Unlike theoretical studies, empirical studies display far more consistency across populations. In regards to the geometry of circumstellar dust – there is *consistent observation of inconsistency* within the inner circumstellar environments of AGB and RSG stars. Common circumstellar descriptions include: elongations, ‘clumpy’ envelopes, ‘blobs’ and generally described inhomogeneity and asymmetry that is often suggested as evidence of episodic mass ejection, which may vary in its morphology on timescales as short as months (Ohnaka et al. 2016; Khouri et al. 2015a; Khouri et al. 2024; Adam et al. 2019; Sacuto et al. 2013; Planquart et al. 2024; Karovicova et al. 2013; Cannon et al. 2021; Jeffers et al. 2012; Richards et al. 1996; Dupree et al. 2022; Munoz-Sanchez et al. 2024b; Scicluna et al. 2015; Wittkowski et al. 2012; Ohnaka et al. 2024; Höfner et al. 2018; O’Gorman et al. 2015; Ohnaka 2014). Ejected dust is observed on a variety of spatial scales, as close as  $1.5 R_*$  from the stellar surface (Munoz-Sanchez et al. 2024b; Scicluna et al. 2015; Wittkowski et al. 2012; Ohnaka et al. 2024; Norris et al. 2012; Haubois et al. 2019). The dust within these regions has been identified as iron-free silicates like forsterite, enstatite or corundum, with typical ‘characteristic’ or ‘averaged’ dust grain sizes between  $0.02$ - $0.7 \mu\text{m}$  (Norris et al. 2012; Scicluna et al. 2015; Ohnaka et al. 2016; Haubois et al. 2019; Karovicova et al. 2013; Dell’Agli et al. 2014; Zhao-Geisler et al. 2012; Gobrecht et al. 2016).

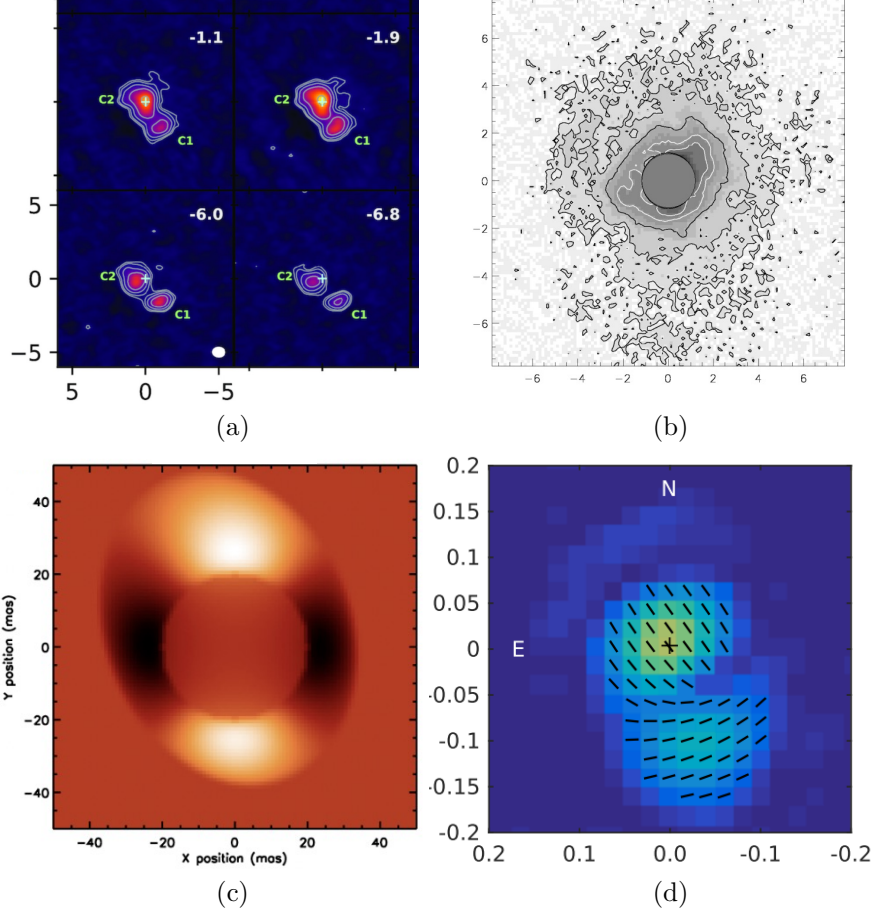
One proposed mechanism to explain mass loss from O-rich evolved stars is a *scattering-driven* stellar wind – whereby photons impart momentum onto circumstellar dust via scattering, which drives the dust grains outward from the star. The dust grains required for this mechanism are large ( $0.1$ - $1 \mu\text{m}$  for the AGB characterised in Höfner 2008) and must be transparent (poorly absorbing) to survive the high temperatures within the wind base (Höfner 2008; Bladh et al. 2012). We refer to these conditions henceforth as the ‘Hofner criteria’ for circumstellar dust grains, or to the dust grains themselves as ‘Hofner grains’. Whilst modelling has demonstrated the theoretical plausibility of this mechanism (Höfner 2008; Bladh et al. 2012), support for this theory relies most heavily on the direct detection of such dust grains, requiring observation of inner circumstellar environments. There is growing observational evidence for these grains – this work is largely based on O-rich AGB stars (Norris et al. 2012; Khouri et al. 2016; Khouri et al. 2020; Ohnaka et al. 2016; Ohnaka et al. 2017a; Adam et al. 2019; Sacuto et al. 2013), whilst the extension of this idea to RSG stars is much newer (Scicluna et al. 2015; Haubois et al. 2019). These studies all used a form of spectro-polarimetry to diagnose the size of circumstellar dust grains. As part of the next generation of spectro-polarimeters, VAMPIRES is optimally positioned to play a critical role in producing future observational evidence of Hofner grains in the inner circumstellar environments of evolved stars.

## 1.5.2 $\mu$ Cephei

Within the present work we will use the VAMPIRES instrument to search for Hofner grains in the circumstellar environment of  $\mu$  Cephei (HD 206936), fondly referred to as “Herschel’s Garnet Star” (LeVan 1989; Russell et al. 1975; Aller et al. 1951).  $\mu$  Cephei is a single M2 Ia O-rich RSG, making it an excellent test for the mechanisms contributing to RSG mass loss without the effects of binary companionship.  $\mu$  Cephei has many interesting features which include a variable light curve ( $\Delta V \sim 1$ ) (Aller et al. 1951; Krisciunas 1986; Sharpless et al. 1966; Polyakova 2003) and high and variable net polarisation ( $\Delta P \sim 2\%$ ,  $\Delta\phi_P \sim 120^\circ$ ) (Larsson-Leander 1964; Coyne et al. 1968; Polyakova 2003) which have been interpreted as evidence of asymmetric and polarised circumstellar material.

As is characteristic of red super-giant stars,  $\mu$  Cephei is undergoing extensive mass loss with estimated rates between  $10^{-7} - 10^{-5} M_\odot \text{ yr}^{-1}$  (Shenoy 2016; Gehrz 1971; Montargès et al. 2019).  $\mu$  Cephei’s mass loss is irregular and complex, with inhomogeneous dust observed from the local circumstellar environment (Safonov et al. 2019) out to scales as broad as  $25''$  (Shenoy et al. 2016). On large scales, direct imaging studies reveal complexity and find strong asymmetry, ‘clumps’ and discrete shells (between  $15$ - $17''$ ) (Mauron 1986; Le Borgne et al. 1989; Mauron 1990; Wit et al. 2008b) (Figure 1.5, (a) and (b)), which are broadly interpreted as evidence of discrete or episodic mass ejections causing variability in the mass loss rate (Mauron 1986; Le Borgne et al. 1989; Shenoy et al. 2016). However, these studies probe large fields with low resolution and have observational reach that is insufficient to study the inner-most dust envelope

which must be characterised in the search for Hofner grains. There are two existing high resolution studies which reveal features of  $\mu$  Cephei’s inner circumstellar environment (Safonov et al. 2019; Norris et al. 2012). Safonov et al. 2019 observed  $\mu$  Cephei at 880 nm with a 2.5 m telescope, unveiling a pronounced asymmetric dust feature at  $\sim 0.1$ -0.15 with a position angle of  $170$ - $220^\circ$  (Figure 1.5 (d)). Within the preliminary thesis work of Norris 2015, VAMPIRES + NRM captured  $\mu$  Cephei with the partially redundant ‘annulus aperture mask’ and modelled a pronounced asymmetry in the inner circumstellar environment at 750 nm (Figure 1.5 (c)).



**Figure 1.5:** Existing images and reconstructed models of  $\mu$  Cephei – Reproduced from: (a) Montargès et al. 2019 (image in CO lines), (b) Wit et al. 2008a (image at  $25\mu\text{m}$ ), (c) Norris 2015 (image of Stokes Q at 750 nm) and (d) Safonov et al. 2019 (image of polarisation P at 800 nm). All axes are in arc seconds except for (d) which is in milli-arcseconds. All images are in celestial coordinates. These existing studies demonstrate the presence of asymmetry and anisotropy in  $\mu$  Cephei’s circumstellar environment on all spatial scales.

There are several questions remaining from these studies. Firstly, Norris 2015 fit without closure phases and to only the inner-most circumstellar environment, leaving opportunity for future analysis to better constrain both broader scale dust features and the full extent of asymmetry. Comparing the images and models of Norris 2015 and Safonov et al. 2019, the orientation of the inner circumstellar environment appears to differ between 2015-2017, suggesting that temporal evolution of the inner environment may occur on the scale of 2 years. Neither study observed  $\mu$  Cephei over the multiple wavelengths required to probe the dust chemistry and grain size, leaving uncertainty as to the plausibility of Hofner’s scattering driven stellar wind from  $\mu$  Cephei. Within the present work we will attempt to constrain these remaining questions on  $\mu$  Cephei’s circumstellar geometry, chemistry and grain size.

## 1.6 Thesis Overview

This thesis aims to advance the use of optical-wavelength spectro-polarimetric interferometry as a method for probing dust in the inner circumstellar environments of evolved stars. We achieve this by pursuing three core goals: (1) the development of our instrumentation – VAMPIRES + NRM, (2) the creation of a suite of modelling and image reconstruction tools which address the limitations of existing techniques and (3) the application of these methods to a scientific target,  $\mu$  Cephei, which we use as a test example to investigate how VAMPIRES + NRM data can constrain circumstellar geometry, chemistry and dust grain size. In particular, we seek to understand if this constraint supports a scattering driven stellar wind from  $\mu$  Cephei (Höfner 2008). Together, these efforts aim to establish VAMPIRES + NRM as a leading and fully-supported instrument for performing optical wavelength spectro-polarimetric interferometry.

We commence our work in Chapter 2 where we place theoretical constraint on the chemistry and dust grain size within  $\mu$  Cephei’s inner circumstellar environment. By rigorously combining: theoretical predictions from the O-rich condensation sequence, Infrared Space Observatory Short Wavelength Spectrometer (ISO-SWS) data and our own radiative transfer simulations for thermodynamic constraint, we provide a renewed perspective on the species of dust which should be tested in these environments. The outcome of this chapter is an extensive and thermodynamically-viable set of dust candidates which are fit to  $\mu$  Cephei data throughout this thesis – providing the most comprehensive pool of dust candidates fit to optical and infrared spectro-polarimetric interferometry data to date.

In Chapter 3 we document the instrumentation work required to fully develop VAMPIRES + NRM for science use. A robust data reduction pipeline is written and integrated into the existing SUB-ARU/SCEXAO facility pipeline `vampires_dpp` – providing the first standardised reduction framework for VAMPIRES + NRM data, a significant step towards increasing its accessibility to the broader community. To the historical treatment of the NRM mode we introduce an empirically fitted Mueller matrix model, permitting the complete disentanglement of non-ideal instrumental polarisation effects from science signal. To benchmark the polarimetric calibration limits of the upgraded VAMPIRES instrument we conduct engineering observations and in doing so correct a computational issue that improves observing efficiency by up to 30%. We review archival VAMPIRES data, commenting on historical trends within polarimetric observations and suggesting methods for making polarimetric corrections where required. With this infrastructure in place we conduct a multi-year observing campaign of  $\mu$  Cephei. We combine our campaign with pre-existing archival data to yield 7 years of optical-wavelength polarised interferometric coverage of  $\mu$  Cephei.

In Chapter 4 we develop our first modelling package – the Parametric Inference Package for Polarimetric Interferometry (PIPPI) – the first publicly available parametric modelling package for optical-wavelength spectro-polarimetric interferometry. PIPPI is a differentiable implementation of the single scattering models used in existing studies (Norris et al. 2012; Haubojs et al. 2019) and permits accelerated inference on circumstellar geometry, chemistry and dust grain size under the assumption that circumstellar dust is optically thin.

In Chapter 5 we apply the PIPPI package to our observational campaign of  $\mu$  Cephei. First, we use a Rayleigh scattering approximation to constrain the circumstellar geometry at  $\lambda$  750,  $760 \pm 25$  nm over a 7 year period. Then, we use a Mie scattering approximation to constrain the circumstellar chemistry and dust grain size using multi-wavelength data obtained in 2023 at  $\lambda$  610, 670, 720,  $760 \pm 25$  nm. We then injection test single scattering models like PIPPI by reproducing them in radiative transfer code MCFOST (Pinte et al. 2006). Our findings provide a new perspective on the utility and physical realism of simple scattering models like PIPPI and recommend a change to how they are used to fit spectro-polarimetric interferometry data.

In Chapter 6 we develop the first image reconstruction algorithm for optical-wavelength polarised interferometry data – Polarimetric Image Reconstruction AI for Tracing Evolved Structures (PIRATES). PIRATES is trained on self-consistent and physically realistic polarised images generated from radiative transfer code and implements a two-stage framework to balance the reconstruction of high fidelity images with the accurate reproduction of interferometric observables. Our paper Lilley et al. 2024 demonstrates that PIRATES can accurately reconstruct images of a wide variety of astronomically relevant polarised

circumstellar environments, and is robust to typical amounts of random observational error.

In Chapter 7 we apply PIRATES to our observing campaign for  $\mu$  Cephei, producing the highest spatio-temporal resolution image reconstructions of  $\mu$  Cephei's inner circumstellar environment to date. We use these image reconstructions as inspiration to parameterise  $\mu$  Cephei's circumstellar geometry and fit our multi-wavelength data with radiative transfer code MCFOST using Bayesian Optimisation package `optuna`. Our results provide a comprehensive description of  $\mu$  Cephei's circumstellar environment, capturing multiple scales of asymmetry. Our constraint of the dust's chemistry and grain size enable us to comment on the plausibility that  $\mu$  Cephei's circumstellar dust can support a scattering driven stellar wind (Höfner 2008).

In Chapter 8 we summarise our findings and make several recommendations for the direction of future work on VAMPIRES + NRM,  $\mu$  Cephei, and modelling of the spectro-polarimetric interferometry data product.

# Chapter 2

## Constraining the properties of $\mu$ Cephei's circumstellar dust

### Chapter Overview

In this chapter we place constraint on the chemical composition and grain size of  $\mu$  Cephei's circumstellar dust. To do so we draw on: a theoretical description of the condensation sequence in O-rich evolved stars, the ISO-SWS spectra of  $\mu$  Cephei, and existing studies which characterise analogous stellar environments. We compile necessary optical constant data and perform radiative transfer simulations to profile the thermodynamic stability of each dust species recommended by our analysis. Combined with condensation temperatures we collate from literature sources, our simulations enable us to rule out iron-rich silicates as candidates for the inner circumstellar environment (between 1-2  $R_*$ ). We then use the methods of Höfner 2008 to re-calculate the grain sizes that would support a scattering driven stellar wind from  $\mu$  Cephei, finding that a broader range of grains (80-1250 nm) will facilitate such a mechanism than in the AGB stars examined historically (Höfner 2008; Norris et al. 2012). The analysis contained within this chapter recommends fitting to a substantially broader and more physically motivated set of iron-poor dust candidates than have been considered within existing spectro-polarimetric interferometry studies. Our suggested pool of dust candidates forms the basis of our modelling of  $\mu$  Cephei in subsequent chapters.

### Contribution Statement

The idea to re-constrain the chemical possibilities of O-rich circumstellar environments was my own. I conducted a review of the relevant literature, collating optical constant data and writing code to derive the Hofner grain size for  $\mu$  Cephei. I performed radiative transfer simulations for each of my candidate dust species and reviewed condensation temperature literature to derive values of the condensation radius. I was assisted by Christophe Pinte in the provision of optical constant data from MCFOST and in the use of MCFOST. I synthesised the final pool of dust candidates from all analysis.

## 2.1 Methods for constraining expectations

The chemical compositions proposed when fitting to the inner circumstellar environments of evolved stars have been relatively limited. Many studies fit only to: amorphous forsterite ( $\text{Mg}_2\text{SiO}_4$ ), enstatite ( $\text{MgSiO}_3$ ) or corundum ( $\text{Al}_2\text{O}_3$ ), finding they cannot be distinguished at optical wavelengths (Adam et al. 2019; Khouri et al. 2020; Scicluna et al. 2015; Ohnaka et al. 2016; Khouri et al. 2016; Ohnaka et al. 2017b; Haubois et al. 2019). To identify that the fitted dust would have the transparency required to support a scattering driven stellar wind, these studies instead posit that the fitted radial location of the dust is consistent with an iron-free (transparent) species that can survive high temperatures (Norris et al. 2012; Haubois et al. 2019). Here we aim to extend on this typical sub-set of dust chemistries to provide a broader, more inclusive and rigorously defined pool of options to fit to the circumstellar

environment of  $\mu$  Cephei, to fully test which species can be distinguished at the operational wavelengths of VAMPIRES. To do so, we first construct a set of chemical possibilities from the condensation sequence of O-rich stars and from Infrared Space Observatory Short Wavelength Spectrometer (ISO-SWS) data of  $\mu$  Cephei. Constrained by the availability of optical constants we propose a final pool of dust candidates for chemical inference. We derive the Hofner grain size criteria for  $\mu$  Cephei and then evaluate the thermodynamic stability of each dust candidate and grain size in radiative transfer code MCFOST.

## 2.2 The condensation sequence of O-rich evolved stars

The chemistry of O-rich stellar environments is complex, with many competing processes occurring and many environmental conditions that may alter the proportions of chemical end products. Here we largely take this complexity for granted and focus on the fundamental processes that suggest a handful of reasonable dust candidates for the inner circumstellar environment. O-rich stars form silicate and oxide dust from the elements: Mg, Fe, Si, Al, Ca and Na (Woolf 1969; Hans-Peter 1941; Asplund et al. 2009). From calculations that yield typical abundances at typical circumstellar pressures ( $\sim 10^{-4} - 10^{-3}$  bar) (Hans et al. 2013; Hans-Peter 1941), it can be derived that the most dominant dust species must be formed from SiO, Mg, Fe and H<sub>2</sub>O, and that the second most dominant species are formed from Ca, Al, AlOH, and Al<sub>2</sub>O (Fig 10.9, 10.12 Hans-Peter 1941). From these molecules and elements, combined with thermodynamic stability limits, the following dust products may condense in solid form: corundum (Al<sub>2</sub>O<sub>3</sub>), forsterite (Mg<sub>2</sub>SiO<sub>4</sub>), olivine (MgFeSiO<sub>4</sub>), enstatite (MgSiO<sub>3</sub>), pyroxene (Mg<sub>0.5</sub>Fe<sub>0.5</sub>SiO<sub>3</sub>), gehlenite (Ca<sub>2</sub>Al<sub>2</sub>SiO<sub>7</sub>), spinel (MgAl<sub>2</sub>O<sub>4</sub>), silica (SiO<sub>2</sub>) and diopside (CaMgSi<sub>2</sub>O<sub>6</sub>) (Hans-Peter 1941).

## 2.3 Spectroscopic analysis of ISO-SWS data of $\mu$ Cephei

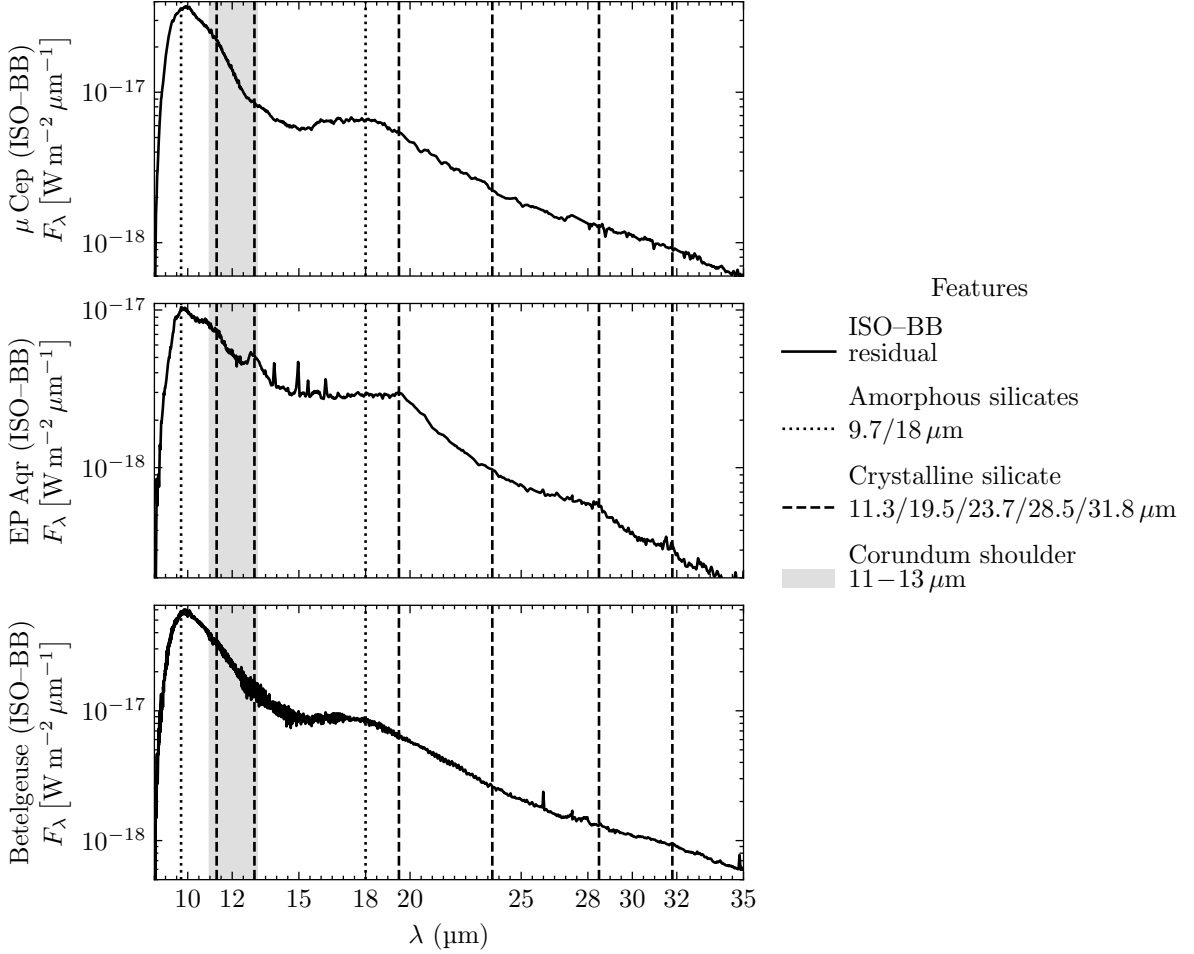
A second method for suggesting the chemical composition of circumstellar grains is mid-IR spectroscopy – here we used data from the ISO-SWS (Heras et al. 1997). At wavelengths of  $\sim 8 \mu\text{m}$  and longer the dominant contribution to emission features is from circumstellar dust (Molster et al. 1999). From the spectra of  $\mu$  Cephei (Figure 2.1) we can infer that the dominant dust species present in the wider field are amorphous silicates like olivine and pyroxene. The sharper emission features expected from crystalline silicates are not present (see EP Aqr for a comparison in Figure 2.1), which suggests that either the dust cools too quickly for the formation of crystal structure or that crystal structure is present in such a small volume that it is dominated by its amorphous counterparts in emission (Gail et al. 1984; Molster 2000). The second significant spectral feature present is corundum, indicated by a plateau between 11-13  $\mu\text{m}$  known as the ‘corundum shoulder’ (Hans-Peter 1941). Our identification of these species is supported by existing spectroscopic study of  $\mu$  Cephei (Gail et al. 2020).

However, care must be taken not to eliminate chemical candidates from spectroscopic analysis alone – for four reasons. Firstly, emission spectra do not reflect how strongly polarising a dust species is – some dust may emit weakly but be strongly polarising at optical wavelengths. Secondly, spectra like those in Figure 2.1 reflect the content and volume of all dust species within the ISO-SWS field of view ( $14'' \times 20''$ ) (Leech et al. 2003). The circumstellar dust around  $\mu$  Cephei extends to scales that encompass this entire field of view (Shenoy et al. 2016). The emission features that define the bulk of the circumstellar dust on broader scales may not accurately reflect the chemical contents or proportions of the inner circumstellar environment probed by VAMPIRES + NRM data ( $\sim 150$  mas from the stellar surface). As dust travels outwards from the star, additional dust species become thermodynamically permitted to condense. Most importantly, iron-rich silicates which would not appear in inner circumstellar regions below  $\sim 2 R_*$  may condense at larger radii of  $\sim 10 R_*$ . Whilst iron is a dominant feature of  $\mu$  Cephei’s emission spectra (Gail et al. 2020), we expect *no* iron-rich species to be thermodynamically stable within the inner circumstellar regions due to its low condensation temperature (we support these statements on iron-rich silicates with simulations in Section 2.6).

Thirdly, we would not expect likely candidates for the inner circumstellar environment (iron-free species) to have large signatures in the IR spectra, as by definition they do not heat strongly due to their transparency (Khouri et al. 2015a; Khouri et al. 2016; Ohnaka et al. 2016). For the same reason we are also under no obligation to provide dust candidates that fully fit mid-IR spectroscopic data, and can

only confirm that dust fitted to the inner environment is broadly *consistent* with spectroscopy of the wider-field, where iron-rich counterparts may condense at lower temperatures.

Fourthly, in addition to chemical evolution, the time required for dust to leave the star and travel outwards allows for structural evolution. Typically, crystalline silicates form slowly at high temperatures (Hans-Peter 1941; Fabian et al. 2000), whilst amorphous silicates form when material is rapidly cooled (Jäger et al. 1994). Likely the rate of mass loss and the speed of the stellar wind will contribute to the initial structure of silicates formed in circumstellar environments. However, conversion between structural states is also possible as amorphous material may re-anneal (Jäger et al. 2003b) and crystalline structure can be re-amorphised by cosmic rays (Jäger et al. 2003c).



**Figure 2.1:** ISO SWS Spectra of  $\mu$  Cephei, EP Aqr and Betelgeuse, data obtained from (Sloan et al. 2003),  $\lambda/\Delta\lambda = 500 - 1000$ . Each spectrum has been subtracted from a ‘continuum’ approximated with a black body at stellar temperatures 3750, 3200 and 3600 K respectively. Amorphous silicate features are broad, seen at 9.7 and 18  $\mu\text{m}$ , FWHM  $\sim 2-4 \mu\text{m}$ . Crystalline silicate features are more narrow, seen at 11.3, 19.5, 23.7, 28.5, 31.8  $\mu\text{m}$ , FWHM  $\sim 0.2 \mu\text{m}$ . There is a corundum ‘shoulder’ (plateau) in the region between 11–13  $\mu\text{m}$ .  $\mu$  Cephei’s spectrum strongly contains broad amorphous silicate emission, and lacks the narrow features expected from crystalline silicates – compared to EP Aqr which has crystalline features at 19.5, 28.5, 31.8  $\mu\text{m}$ .  $\mu$  Cephei has a corundum shoulder, as does Betelgeuse.

Whilst amorphous silicate features are predominantly found in the spectra of O-rich evolved stars, some analysis suggests that the relative strength of amorphous and crystalline features is correlated with geometry – objects with ‘disk-line’ geometry have stronger crystalline bands than objects with ‘dust shell’ outflows (Molster et al. 2002a; Molster et al. 2002b). However, it may be possible that crystalline

grains form in hot inner circumstellar environments and then rapidly amorphise via interaction with cosmic rays. If this is the case then we would detect crystalline silicates in the inner circumstellar environment where they are first formed and then observe amorphous silicate features on the broad scales that dominate emission spectra (eg. Figure 2.1). To test this hypothesis we include crystalline forms of the iron-free silicates and alumina expected from the O-rich condensation sequence within our chemical inference (forsterite, enstatite and corundum).

Existing spectroscopic analysis of  $\mu$  Cephei has identified iron as an essential component of circumstellar dust, causing strong extinction within the UV portion of the emission spectrum (Gail et al. 2020). Whilst most iron-rich species are unable to condense in proximity to the stellar surface (Norris et al. 2012), it may be thermodynamically possible for *iron-poor* silicates to exist in inner circumstellar environments. As such, we include iron poor versions of olivine and pyroxene (with Mg/Fe ratios of: 95/5, 80/20, 70/30 and 60/40) within our pool of dust candidates for thermodynamic assessment.

## 2.4 On the availability of optical constant data

The condensation sequence of O-rich stellar environments and  $\mu$  Cephei’s spectroscopic features have so far suggested a pool of chemical candidates for  $\mu$  Cephei’s circumstellar dust. To properly test each species the complex refractive index  $m(\lambda) = n + ik$  is required over a) the visible observational wavelength bands of VAMPIRES, and b) over the majority of the stellar SED. Optical data is required over a) to fit to observed data, and b) to evaluate the thermodynamic stability of each material via integration over the stellar spectrum (which we approximate with a blackbody of  $T_{\star} = 3750$  K (Josselin et al. 2007)).

For some dust candidates  $m(\lambda)$  values were available as input files to MCFOST (Pinte et al. 2006; Pinte et al. 2009). However, for most species we were required to generate  $m(\lambda)$  files using data sourced from literature. Of significant help were the Jena Database of Optical Constants for Cosmic Dust (Harald Mutschke 2020), or the Heidelberg-Jena Database of Optical Constants (HEJDOC) (Heidelberg-Jena Database 2024), which collate  $m(\lambda)$  for a vast number of astrophysically meaningful dust species. However, we found it common for  $m(\lambda)$  data to be tabulated for wavelengths larger than  $\sim 1 \mu\text{m}$  (Harald Mutschke 2020), omitting the shorter optical wavelength values we also require.

For some dust candidates it was possible to construct the required range of  $m(\lambda)$  from multiple sources (eg. crystalline enstatite and amorphous spinel). Some dust species had no available optical wavelength values for  $m(\lambda)$ , which prohibited us from testing them entirely (eg. gehlenite and diopside). Crystalline data was often available as  $m(\lambda)$  values along each principal axis of the crystal. It is typical to ‘mix’ the polarisation properties of each axis in dielectric space (Bohren et al. 1998) – ie, to assume there is no preferential alignment of grains along a particular optic axis. To derive (n,k) for iron-poor silicate mixtures of olivine (which were not available as pre-computed values), mixing in dielectric space between forsterite and olivine was required (optical coefficients for iron rich fayalite were unavailable).

All of these restrictions yield the dust candidates detailed in Table 2.1 as the starting point for our chemical inference. The optical coefficients (n,k) used for each species are plotted in Figures A.1 - A.3. The assumption of optically thin dust permits a tremendous number of combinatorial options during fitting; as all photon-grain interactions are singular and independent, mixtures of grain materials may be composed by linearly averaging the scattering cross sections. We will return to this idea during chemical inference in Chapter 5.

## 2.5 Deriving the Hofner grain size criteria for $\mu$ Cephei

The size of astronomical dust grains is a function of environmental and stellar conditions (Hutchison et al. 2016). To inform the range of dust grain sizes to test in our inference we draw on existing observational and laboratory studies which suggest physically realistic grain sizes, and derive the grain sizes required to facilitate a scattering driven stellar wind from  $\mu$  Cephei by re-performing the calculations from Höfner 2008 for  $\mu$  Cephei.

As reviewed in Chapter 1, observational studies have found circumstellar dust around evolved stars at

distances as close as  $1.5R_*$  with grain radii between  $0.02$ - $0.7\mu\text{m}$  (Munoz-Sanchez et al. 2024b; Scicluna et al. 2015; Wittkowski et al. 2012; Ohnaka et al. 2024; Norris et al. 2012; Ohnaka et al. 2016; Karovicova et al. 2013). We may also look to the interstellar medium as much of the dust permeating the ISM originates from evolved stars (Gehrz 1988). Laboratory studies of interstellar grains suggest a grain size distribution of  $0.1$ - $1\mu\text{m}$  with a peak at  $0.3\mu\text{m}$  (Srama et al. 2009; Mathis et al. 1977a), and that power law distributions with exponents between  $-2.5$  and  $-4$  are typical (Mathis et al. 1977a; Weingartner et al. 2001). Of course, it is also understood that grain size evolution occurs after dust is lost from an evolved star (Asano et al. 2013), so caution must be taken not to exclude grain sizes based only on representative values observed in the broader ISM.

The dust grain sizes which will facilitate a scattering driven stellar wind depend on the dust and stellar properties. The sizes of these grains derived in Höfner 2008 ( $125$ - $1250\text{nm}$ ) were for a  $1M_\odot$ ,  $7000L_\odot$  and  $2700\text{K}$  star. Here we re-derive these grain sizes for  $\mu$  Cephei using stellar parameters:  $M_* = 25M_\odot$  (Ariste 2023),  $T_* = 3750\text{K}$  (Josselin et al. 2007),  $L_* = 269\,000L_\odot$  (Kravchenko et al. 2019). For a full description of this derivation, see Höfner 2008 – here we summarise the main equations. To drive a stellar wind, dust grains must have a flux mean opacity per mass that is higher than the critical threshold given by  $\kappa_{crit}$  – at which point radiation pressure overcomes the inward gravitational pull.  $\kappa_{crit}$  is given by Equation 2.2 and is proportional to the ratio between the stellar mass  $M_*$  and luminosity  $L_*$  where  $c$  and  $G$  are the speed of light and the gravitational constant.

$$\kappa_{crit} = \frac{4\pi cGM_*}{L_*} \quad (2.1)$$

The total cross section per unit mass of dust grains with radius  $a_{gr}$ , density  $\rho_{grain}$  and efficiency  $Q_{rp}$  is given by Equation 2.2, where we use the constants as defined in Höfner 2008:  $\epsilon_c = 3.55 \times 10^{-5}$ ,  $\epsilon_{He} = 0.1$ ,  $A_{mon} = 140$  and  $f_c = 0.30$ . The efficiency  $Q_{rp}$  is calculated from Mie theory, we used python package `miepython` (Prah 2025) and assumed Forsterite grains with  $\rho = 3.3\text{g cm}^{-3}$  as in Höfner 2008.


$$\kappa_{rp}(\lambda, a_{gr}) = \frac{3}{4} \frac{A_{mon}}{\rho_{grain}} \frac{Q_{rp}(\lambda, a_{gr})}{a_{gr}} \frac{\epsilon_c}{1 + 4\epsilon_{He}} f_c \quad (2.2)$$

The grains capable of driving the stellar wind must have a flux weighted opacity  $\tilde{\kappa}_{rp}(\lambda, a_{gr}, T_*)$  greater than the critical opacity  $\kappa_c$ . For each wavelength  $\lambda$  the stellar intensity is approximated as a blackbody using the Planck function  $B_\lambda(T_*)$  (Equation 2.3).

$$\tilde{\kappa}_{rp}(\lambda, a_{gr}, T_*) = \kappa_{rp}(\lambda, a_{gr}) \frac{\lambda B_\lambda(T_*)}{T_*^4 \sigma / \pi} \quad (2.3)$$

The results of our calculations are displayed within Figure 2.2. The top panel depicts our re-computation of Figure 1 of Höfner 2008 for a star of  $1M_\odot$ ,  $7000L_\odot$  and  $2700\text{K}$  for comparison and the bottom panel depicts our new application of this formalism to  $\mu$  Cephei. The black circles in each figure depict the regions for which  $\tilde{\kappa}_{rp}(\lambda, a_{gr}, T_*) > \kappa_{crit}$ , which defines where grains are large enough to facilitate a scattering driven stellar wind. When compared to the star described in Höfner 2008 (Figure 2.2, top panel) which is used as the standard for Mira variables in Norris et al. 2012, our results demonstrate that the population of grain sizes capable of facilitating a scattering driven wind from  $\mu$  Cephei is *larger* and includes smaller grains (Figure 2.2, bottom panel).

The dust grains capable of driving a scattering driven stellar wind from  $\mu$  Cephei have radii between  $80$ - $1250\text{nm}$  (Figure 2.2, bottom panel). An interesting point of difference between these values derived for  $\mu$  Cephei and those of Höfner 2008 for a prototypical AGB star is that they include grains small enough to be described by the Rayleigh scattering approximation (smaller than  $\sim 100\text{nm}$  at  $\sim \lambda 700\text{nm}$ ). Thus, to discriminate between grains that can drive a scattering driven wind and those that cannot, we fit  $\mu$  Cephei with grain sizes on either side of this critical region and test grain radii between  $0.001$ - $2\mu\text{m}$ .



**Figure 2.2:** Comparison between scattering driven wind grain radii derived in Höfner 2008 for a star of  $1 M_{\odot}$ ,  $7000 L_{\odot}$  and  $2700 \text{ K}$  (top panel) and our new derivation for  $\mu$  Cephei with  $M_{\star} = 25 M_{\odot}$ ,  $T_{\star} = 3750 \text{ K}$ ,  $L_{\star} = 269\,000 L_{\odot}$  (bottom panel). The heavy black contour constrains the dust grain sizes for which  $\tilde{\kappa}_{rp}(\lambda, a_{gr}, T_{\star}) > \kappa_{crit}$ , within which grains are of appropriate size to facilitate a scattering driven stellar wind. The grain sizes which facilitate a scattering driven stellar wind from the prototypical AGB modelled in Höfner 2008 are between  $0.1\text{--}1.25 \mu\text{m}$ , whilst the population that supports this mechanism from  $\mu$  Cephei are  $0.08\text{--}1.25 \mu\text{m}$ . Interestingly,  $\mu$  Cephei’s population includes grains small enough to satisfy the Rayleigh scattering approximation at VAMPIRES observational wavelengths ( $\sim 700 \text{ nm}$ ).

## 2.6 Thermodynamic constraint via radiative transfer simulation

Finally, for our pool of dust candidates we derive thermodynamic stability limits that constrain the geometric regions in which each dust species may exist in solid form. To do so we simulate  $\mu$  Cephei in radiative transfer code MCFOST ( $M_\star = 25 M_\odot$ ,  $R_\star = 2257 R_\odot$ ,  $T_\star = 3750\text{K}$ ).<sup>1</sup> For each dust species and grain size we inject a uniform field of (optically thin) dust into the model domain and compute the radial temperature profile as a function of distance from the stellar surface. Radial temperature profiles are then compared to the condensation temperatures of each species to derive a ‘condensation radius’ – the inner-most distance at which each species of dust can form as a solid. Condensation and sublimation temperatures vary significantly across the literature, reported values depend on circumstellar pressure. We collated values relevant to typical circumstellar pressures of RSGs  $\sim 10^{-4}$  -  $10^{-3}$  bar (Table 2.1) (Hans-Peter 1941).

We used condensation temperature (gas  $\rightarrow$  solid) and not sublimation temperature (solid  $\rightarrow$  gas) to be conservative – of course it is possible that formed dust may infall and exist as a solid until crossing the sublimation radius, which will be closer to the stellar surface than the condensation radius. In-falling material is more likely to be a feature of the periods between episodic mass loss events than a feature of persistent mass loss. The mass loss rate of  $\mu$  Cephei is at times very high (up to  $10^{-5} M_\odot \text{ yr}^{-1}$ ) (Gehrz 1971), and whilst there are defined shells and anisotropic and ‘clumpy’ features that suggest the role of episodic mass loss events (Mauron 1986; Le Borgne et al. 1989; Mauron 1986), there is also a large amount of dust distributed continuously out towards wider scales (Wit et al. 2008a). This suggests that the mass loss mechanism (whether it be persistent, episodic, or a combination of both) is reasonably active, that it is appropriate to assume the bulk of dust movement is radially outward and to use condensation temperature to derive radial thermodynamic stability limits.

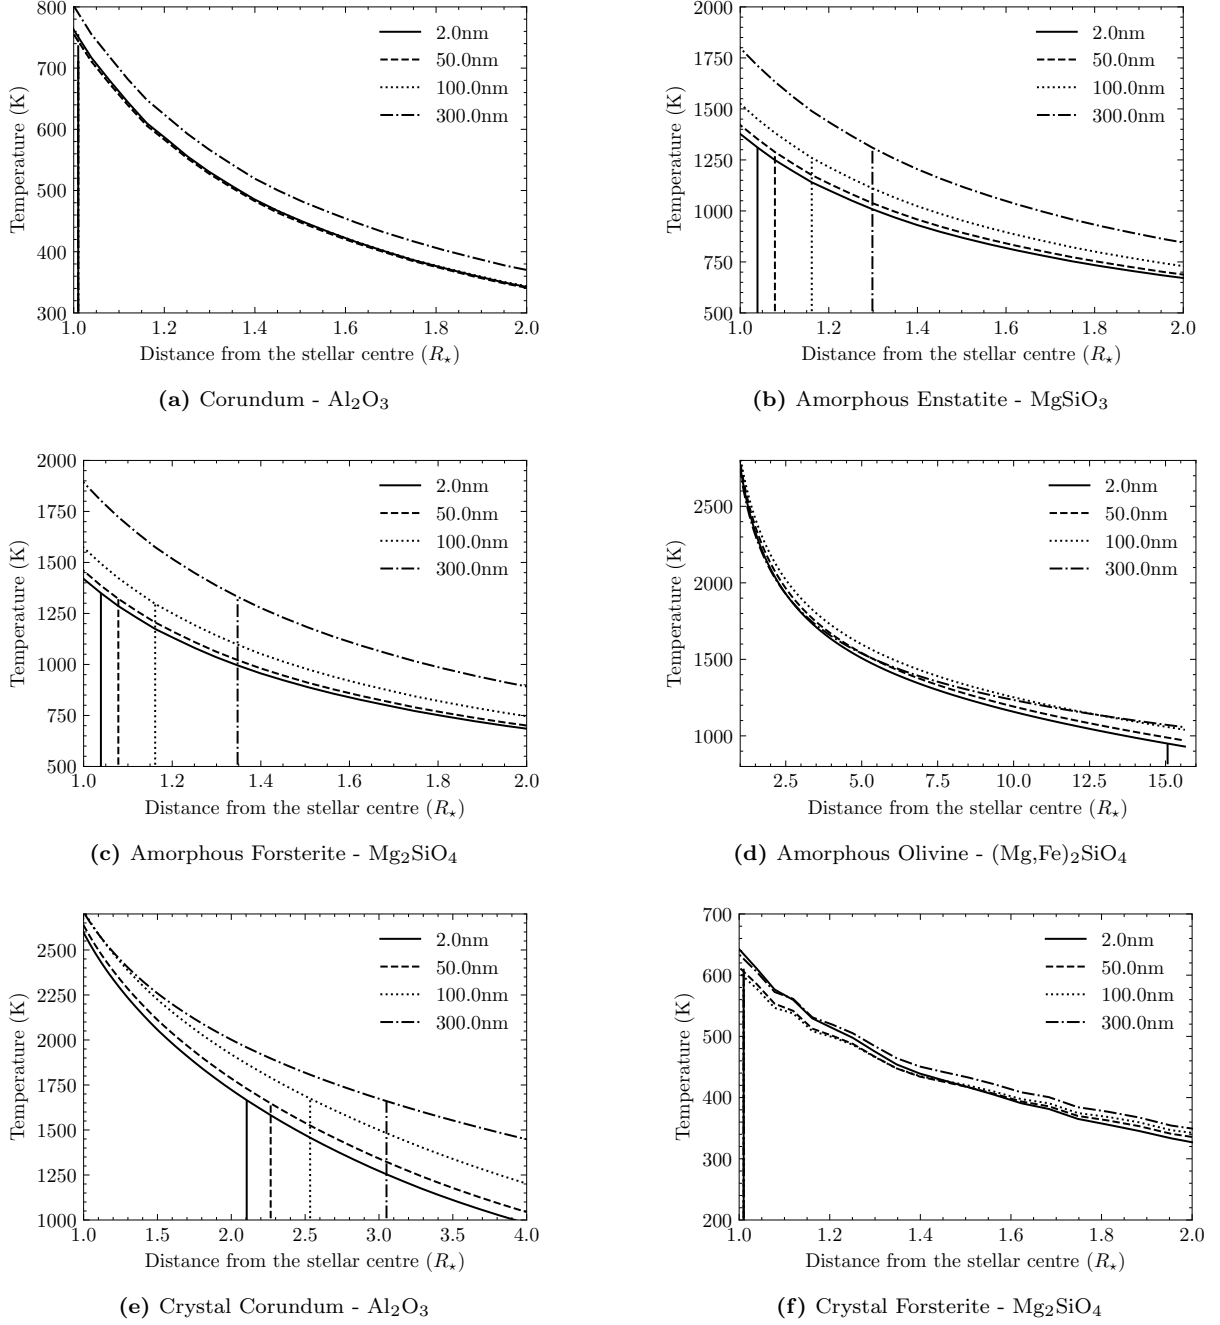
Thermodynamic profiles of dust temperature as a function of radius are depicted in Figures 2.3-2.5 and demonstrate which species may exist stably in solid form between  $1 - 2 R_\star$ . These are: enstatite, forsterite (all in amorphous and crystalline forms), amorphous corundum, spinel, silica, and iron poor (Mg/Fe = 95/5) pyroxene. Crystalline corundum and species with heavier iron content, including: olivine, pyroxene, most iron poor pyroxenes and all iron rich olivines are not thermodynamically stable within these regions and thus may be excluded from chemical inference concerning the inner-most circumstellar environment.

Our thermodynamic simulations also demonstrate that the condensation radius  $R_C$  depends non monotonically on grain size, reflecting the non-linear relationship between grain size (given by the size parameter  $x = 2\pi a/\lambda$ ) and the absorption efficiency predicted by Mie theory. Very small grains ( $x \ll 1$ ) absorb inefficiently, then as grains grow to intermediate sizes ( $x \sim 1$ ) their emission efficiency decreases. As grains become larger ( $x \gg 1$ ) they approach ‘grey’ absorbers with efficient emission, decreasing the condensation radius again (Ysard et al. 2018) (Figure 2.6).

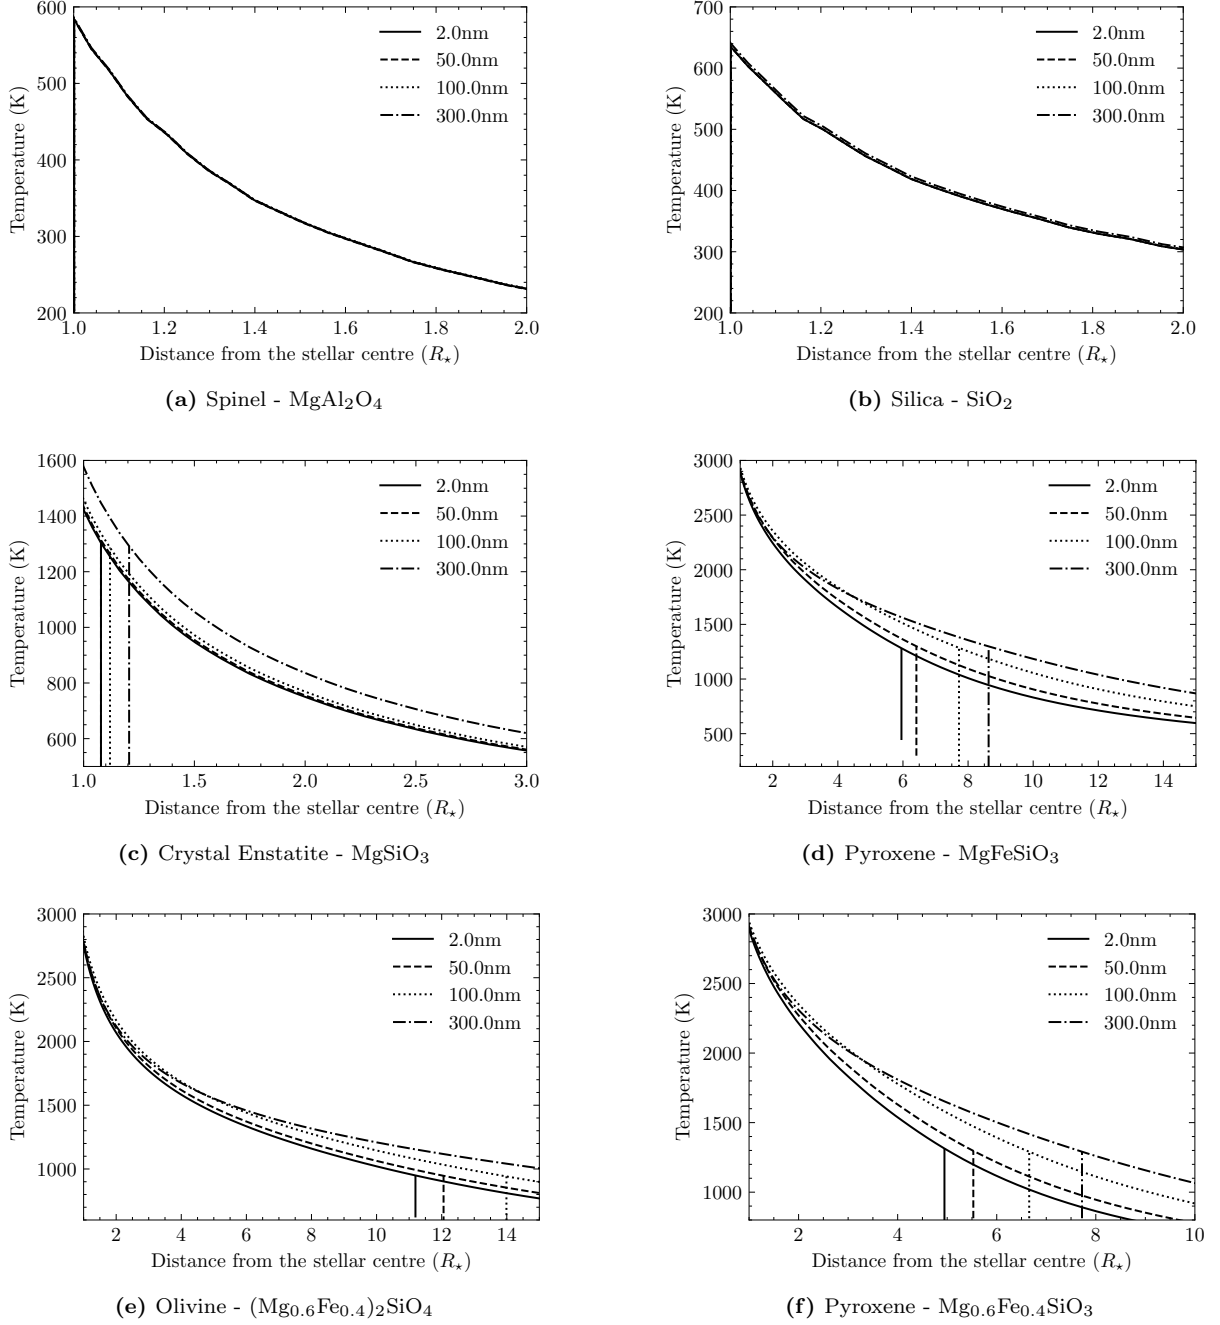
Figure 2.6 reveals that for some dust species the condensation radius can vary by large values (up to  $4 R_\star$  for species like olivine) as the grain size is changed, which would theoretically permit some constraint of dust grain size from the dust’s radial location. We note that for iron-poor candidates there is a weaker (fractional) variation in condensation radius with grain size, and so stochastic perturbations present in simulations for all species (from finite sampling and a discrete number of photons) are more evident in plots for iron-free species (bottom panel) than they are for iron-rich species (top panel).

In Table 2.1 we report the smallest possible condensation radius  $R_C [R_\star]$  for each grain species, which are typically for the smallest grain size. From our thermodynamic analysis we are able to eliminate iron-rich species and crystalline corundum as candidates for the inner-most circumstellar environment ( $1 - 2 R_\star$ ). Dust species which are thermodynamically permitted to exist in the inner circumstellar environment and will be tested in our subsequent analysis are marked with a  $\checkmark$ , whilst those that we can exclude are marked with a  $\times$ .

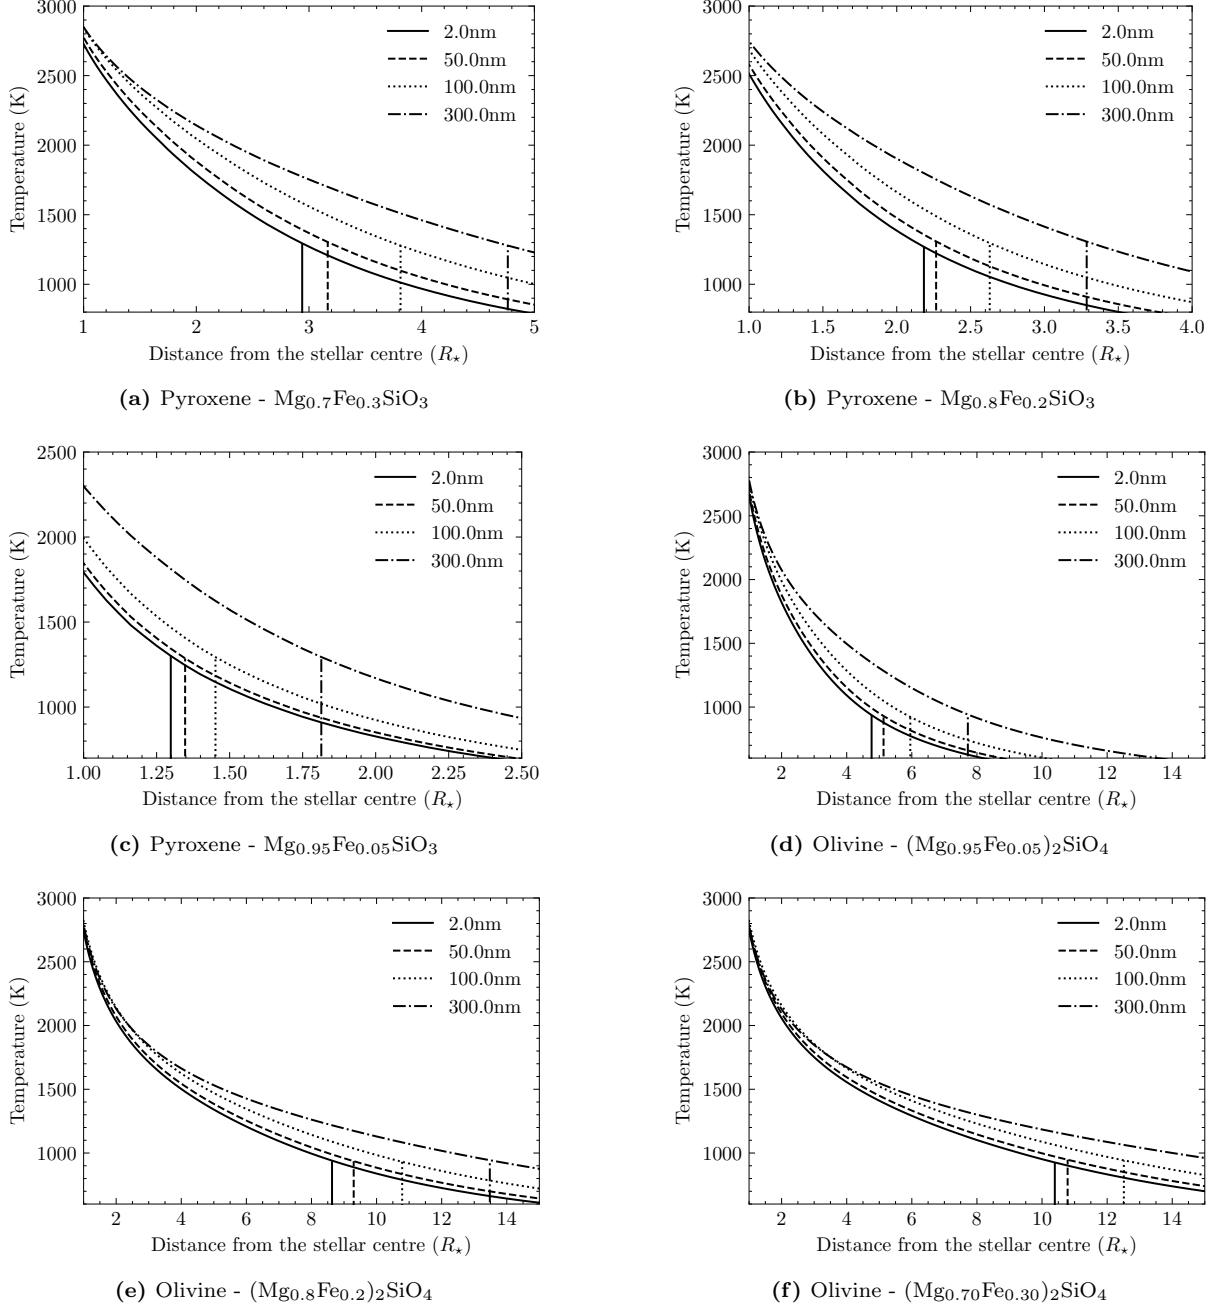
<sup>1</sup>To obtain this value for stellar radius we used a literature value for the angular size obtained in the continuum at 800 nm –  $R_\star = 10 \text{ mas}$  (Quirrenbach et al. 1993). We then used a distance of 1000 pc which yields a stellar size of  $2257 R_\odot$ . The distance to  $\mu$  Cephei carries uncertainty – standard distance catalogues like Hipparcos or Gaia yield unreliable results for red supergiants (Chiavassa et al. 2022). As such, we adopt a value of 1000 pc which is common for studies on  $\mu$  Cephei (Davies et al. 2020).




**Figure 2.3:** Condensation Radii Derivations – Part 1. Dust composed of each chemical and grain size is injected into a radiative transfer simulation surrounding  $\mu$  Cephei. The temperature profile is computed as a function of distance from the stellar surface in units of  $R_* = 2257R_\odot$ . Vertical lines from each temperature profile indicate the location of the condensation radius for each dust grain size, which is calculated by comparison of the temperature profile to the condensation temperature. The condensation radius defines how close to the stellar surface each dust species may exist in solid form.



**Figure 2.4:** Condensation Radii Derivations – Part 2. Dust composed of each chemical and grain size is injected into a radiative transfer simulation surrounding  $\mu$  Cephei. The temperature profile is computed as a function of distance from the stellar surface in units of  $R_* = 2257R_\odot$ . Vertical lines from each temperature profile indicate the location of the condensation radius for each dust grain size, which is calculated by comparison of the temperature profile to the condensation temperature. The condensation radius defines how close to the stellar surface each dust species may exist in solid form.



**Figure 2.5:** Condensation Radii Derivations – Part 3. Dust composed of each chemical and grain size is injected into a radiative transfer simulation surrounding  $\mu$  Cephei. The temperature profile is computed as a function of distance from the stellar surface in units of  $R_* = 2257R_\odot$ . Vertical lines from each temperature profile indicate the location of the condensation radius for each dust grain size, which is calculated by comparison of the temperature profile to the condensation temperature. The condensation radius defines how close to the stellar surface each dust species may exist in solid form.



**Figure 2.6:** Condensation radii ( $R_*$ ) as a function of grain size (radius in microns) for different dust species. Row 1: iron-rich dust; Row 2: iron-poor and iron-free dust.  $\mu$  Cephei was simulated in radiative transfer code ( $M_* = 25_\odot$ ,  $T_* = 3750$  K,  $L_* = 269\,000 L_\odot$ ,  $R_* = 2257R_\odot$ ) and each dust species was injected into the model domain. Condensation radii were computed by comparing the dust’s temperature to its condensation temperature.

Mineral	Formula	$T_{\text{cond}}$ [K]	(n,k)	$R_C$ [ $R_\star$ ]	Candidate?
A. Olivine	$(\text{Mg, Fe})_2\text{SiO}_4$	950 [o]	[a]	10.0	✗
A. Forsterite	$\text{Mg}_2\text{SiO}_4$	1354 [m]	[b]	1.0	✓
A. Pyroxene	$(\text{Mg, Fe})\text{SiO}_3$	1316 [m]	[b]	5.1	✗
A. Enstatite	$\text{MgSiO}_3$	1316 [m]	[b]	1.1	✓
A. Corundum	$\text{Al}_2\text{O}_3$	1677 [m]	[c]	1.0	✓
C. Enstatite	$\text{MgSiO}_3$	1316 [m]	[d],[l]	1.1	✓
C. Forsterite	$\text{Mg}_2\text{SiO}_4$	1354 [m]	[e]	1.0	✓
C. Corundum	$\text{Al}_2\text{O}_3$	1677 [m]	[f]	2.1	✗
A. Silica	$\text{SiO}_2$	950 [o]	[g]	1.0	✓
C. Spinel	$\text{MgAl}_2\text{O}_4$	1397 [m]	[h],[k]	1.0	✓
A. Pyroxene Mg 95	$\text{Mg}_{0.95}\text{Fe}_{0.05}\text{SiO}_3$	1316 [m]	[b]	1.3	✓
A. Pyroxene Mg 80	$\text{Mg}_{0.80}\text{Fe}_{0.20}\text{SiO}_3$	1316 [m]	[b]	2.2	✗
A. Pyroxene Mg 70	$\text{Mg}_{0.70}\text{Fe}_{0.30}\text{SiO}_3$	1316 [m]	[b]	3.0	✗
A. Pyroxene Mg 60	$\text{Mg}_{0.60}\text{Fe}_{0.40}\text{SiO}_3$	1316 [m]	[b]	5.0	✗
A. Olivine Mg 95	$(\text{Mg}_{0.95}\text{Fe}_{0.05})_2\text{SiO}_4$	950 [o]	[p]	5.0	✗
A. Olivine Mg 80	$(\text{Mg}_{0.80}\text{Fe}_{0.20})_2\text{SiO}_4$	950 [o]	[p]	8.6	✗
A. Olivine Mg 70	$(\text{Mg}_{0.70}\text{Fe}_{0.30})_2\text{SiO}_4$	950 [o]	[p]	9.7	✗
A. Olivine Mg 60	$(\text{Mg}_{0.60}\text{Fe}_{0.40})_2\text{SiO}_4$	950 [o]	[p]	10.0	✗

**Table 2.1:** Chemical formulae, condensation temperatures ( $T_{\text{cond}}$ ), optical-constant (n,k) references and computed condensation radii  $R_C$  for candidate dust species. A/C denote amorphous or crystalline structural form of each dust species. References for optical constants: [a] (Kama et al. 2009); [b] (Jäger et al. 2003a); [c] (Kobayashi et al. 2011); [d] (Rucks et al. 2022); [e] (Pitman et al. 2013); [f] (Querry 1985); [g] (Franta et al. 2016); [h] (Chernova et al. 2017); [k] (Zeidler et al. 2013); [l] (Jaeger et al. 1998); [m] (Lodders 2003); [o] (Gail et al. 1999); [p] (Jager et al. 1994)

## 2.7 Chapter Summary

This chapter derived a comprehensive subset of candidates for the composition of circumstellar dust around RSG  $\mu$  Cephei, extending on the pool of candidates explored in prior studies of the inner circumstellar environments of evolved stars, in particular to those fitted via spectro-polarimetric interferometry (Norris et al. 2012; Haubois et al. 2019; Adam et al. 2019; Khouri et al. 2020; Scicluna et al. 2015; Ohnaka et al. 2016; Khouri et al. 2016; Ohnaka et al. 2017b). We derived the grain sizes sufficient to support a scattering driven stellar stellar wind from  $\mu$  Cephei, finding that grain radii between 80-1250 nm are required. This range of dust grain sizes includes grains *smaller* than those quoted in Höfner 2008, most notably including grains small enough to be described by the Rayleigh scattering approximation. By simulating each dust species within radiative transfer code MCFOST, we evaluated the regions in which each species is thermodynamically stable. This has confirmed that crystalline corundum and iron-rich silicates are not candidates for the inner circumstellar environment (1-2  $R_\star$ ).

This same thermodynamic reasoning yields that all our proposed iron-free dust species (except crystalline corundum) may exist within the inner circumstellar environment over a diverse range of grain sizes. Existing studies which have characterised the inner circumstellar environments of evolved stars have found that species like forsterite, enstatite and corundum are indistinguishable at optical and mid-IR wavelengths (Haubois et al. 2019; Adam et al. 2019; Khouri et al. 2020; Scicluna et al. 2015; Ohnaka et al. 2016; Khouri et al. 2016; Ohnaka et al. 2017b). Here, we provide our pool of candidate dust grains to ascertain what discrimination the optical wavelengths of VAMPIRES may have over a broader and more exhaustive range of plausible dust species. We will use these derived candidates for dust chemistry and grain size in our subsequent modelling of  $\mu$  Cephei.

# Chapter 3

## Developing the NRM mode of the VAMPIRES Instrument

*“An expert is a person who has made all the mistakes that can be made in a very narrow field.”*

---

Niels Bohr

### Overview of Chapter

This chapter describes the instrumentation and engineering work necessary to develop the NRM mode of the VAMPIRES Instrument for science use. This consisted of software development, instrumental characterisation, engineering activities, archival data analysis and observation of  $\mu$  Cephei. We first introduce the VAMPIRES instrument, reviewing the instrumental Mueller matrix model and integrating it into our treatment of the NRM data product. We then develop a robust pipeline for reducing VAMPIRES + NRM data which is now integrated into the SCEXAO facility `vampires_dpp` pipeline for public use. We conducted engineering observations to measure the intrinsic polarimetric calibration limit of new VAMPIRES – during which we resolved a computing issue that improved the observing efficiency of the NRM mode by up to  $\sim 30\%$ . We present our observing campaign of  $\mu$  Cephei – composed of existing archival observations and our own new observations. Our campaign provides the highest spatio-temporal coverage of  $\mu$  Cephei at optical wavelengths to date, spanning a 7 year period. Lastly, we review archival VAMPIRES data to identify historical trends and provide methods for making polarimetric corrections to archival data where required.

### Contribution Statement

I created a pipeline for VAMPIRES + NRM data which was integrated into the SCEXAO facility code `vampires_dpp` in collaboration with Miles Lucas. The VAMPIRES Mueller matrix instrument model and empirically fitted matrix coefficients are the work of (Zhang et al. 2023). I worked closely with the VAMPIRES polarimetry team to implement their model within the NRM framework. Engineering observations conducted are my own but were greatly supported by the SCEXAO team. I identified the issue with observing inefficiency and worked in collaboration with the SCEXAO team to identify its source. In collaboration with my supervisors I wrote three successful observing proposals for the SUBARU telescope. Due to technical issues experienced in 2023 I was generously gifted additional observation time by Miles Lucas. The review of archival data is my own work, as were the identification of and methods for correcting polarimetric errors.

### 3.1 Instrumental Mueller Matrix Model of VAMPIRES

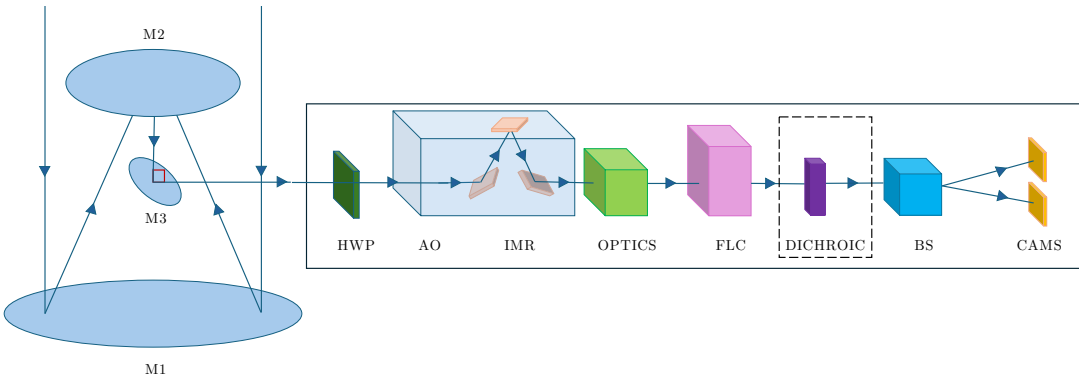
The VAMPIRES Instrument (Visible Aperture Masking Polarimeter for Resolving Exoplanetary / Evolved Star Signatures) is a differential imaging polarimeter mounted on the SCEXAO bench of the SUBARU telescope on Mauna Kea, Hawaii. In addition to its functionality as a polarimetric imager, VAMPIRES is equipped with a suite of aperture masks which enable *differential, spectro-polarimetric interferometry*. VAMPIRES + NRM is a prototypical example of polarised interferometry – preceded by the now decommissioned NACO + SAMPOL/VLT (Lenzen et al. 2003; Tuthill et al. 2010b), and SUSI instruments (Davis et al. 1999). Operating at optical wavelengths (600–800 nm) where polarisation signatures induced by dust-scattered starlight are strong, VAMPIRES + NRM performs multi-wavelength spectro-polarimetric interferometry of linear polarisation states Stokes I, Q and U, providing a sub-diffraction limited spatio-spectral probe of dusty inner-circumstellar environments (Norris et al. 2015a; Norris 2015).

Two distinct hardware versions of the VAMPIRES instrument contributed data to this thesis – we refer to them as the *original* (Norris et al. 2015a) and *upgraded* (Lucas et al. 2024) instrument versions. All instrument analysis is common to both unless specified. VAMPIRES is complex and contains a long optical train with many moving parts – for a full review see Norris et al. 2015a and Lucas et al. 2024. The components required to understand and model the polarimetric data products are those explicitly defined in the instrument model – outlined schematically in Figure 3.1. In the context of polarimetry an instrument model is typically a Mueller matrix model where each optic is defined by a (4 × 4) matrix (Equation 3.1) (Bickel et al. 1985; Snik et al. 2013). Each matrix element quantifies the materially-induced modification made by optic M to incident Stokes vector  $S_{in}$ , yielding output Stokes vector  $S_{out}$  (Equation 3.2).

$$\mathbf{M} = \begin{bmatrix} I \rightarrow I & Q \rightarrow I & U \rightarrow I & V \rightarrow I \\ I \rightarrow Q & Q \rightarrow Q & U \rightarrow Q & V \rightarrow Q \\ I \rightarrow U & Q \rightarrow U & U \rightarrow U & V \rightarrow U \\ I \rightarrow V & Q \rightarrow V & U \rightarrow V & V \rightarrow V \end{bmatrix} \quad (3.1)$$

$$\mathbf{S}_{out} = \mathbf{M} \cdot \mathbf{S}_{in} \quad (3.2)$$

The instrument model is in many ways the most important part of any polarimetric study – it is a deterministic function that specifies how the polarisation of light is influenced by propagation through the instrument. Provision of an instrument model enables the disentanglement of astrophysical and instrumental origins of measured polarisation, causes which are otherwise degenerate.



**Figure 3.1:** Schematic of key Subaru telescope and VAMPIRES Instrument hardware components necessary for modelling the polarimetric data product. Starlight arrives at the SUBARU Telescope mirrors (primary (M1), secondary (M2), and tertiary (M3)). Components defining the VAMPIRES instrument are then depicted – five elements are common to both the original and upgraded versions of VAMPIRES: a half wave plate (HWP), an image rotator (IMR), an optics term (OPTICS), a ferro-electric liquid crystal (FLC), a polarising Beam splitter (BS) and two cameras (CAMS). Upgraded VAMPIRES features a dichroic stack (DICHROIC, dotted line) which enables simultaneous multi-wavelength observation.

Mueller matrices are multiplicative but do not commute; as such, the interaction of light with  $n$  instrument elements can be written as Equation 3.3 where sub-indices are ordered from upstream (1)  $\rightarrow$  downstream ( $n$ ). This formalism produces one aggregate Mueller matrix for the instrument, a dynamic model that is a function of instrumental telemetry and will change throughout an observation as optics rotate.

$$\mathbf{S}_{\text{out}} = \mathbf{M}_n \mathbf{M}_{n-1} \dots \mathbf{M}_2 \mathbf{M}_1 \mathbf{S}_{\text{in}} = \mathbf{M}_{\text{total}} \cdot \mathbf{S}_{\text{in}} \quad (3.3)$$

The terms of a Mueller matrix are either approximated with idealised values or fitted empirically using a combination of internal calibration measurements and on-sky polarisation standard star measurements. In this work we use empirically fitted Mueller matrix coefficients fitted by Zhang et al. 2023.<sup>1</sup>

Subaru is an alt-az telescope. The SCEXAO bench operates in ‘pupil tracking mode’ where the orientation of the telescope pupil is fixed while the sky field of view rotates clockwise with the parallactic angle (pa) throughout observation. From altitude (alt) and azimuth angles (az) all other instrument angles can be calculated. For convenience we first define several Mueller matrix expressions common to multiple optical components.

### Common Mueller matrices

The Muller matrix for an optic which rotates the polarisation reference frame by angle  $\theta$ , also referred to as a ‘rotator array’ is given by Equation 3.4.

$$M_{\text{rotator}}(\theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(2\theta) & \sin(2\theta) & 0 \\ 0 & -\sin(2\theta) & \cos(2\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.4)$$

The Muller matrix for a general optical component oriented at an angle  $\theta$  with a diattenuation  $\epsilon$  and a phase retardance  $\delta$  is given by Equation 3.5.

$$M_{\text{generic}}(\theta, \epsilon, \delta) = \begin{bmatrix} 1 & \epsilon \cos(2\theta) & \epsilon \sin(2\theta) & 0 \\ \epsilon \cos(2\theta) & \cos^2(2\theta) + \sin^2(2\theta) & -\sin(2\delta) & 0 \\ 0 & \sin(2\delta) & \cos(2\delta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.5)$$

The Muller matrix for a wave-plate at angle  $\theta$  with retardance  $\delta$  is given by Equation 3.6.

$$M_{\text{waveplate}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos^2(2\theta) + \cos(\delta) \sin^2(2\theta) & (1 - \cos(\delta)) \sin(2\theta) \cos(2\theta) & -\sin(\delta) \sin(2\theta) \\ 0 & (1 - \cos(\delta)) \sin(2\theta) \cos(2\theta) & \sin^2(2\theta) + \cos(\delta) \cos^2(2\theta) & \sin(\delta) \cos(2\theta) \\ 0 & \sin(\delta) \sin(2\theta) & -\sin(\delta) \cos(2\theta) & \cos(\delta) \end{bmatrix} \quad (3.6)$$

The Mueller matrix for a linear attenuator with retardance  $\phi$  and diattenuation  $\epsilon$  is given by Equation 3.7.

$$M_{\text{att}} = \begin{bmatrix} 1 & \epsilon & 0 & 0 \\ \epsilon & 1 & 0 & 0 \\ 0 & 0 & \sqrt{1 - \epsilon^2} \cos \phi & \sqrt{1 - \epsilon^2} \sin \phi \\ 0 & 0 & -\sqrt{1 - \epsilon^2} \sin \phi & \sqrt{1 - \epsilon^2} \cos \phi \end{bmatrix} \quad (3.7)$$

The Muller matrix for a polarising beam-splitter or Wollaston prism with diattenuation  $\eta$  is given by Equation 3.8 where the direction of  $\eta$  is reversed to switch between the ordinary and extraordinary ray states.

<sup>1</sup>We used the most updated coefficients available at the time of modelling our data, which are not the same values in the original paper by Zhang et al. 2023. We obtained updated coefficients via private communication with the authors of Zhang et al. 2023.

$$M_{\text{PBS}} = 0.5 \times \begin{bmatrix} 1 & \eta & 0 & 0 \\ \eta & 1 & 0 & 0 \\ 0 & 0 & \sqrt{(1-\eta)(1+\eta)} & 0 \\ 0 & 0 & 0 & \sqrt{(1-\eta)(1+\eta)} \end{bmatrix} \quad (3.8)$$

With these general expressions defined we now employ them to define the core components of the VAMPIRES Mueller matrix instrument model – following the optical train from sky to detector as outlined in Figure 3.1.

### Telescope Mirrors

The primary (M1), secondary (M2) and tertiary (M3) mirrors collect and focus light into the SCExAO bench. The reflection angle of M3 is maintained at  $45^\circ$  while its plane of incidence is rotated relative to the Nasmyth platform. Subaru has two tertiary mirrors, each sending light to the optical and infrared Nasmyth platforms. Counter-intuitively (given VAMPIRES is an optical-wavelength instrument) SCExAO/VAMPIRES is located on the *infrared* bench – SCExAO is a near-IR instrument and VAMPIRES operates simultaneously in ‘hitch-hiker mode’, collecting unused optical-wavelength photons. The IR M3 is uncoated silver (NAOJ 2024). The Muller matrix for M3 is given by Equation 3.9, where  $M_{\text{rot}}$  and  $M_{\text{generic}}$  were defined in Equations 3.4 and 3.5.

$$M_{\text{M3}} = M_{\text{rotator}}(-\text{alt}) \cdot M_{\text{generic}}(\theta_{\text{M3}}, \epsilon_{\text{M3}}, \delta_{\text{M3}}) \cdot M_{\text{rotator}}(\text{pa}) \quad (3.9)$$

The current instrument model for VAMPIRES uses idealised coefficients for M3. We tested the use of material coefficients for un-coated silver but found this overestimated the amount of instrumental polarisation induced by the tertiary mirror and created an instrument model which calibrated poorly. The inadequacy of using material coefficients for deriving instrumental polarisation in this context has been addressed previously (Hart et al. 2021). The use of on-sky standards to empirically calibrate the coefficients of M3 is part of ongoing instrumental work (Zhang et al. 2023).

### VAMPIRES Components

#### Half Wave Plate (HWP)

A half wave plate is rotated to probe four orthogonal polarisation states (H, V, H45, V45) which are combined to measure Stokes Q (H-V) and Stokes U (H45-V45). There are two modes of operation for the HWP: ADI sync ‘ON’ and ‘OFF’. In ADI sync OFF, HWP angles are fixed relative to the Nasmyth platform at  $0^\circ, 22.5^\circ, 45^\circ$ , and  $67.5^\circ$ . In ADI sync ON, these HWP angles are offset by a term that depends on telemetry and tracks the parallactic angle (Equation 3.11). The Muller Matrix for the HWP is given by Equation 3.10, where  $\Delta\theta_{\text{HWP}} = 0$  in ADI sync ‘OFF’ mode, and ‘lat’ denotes the observing latitude.

$$M_{\text{HWP}} = M_{\text{waveplate}}(\theta_{\text{HWP}} + \Delta\theta_{\text{HWP}}, \delta_{\text{HWP}}) \quad (3.10)$$

$$\begin{aligned} \Delta\theta_{\text{HWP}} &= \frac{1}{2} \arctan \frac{\alpha}{\beta} + \text{alt} \\ \alpha &= \sin(\text{az}) \\ \beta &= \sin(\text{alt}) \cos(\text{az}) + \cos(\text{alt}) \tan(\text{lat}) \end{aligned} \quad (3.11)$$

#### Image Rotator (IMR)

The image rotator (also known as a ‘K mirror’) is a rotating assembly of three inclined mirrors. For simplicity it is modelled here as a non-ideal half wave-plate (Equation 3.4, 3.12) which rotates by angle  $\theta_{\text{IMR}}$  (Equation 3.13).  $\text{IMR}_{\text{pap}}$  is a static instrument offset (Lucas et al. 2024).

$$M_{\text{IMR}} = M_{\text{generic}}(\theta_{\text{IMR}}) \quad (3.12)$$

$$\begin{aligned}\theta_{\text{IMR}} &= \frac{1}{2}(\text{alt} + 90 - \text{IMR}_{\text{pap}}) \\ \text{IMR}_{\text{pap}} &= -39\end{aligned}\tag{3.13}$$

### Optics

The ‘optics’ term aggregates all non-individually modelled optical components (Equation 3.14). The main role of the ‘optics’ term here is to model any residual instrumental polarisation that is not corrected by two quarter wave-plates (QWPs) which are tuned to minimise instrumental polarisation predominantly introduced by a periscope that moves light vertically between the ‘top’ and ‘bottom’ benches of VAMPIRES. In original VAMPIRES the QWPs were tuned to optimal values for each wavelength and were moved to those positions when changing the wavelength filter during observation. In upgraded VAMPIRES the new ‘multi-band imaging mode’ uses a dichroic stack to capture multi-wavelength data, requiring use of averaged QWP values. Thus particularly in upgraded VAMPIRES, the optics term is an important part of compensating for residual instrumental polarisation effects.

$$M_{\text{OPT}} = M_{\text{generic}}(\theta_{\text{OPT}}, \delta_{\text{OPT}}, \epsilon_{\text{OPT}})\tag{3.14}$$

### Ferroelectric Liquid Crystal (FLC)

The FLC is modelled with Equation 3.16 where  $S_{\text{o,e}}$  defines whether the ordinary or extraordinary crystal axis is being used. Each transmits a different polarisation component – these two FLC states are referred to as ‘A’ and ‘B’ within VAMPIRES meta-data.

$$M_{\text{FLC}} = M_{\text{waveplate}}(S_{\text{o,e}}, \delta_{\text{FLC}})\tag{3.15}$$

### Dichroic Stack – Upgraded VAMPIRES only

The dichroic stack is a new instrumental component that was introduced as part of upgraded VAMPIRES. Rather than requiring sequential changing of wavelength filters, the dichroic stack enables simultaneous collection of multi-wavelength data (Lucas et al. 2024). The dichroic stack is modelled with a general expression for a retarder (Equation 3.7).

$$M_{\text{DICHROIC}} = M_{\text{rotator}}(\theta_{\text{DICH}}) \cdot M_{\text{att}}(0, \epsilon_{\text{DICH}}, \phi_{\text{DICH}}) \cdot M_{\text{rotator}}(-\theta_{\text{DICH}})\tag{3.16}$$

### Polarising Beam Splitter

The beam splitter is modelled with Equation 3.17 and sends light to each camera.

$$M_{\text{PBS}} = M_{\text{wollaston}}(\theta_{\text{FLC}}, \delta_{\text{FLC}})\tag{3.17}$$

### Cameras

Two cameras are oriented orthogonally with each viewing one of the two output beams of the polarising beam splitter. There is no explicit term in the Muller matrix for each camera – the model assumes the image intensity measured corresponds to each beam splitter state. The cameras are numbered ‘1’ and ‘2’ within meta-data.

## Net Instrument Expressions

The total Mueller matrix models for original and upgraded VAMPIRES are given by Equations 3.18 and 3.19.

$$M_{\text{ORIGINAL}} = M_{\text{PBS}} \cdot M_{\text{FLC}} \cdot M_{\text{OPT}} \cdot M_{\text{IMR}} \cdot M_{\text{HWP}} \cdot M_{\text{M3}}\tag{3.18}$$

$$M_{\text{UPGRADED}} = M_{\text{PBS}} \cdot M_{\text{DICHROIC}} \cdot M_{\text{FLC}} \cdot M_{\text{OPT}} \cdot M_{\text{IMR}} \cdot M_{\text{HWP}} \cdot M_{\text{M3}}\tag{3.19}$$

### Net image transformations

The wavefront incident on M1 undergoes a spatial transformation as it propagates through the instrument to the detector. Upgraded VAMPIRES has a parity flip relative to original VAMPIRES, which results in a different relation for this transformation between instrument versions. In both instrument versions there is a relative mirror inversion in the images for camera 2 due to the staging of the cameras on orthogonal sight-lines: one sees a reflected beam while the other a transmitted beam. The net image rotations  $I_{\text{rot}}$  (counter clockwise) for original VAMPIRES (Equation 3.20) and upgraded VAMPIRES (Equation 3.21) define the rotation of the camera 1 image with respect to the parallactic angle.

$$I_{\text{rot, orig}} = \text{pa} - 78.6 \quad (3.20)$$

$$I_{\text{rot, upgraded}} = \text{pa} + 102.25 \quad (3.21)$$

### Empirically fitted Mueller matrix coefficients

Empirically fitted Mueller matrix coefficients for each instrumental version are tabulated in Tables 3.1 and 3.2.

Parameter	Symbol	unit	625 nm	675 nm	725 nm	750 nm	775 nm	Ideal
HWP retardance	$\delta_{\text{HWP}}$	$\lambda$	0.432	0.451	0.465	0.479	0.5	0.5
HWP offset	$\theta_{\text{HWP}}$	$^\circ$	-0.002	-2.642	-4.472	-2.316	2.508	0
IMR retardance	$\delta_{\text{IMR}}$	$\lambda$	0.223	0.32	0.446	0.479	0.497	0.5
IMR offset	$\theta_{\text{IMR}}$	$^\circ$	-0.118	-0.011	0.192	0.127	-0.133	0
OPTICS retardance	$\delta_{\text{OPTICS}}$	$\lambda$	0.13	-0.163	-0.104	-0.155	-0.272	0
OPTICS diattenuation	$\epsilon_{\text{OPTICS}}$	-	0.005	0.036	0.011	0.001	0.002	0
OPTICS offset angle	$\theta_{\text{OPTICS}}$	$^\circ$	60.749	-7.151	-31.142	-26.709	-17.222	0
FLC retardance	$\delta_{\text{FLC}}$	$\lambda$	0.3	0.302	0.281	0.242	0.389	0.5
FLC offset angle	$\theta_{\text{FLC}}$	$^\circ$	-2.718	0.256	4.238	-1.257	2.373	0

**Table 3.1:** Empirically fitted Mueller matrix coefficients for original VAMPIRES, obtained by private communication with the authors of Zhang et al. 2023 who obtained these coefficients by fitting internal polarised source data with an MCMC and the model defined in Equation 3.18.

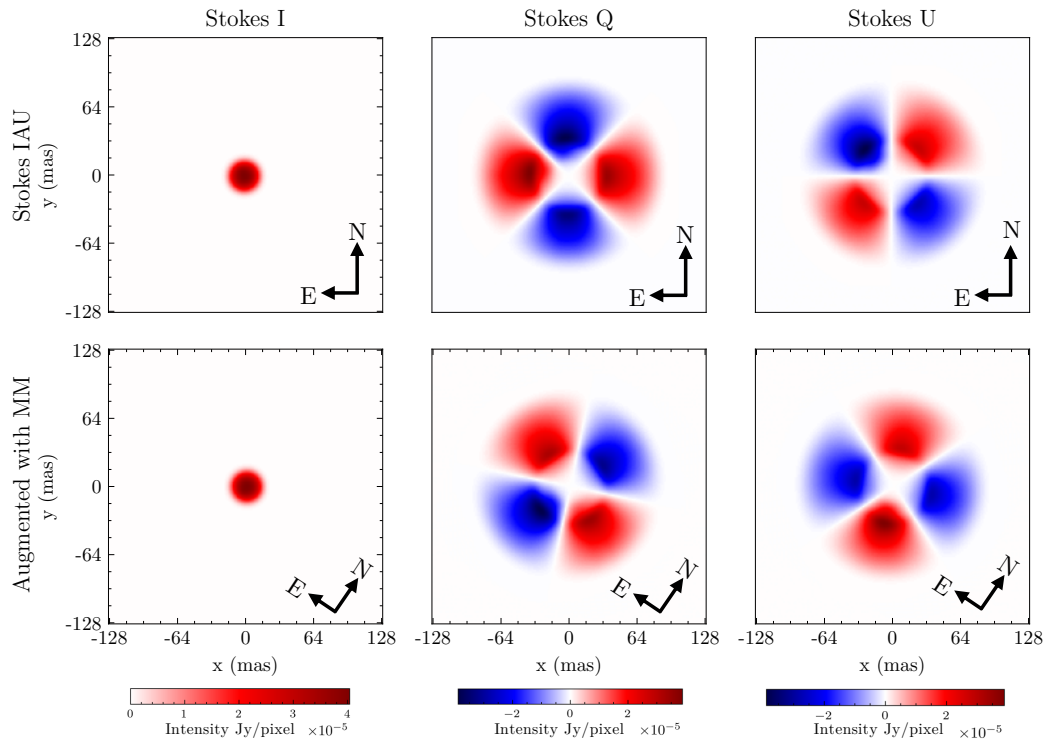
Parameter	Symbol	unit	760 nm	720 nm	670 nm	610 nm	Ideal
HWP retardance	$\delta_{\text{HWP}}$	$\lambda$	0.5	0.48	0.46	0.45	0.5
HWP offset	$\theta_{\text{HWP}}$	$^\circ$	-4.8	-2.4	-0.5	5*	0
IMR retardance	$\delta_{\text{IMR}}$	$\lambda$	0.53	0.42	0.32	0.22	0.5
IMR offset	$\theta_{\text{IMR}}$	$^\circ$	0	0	0	0	0
optics retardance	$\delta_{\text{OPTICS}}$	$\lambda$	0.07	0.04	0.86	0.18	0
optics diattenuation	$\epsilon_{\text{OPTICS}}$	-	0.17	0	0	0.02	0
optics offset angle	$\theta_{\text{OPTICS}}$	$^\circ$	-85	-2	-82	4.4	0
FLC retardance	$\delta_{\text{FLC}}$	$\lambda$	0.5	0.52	0.48	0.49	0.5
FLC offset angle	$\theta_{\text{FLC}}$	$^\circ$	1.3	1.1	-0.9	-4.8	0
dichroic retardance	$\delta_{\text{DICH}}$	$\lambda$	0	0	0.16	0.03	0
dichroic diattenuation	$\epsilon_{\text{DICH}}$	-	0.36	0.3	0	0	0
dichroic offset angle	$\theta_{\text{DICH}}$	$^\circ$	-45	-45	86	63	0

**Table 3.2:** Empirically fitted Mueller matrix coefficients for upgraded VAMPIRES, obtained by private communication with the authors of Zhang et al. 2023 who obtained these coefficients by fitting internal polarised source data with an MCMC and the model defined in Equation 3.19. \*A 2.5 degree adjustment was added to this value in consultation with the wider polarimetry team on account of a miscalibration present within the 610 nm matrix, which is the subject of ongoing calibration work.

It is worth noting that each coefficient on its own may not be interpreted as physically meaningful – as the entire optical train was fitted simultaneously. A good example of this is the variable ‘optics offset angle’ for upgraded VAMPIRES. It is highly implausible that a physical optic has an offset that varies by  $80^\circ$  over 40 nm of wavelength change (eg. compare 625 and 675 nm). Instead, as there are many degenerate ways to apportion non-ideal behaviour among modelled components, refitting the coefficients for each wavelength can result in different allocations of the non-ideal behaviour – even though the underlying optical system is the same and the outcome of applying the fitted Mueller matrix is unaffected. Additionally, although cursory inspection of some empirically fit values might suggest that the instrument behaves *very* non-ideally – this is also not the case. The idealised instrument model provides an excellent approximation for all wavelengths except the shortest (625 nm and 610 nm), which suffer most from a dichroic that splits light between VAMPIRES and a visible wavefront sensor. We used these empirically fitted coefficients (Tables 3.1 and 3.2) to build a forwards model for the VAMPIRES instrument, creating a function that produces the polarimetric coordinates on the detector at the time of each observation. A clear description of the polarimetric coordinates on the detector is essential for our approach to modelling polarimetric NRM data.

### NRM’s ‘detector’ polarisation coordinate system

When performing PDI, the instrument model and image transformations are applied to the raw images, returning them to the Stokes IAU celestial coordinate system for inspection. However, when polarimetric NRM is performed, there are no images to apply the instrument model to until a model has been fitted to the polarised interferometric observables. To fit a model to the polarised interferometric observables in the first place, the polarimetric coordinate system on the detector must be known so that models can be fitted within that local polarimetric coordinate system (Figure 3.2).



**Figure 3.2:** Model of a thick circumstellar envelope in the Stokes IAU celestial coordinate system (top row) and in the local detector coordinate system for the instrument configuration in a given observing epoch (bottom row). Here we fit observables within the local ‘detector’ coordinate system (bottom row) which is obtained by applying the instrument Mueller matrix model in the forwards direction.

Defining a coordinate system with respect to a detector needs care. Here we adopt a polarimetric

coordinate system defined relative to the perpendicular detector edges which are fixed to the instrumental bench and have a known orientation relative to the telescope’s field of view on sky. Within the present work we generate one instrument model (Mueller matrix) for every observation, which is used during model fitting. The images generated from fitted models are then de-rotated into the celestial and Stokes IAU coordinate system for examination.

## 3.2 A facility data reduction pipeline for VAMPIRES + NRM

I wrote a data reduction pipeline for VAMPIRES + NRM and integrated it into the SCEXAO facility supported pipeline `vampires_dpp`<sup>2</sup>, which historically has only provided PDI support. The addition of an NRM pipeline into `vampires_dpp` is a significant step towards making VAMPIRES + NRM more accessible to the broader scientific community, as VAMPIRES matures towards becoming a facility instrument in the coming years. A schematic of the NRM pipeline is outlined in Figure 3.3. Our pipeline features four main components: `nrm`: which performs raw frame pre-processing and cleaning, `mask alignment`: where the mask is aligned to sample the power spectra, `run amical`: where the AMICAL package (Soulain et al. 2020) is used to extract interferometric observables from the data, and `triple diff`: where triple differential interferometric quantities are computed. Each of these components are now detailed in turn.

### Component 1 – nrm

The following pre-processing steps are required prior to extracting interferometric observables.

#### FLC de-interleaving

Raw data cubes output from VAMPIRES each have a unique HWP angle and camera but have FLC frames ‘interleaved’ – ie, each frame switches between FLC state ‘A’ and ‘B’. Frames must first be ‘de-interleaved’ to create data cubes that each have a unique HWP, camera and FLC state. To perform de-interleaving a script was already available in `vampires_dpp` (Lucas et al. 2024).

#### Dark Frame subtraction

Dark frames are obtained on the same night of observation with the same instrumental configuration as science measurements. Dark frame statistics (mean, mode, std) are examined for obvious artefacts or structural noise prior to compilation of master dark frames which are subtracted from every frame.

#### Frame re-centring

Each frame is re-centred on the brightest pixel to ensure that windowing applied subsequently does not remove information from the edges of the interferogram. The frames are padded with zeros after cropping. Re-centring employed within `vampires_dpp` uses an automated method to centre the PSF.

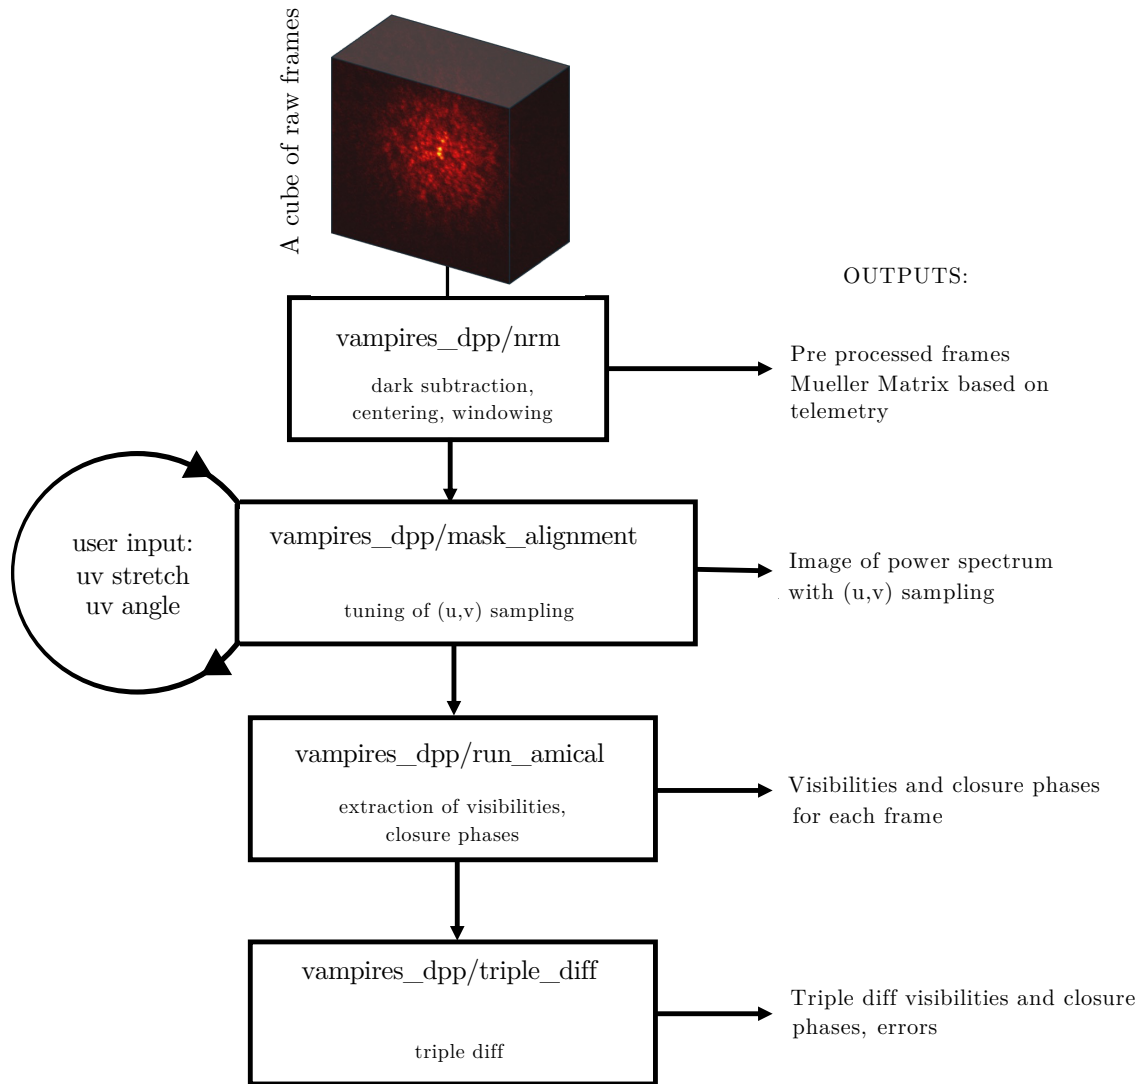
#### Windowing

A Gaussian window function is applied to every frame to ensure that the flux is sufficiently attenuated towards the edges of each frame – this avoids the introduction of classical ‘cross artefacts’ into the power spectra from non-zero image edges. Care must be taken to ensure the window size does not attenuate visibilities from short baselines which are closer to the data array’s edge – to facilitate this, `vampires_dpp` displays the interferogram and window function on a log stretch.

#### Sense checks

To confirm there were no significant disruptions to an observation several ‘sense check’ statistics are plotted as a function of observing time (the average and mode frame flux). It is possible to use these statistics to diagnose significant changes to the conditions or to the AO performance, both of which may require data to be trimmed from the total integration time to yield observationally consistent measurements.

<sup>2</sup>[https://github.com/scexao-org/vampires\\_dpp](https://github.com/scexao-org/vampires_dpp)




**Figure 3.3:** Flowchart of our pipeline for processing VAMPIRES + NRM data, which is now integrated into the `vampires.dpp` SExAO facility pipeline. Our pipeline has four main components: `nrm`, `mask_alignment`, `run_amical` and `triple_diff`. The pipeline outputs are the raw and reduced polarised interferometric observables – visibilities and closure phases.

## Component 2 – mask alignment

AMICAL has hard coded parameters for the g18 mask stored in `ami_function/get_mask`. These coordinates are then used to generate sampling regions in (u,v) space. However, small changes in staging and focus require the user to modify two parameters that adjust the (u,v) sampling: `theta_detector` – a radial stretch, and `scaling_uv` – a rotation angle. AMICAL function `amical/_show_ft_arr_peak` generates a figure that overlays the (u,v) sampling coordinates on the log stretched power spectrum so that these parameters are visually easy to update (Figure 3.4).

At present this adjustment is manual. Automation of (u,v) alignment will be introduced in a future version of the pipeline. Mask parameters `theta_detector` and `scaling_uv` may require some adjustment for each observation and should be confirmed before every data reduction. There are three different visibility sampling methods available within the AMICAL package – at the time of writing this package the `square` method was the most thoroughly tested and so is used here.



**Figure 3.4:** A power spectrum of a VAMPIRES + NRM frame. The splodges (purple) are sampled with the (u,v) coordinates (overlaid points in various colours). The left panel shows the full power spectra, the right panel shows a section from the lower right quadrant of the same power spectra. The power spectra are log stretched to maximise contrast, revealing power on all baselines despite a large intrinsic dynamic range in signal. This ensures that alignment can be optimally tuned by the user for all baselines.

## Component 3 – run amical

The AMICAL pipeline (Soulain et al. 2020) is a publicly available aperture masking pipeline for which there is a VAMPIRES branch.<sup>3</sup> We were involved in the development of this branch and used its initial release to identify bugs and recommend improvements to the authors. Our `vampires_dpp` NRM pipeline deploys AMICAL function `extract_bs.py` to sample interferometric observables from power spectra. Each observational data cube (each `.fits` file) is passed to this function and the output is saved as a `.h5` file. We use the following variables from each `.h5` file in subsequent reductions – here `file` denotes the loaded AMICAL `.h5` object,  $N_f$  denotes the number of temporal frames, and  $N_{\text{pix}}$  denotes the number of spatial pixels:

- `file['matrix'].v2_arr` –  $[N_f, N_{\text{pix}}, N_{\text{pix}}]$ : visibilities computed for every frame
- `file['matrix'].cp_arr` –  $[N_f, N_{\text{pix}}, N_{\text{pix}}]$ : closure phases computed for every frame
- `file.mask.bs2bl_ix` –  $[3, 816]$ : closure phase closing triangle indices
- `file['u']` –  $[153,]$ : u coordinates used for sampling interferometric quantities
- `file['v']` –  $[153,]$ : v coordinates used for sampling interferometric quantities

<sup>3</sup>[https://github.com/SAIL-Labs/AMICAL/tree/vampires\\_dev](https://github.com/SAIL-Labs/AMICAL/tree/vampires_dev)

*A word of caution about (u,v) coordinates*

As previously described, VAMPIRES has two cameras mounted to view opposing beams emerging from the polarising beam-splitter. This means that the (u,v) coverage is flipped between cameras. It can be confusing if the (u,v) coverage used is not consistent in plotting and modelling – within the present work we always use the (u,v) coverage of camera 2 as the reference. The choice is arbitrary so long as it is consistent. It is essential that the same (u,v) coverage is also used when making comparisons to the instrument model, as using an alternative camera’s (u,v) coverage will appear as a flip in polarimetric coordinates at first inspection.

**Component 4 – triple diff**

With visibilities and closure phases extracted for each frame, triple differential visibilities and closure phases are now computed. First, for every instrumental configuration (of HWP, camera and FLC), all visibilities and closure phases are concatenated together. This implicitly assumes the entire observing campaign has experienced minimal rotation of the polarimetric coordinates (we found that up to  $\sim 3^\circ$  was acceptable). The information for each instrumental state is now contained in an array (M, 153) for visibilities, and (M, 816) for closure phases, where M denotes the number of frames (and therefore number of measurements). To compute the final triple differential measurements and to quantify the amount of random error present, we use bootstrapping (Efron 1979). The procedure for bootstrapping is outlined below and the outputs are summarised in Table 3.3.

Output	Dimensions	Meaning	Variable Name
$Q_{\text{all, vis}}$	[153, N]	All Q vis triple-diff bootstrap samples	qvis_all
$U_{\text{all, vis}}$	[153, N]	All U vis triple-diff bootstrap samples	uvis_all
$Q_{\text{all, CP}}$	[816, N]	All Q CP triple-diff bootstrap samples	qcp_all
$U_{\text{all, CP}}$	[816, N]	All U CP triple-diff bootstrap samples	ucp_all
$Q_{\text{final, vis}}$	[153,]	Q vis triple-diff bootstrap average	qvis_final
$U_{\text{final, vis}}$	[153,]	U vis triple-diff bootstrap average	uvis_final
$Q_{\text{final, CP}}$	[816,]	Q CP triple-diff bootstrap average	qcp_final
$U_{\text{final, CP}}$	[816,]	U CP triple-diff bootstrap average	ucp_final
$\Delta Q_{\text{final, vis}}$	[153,]	Q vis triple-diff bootstrap standard deviation	qvis_err
$\Delta U_{\text{final, vis}}$	[153,]	U vis triple-diff bootstrap standard deviation	uvis_err
$\Delta Q_{\text{final, CP}}$	[816,]	Q CP triple-diff bootstrap standard deviation	qcp_err
$\Delta U_{\text{final, CP}}$	[816,]	U CP triple-diff bootstrap standard deviation	ucp_err
$M_{\text{total}}$	[4,4]	Instrumental Mueller matrix	M

**Table 3.3:** Outputs of the `vampires_dpp` NRM pipeline for the g18 mask — where there are 153 baseline vectors and 816 closure phase quantities. All bootstrapping samples are stored along with the bootstrapped mean and standard deviation of each observable. The instrumental Mueller matrix is also supplied, to be used within any subsequent modelling of the data.

For N bootstraps (N chosen when the observed quantities are not modified by adding new bootstraps):

- Randomly sample visibilities for each state – M times from [M, 153] with replacement. Average the visibility populations for each state.
- Randomly sample closure phases for each state – M times from [M, 816] with replacement. Average the closure phase populations for each state.
- Compute  $Q_{\text{final,vis}}$  and  $U_{\text{final,vis}}$  (Equations A.1, A.2, A.3, A.4) and store
- Compute  $Q_{\text{final,CP}}$  and  $U_{\text{final,CP}}$  (Equations A.5, A.6, A.7, A.8) and store

After  $N$  bootstraps:


- Final values of each quantity:  $Q_{\text{final}_{\text{vis}}}$ ,  $U_{\text{final}_{\text{vis}}}$ ,  $Q_{\text{final}_{\text{CP}}}$ ,  $U_{\text{final}_{\text{CP}}}$  are obtained by taking the mean over all  $N$  bootstraps.
- Final errors of each quantity:  $\Delta Q_{\text{final}_{\text{vis}}}$ ,  $\Delta U_{\text{final}_{\text{vis}}}$ ,  $\Delta Q_{\text{final}_{\text{CP}}}$ ,  $\Delta U_{\text{final}_{\text{CP}}}$  are obtained by taking the standard deviation over all  $N$  bootstraps.

### 3.3 Polarimetric calibration limit of upgraded VAMPIRES

We made many attempts to characterise the intrinsic polarisation calibration limits of upgraded VAMPIRES, to facilitate comparison with original VAMPIRES. To do so we observed bright un-polarised calibrators like Alioth and Vega between 2023-2025. However, persistent issues arose in reaching the noise floor, with the total integration time devoted appearing systematically insufficient despite prior calculation of the required integration time. This discrepancy was eventually traced to a computational inefficiency arising from the control electronics: details of the problem and its resolution are outlined in Section 3.4. The best measurement we obtained for the polarimetric calibration limit of upgraded VAMPIRES was in 2025 (Figure 3.5).

Here we plot Stokes Q which has a residual error of  $\sim 0.44\%$ , similar to the  $\sim 0.42\%$  residual error quoted for the polarimetric calibration of the original VAMPIRES instrument (Norris et al. 2015a). It was noted in Norris et al. 2015a (alongside equivalent measurements for original VAMPIRES) that their estimation of the polarimetric calibration limit would benefit from additional integration time – as would ours, as evidenced by the left panel of Figure 3.5. Following the procedure in Norris 2015, we use this polarimetric calibration to estimate the magnitude of systematic errors, by increasing the calibrator’s error bars until a  $\chi_{\text{red}}^2 = 1$  is obtained (representing an ideal polarimetric calibration). From this data we found a factor of 2.5 was required and adopt this value when modelling  $\mu$  Cephei.

Work to characterise polarisation aberrations (non-uniformity in optics) is forthcoming and may be used in future to provide more accurate constraint of systematic errors, which will permit recovery of a better polarimetric calibration in post processing (Ashcraft et al. 2025). Given our (and Norris et al. 2015a’s) limited total integration time, we cannot draw firm conclusions on the relative performance of the original and upgraded VAMPIRES instruments. However, our analysis suggests they have similar performance and provides an upper calibration limit of  $\sim 0.44\%$  for upgraded VAMPIRES.



**Figure 3.5:** Measuring the intrinsic polarisation calibration limit of upgraded VAMPIRES – observations of Vega were obtained on 2025/05/09 using the g18 aperture mask. The polarimetric calibration limit of the instrument has not yet been reached within the requested integration time. This is indicated by the black curve (left) which has not yet plateaued. A slope of zero indicates that additional measurements will not improve the polarimetric calibration and that data are limited only by systematics. The best polarimetric calibration limit we obtained is plotted on the right – our error is  $\sim 0.44\%$  which is similar to  $\sim 0.42\%$  quoted for the original VAMPIRES instrument. Errors plotted are obtained from bootstrapping and represent the amount of random error within the measurements.

### 3.4 A 30% improvement to observing efficiency


In the process of obtaining data for Section 3.3 we found that the total integration time required to reach the polarimetric calibration limit was systematically under-estimated, despite use of deterministic relations that predict theoretical observing efficiency. Our final attempt at this measurement (obtained in May 2025) is displayed in Figure 3.5 where the left panel demonstrates that the noise floor has not been reached. During this observation we found that for 126 minutes of total ‘on sky’ time only  $\sim 13$  minutes of data were obtained per camera (Figure 3.6). Figure 3.6 displays 100 seconds of our observation from Section 3.3. Time acquiring data is plotted in blue against the time spent ‘on sky’, revealing that significant amounts of ‘down time’ occur where no data is acquired.

We could not explain these latent periods with known instrumental overheads (like HWP switching or camera triggering). In collaboration with the VAMPIRES/SCEXAO instrument team, further investigation of this issue traced the cause to the data acquisition system’s management of parallel processing during ‘fast’ FLC read-out observations (in which all of our NRM datasets are obtained). The upgraded VAMPIRES data acquisition process involves two continuous streams of ‘video’ on each camera and the data is then ‘synchronised’ into files of each FLC state as part of post processing. This system involves two key processes per camera:

- A bufferiser, which handles acquisition timing and temporarily stores frames in memory
- A FITS writer, which uses the `cfitsio` library to write the buffered frames to disk

These two processes manage the transition from high-frame-rate video acquisition to the output FITS data cubes. With two cameras in operation the system runs a total of four processes during each observation. The root cause of the observing inefficiency we identified was traced to these four processes being scheduled on a single CPU core. The operating system’s scheduler was rapidly switching between processes sharing the same core, undermining the true parallelism which the VAMPIRES computers rely on to process and save data in real time. The solution implemented by the SCEXAO team was to reconfigure software to allocate four dedicated CPU cores during observation.

Fixing this issue improved observing efficiency by rates of up to  $\sim 30\%$ .<sup>4</sup> This discovery explained why we had struggled to reach the noise floor in our earlier engineering observations (Section 3.3). This issue appears to have been present in the computational software from VAMPIRES’ upgrade in mid-2023 until its correction in mid-2025. Rectification of this issue will significantly lower observational overheads for future users of the fast-FLC mode (predominantly NRM users) and will re-enable reliable budgeting of integration time.



**Figure 3.6:** A representative time series selected from our observation of Vega on 2025/05/09. Between sparse periods of data acquisition (blue), there is a significant amount of ‘down time’ (white). Despite a heavy penalty expected in FLC ‘fast’ mode and the compulsory periods of switching time intrinsic to instrumental overheads, the length of the down time periods could not be explained with understood sources.

<sup>4</sup>The exact improvement yielded by our resolution depends on the frame size and exposure time. We calculated an improvement of 30% for FAST-mode FLC data obtained with an exposure time of 1 ms, which is representative of our Vega engineering observations depicted in Figure 3.5.

### 3.5 Observing Campaign of $\mu$ Cephei

This section describes the VAMPIRES + NRM data of  $\mu$  Cephei used within this thesis. Our data (summarised in Table 3.4) is a combination of existing archival observations (datasets 1-4), and newly obtained datasets (datasets 5-9). In collaboration with my supervisors I wrote three successful observing proposals for open-use time on the Subaru Telescope and observed with VAMPIRES in Semesters S23A, S24B and S25A.

Our observations in S23A (dataset 5) suffered from an FLC malfunction which rendered this dataset unusable. However, we were able to use  $\mu$  Cephei as a dual-purpose engineering target within the following month (dataset 6). Observations obtained in 2024 (datasets 7 and 8) await a polarimetric calibration model from the instrument team as it appears there was internal re-alignment after a long period of down time in late 2023.<sup>5</sup> Our 2025 observations (dataset 9) suffer from a significant polarimetric miscalibration issue that is common to multiple observations within this period and is the subject of an ongoing observation. As a result we were only able to contribute our 2023 (dataset 6) to existing archival datasets, yielding a total observing campaign that spans a 7 year period.

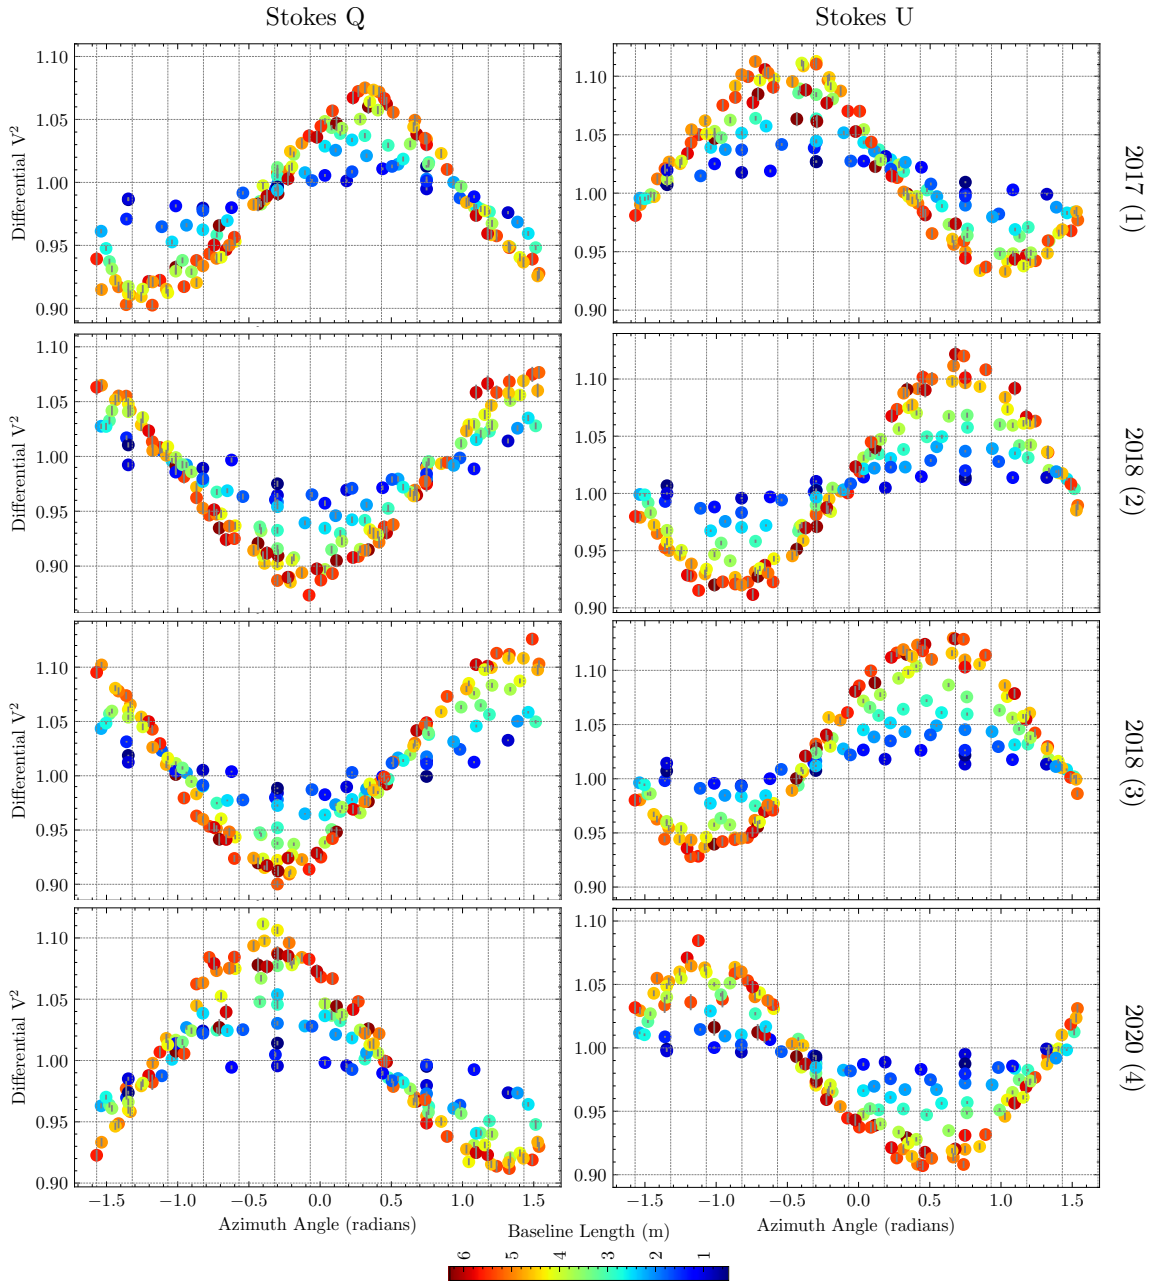
Target	Code	Date	Observer	Outcome	Modelled?	$\lambda$
$\mu$ Cephei	1	19/06/2017	B. Norris	Good	✓	S
$\mu$ Cephei	2	23/10/2018	B. Norris	Good	✓	S
$\mu$ Cephei	3	23/10/2018	B. Norris	Good	✓	S
$\mu$ Cephei	4	07/12/2020	B. Norris	Good	✓	S
$\mu$ Cephei	5	27/06/2023	L. Lilley	FLC Issue	✗	M
$\mu$ Cephei	6	07/07/2023	L. Lilley, M. Lucas	Good	✓	MBI
$\mu$ Cephei	7	26/07/2024	L. Lilley, M. Lucas	IP Cal	✗	MBI
$\mu$ Cephei	8	26/07/2024	L. Lilley, M. Lucas	IP Cal	✗	MBI
$\mu$ Cephei	9	09/05/2025	L. Lilley	Pol Miscal	✗	MBI

**Table 3.4:** Summary of our observing campaign for  $\mu$  Cephei. The outcome of each observation is listed in column 5 – here ‘FLC Issue’ denotes an FLC malfunction, ‘IP Cal’ denotes that the observations await an instrumental polarisation model, and ‘Pol Miscal’ denotes that there is a polarisation miscalibration issue which is the subject of ongoing work. The datasets which could be modelled (column ‘Modelled?’) in this thesis are marked with a ✓ – those that could not are marked with a ✗. The wavelength range of each observation is marked as either S – single-band ( $750 \pm 25$  nm), M – multiple single-band measurements or MBI – multi-band measurements ( $610, 670, 720, 760 \pm 25$  nm).

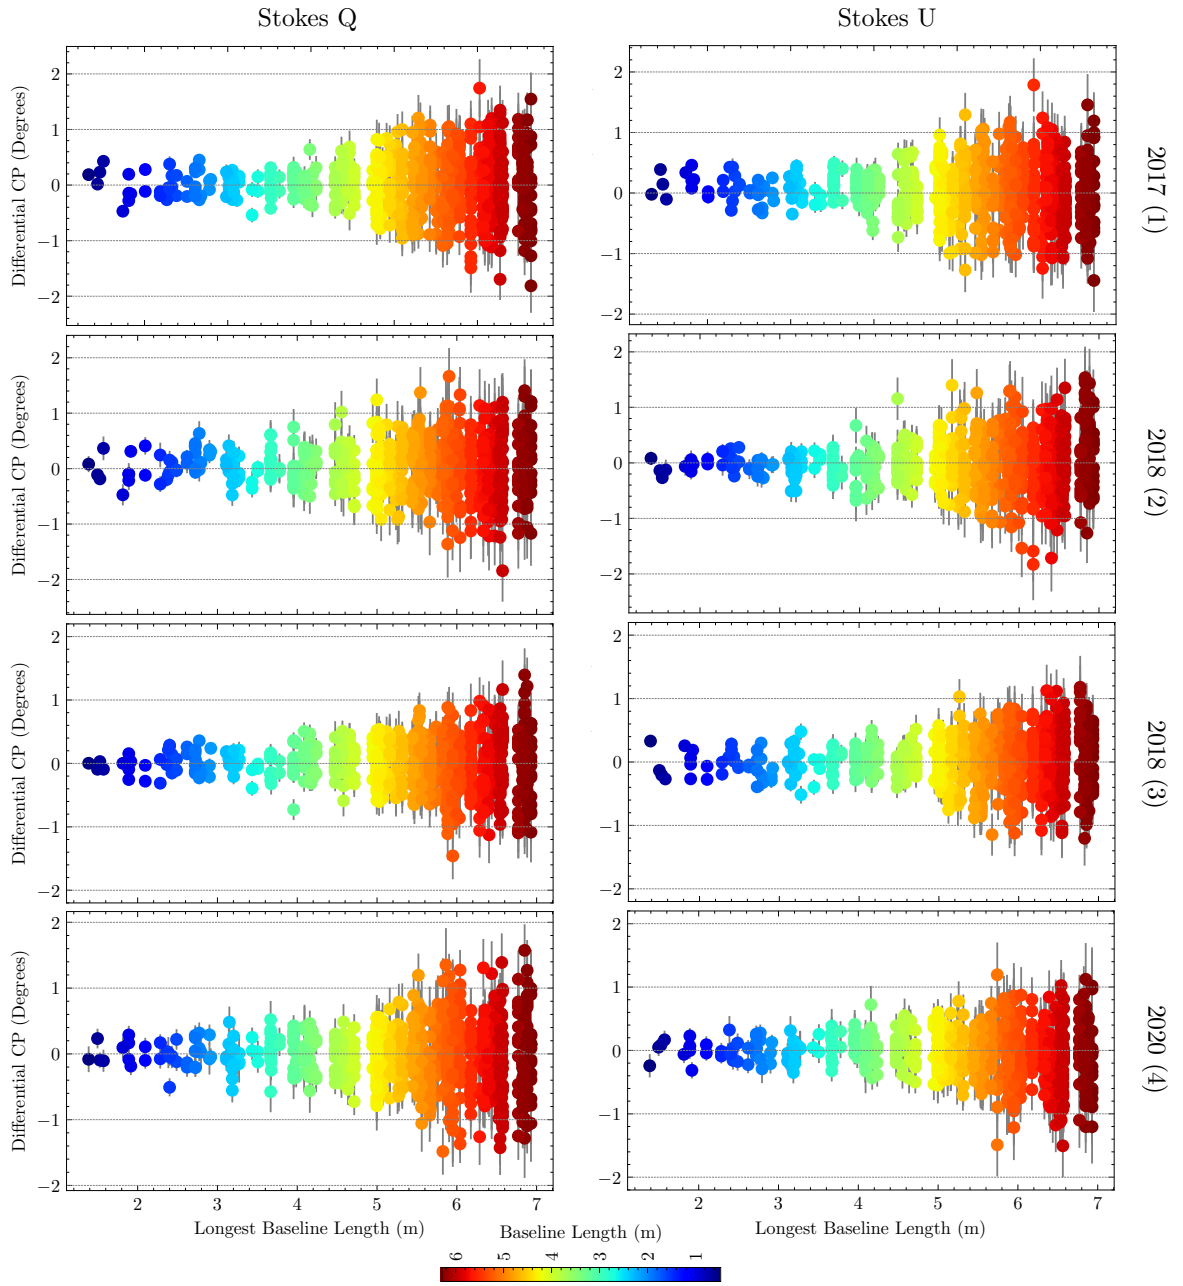
### 3.6 Reduced Data Products of $\mu$ Cephei

Reduced datasets for  $\mu$  Cephei – which are polarised interferometric visibilities and closure phases – are displayed within this section. Figures 3.7 and 3.8 depict observables from our archival data on  $\mu$  Cephei, spanning 2017-2020 (datasets 1-4). Figures 3.9 and 3.10 depict observables from our newly obtained data on  $\mu$  Cephei in 2023 (dataset 6). Errors plotted are obtained by bootstrapping and do not include the systematic error factor of 2.5 derived in Section 3.3, for ease of visualisation.

<sup>5</sup>A guard rail fell onto Subaru’s primary mirror in late 2023 and the telescope was offline whilst the damage was repaired <https://subarutelescope.org/en/news/announcements/2023/10/16/3312.html>



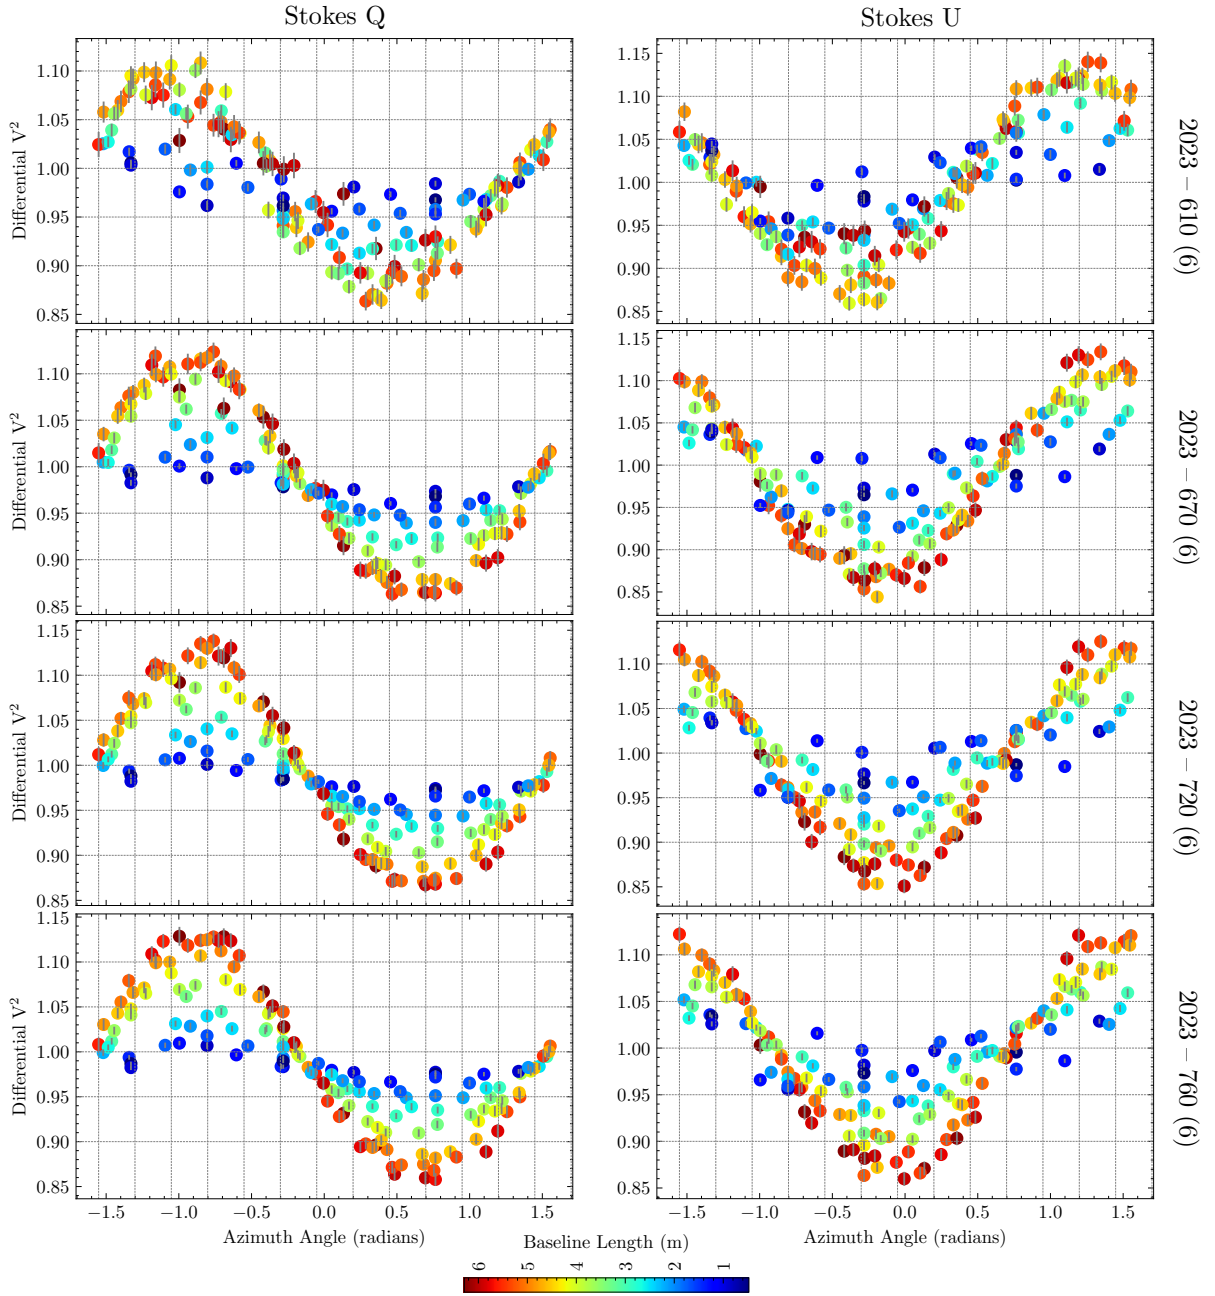
**Figure 3.7:** Polarised differential visibilities for archival datasets 1-4 (2017-2020), obtained with the original VAMPIRES instrument. Each dataset is obtained in a unique polarimetric coordinate system. Data are obtained at  $\lambda = 750 \pm 25$  nm. Dataset codes denoting observation date are given on the right of each row.



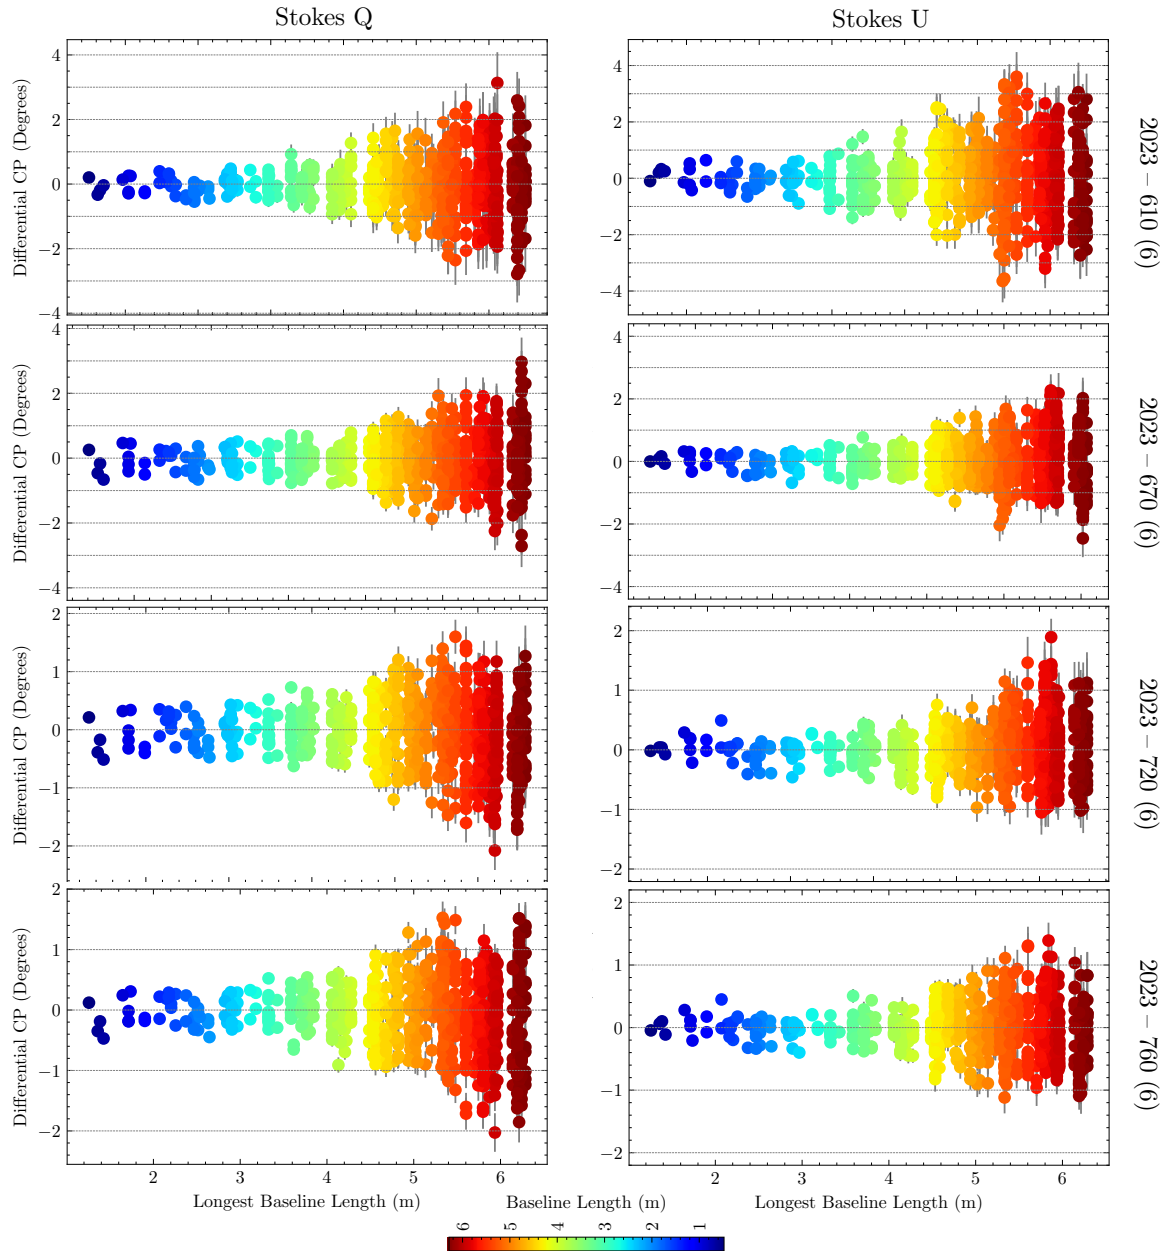
**Figure 3.8:** Polarised differential closure phases for original datasets 1-4 (2017-2020), obtained with the original VAMPIRES instrument. Each dataset is obtained in a unique polarimetric coordinate system. Data are obtained at  $\lambda = 750 \pm 25$  nm. Dataset codes denoting observation date are given on the right of each row.

Observing logs from archival data:

- Dataset 1 (2017): Target at very low airmass.
- Dataset 2 (2018): PWFS loop open, AO188 only.
- Dataset 3 (2018): Both loops closed.
- Dataset 4 (2020): Seeing is  $0.6''$ .



**Figure 3.9:** Polarised differential visibilities for our newly obtained dataset 6 (2023). Data are obtained at  $\lambda = 610, 670, 720, 760 \pm 25$  nm in the ‘multi-band’ mode of the upgraded VAMPIRES instrument. Dataset codes denoting observation date are given on the right of each row.




**Figure 3.10:** Polarised differential closure phases for our newly obtained dataset 6 (2023). Data are obtained at  $\lambda = 610, 670, 720, 760 \pm 25$  nm in the ‘multi-band’ mode of the upgraded VAMPIRES instrument. Dataset codes denoting observation date are given on the right of each row.

Observing log from new data:

Dataset 6 (2023): Good seeing.

### 3.7 Diagnosis and correction of faults in archival VAMPIRES data

When processing archival data some anomalies were detected that require discussion and correction. As described in Section 3.1 each observing epoch has a single instrumental Mueller matrix model. It is useful to propagate a synthetic polarised image through the instrument model – we recommend using a spherically symmetric circumstellar shell with size that matches the highest visibility amplitude of the data (here, red,  $\sim 6$  m). To first order, comparison of the synthetic image’s polarised visibilities with those of the science signal will confirm that the Mueller matrix is an accurate description of the detector polarimetric coordinates during observation (allowing for some small perturbations due to asymmetry and non-azimuthal scattering). When performing this sanity check for our observing campaign we found that the Mueller matrices for 2017, 2020 and 2023 were consistent with our data but that in 2018 (datasets 2 and 3) there was a significant discrepancy (Figure 3.11).




**Figure 3.11:** Comparison of polarised differential visibilities for  $\mu$  Cephei dataset 2 (2018/10/23) (bottom row), with a model envelope propagated through the Muller matrix instrument model (top row). The polarimetric coordinate system, as indicated by the phase of the sinusoidal and cosinusoidal polarised visibilities, should be approximately matched between the real data and model data (allowing for some minor distortions due to asymmetry and non-azimuthal scattering signal). Instead we found a large discrepancy which persists over both visits to  $\mu$  Cephei on 2018/10/23 (datasets 2 and 3).

To gain clarity on the cause of this discrepancy and to ascertain whether it was persistent, we performed an archival review of PDI data of AB Aurigae, a Herbig Ae/Be star surrounded by a large proto-planetary disk (Tang et al. 2017). The choice of AB Aurigae was twofold – firstly, it is a source dominated by azimuthal scattering and so is a reasonable proxy for a ‘resolved polarisation calibrator’ (Perrin et al. 2009b). Secondly, AB Aurigae has more than a decade of observational coverage within our archive with multiple epochs obtained near to our 2018 observations, making it best poised to ascertain the general behaviour of the instrument during this period of time.

Each archival dataset of AB Aurigae was compared to an instrument model (Section 3.1). We identified several trends from these comparisons. Firstly, we found that undocumented changes to the definitions of H and V (or H45 and V45) are common, a result of an error in labelling the FLC or camera states (which cannot be of physical origin during an observation). Resolving these issues is usually simple as

the instrument model requires a simple switch of H and V (and H45 and V45) to restore the true detector coordinates. A clear example of this is illustrated within Figure 3.12.

Conveniently there was archival PDI data of AB Aurigae from 2018/10/17 and 2018/10/22 in the nights prior to our 2018 observations on 2018/10/23. Reduction of these AB Aurigae datasets revealed a polarimetric mis-calibration preceding our observations that could not be resolved with the simple mislabelling of states described above. Instead, the error appears as a rotation of (almost)  $H \rightarrow H45$ ,  $V \rightarrow V45$ , and  $H45 \rightarrow V$ ,  $V45 \rightarrow H$  (Figure 3.13). According to instrumental telemetry from this evening, all wave plates were configured ‘typically’ in the positions implicitly assumed by the Mueller matrix, indicating that this rotation was not intentional.

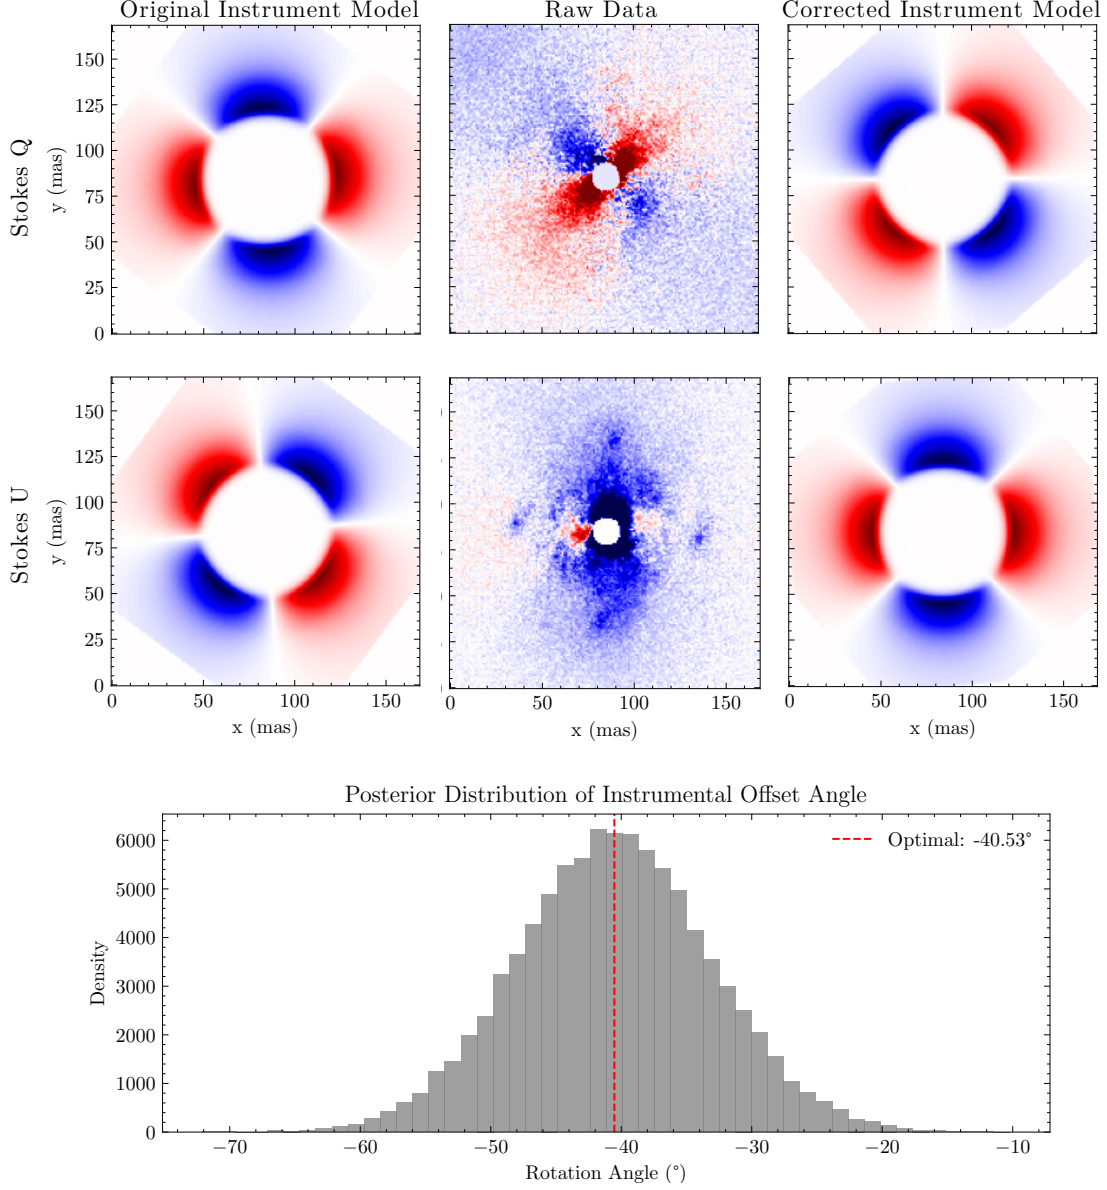


**Figure 3.12:** PDI data of AB Aurigae (Stokes Q is displayed in the bottom row) obtained 2017/03/13, compared to the predicted instrument state (top row). Each column corresponds to a different parallactic angle. The observed data are inconsistent with the instrumental model but are easily reconciled with a simple flip in convention of H/V and H45/V45. Within polarimetric images, red denotes positive and blue denotes negative.

Instrumental metadata from this era recorded the positions ‘requested’ of optics but not their true position, unless this information was specially requested prior to an observation (which for this data, it was not). The most likely candidates for a static rotational offset to the expected polarimetric coordinates are the homing units of the half wave plate and quarter wave plates (QWP). Observing logs from nights prior to these observations indicate that a series of QWP tests took place, wherein QWPs were moved and then (allegedly) returned to their usual positions. However, we suspect that the QWP homing unit failed during this last request. Regardless of the the exact cause, to fit to our 2018 data we were required to modify the Mueller matrices for datasets 2 and 3. To do so we solved for a single polarimetric rotation by fitting to our 2018/10/22 PDI data of AB Aurigae using an MCMC, yielding a waveplate offset of  $\sim 20^\circ$  (Figure 3.13). Applying this updated Mueller matrix model to our  $\mu$  Cephei data from 2018 produced excellent alignment for both datasets 2 and 3 (Figure 3.14).

We were also able to confirm that this offset was no longer present by 2019/03/20 where the un-modified instrument model again matches the polarimetric orientation of AB Aur. Our analysis highlights the importance of observing polarisation calibrators and that observing resolved polarimetric ‘coordinate standards’ at the beginning of any major observing campaign is highly advisable as an easy sanity check.

Our ability to diagnose and trace the polarimetric miscalibration issue afflicting our 2018 data was luck as there had been a simultaneous observation campaign for AB Aurigae.

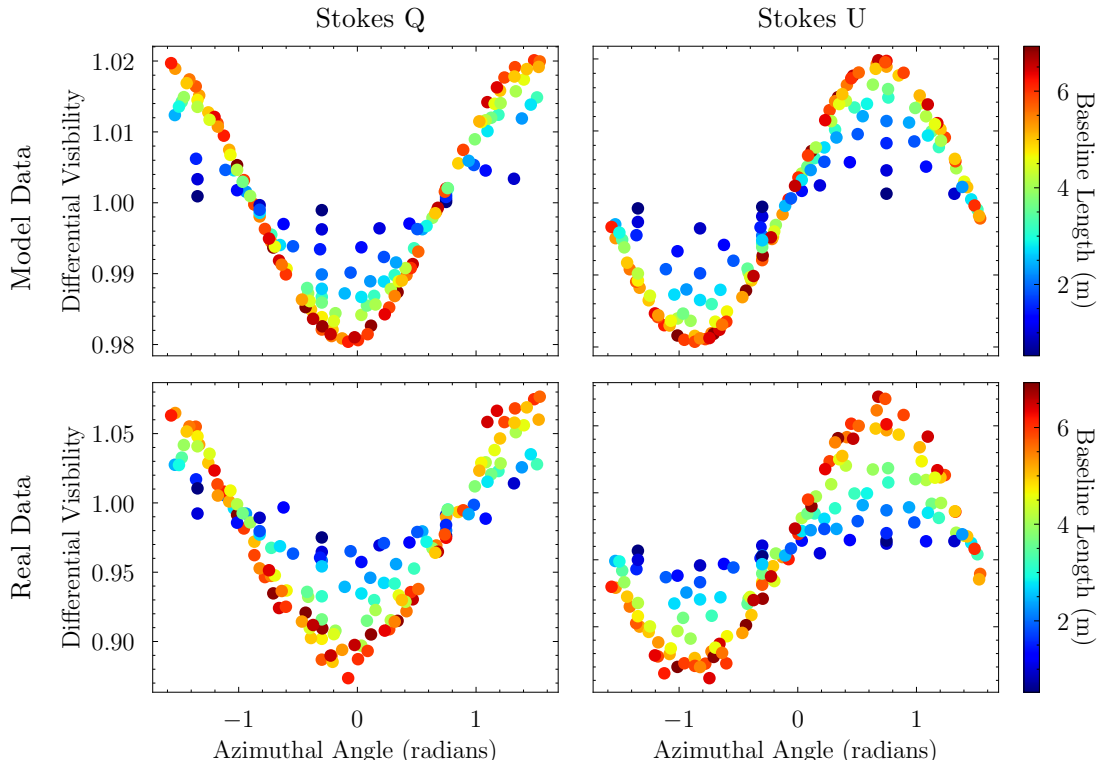


**Figure 3.13:** Fitting AB Aurigae PDI data from 2018/10/22 to solve for a correction to the Mueller matrix instrument model. We fitted to an additional image rotation (of the polarimetric coordinates) using the PDI images of AB Aurigae. We ran an MCMC for 50,000 epochs with 4000 discards (posterior distribution is displayed in the bottom panel). An image rotation of  $\sim 40^\circ$  was fitted which is equivalent to a HWP angle offset of  $-20^\circ$ , which we added to the HWP term of our Mueller matrix. This correction was applied to our  $\mu$  Cephei datasets 2 and 3. Within polarimetric images, red denotes positive and blue denotes negative.

We note more generally that VAMPIRES data pre-2019 has been described by the broader instrument team as ‘difficult’ to work with – the Mueller matrices do not always return the correct polarisation vectors for standard stars. <sup>6</sup> In our archival data review of AB Aurigae alone we found multiple instances of either state mislabelling or probable wave-plate re-homing issues. If these events are equally persistent

<sup>6</sup>Private communication with VAMPIRES Instrument polarimetry team.

throughout archival VAMPIRES data of this era then this would explain the difficulty of working with data from this period. Our method for applying polarimetric corrections to our 2018 data is simple yet effective – more nuanced methods may be required to unlock other datasets from this period of time.



**Figure 3.14:**  $\mu$  Cephei data (bottom row) and the now corrected instrument model (top row), with correction of  $-20.5$  degrees to the half wave plate angle applied as derived in Figure 3.13. The polarimetric coordinate systems now appear aligned and the Mueller Matrix is ready to be used in modelling of  $\mu$  Cephei datasets 2 and 3.

## 3.8 Chapter Summary

In this chapter we presented the instrumentation and engineering work required to develop the NRM mode of VAMPIRES for science use. To the SCEXAO facility code `vampires_dpp` we have now contributed a robust data reduction pipeline for VAMPIRES + NRM data. Our integration of an empirically fitted Mueller matrix model has improved upon the pre-existing treatment of VAMPIRES + NRM as it allows for complete disentanglement of astrophysical non-azimuthal polarisation signatures from instrumental polarisation effects. We then conducted engineering observations to characterise the polarimetric calibration limits of the upgraded VAMPIRES instrument, demonstrating it has a similar performance to the original VAMPIRES instrument. During this engineering we diagnosed a significant issue with VAMPIRES’ computational infrastructure, correction of which improved observing efficiency by up to 30%. We then conducted a multi-epoch observing campaign of  $\mu$  Cephei which was combined with archival datasets to yield a 7 year observation campaign on  $\mu$  Cephei. We then reviewed archival data, uncovering several trends that may explain the difficulty of working with data pre-dating 2019, from which we recommended methods to make polarimetric corrections where required.

# Chapter 4

## PIPPI

*“I have never tried that before, so I think I should definitely be able to do that.”*

---

— Astrid Lindgren,  
Pippi Longstocking

### Chapter Overview

A number of parametric modelling pipelines exist for interferometric and spectro-interferometric data (Baron 2020; JMMC n.d.; Mérand 2022; Martí-Vidal et al. 2014; Tazzari et al. 2018; Jennings et al. 2020; Chael et al. 2018). However, to date those that treat spectro-*polarimetric* interferometry have only been built for radio or sub-millimeter wavelengths (The CASA Team et al. 2022; Pesce 2021). At the time of writing there is no publicly available parametric modelling package for optical wavelength spectro-polarimetric interferometry data. Here we present the first, PIPPI – the ‘Parametric Inference Package for Polarimetric Interferometry’. The PIPPI package is a *differentiable* implementation of the simple scattering models which have been used historically to fit spectro-polarimetric interferometry data (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019). PIPPI’s creation has two purposes – to provide an accelerated inference framework to fit VAMPIRES + NRM data, and to permit reanalysis of the utility and physical realism of single-scattering models via comparison with radiative transfer.

### Contribution statement

The idea to fit parametric Rayleigh scattering models to VAMPIRES + NRM data in a Bayesian framework was originally my supervisors. The idea to expand PIPPI to include Mie scattering was my own. I designed PIPPI and implemented it in `jax`, with assistance from Louis Desdoigts in implementing complex `jax` functionalities. The design of models in PIPPI are my own.

## 4.1 A brief introduction to Parametric Modelling

A *parametric* model is one with a finite number of continuous variables, or ‘parameters’ (Bickel et al. 2015). At the outset of a parametric modelling study there is a decision to be made on the level of physical complexity required to model the system of interest. Within this decision there is a tension between computational efficiency and physical realism – typically motivating an initial choice of ‘first order’ models which are physically simple and computationally inexpensive. Even if the initial choice of physical complexity is unrealistically simple, quick inference enables this to be rapidly uncovered. In many ways this is an implementation of Occam’s Razor, a philosophy for which William of Occam is credited but the foundational ideas of which can be dated as far back as Aristotle’s (384–322 BC) Posterior Analytics – ‘*We may assume the superiority ceteris paribus (other things being equal) of the demonstration which*

*derives from fewer postulates or hypotheses*' (Aristotle 1981).<sup>1</sup> The simplest possible models should be fitted to data and model complexity should only be increased with quantitative justification. These principles drove the creation and use of single scattering models as the original parametric models used to fit optical wavelength spectro-polarised interferometry data (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019) – of which PIPPI is a differentiable implementation.

There are many algorithms for parametric model fitting. Here we employ a Markov Chain Monte Carlo (MCMC) method, a class of algorithms used to sample from a posterior distribution (Hogg et al. 2018). Use of an MCMC permits Bayesian inference, where model parameters can be expressed as probability distributions and uncertainty in their values quantified with respect to this distribution (Silvia D 2008). Bayesian inference is underpinned by Bayes' theorem (Equation 4.1) which defines the probability of an outcome – the *posterior*  $P(\theta_k | D, M_k)$  (here, of a model  $M_k(\theta_k)$  explaining the data  $D$ ), in terms of what is already known – the *prior*  $P(\theta_k | M_k)$ , how well the model fits the data – the *likelihood*  $P(D | \theta_k, M_k)$ , and a normalisation factor – the *Bayesian evidence*  $P(D | M_k)$  (Equation 4.2).

$$P(\theta_k | D, M_k) = \frac{P(D | \theta_k, M_k) P(\theta_k | M_k)}{P(D | M_k)} \quad (4.1)$$

$$P(D | M_k) = \int P(D | \theta_k, M_k) P(\theta_k | M_k) d\theta_k \quad (4.2)$$

The most probable value of each model parameter is the mean or mode of the marginalised posterior distribution, which provides a *maximum likelihood estimate* (MLE) (Equation 4.3) with uncertainty given by the posterior's width. PIPPI uses a chi-squared statistic ( $\chi^2$ ) as its likelihood function, which is the summed squared difference between the observed data  $O_i$  and model data  $M_i$ , scaled by observed errors  $\sigma_i$  (Equation 4.4). By default PIPPI uses uniform prior distributions for all parameters.

$$\hat{\theta}_{MLE,k} = \arg \max_{\theta_k} P(D | M_k) \quad (4.3)$$

$$P(D | \theta_k, M_k) = \chi^2 = \sum \frac{(O_i - M_i)^2}{\sigma_i^2} \quad (4.4)$$

After sampling the posterior distributions of each model, two models can be compared by using the ratio of their Bayesian evidences – the *Bayes factor*. Computing the Bayesian evidence  $P(D | M_k)$  requires marginalisation over all possible parameter values  $\theta_k$  (Equation 4.5).

$$P(D | M_k) = \int P(D | \theta_k, M_k) P(\theta_k | M_k) d\theta_k \quad (4.5)$$

However, this integral is intractable in high dimensions and may be approximated with the Laplace approximation, which expands the posterior around the maximum a posteriori (MAP) estimate under the assumption of a Gaussian distribution (Equation 4.6).

$$\theta_{MAP,k} = \arg \max_{\theta_k} P(\theta_k | D, M_k) = \arg \max_{\theta_k} \left[ P(D | \theta_k, M_k) \cdot P(\theta_k | M_k) \right] \quad (4.6)$$

The posterior is then approximated as Equation 4.7, where  $H_k$  is the Hessian of the negative log posterior at  $\theta_{MAP,k}$ .

$$\begin{aligned} \log \left[ P(D | \theta_k, M_k) P(\theta_k | M_k) \right] &\approx \log P(D | \theta_{MAP,k}, M_k) + \log P(\theta_{MAP,k} | M_k) \\ &\quad - \frac{1}{2} (\theta_k - \theta_{MAP,k})^T H_k (\theta_k - \theta_{MAP,k}) \end{aligned} \quad (4.7)$$

---

<sup>1</sup>This date is of the translation.

The log of the Bayesian evidence for model  $M_k$  is then given by Equation 4.8, where  $H_k^{-1}$  is estimated from a covariance matrix of posterior samples (within PIPPI, posterior samples are obtained from MCMC chains – explained further in Section 4.7).

$$\log P(D | M_k) \approx \log P(D | \theta_{\text{MAP},k}, M_k) + \log P(\theta_{\text{MAP},k} | M_k) - \frac{d}{2} \log(2\pi) - \frac{1}{2} \log |H_k^{-1}| \quad (4.8)$$

Computing the Bayes factor for two models  $M_1$  and  $M_2$  then enables Bayesian hypothesis testing via comparison of competing models (Equation 4.9) (Ando 2010).

$$B_{12} \equiv \frac{P(D | M_1)}{P(D | M_2)} = \frac{\int P(D | \theta_1, M_1) P(\theta_1 | M_1) d\theta_1}{\int P(D | \theta_2, M_2) P(\theta_2 | M_2) d\theta_2}. \quad (4.9)$$

To interpret the value of the ratio  $B_{12}$  we use the Jeffrey’s scale (Jeffreys 1998). For ease of interpretation within the present work we record the natural logarithm of the Bayes factor,  $\ln(B_{12})$ . The Jeffrey’s scale gives that values  $\ln(B_{12}) < 0$  support  $M_2$ , values  $0 < \ln(B_{12}) < 1$  are weak evidence for  $M_1$ , values  $1 < \ln(B_{12}) < 2.5$  are substantial evidence for  $M_1$ , values  $2.5 < \ln(B_{12}) < 5$  are strong evidence for  $M_1$  and values  $5 < \ln(B_{12})$  are very strong evidence for  $M_1$  (Jeffreys 1998).

One important limitation of the parametric modelling process is that it is intrinsically non-exhaustive. For each class of model, the question ‘which model within *this class* best fits my data’ may be posed and answered, but not ‘what is the *best possible* model to describe my data’ – this latter question is never truly answered. Whilst more complex models can be designed and their use justified with metrics like the Bayes factor, the exact form of additional complexities can be ambiguous. The parametric modelling process is typically ceased when a model is found that fits data to below  $\chi_{red}^2 < 1$ , where  $p$  is the number of model parameters and  $N$  is the number of observables (Equation 4.10). However, as the  $\chi_{red}^2$  statistic does not penalise for model complexity it must be accompanied with computation of the Bayes factor when justifying an increase in model complexity.

$$\chi_{red}^2 = \frac{1}{N - p} \sum \frac{(O_i - M_i)^2}{\sigma_i^2} \quad (4.10)$$

## 4.2 PIPPI – Parametric Inference Package for Polarimetric Interferometry

PIPPI (Parametric Inference Package for Polarimetric Interferometry) is a differentiable parametric modelling framework for optical wavelength polarimetric interferometry data. PIPPI is designed to provide a fast and statistically principled method for rapidly assessing the geometry, chemical composition and grain size of the dusty circumstellar environments of evolved stars. We developed PIPPI with the VAMPIRES Instrument in mind but maintained flexibility where possible to facilitate use of PIPPI on alternative instruments. PIPPI is available for public download via the SAIL Github. <sup>2</sup> PIPPI has five core components:

- I: *Astronomical Scenes*
- II: *Instrument Model*
- III: *Non Redundant Masking Simulation*
- IV: *Graphical User Interface*
- V: *Inference Framework*

---

<sup>2</sup><https://github.com/SAIL-Labs/VAMPIRES>

The core components and flow of information through PIPPI are outlined in Figure 4.1 and Table 4.1. Parts I-III include a complete forwards model of the VAMPIRES instrument – from the astronomical scene to the interferometric observables computed from the interferogram on the detector. Parts I-III are designed to be used iteratively as input to the Graphical User Interface (GUI) (Part IV) to facilitate the user’s design and understanding of candidate models and then as input to the Inference Framework (Part V) where models are fit to data using an MCMC framework.

PIPPI is a Python code written in `jax` (Bradbury 2018), a Python library providing code acceleration via the XLA compiler (Accelerated Linear Algebra). Once a function has been ‘just-in-time’ (JIT) compiled, subsequent function calls are substantially accelerated – providing a clear advantage when parallelised codes like MCMCs are used iteratively after a single compilation. Further acceleration is obtainable by running `jax` code on a GPU. Within this work we used a GeForce RTX 4090 and were able to accelerate each PIPPI model evaluation from  $\sim 3$  s to  $\sim 33$  ms. We now derive each component of PIPPI.

Function	Part	Inputs	Outputs
<code>make_dust()</code>	I	Model Parameters	Dust Density Grid
<code>make_scattering_grid()</code>	I	Dust Grain Size Dust Chemistry Star Size	Scattering Grids in IAU, NE
<code>apply_instrument_model()</code>	II	Altitude Azimuth	Scattering Grids on Detector
<code>simulate_nrm()</code>	III	(u,v) coordinates Indices of CP	Stokes I: $V^2$ , CP Stokes Q: $V^2$ , CP Stokes U: $V^2$ , CP
<code>gui_nrm()</code>	IV	Model Class (u,v) coordinates Indices of CP	GUI
<code>model_fitting_nrm()</code>	V	Model Class MCMC Parameters Real Data (Q,U: $V^2$ , CP)	MLE Diagnostic Plots $\chi^2$ , $\chi_{red}^2$

**Table 4.1:** Outline of PIPPI’s core components – inputs and outputs of each function are summarised. PIPPI Parts I-III define a complete forwards model for simulating observation of dusty circumstellar scenes with the NRM mode of the VAMPIRES Instrument. A circumstellar dust geometry (I) is propagated through an empirically fitted instrument model to correct for instrumental polarisation (II), after which, observation with an aperture mask is simulated (III). PIPPI Parts I-III are parsed iteratively to either the GUI (IV) for model development, or to the model fitting framework (V).

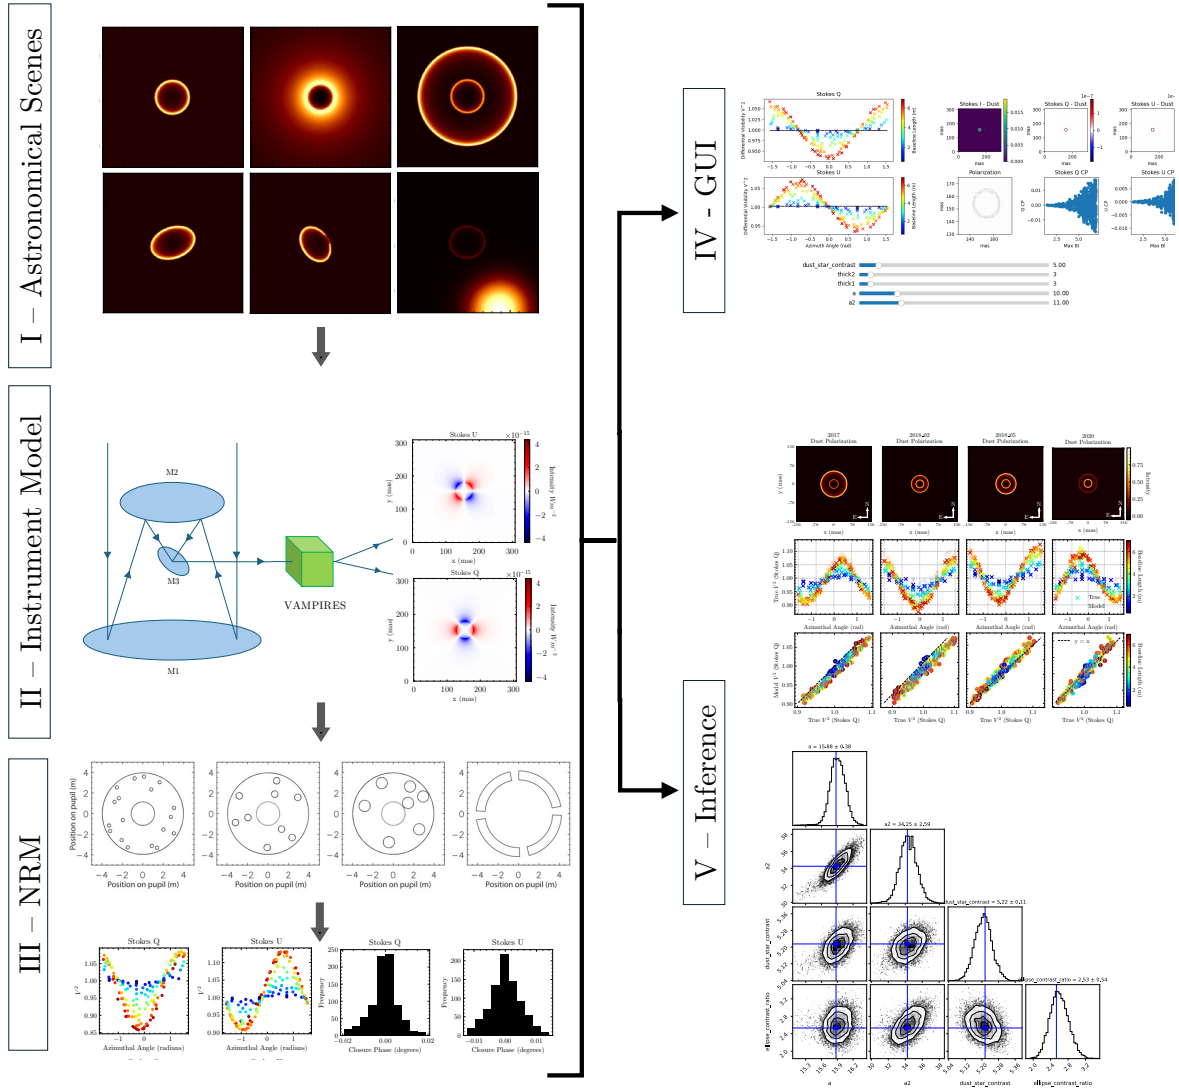
## 4.3 Part I – Astronomical Scenes

Part I of PIPPI generates dusty circumstellar environments and models single scattering through a simulated dust media to generate polarised images of Stokes I, Q and U. To do so, three dimensional dust density grids of the desired circumstellar morphology are first generated.

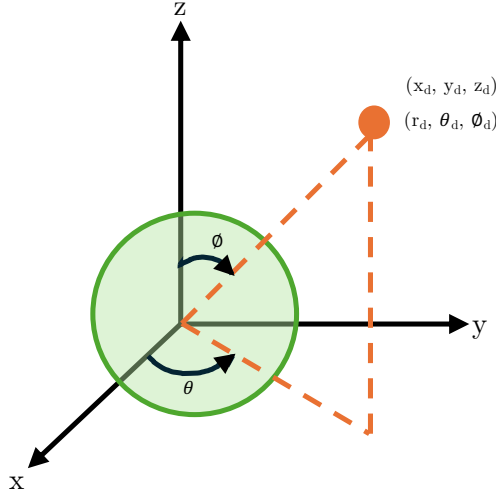
### Generation of Dust Morphologies

PIPPI function: `make_dust.py`

The equations used to define circumstellar dust geometries are defined below in the coordinate system of Figure 4.2. To generate dust volumes a Gaussian function  $G$  is used (Equation 4.14). From this, thin shells  $S$  (Equations 4.13, 4.15) and power law outflows  $S_p$  (Equation 4.16) may be generated.



**Figure 4.1:** An overview of the PIPPI package – Parametric Inference Package for Polarimetric Interferometry. PIPPI has five core components: Part I: Astronomical Scenes – dusty circumstellar environments are designed by the user, Part II: Instrument Model – an empirically fitted Muller Matrix model is generated to determine the polarisation coordinates at the time of observation and correct for non-ideal instrumental polarisation, Part III: Simulation of NRM – where NRM measurement is simulated using the g18 mask. Parts I-III are used iteratively as input to Parts IV and V. Part IV: GUI – where the user can interactively modify model parameters to build intuition for model candidates and Part V: Inference – where the models are fit to data using an MCMC framework.



**Figure 4.2:** Coordinate system for PIPPI (right handed). The green circle denotes the star centred at the origin. The orange circle denotes a dust grain located at coordinates  $(x_d, y_d, z_d)$  and  $(r_d, \theta_d, \phi_d)$ . Angles are defined as  $\theta = \arctan(y/x)$  and  $\phi = \arcsin(R_*/R_{\text{dust}})$ . The observer sits at  $z = \infty$ .

$$\begin{aligned} x' &= x \cos(\alpha) - y \sin(\alpha), \\ y' &= x \sin(\alpha) + y \cos(\alpha), \\ z' &= z. \end{aligned} \quad (4.11)$$

$$r = \sqrt{x^2 + y^2 + z^2} \quad (4.12)$$

$$E(x, y, z) = \exp\left(-\frac{1}{2} \left( \frac{(x-h)^2}{a^2} + \frac{(y-k)^2}{b^2} + \frac{(z-n)^2}{c^2} \right)^{\frac{1}{\beta}}\right) \quad (4.13)$$


$$G(x, y, z) = 1 - E(x, y, z) \quad (4.14)$$

$$S(x, y, z) = G(x, y, z) * E(x, y, z) \quad (4.15)$$

$$S_p(x, y, z) = G(x, y, z) * r^{pl} \quad (4.16)$$

- $h, k, n$  are the centre coordinates of the Gaussian or shell in the  $x, y,$  and  $z$  directions, respectively.
- $a, b, c$  are the radii of the Gaussian or shell in the  $x, y,$  and  $z$  directions, respectively.
- $\beta$  is the shape parameter, a value of 0.1 was found to create appropriately thin and smooth shells.
- $pl$  is a power law exponent
- $\alpha$  is a rotation angle defined in the plane of the sky

A multitude of circumstellar morphologies can be generated using these equations, six examples are depicted in Figure 4.3. The set of basis functions available in PIPPI are inspired by observational studies which indicate the prevalence of shells and ‘clumpy’ dust structures in the circumstellar environments of evolved stars (see Chapter 1 for relevant studies). Additional parametrisations can be easily introduced by the user if required. The pixel scale of the model (in units of mas) is a free parameter – oversampling is recommended to prevent pixel-artefact noise (we used 1 mas/pixel).



**Figure 4.3:** Examples of dust density distributions generated by PIPPI using combinations of shells and blob structures identified in observational studies on evolved star dust morphology. Images are of 3D density distributions summed over the line of sight. Models include the simple thin spherical shell models (top left) fitted to polarimetric NRM data in existing studies (Norris et al. 2012; Haubois et al. 2019), as well as double-shell (top right) and elliptical shell morphologies (bottom right). These parametrisations can be used independently or in combination.

### Stellar Function

A star is inserted at the centre of the simulated dust. With  $\mu$  Cephei in mind, PIPPI is designed to model late-type evolved stars which typically have low surface gravity and strong limb darkening (Hofmann et al. 1998). As such, a fully darkened disk (FDD) is used as the default centre-to-limb (CLV) function for stars generated by PIPPI (Equation 4.17): where  $R$  defines the radius of the solar disk and  $r$  is the radial distance from the star (Equation 4.12).

$$I(r) = I(0) \left( 1 - \left( \frac{r}{R} \right)^2 \right) \quad (4.17)$$

### Simulation of Scattering – generating ‘scattering grids’

PIPPI function: `make_scattering_grid.py`

PIPPI simulates the scattering of starlight off the constructed dust morphology, generating polarised images that are the input to the Mueller matrix model of the VAMPIRES instrument. Both Rayleigh scattering and Mie scattering functionalities are supported – the user may specify which approximation they wish to utilise. Each is used in our subsequent analysis for a different purpose (see Chapter 5) – here we simply summarise their implementation in PIPPI. PIPPI is limited to modelling a single scattering process, which implicitly assumes that the underlying astrophysical scene is optically thin. This assumption and its utility is evaluated in detail within our later analysis (see Chapter 5).

In both Rayleigh and Mie approximations, equations which define the spatial dependence of scattering induced polarisation are used to pre-compute 3D cubes (representing our model domain) that we refer to as ‘scattering grids’. To generate polarised ‘images’ these scattering grids are multiplied by the dust

morphologies generated with `make_dust.py` and are then summed over the line of sight to form polarised images. However, even with accelerated code, calculation of these scattering grids is the slowest part of model construction and computation of a single scattering grid may take several minutes. Conveniently, a scattering grid only needs to be computed once for each stellar size and dust composition (grain size, chemistry). As such, scattering grids can be pre-computed and parsed as constants to the GUI (Part IV) and the inference (Part V) components of PIPPI.

We now detail the mathematics involved in computing a scattering grid. Each volumetric pixel in the scattering grid is approximated as a single dust grain (Figure 4.4 (a), orange circle). For each of these dust grains, computing the total outgoing scattering induced polarisation requires integration of incident flux from the visible stellar surface  $\theta_{max}$  – each patch of stellar surface is parameterised by angles  $\theta_i$  and  $\phi_i$  (Figure 4.4, (b)) (Equations 4.18, 4.19). Below we describe the computation of this integral for a single dust grain. The same process is repeated for all dust grains in the model domain.

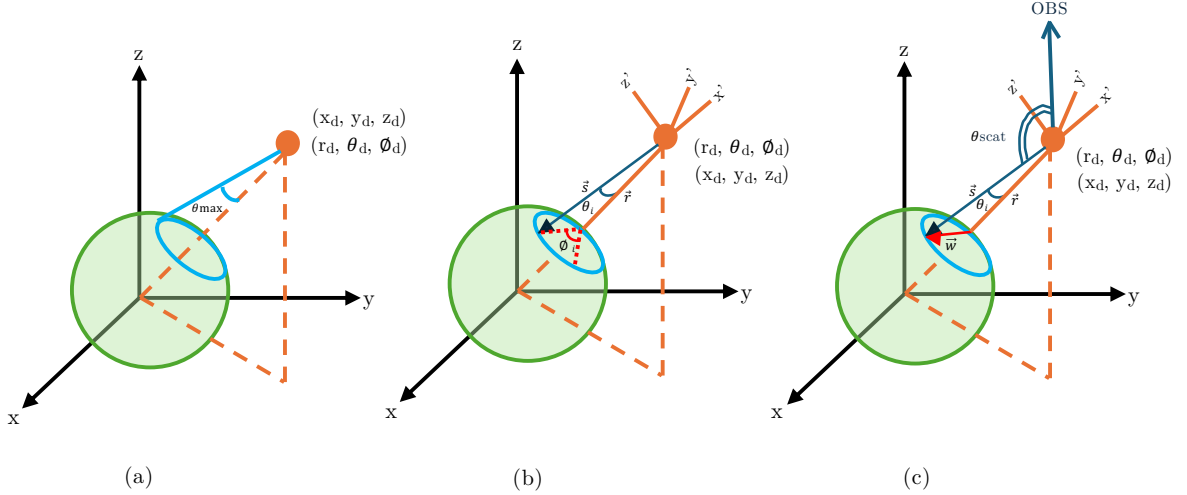
$$\theta = \arctan\left(\frac{y}{x}\right) \quad (4.18)$$

$$\theta_{max} = \arcsin\frac{R_*}{R_{dust}} \quad (4.19)$$

For our dust grain at  $(x_d, y_d, z_d)$ , integrating over  $\theta_i$  and  $\phi_i$  requires definition of two vectors –  $\vec{r}$  and  $\vec{s}$  (Figure 4.4, (b)). The final scattered flux from our dust grain will be the integral of all scattering contributions from all  $\vec{s}$  vectors (over all  $\theta_i$  and  $\phi_i$ ). Here we outline the process involved in computing the polarised scattering contribution from a single  $\vec{s}$  vector (ie. for a given  $\theta_i$  and  $\phi_i$ ). Each vector  $\vec{s}$  (Equation 4.21) is defined in terms of an intermediate vector  $\vec{w}$  (Figure 4.4, (c)) and the vector  $\vec{r}$  (Equation 4.20).

$$\vec{r} = -x_d\hat{x} - y_d\hat{y} - z_d\hat{z} \quad (4.20)$$

$$\vec{s} = \vec{r} + \vec{w} \quad (4.21)$$



**Figure 4.4:** The angles and vectors required to pre-compute scattering grids, as defined in Equations 4.18-4.26. The star (green), and an example dust grain (orange) are depicted. The star is centred on the global observational coordinate system  $(x,y,z)$  whilst the dust grain is offset and has its own local coordinate system  $(x',y',z')$ . The observer is located at  $z = \infty$  along a vector indicated by ‘OBS’ in (c).

To derive intermediate vector  $\vec{w}$  a second coordinate system centred on the dust grain is defined, where vector  $\vec{r}$  is the  $-x'$  axis (Figure 4.4, (b)) and the  $-y'$  axis is  $\vec{w}$  when  $\phi_i = 0$  (Figure 4.4, (c)). We thus now have two coordinate systems, the ‘stellar’ coordinate system defined by axes  $(x, y, z)$  or  $(r, \theta, \phi)$  and the

‘dust’ coordinate system defined by axes with primes ( $x', y', z'$ ) or ( $r', \theta', \phi'$ ). To construct vector  $\vec{w}$  a unit vector along the  $y'$  axis is defined as  $\vec{p}'$  (Equation 4.22).

$$\vec{p}' = \hat{y}' \quad (4.22)$$


To parameterise  $d\phi_i$ ,  $\vec{p}'$  is rotated about the  $x'$  axis using an Euler rotation matrix (Equation 4.23).

$$E_{x'} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\phi_i) & -\sin(\phi_i) \\ 0 & \sin(\phi_i) & \cos(\phi_i) \end{pmatrix} \quad (4.23)$$

To parameterise  $d\theta_i$ , the length of  $\vec{p}'$  is varied to produce vector  $\vec{w}'$  – for an angle  $\theta_i$  at radial distance  $R_{rad}$  the length of  $\vec{w}'$  is given by Equation 4.24.

$$|\vec{w}'| = R_{rad} \times \tan(\theta_i) \quad (4.24)$$

These calculations remain in the dust grain coordinate system. The relationship between the dust coordinate system and the stellar coordinate system is given by two Euler rotation matrices (Figure 4.5). Stellar coordinates are first rotated by an angle  $\alpha$  about the  $z$  axis to arrive at a system described by ( $x'', y'', z''$ ). This intermediate coordinate system is then rotated by an angle  $\beta$  (defined as a rotation around intermediate dust axis  $y''$ ) to arrive at the dust system ( $x', y', z'$ ). Angles  $\alpha$  and  $\beta$  may be defined in the stellar coordinate system for convenience. Angle  $\alpha$  is simply angle  $\theta \in [0, 2\pi]$ . In the intermediate coordinate system ( $x'', y'', z''$ ) angle  $\beta$  is defined in the clockwise direction and so it is convenient to define a second angle  $\gamma = 2\pi - \beta \in [0, 2\pi]$  to maintain the consistency of anti-clockwise rotations. The compound rotation by  $\alpha$  and  $\gamma$  is defined in Equation 4.25 and is applied to  $\vec{s}'$  (Equations 4.25, 4.26).



**Figure 4.5:** Derivation of local dust grain coordinate system ( $x', y', z'$ ) – two rotations by  $\alpha$  and  $\beta$  are required as depicted. The star is denoted as the green circle at the global origin and the dust grain is assumed to be at (0,0,0) within the ( $x', y', z'$ ) system.

$$\vec{s} = E_{tot} \cdot \vec{s}' = E_{y'}(\gamma) \cdot E_z(\alpha) \cdot \vec{s}' = \begin{pmatrix} \cos(\gamma) \cos(\alpha) & -\sin(\alpha) \cos(\gamma) & \sin(\gamma) \\ \sin(\alpha) & \cos(\alpha) & 0 \\ -\sin(\gamma) \cos(\alpha) & \sin(\alpha) \sin(\gamma) & \cos(\gamma) \end{pmatrix} \vec{s}' \quad (4.25)$$

$$\vec{s} = \begin{pmatrix} \cos(\gamma) \cos(\alpha) & -\sin(\alpha) \cos(\gamma) & \sin(\gamma) \\ \sin(\alpha) & \cos(\alpha) & 0 \\ -\sin(\gamma) \cos(\alpha) & \sin(\alpha) \sin(\gamma) & \cos(\gamma) \end{pmatrix} \begin{pmatrix} s'_x \\ s'_y \\ s'_z \end{pmatrix} \quad (4.26)$$

The angle between  $\vec{s}$  and the observing vector is the scattering angle  $\theta_{scat}$  (Figure 4.4 (c)). The scattering angle  $\theta_{scat}$  is used as input into scattering equations for either the Mie or Rayleigh approximation. The details of these scattering computations are outlined in the next section; for now we assume the horizontal ( $\sigma_{HH}$ ) and vertical ( $\sigma_{VV}$ ) components of the polarised scattering cross section have been computed in

whatever approximation was specified by the user. Equation 4.26 gives the scattering contribution from a given  $d\theta_i$  and  $d\phi_i$  of stellar surface. The total polarised scattered flux in the horizontal and vertical directions are the integral over all  $d\theta_i$  and  $d\phi_i$  (Equations 4.27 and 4.28).

$$I_{H,d} = \int_{\phi=0}^{\phi=2\pi} \int_{\theta=-\theta_{max}}^{\theta=\theta_{max}} \sin(\theta_{max}) I_{H,d,i} d\theta d\phi \quad (4.27)$$

$$I_{V,d} = \int_{\phi=0}^{\phi=2\pi} \int_{\theta=-\theta_{max}}^{\theta=\theta_{max}} \sin(\theta_{max}) I_{V,d,i} d\theta d\phi \quad (4.28)$$

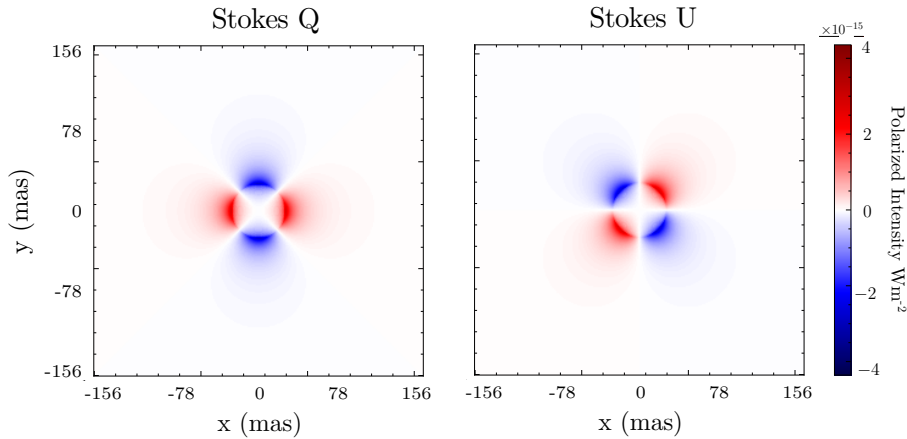
The final Stokes vector emerging from each dust grain is given by Equation 4.29, where Stokes V is set to 0 as we do not measure it with VAMPIRES. However, this Stokes vector is defined in terms of  $I_{H,d}$  and  $I_{V,d}$  where the horizontal and vertical planes are defined in the dust grain coordinate system.

$$s^{scat'} = \begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} = \begin{pmatrix} I_{H,d} + I_{V,d} \\ I_{H,d} - I_{V,d} \\ 0 \\ 0 \end{pmatrix} \quad (4.29)$$

A final transformation is required to measure the Stokes vector in the global stellar coordinate system, achieved by application of Mueller rotation matrix  $E_R$ , where  $\theta_r$  is the angle between the  $x, y$  projection of our vector  $\vec{s}$  and the star coordinate's x axis (Equation 4.30). This yields the Stokes vector in the stellar coordinates ( $I_s, Q_s, U_s, V_s$ ), equivalent to what we measure as the observer.

$$E_R \cdot s^{scat'} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(2\theta_r) & \sin(2\theta_r) & 0 \\ 0 & -\sin(2\theta_r) & \cos(2\theta_r) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} I_{HH,i} + I_{VV,i} \\ I_{HH,i} - I_{VV,i} \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} I_s \\ Q_s \\ U_s \\ 0 \end{pmatrix} = s^{scat} \quad (4.30)$$

Performing this for every volumetric pixel in the model domain [ $N^3$ ] creates our ‘scattering grid’ [ $4 \times N^3$ ]. The scattering grids are by default produced in the Stokes IAU NE coordinate system (Hamaker et al. 1996). Figure 4.6 depicts the sum over the grid in the observing direction after multiplication by a spherically symmetric dust density grid.



**Figure 4.6:** The output of multiplying a Rayleigh scattering grid (10 nm grains of Forsterite) with a spherically symmetric dust envelope – polarised differential images of Stokes Q (left) and Stokes U (right). Models were summed along the line of sight to produce polarised images.

### Scattering Cross Section Computation

We use a different method for computing polarised scattering cross sections for the Rayleigh and Mie approximations, to capitalise on different computational acceleration limitations for each method. For both formalisms the angle  $\theta_{scat}$  (Figure 4.4 (c)) is used as input, along with the grain size and optical constant data that specifies the dust chemistry.

#### *Rayleigh scattering*

For Rayleigh scattering Equations 4.31 and 4.32 are used to compute the polarised scattering cross sections –  $\lambda$  is  $\frac{\lambda_0}{m_0}$  where  $\lambda_0$  is the incident wavelength and  $m_0$  is the refractive index of the surrounding medium. The dust’s refractive index is  $m = n - ik$  where  $n$  and  $k$  are the real and complex coefficients. The approximation of no absorption is made ( $m = n, k = 0$ ). The dimensionless size parameter is defined as  $\alpha = \frac{2\pi a}{\lambda}$  where  $a$  is the grain radius. The Rayleigh scattering functionality of PIPPI is designed for use simulating single wavelength data (over an observational bandwidth). Whilst equations for Rayleigh scattering reference dust size and chemical-specific refractive indices  $n$  and  $k$ , dust chemistry and size are entirely degenerate at a single wavelength.

$$\sigma_{VV} = \frac{\lambda^2}{4\pi^2} \alpha^6 \left| \frac{\bar{m}^2 - 1}{\bar{m}^2 + 2} \right|^2 \quad (4.31)$$

$$\sigma_{HH} = \sigma_{VV} \cos^2(\theta_{scat}) \quad (4.32)$$

Here we use a grain size of 10 nm and assume forsterite ( $\text{Mg}_2\text{SiO}_4$ ) grains with  $n = 1.636$ . Like the simple models employed by prior studies, PIPPI has a ‘star-dust-contrast’ parameter which permits the polarised scattering cross section to be modified in a spatially uniform manner (as a proxy for increasing the amount of dust). As such, the scattering is not restricted to emulating 10 nm grains of Forsterite, as the ‘star-dust-contrast’ is effectively able to modify the polarised scattering cross section. As the goal of applying a Rayleigh scattering approximation is simply to fit to an isotropic scattering regime, any small grains (with an isotropic scattering function) will suffice as input. To accelerate pre-computation of Rayleigh scattering grids our code is implemented in `jax` and accelerated by using a GPU (NVIDIA GeForce RTX 4090) – taking  $\sim 2$  minutes to compute a single set of scattering grids.

#### *Mie scattering*

In contrast to the simple analytic solutions for Rayleigh polarised scattering cross sections, Mie scattering equations are numerical approximations to infinite series of spherical harmonics (Max Born 1975). PIPPI implements Mie scattering via `miepython` (Prahl 2025) which accepts  $\theta_{scat}$  and the physical dust properties as input. The Mie functionality of PIPPI is designed to fit to multi-wavelength data. Scattering grids need to be computed at all observational wavelengths (over observational bandwidth) for each dust grain chemistry and size. To encode the dust grain chemistry, the complex refractive index  $m(\lambda)$  as function of wavelength is required. The construction of  $m(\lambda)$  data and choice of dust grain materials are motivated by the modelled target and were outlined in Chapter 2. `miepython` is not differentiable but may be accelerated using `numba` on a high performance CPU. With acceleration we were able to generate each set of multi-wavelength scattering grids (for a given dust grain size and chemistry) within  $\sim 8$  minutes.

The outputs of PIPPI Part I are scattering grids and dust density distributions. By multiplying these together, polarised images are generated (Figure 4.6) which simulate the astronomical input to the Mueller Matrix instrument model of the Subaru Telescope and VAMPIRES instrument.

## 4.4 Part II – Instrumental Mueller Matrix Model

PIPPI function: `apply_instrument_model.py`

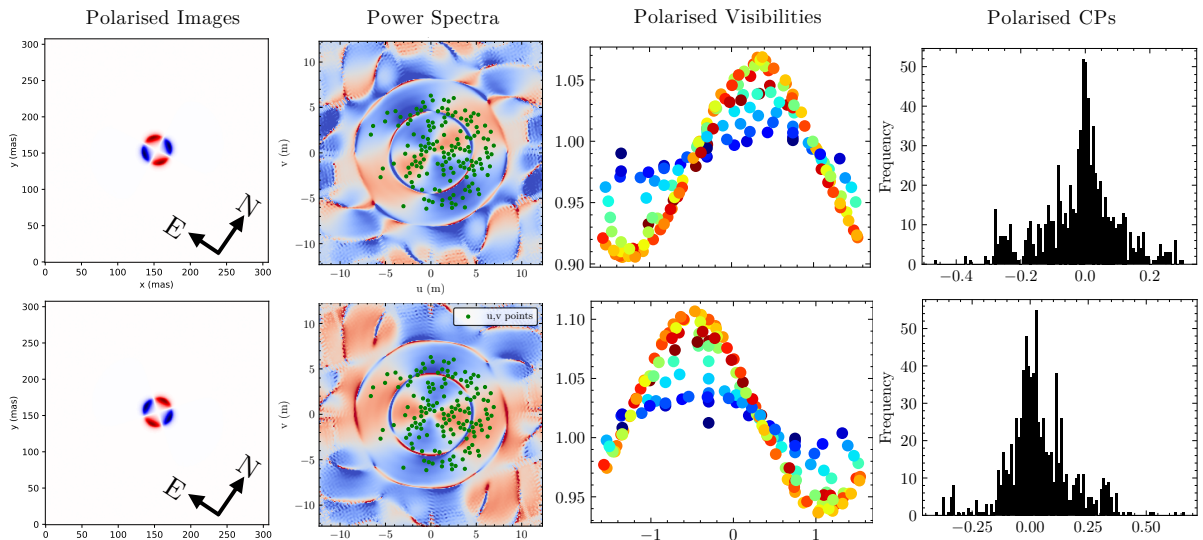
Part II of PIPPI consists of the instrumental forwards model outlined in Chapter 3. For each observing epoch the `vampires_dpp` NRM pipeline returns a Mueller matrix and net image transformation. These are parsed as inputs to PIPPI function `apply_instrument_model.py` to augment the scattering grids

such that they reflect the detector’s polarimetric coordinates at the time of observation (Figure 3.2). Models are always constructed and fit in this local detector polarimetric coordinate system. Polarised images generated from the fitted models can later be de-rotated back into the Stokes IAU system like PDI images. PIPPI is agnostic to the instrument used, only applying an instrumental Mueller Matrix and a net image transformation. With appropriate equivalent metadata, data from alternative instruments can be easily parsed instead.

## 4.5 Part III – Simulating Non Redundant Masking

PIPPI function: `simulate_nrm.py`

Part III of PIPPI simulates NRM measurement – computing visibilities and closure phases in Stokes I, and differential polarised visibilities and closure phases in Stokes Q and Stokes U. All quantities are computed in the local detector coordinate system at the time of observation (Figure 4.7). Rather than hard coding the pupil plane mask coordinates in  $(x,y)$ , PIPPI reads in the  $(u,v)$  coordinates output from `vampires_dpp_nrm`. This permits flexible use of alternative masks and accounts for small changes to the  $(u,v)$  coordinates on account of variability in focus within optics. The  $(u,v)$  coordinates are used to compute a discrete Fourier transform matrix which samples the complex coherence in the Fourier domain to produce visibilities and closure phases. PIPPI treats the  $(u,v)$  coordinates as non-rotating, instead implementing rotation of the underlying astronomical scene. This is largely for convenience – of course, as the rotating sky is tracked by an alt-az telescope the relative orientation of the mask and sky varies. If projected onto sky coordinates, the  $(u,v)$  coordinates would rotate in circles (not the typical elliptical arcs) due to SCEXAO’s pupil-tracking mode of operation (Millour 2008; Millour 2014).



**Figure 4.7:** Polarised images, power spectra and observables produced by PIPPI Part (II), all measured within the detector coordinate system. Column 1 – differential polarised images for Stokes Q and U. Column 2 – differential polarised power spectra from which interferometric observables are sampled. Green points overlaid are the  $(u,v)$  coordinates of the g18 NRM mask. Column 3 – differential polarised visibilities. Column 4 – differential polarised closure phases.

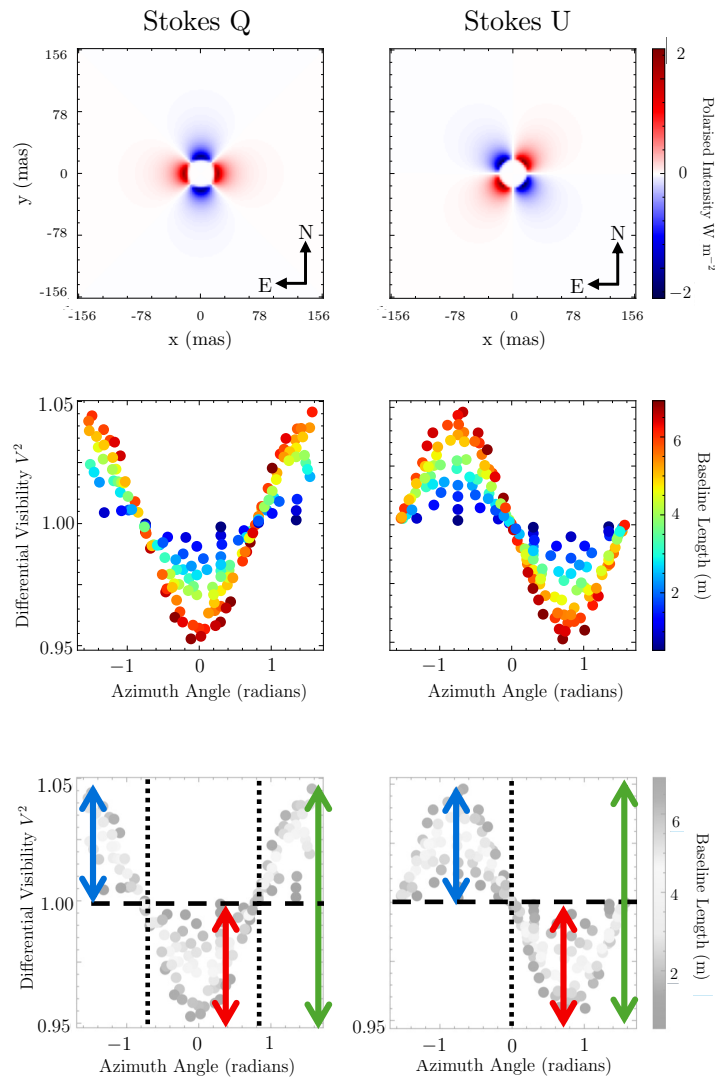
## 4.6 Part IV – Graphical User Interface (GUI)

Optical wavelength polarised interferometric observables are a largely unexplored and under-utilised data type – most of the geometric models which have been fitted in prior works have been simple and symmetric (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019). To aid in developing intuition and understanding of how astrophysical scenes are embedded in polarised interferometric observables, PIPPI has a Graphical User Interface (GUI). The GUI enables the user to change a model and simultaneously

observe the influence of these changes on the polarised images and observables. We found our GUI to be a very effective way to quickly build intuition on this data product and recommend its use during modelling. Within the following section the PIPPI GUI is used to generate astrophysical scenes and compare polarised interferometric observables (Stokes Q and U) to their unpolarised interferometric counterparts (Stokes I). By generating models and experimenting with scene rotation we will explore how asymmetry is embedded in polarised observables. This analysis will guide our design of parametric models in Chapter 5.

### Terminology

For consistency we define several new pieces of terminology for different features of polarised interferometric visibilities (for each linear polarisation state Stokes Q and U) – represented diagrammatically within Figure 4.8. We will refer to these terms within subsequent discussions.




**Figure 4.8:** A modelled spherical circumstellar envelope with simulated polarised visibilities from observation with VAMPIRES and the ‘g18’ mask. Row 1 depicts the polarised differential images of Stokes Q (left) and U (right). Row 2 gives the corresponding differential polarised visibilities. Row 3 is a duplicate of row 2, faded to illustrate the terminology defined in the main text. The horizontal dashed black line defines the ‘mean of differential visibilities’. The vertical dotted black line denotes the ‘ $\frac{V_H}{V_V} = 1$  crossing’. The blue and red arrows denote the ‘upwards and downwards lobes’ (respectively), the amplitude of both is indicated by the green arrow. Illustrations are drawn onto both Stokes Q and U visibilities, as each polarisation state may have different values of these features.

- *Mean of differential visibilities* – represented by black dashed horizontal lines. These lines define the average of each signal (over azimuth angle). The mean of differential visibilities can be shifted away from  $\frac{V_H}{V_V} = 1$  as a result of asymmetry in the dust.
- $\frac{V_H}{V_V} = 1$  *crossing (radians)* – represented with black dotted vertical lines. These values are used as a proxy for understanding how the inflexion points of the differential visibilities are affected by asymmetry. There may be different values of this term for different baseline lengths, depending on the distribution of asymmetry across spatial scales.
- *Upwards Lobe* – represented with blue arrows. These arrows denotes the amplitude of the polarised visibility lobe above the line  $\frac{V_H}{V_V} = 1$ .
- *Downwards Lobe* – represented with red arrows. These arrows denote the amplitude of the polarised visibility lobe below the line  $\frac{V_H}{V_V} = 1$ .
- *Amplitude of Differential Visibilities* – represented with green arrows. This term describes global variation in the differential visibility amplitudes, caused by changing polarisation strength.

**Polarimetric visibilities give us higher sensitivity in high contrast regions than unpolarised visibilities.**

The observational advantages of polarimetry and interferometry were summarised in Chapter 1. The combination of polarimetry and interferometry provides a high resolution measurement of polarised circumstellar dust that would be too faint to measure in Stokes I. Contrasts typical of circumstellar environments lie between  $10^{-3}$  and  $10^{-6}$  (Avenhaus et al. 2018b; Garufi et al. 2020; Esposito et al. 2020). Stars with circumstellar dust envelopes lying towards the higher range of these contrasts (between  $10^{-4}$  and  $10^{-6}$ ) have Stokes I visibilities dominated by the stellar flux – in the presence of realistic random error the signal from the dust will likely lie below the noise floor (Figure 4.9). Only prominent dust shells at low contrast ( $\sim 10^{-3}$ ) will have measurable impact on visibilities (Figure 4.9, red).



**Figure 4.9:** The impact of a spherically symmetric circumstellar dust envelope at  $1.5 R_\star$  on Stokes I visibilities – plotted for typical contrast ratios given in the legend. Visibilities for high contrast dust shells lying between  $10^{-4}$  and  $10^{-6}$  are dominated by the stellar signal – information on the dust is indistinguishable from that of the star. Shells with these high contrasts are indistinguishable from one another in the presence of realistic error from seeing and/or calibration precision. In this regime little information about the dust may be recovered. Where scattered light is more prominent at lower contrasts (here,  $10^{-3}$ ), the impact of dust on the signal becomes clearer.

Contrastingly, differential polarimetry yields measurements of Stokes Q and U which show only the polarised circumstellar dust, removing the unpolarised stellar contribution (Figure 4.8, top row). Furthermore, without averaging over azimuthal angle, polarised interferometric visibilities encode information on the high contrast dust over the range of spatial frequencies probed by the available baselines (Figure

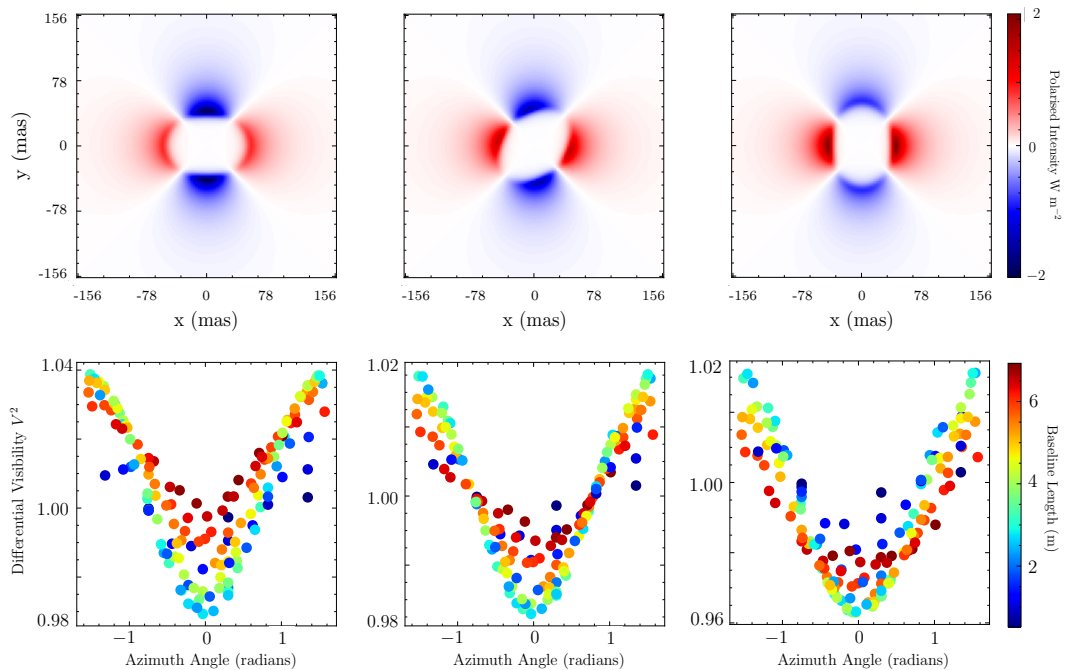
4.8, middle row). While the absolute contrast ratio between the star and dust is inherently embedded in these measurements, by using differential polarimetric interferometry we have gained access to a new regime of measurement where signal strength is dominated by intrinsically faint polarised flux resulting from scattered light off circumstellar dust.

**Polarimetric visibilities give us detailed information on asymmetric circumstellar structure.**

Asymmetric objects have polarised visibilities which depart from the sinusoidal and cosinusoidal shapes characteristic of symmetric objects like those depicted in Figure 4.9. To illustrate how signal can vary on account of asymmetry in the dust an elliptical dust morphology is plotted in Figure 4.10 at various orientations relative to the polarimetric coordinate system. Asymmetry can vary the shape of the lobes – causing them to be more or less rounded, changing the relative amplitude of the upwards and downwards lobes and causing the mean of the polarised visibilities to depart from 1 (these behaviours exhibited in Figure 4.10 – compare the left and middle columns).


When an asymmetric scene is rotated with respect to the polarimetric coordinates, the polarised visibility signal is modulated. The nature of this modulation depends on the distribution of asymmetry along each polarimetric axis. In Figure 4.10 an asymmetric elliptical shell is plotted at three different orientations. When the ellipse is oriented such that flux is most evenly distributed between the positive and negative Stokes quadrants, the distortion of the polarised visibilities is minimised and visibilities are most co-aligned across baseline lengths (Figure 4.10, middle column).

The most significant distortions are observed when there is a large imbalance in flux (and symmetry) between the two Stokes Q quadrants (Figure 4.10 – compare the left and right columns). This imbalance causes the polarised visibilities to ‘fan out’ at azimuthal angles of  $-1$  and  $1$  radians causing the  $\frac{V_H}{V_V} = 1$  crossing points of all baseline lengths to no longer be co-aligned. The variation in this co-alignment with scene orientation can be predicted for a given dust morphology by grouping the triple differential visibilities into baseline lengths and tracking the azimuthal location of the  $\frac{V_H}{V_V} = 1$  crossing of each baseline length. Figure 4.11 displays this effect for the elliptical geometry depicted in Figure 4.10.



**Figure 4.10:** Model of an elliptical circumstellar envelope, captured with Stokes Q polarimetric images (top row) and Stokes Q polarimetric visibilities (bottom row). The dust morphology is an asymmetric ellipse and distorts the polarised visibilities. As the ellipse is rotated the polarised visibilities are modulated – the extent and nature of this modulation depends on the orientation of the ellipse with respect to the polarisation coordinates.

Many of these markers of structural asymmetry are present in our polarised interferometric observables of  $\mu$  Cephei (see Section 3.6 of Chapter 3). These features include: uneven sizes of the upwards and downwards polarimetric lobes, a mean of polarised visibilities that is not equal to 1, and distortion to the polarised visibilities such that they are skewed from the sinusoidal and cosinusoidal form of a symmetric object. The importance of asymmetry is also supported by  $\mu$  Cephei’s polarised differential closure phases which diverge from zero as is expected from asymmetric structure.



**Figure 4.11:** The location of Stokes  $Q$ ,  $\frac{V_H}{V_V} = 1$  ( $\frac{V_H}{V_V} = 1$  crossing) as a function of ellipsoid rotation, as simulated with the g18 mask. The baseline lengths are grouped into 4 bins and a sine function was fitted to the Stokes  $Q$  visibilities for each bin. The Stokes  $Q$ ,  $\frac{V_H}{V_V} = 1$  location in azimuth is plotted as a function of the elliptical rotation  $\alpha$ . As the ellipse rotates, it is able to distort the location of the Stokes  $Q$   $\frac{V_H}{V_V} = 1$  crossing by as much as  $\sim 0.12$  rad.

## 4.7 Part V – Inference

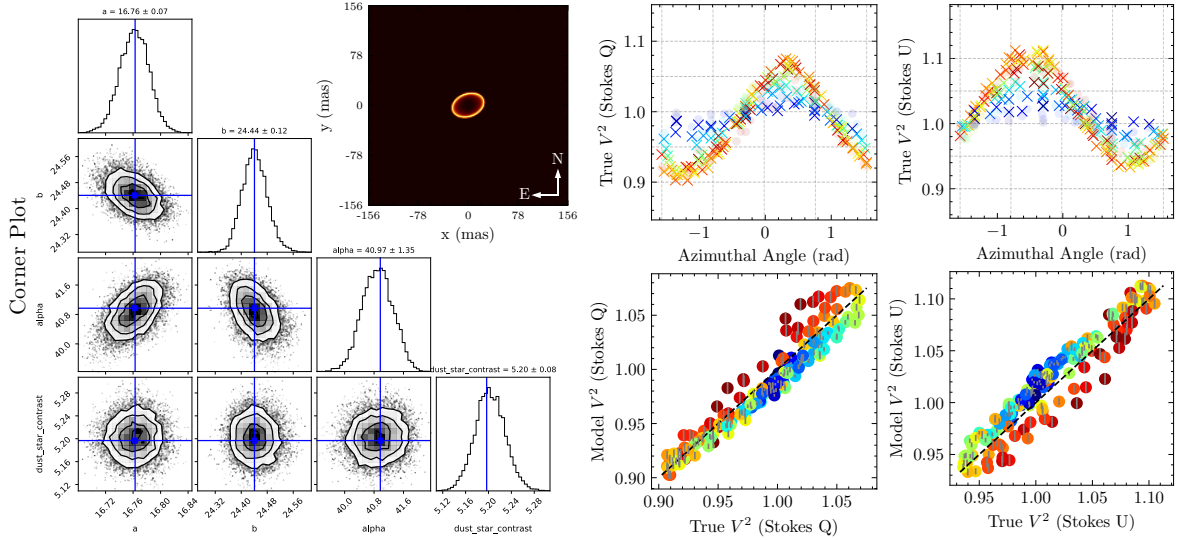
PIPPI function: `model_fitting_nrm.py`

PIPPI uses python package `emcee` (Foreman-Mackey et al. 2013) to implement the affine invariant MCMC to fit VAMPIRES + NRM data. Affine invariant MCMCs exploit the idea that an affine transformation can be used to make convergence less sensitive to covariances among parameters (Goodman et al. 2010). To accelerate inference further than already described by writing PIPPI in `jax`, PIPPI also has a `jit` compiled log likelihood function which is passed to `emcee`, allowing the inference within Part V to be accelerated by use of a GPU. The inputs to PIPPI Part V were outlined in Table 4.1. The user specifies a model class to fit – PIPPI provides 17 pre-defined classes inspired by evolved circumstellar environments and additional classes are easily defined with combinations of the pre-existing classes. As is customary with any MCMC code the user also specifies priors, initial conditions and hyper-parameters specifying the number of walkers and the number of walker steps.

By default, uniform priors with generous ranges are used, accounting for geometric degeneracies in parameters. For initialisation and tuning of each MCMC PIPPI adopts several recommendations made by Foreman-Mackey et al. 2013; that the walkers are initialised in a tight  $N$  dimensional ball around initial values and that the acceptance fraction parameter ‘ $a$ ’ is set to 2. PIPPI uses the auto-correlation time to assess convergence, which quantifies how efficiently the sample mean of a function converges to its expectation value (Thompson 2010). Several diagnostic plots are then produced to assess the features and goodness of each model fit. The `corner` package is used to visualise marginalised posteriors and plots of the maximum likelihood model and data are produced (Figure 4.12).

PIPPI’s simplest model is that which has been historically fitted to optical polarimetric interferometry

data – a symmetric ‘thin shell’ model (Norris et al. 2012). When accelerated on a GeForce RTX 4090 this model satisfies autocorrelation time requirements within  $\sim 3$  minutes. As a result of this accelerated inference PIPPI can rapidly assess a variety of dust grain sizes and compositions, enabling finer sampling of dust grain sizes and testing of a wider variety of chemical possibilities than were computationally tractable to test in prior studies (Haubois et al. 2019; Norris et al. 2012).



**Figure 4.12:** Outputs of PIPPI: A corner plot assesses the convergence and fit of each model candidate. The Maximum Likelihood Estimate (MLE) solution for each model class is plotted – the dust density, observables and correlation of observables with the fitted data. Due to its implementation in `jax` PIPPI can be accelerated to run simple models in as little as 3 minutes. PIPPI returns a  $\chi^2_{red}$  value and the Bayesian Evidence for the model class, which should be used in combination to assess whether greater model complexity is required to fit the data.

## 4.8 Chapter Summary

This chapter presented PIPPI, the first publicly available parametric modelling package for optical wavelength spectro-polarimetric interferometry data. PIPPI’s differentiable design permits accelerated inference for both single-wavelength and multi-wavelength datasets, permitting constraint of circumstellar geometry, chemistry and grain size. Here PIPPI was discussed in the context of VAMPIRES + NRM data, however with small adjustments to metadata PIPPI is straightforwardly adapted to alternative polarised interferometry instruments. PIPPI is best deployed with a combination of high performance computing resources – both GPU and CPU, to accelerate the creation of input scattering grids and the inference process. PIPPI makes several simplifying assumptions – that absorption is negligible and that the dust is optically thin, as is customary of single scattering codes (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019)). The following Chapter 5 will apply both the Rayleigh and Mie approximations of PIPPI to our VAMPIRES + NRM datasets for  $\mu$  Cephei.

# Chapter 5

## Application of PIPPI to $\mu$ Cephei

### Overview of Chapter

In this chapter we apply the PIPPI package to our observing campaign of  $\mu$  Cephei. We first assume Rayleigh scattering and fit the circumstellar geometry at a single wavelength ( $\lambda 750, 760 \pm 25$  nm). We detect strong asymmetry in the inner circumstellar environment and in extended material, confirming the presence of a previously reported ‘clump’ of dust in the South-Western quadrant (Safonov et al. 2019). Within our 7 year observing campaign we observe temporal evolution of the circumstellar dust, which we attribute to either moving convection illumination effects or the evolution of an inhomogeneous dust distribution at the speeds of typical stellar winds. We then use Mie scattering to fit multi-wavelength data from 2023 ( $\lambda 610, 670, 720, 760 \pm 25$  nm), constraining the chemistry and dust grain size of  $\mu$  Cephei’s circumstellar material. Our chemical inference indicates that our data have weak ability to distinguish between the chemical options postulated in Chapter 2 and that many iron-free dust species all provide excellent fits to our data. We find three degenerate regimes of dust grain size that all equally fit our data – grains with radius  $\sim 250, 750$  and above  $1200$  nm. Whilst grains of  $250$  or  $750$  nm lend support to Hofner’s scattering driven stellar wind, grains above  $1200$  nm are too large, leaving our inference as to whether  $\mu$  Cephei’s grains support Hofner’s theory, inconclusive. We then evaluate the physical realism of single scattering models by injecting them into the radiative transfer code MCFOST. Our analysis reveals that the vast majority of the PIPPI models used in our chemical inference do not represent physically valid solutions – thin shells composed of these grains are no longer optically thin when they are made dense enough to produce polarised visibility amplitudes as high as  $\mu$  Cephei’s. Our findings demonstrate that caution and rigorous injection testing in radiative transfer are necessary to determine whether simple scattering codes like PIPPI should be used for chemical inference.

### Contribution Statement

The idea of a simple scattering code was originally my supervisors – and has been used in their pioneering of spectro-polarised interferometry data (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019). I built PIPPI (Chapter 4) and applied it to our observing campaign of  $\mu$  Cephei, designing models and guiding the direction of inference and subsequent analysis. The re-analysis of these simple models and their injection testing in radiative transfer code was my own idea and work, as is subsequent analysis on their suitability for use in chemical inference.

## 5.1 Constraining $\mu$ Cephei’s circumstellar geometry with Rayleigh scattering

We commence by using PIPPI’s Rayleigh scattering mode to constrain  $\mu$  Cephei’s circumstellar dust geometry. We fit to all four of our single-wavelength archival VAMPIRES datasets (datasets 1-4,  $\lambda 750$  nm) and to a single wavelength of our upgraded VAMPIRES dataset (dataset 6,  $\lambda 760$  nm), which collectively span a 7 year period (2017-2023). Existing literature on the circumstellar geometry of  $\mu$  Cephei and analogous evolved stars was reviewed in Chapter 1 – here we draw inspiration from these descriptions when designing candidate models. Commencing from the simple models used in Ireland et al. 2005, Norris et al. 2012 and Haubois et al. 2019, we first explore how well the historically employed symmetric ‘thin

shell’ model is able to fit our  $\mu$  Cephei VAMPIRES + NRM data. Then, as required, we deploy increasingly complex models by adding ‘shell’ and ‘blob’ structures. We expect  $\mu$  Cephei’s circumstellar dust to create a complex scene spanning multiple spatial scales, as our data suggests the scene is well resolved and has asymmetric contributions on most baseline lengths (Figures 3.7-3.9). Asymmetric structure has been previously identified on both broad scales and within the inner circumstellar environment (Norris 2015; Safonov et al. 2019), however our data should capture both scales simultaneously and with higher resolution. Given the high rates of mass loss exhibited by  $\mu$  Cephei ( $10^{-5} M_{\odot} \text{yr}^{-1}$ ) (Gehrz 1971) we anticipate that circumstellar material may evolve within our observing campaign – as has been observed to occur on scales as short as months in analogous circumstellar environments (Ohnaka et al. 2017b; Khouri et al. 2016). We aim to probe the following features of  $\mu$  Cephei’s circumstellar environment:


1. The outflow properties of the inner circumstellar environment – whether it is a discrete shell or outflow and what this may suggest about the persistence or episodic nature of the mass loss occurring,
2. The role of asymmetry – in both the inner circumstellar environment and on broad scales, and whether we can model these previously separately identified phenomena simultaneously,
3. The proximity of dust to the (nominal) stellar surface,
4. Whether temporal evolution is observed over a 7 year observing campaign.

Our candidate models are formulated in pairs to facilitate the hypothesis testing of geometric features that probe these questions. The geometric parameters used to construct our models are defined in Figure 5.1 – all model parameters are depicted on the same axes for illustration purposes only and Table 5.1 indicates which subset of model parameters is fitted to each model class.

**Table 5.1:** Overview of all parametric models fitted to  $\mu$  Cephei with PIPPI. Each model class has a label defined in column ‘#’, a description, number of parameters  $N_p$ , and the modelled parameters as defined in Figure 5.1.

#	Description	$N_p$	Parameters*
A	Thin Spherical Shell	2	a, $C_D$
B	Thick Spherical Shell	3	a, $C_D$ , $S_T$
C	Power Law Spherical Shell	3	a, $C_D$ , pl
D	Two thin Spherical Shells	4	a, a2, $C_D$ , $S_T$
E	Thin Elliptical Shell	4	a, $C_D$ , b, $\alpha$
F	Thick Elliptical Shell	5	a, $C_D$ , b, $\alpha$ , $S_T$
G	Power Law Elliptical Shell	5	a, $C_D$ , b, $\alpha$ , pl
H	Thin Offset Circle	4	a, $C_D$ , h, k
I	Thin Offset Ellipse	6	a, b, $\alpha$ , $C_D$ , h, k
J	Thin Spherical Shell + Blob	5	a, $C_D$ , $C_B$ , $r_b$ , $\theta_b$
K	Thin Elliptical Shell + Blob	7	a, b, $\alpha$ , $C_D$ , $C_B$ , $r_b$ , $\theta_b$
L	Thick Spherical Shell + Blob	6	a, $S_T$ , $C_D$ , $C_B$ , $r_b$ , $\theta_b$
M	Thick Elliptical Shell + Blob	8	a, b, $\alpha$ , $S_T$ , $C_D$ , $C_B$ , $r_b$ , $\theta_b$
N	Power Law Spherical Shell + Blob	6	a, $C_D$ , pl, $C_B$ , $r_b$ , $\theta_b$
O	Power Law Elliptical Shell + Blob	8	a, b, $\alpha$ , $C_D$ , pl, $C_B$ , $r_b$ , $\theta_b$
P	Offset Power Law Spherical Shell + Blob	8	a, $C_D$ , pl, $C_B$ , $r_b$ , $\theta_b$ , h, k
Q	Offset Power Law Elliptical Shell + Blob	10	a, b, $\alpha$ , $C_D$ , pl, $C_B$ , $r_b$ , $\theta_b$ , h, k
R	Enhanced Power Law Spherical Shell + Blob	8	a, $C_D$ , pl, $C_B$ , $r_b$ , $\theta_b$ , $E_a$ , $E_l$
S	Enhanced Power Law Elliptical Shell + Blob	10	a, b, $\alpha$ , $C_D$ , pl, $C_B$ , $r_b$ , $\theta_b$ , $E_a$ , $E_l$

\*: All parameters fitted are defined in Figure 5.1.



**Figure 5.1:** Geometric parameters fitted within PIPPI. All parameters are displayed on the same axes for illustration purposes only – the subsets of parameters used within each model class are outlined in Table 5.1. The parameters inscribed define the geometric configuration of a dust density distribution around a star centred at the origin (denoted with a  $\star$ ).  $a$  and  $b$  are the semi-minor and semi-major axes of the circumstellar shell,  $\alpha$  is the clocking angle of the (elliptical) semi-major axis. The third dimension of the shell (in  $z$ ) is assumed to be the same as  $a$ . The centre of the shell may have an offset from the stellar location in  $(x, y)$ , denoted  $(h, k)$ , and the star dust contrast is denoted  $C_D$ . A second shell (only in full spherical symmetry for the present work) is defined with radius  $a_2$  and shell contrast  $S_C$  (where this contrast is between shell 1 and shell 2). A Gaussian blob is parameterised by its radial distance  $r_b$ , angular location  $\theta_b$  and contrast  $C_B$ . The internal structure of the circumstellar shell (blue) can take three forms: a finite thin shell (2 mas), a finite thick shell (thickness  $S_T$ ) or a power-law outflow (exponent denoted  $pl$ ). A density enhancement along the circumstellar shell’s limb (marked in green) may also be added at an angular location  $E_l$  with a contrast  $E_a$ .

### 5.1.1 Results of Rayleigh scattering inference

We fit 19 model classes to  $\mu$  Cephei – for each model the parameter MLEs,  $\chi_{red}^2$  and  $\ln(B)$  values are provided in Tables 5.2 - 5.6. All angles are in celestial coordinates. The priors, hyper-parameters and corner plots for each class of model are included in Appendix A.3.

**Table 5.2:** Summary of Model Fits and Parameter MLEs (I). Superscript \* denotes that in the model's original form the dust converged on the nominal stellar surface – and that the model was re-run to enable the dust to converge below the nominal stellar surface. All parameters stated within the present table have been transformed into celestial NE coordinates – angles are defined in the direction East of North. Parameter values stated are the maximum likelihood estimates, with uncertainties given by the posterior distribution widths  $1\sigma$ .  $\ln(B)$  denotes the natural logarithm of the Bayesian Evidence. Larger values (less negative) denote more probable models.

<b>Model</b>	$a$ [mas]	$a_2$ [mas]	$S_C$	$C_D$	$S_T$ [mas]	$pl$	$\chi_{\text{red}}^2$	$\ln(B)$
$A_{2017}$	$18.4 \pm 0.3$	-	-	$5.8 \pm 0.3$	-	-	1.97	-1843.01
$A_{201802}$	$16.7 \pm 0.5$	-	-	$5.8 \pm 0.2$	-	-	1.06	-1399.16
$A_{201805}$	$17.1 \pm 0.3$	-	-	$5.6 \pm 0.3$	-	-	3.38	-2526.86
$A_{2020}$	$17.0 \pm 0.4$	-	-	$6.4 \pm 0.2$	-	-	1.23	-1483.83
$A_{2023}$	$16.3 \pm 0.6$	-	-	$6.7 \pm 0.2$	-	-	1.67	-1695.85
$B_{2017}$	$12.6 \pm 1.3$	-	-	$5.1 \pm 0.4$	$17.0 \pm 4.6$	-	1.96	-1834.35
$B_{201802}^*$	$9.5 \pm 1.2$	-	-	$4.5 \pm 0.4$	$23.9 \pm 5.7$	-	1.02	-1379.85
$B_{201805}^*$	$10.0 \pm 0.5$	-	-	$4.4 \pm 0.3$	$24.2 \pm 4.3$	-	3.27	-2476.89
$B_{2020}$	$12.0 \pm 2.5$	-	-	$5.8 \pm 0.7$	$13.2 \pm 8.1$	-	1.23	-1466.04
$B_{2023}^*$	$1.7 \pm 8.4$	-	-	$5.9 \pm 0.4$	$19.2 \pm 8.6$	-	1.67	-1691.61
$C_{2017}$	$14.4 \pm 0.6$	-	-	$5.0 \pm 0.3$	-	$2.3 \pm 0.7$	1.94	-1824.04
$C_{201802}$	$12.1 \pm 1.2$	-	-	$4.7 \pm 0.4$	-	$1.6 \pm 0.8$	1.01	-1372.03
$C_{201805}$	$12.0 \pm 0.4$	-	-	$4.4 \pm 0.2$	-	$1.2 \pm 0.3$	3.22	-2445.19
$C_{2020}$	$14.8 \pm 1.3$	-	-	$6.0 \pm 0.4$	-	$4.4 \pm 2.4$	1.22	-1478.20
$C_{2023}^*$	$7.8 \pm 2.1$	-	-	$4.1 \pm 0.8$	-	$4.5 \pm 0.4$	1.62	-1673.42
$D_{2017}$	$17.9 \pm 0.6$	$52.7 \pm 13.1$	$0.1 \pm 0.2$	$5.3 \pm 0.4$	-	-	1.93	-1823.88
$D_{201802}$	$15.9 \pm 1.0$	$34.3 \pm 6.5$	$2.5 \pm 1.4$	$5.2 \pm 0.3$	-	-	1.01	-1376.32
$D_{201805}$	$16.9 \pm 2.5$	$41.0 \pm 13.7$	$0.3 \pm 0.2$	$5.1 \pm 0.3$	-	-	3.24	-2462.10
$D_{2020}$	$16.8 \pm 0.5$	$38.8 \pm 12.1$	$1.2 \pm 1.7$	$6.2 \pm 0.3$	-	-	1.22	-1480.10
$D_{2023}$	$15.3 \pm 1.0$	$41.1 \pm 4.7$	$6.4 \pm 1.7$	$5.8 \pm 0.3$	-	-	1.63	-1680.75

**Table 5.3:** Summary of Model Fits and Parameter MLEs (II). Superscript \* denotes that in the model's original form the dust converged on the nominal stellar surface – and that the model was re-run to enable the dust to converge below the nominal stellar surface. All parameters stated within the present table have been transformed into celestial NE coordinates – angles are defined in the direction East of North. Parameter values stated are the maximum likelihood estimates, with uncertainties given by the posterior distribution widths  $1\sigma$ .  $\ln(B)$  denotes the natural logarithm of the Bayesian Evidence. Larger values (less negative) denote more probable models.

<b>Model</b>	$a$ [mas]	$b$ [mas]	$\alpha$ [°]	$C_D$	$S_T$ [mas]	$h$ [mas]	$k$ [mas]	$\chi_{\text{red}}^2$	$\ln(B)$
$E_{2017}$	$16.8 \pm 0.1$	$24.4 \pm 0.2$	$109.3 \pm 2.8$	$5.2 \pm 0.2$	-	-	-	1.67	-1694.09
$E_{201802}$	$14.9 \pm 0.1$	$20.4 \pm 0.2$	$98.1 \pm 1.4$	$6.2 \pm 0.2$	-	-	-	0.85	-1292.65
$E_{201805}$	$14.5 \pm 0.0$	$21.7 \pm 0.2$	$93.2 \pm 1.4$	$6.1 \pm 0.2$	-	-	-	2.76	-2222.90
$E_{2020}$	$15.1 \pm 0.1$	$19.9 \pm 0.2$	$77.4 \pm 4.2$	$6.8 \pm 0.2$	-	-	-	0.99	-1364.03
$E_{2023}$	$15.2 \pm 0.4$	$17.6 \pm 0.5$	$275.0 \pm 3.8$	$6.8 \pm 0.2$	-	-	-	1.58	-1659.70
$F_{2017}$	$16.9 \pm 0.0$	$24.4 \pm 0.2$	$109.0 \pm 0.9$	$5.2 \pm 0.2$	$1.1 \pm 0.1$	-	-	1.64	-1656.42
$F_{201802}$	$15.0 \pm 0.0$	$20.6 \pm 0.2$	$95.9 \pm 1.2$	$6.3 \pm 0.2$	$1.3 \pm 0.1$	-	-	0.81	-1289.15
$F_{201805}$	$13.2 \pm 0.1$	$19.5 \pm 0.2$	$93.4 \pm 1.4$	$6.0 \pm 0.2$	$5.3 \pm 0.3$	-	-	2.75	-2226.53
$F_{2020}$	$15.5 \pm 0.0$	$19.7 \pm 0.1$	$74.1 \pm 1.2$	$6.7 \pm 0.1$	$1.6 \pm 0.1$	-	-	0.97	-1360.23
$F_{2023}^*$	$1.2 \pm 6.6$	$4.0 \pm 6.6$	$274.0 \pm 3.8$	$6.0 \pm 0.3$	$18.0 \pm 6.8$	-	-	1.59	-1626.60
$G_{2017}$	$12.8 \pm 0.2$	$16.0 \pm 0.3$	$111.8 \pm 2.2$	$8.1 \pm 0.2$	-	-	-	1.18	-1436.32
$G_{201802}$	$11.4 \pm 0.1$	$14.1 \pm 0.1$	$97.8 \pm 1.4$	$8.1 \pm 0.1$	-	-	-	0.63	-1203.24
$G_{201805}$	$11.7 \pm 0.1$	$15.0 \pm 0.2$	$95.8 \pm 1.4$	$7.9 \pm 0.2$	-	-	-	1.87	-1801.30
$G_{2020}$	$11.7 \pm 0.1$	$13.6 \pm 0.1$	$79.1 \pm 0.2$	$8.9 \pm 0.2$	-	-	-	0.75	-1255.96
$G_{2023}^*$	$8.3 \pm 0.6$	$11.6 \pm 0.5$	$269.3 \pm 5.4$	$5.2 \pm 0.2$	-	-	-	1.58	-1660.81
$H_{2017}$	$20.6 \pm 1.0$	-	-	$5.0 \pm 0.3$	-	$-4.8 \pm 0.5$	$-9.3 \pm 2.3$	1.77	-1744.34
$H_{201802}$	$18.3 \pm 0.6$	-	-	$5.0 \pm 0.3$	-	$7.7 \pm 1.4$	$1.8 \pm 1.1$	0.91	-1328.97
$H_{201805}$	$19.9 \pm 1.0$	-	-	$4.4 \pm 0.2$	-	$10.6 \pm 0.8$	$-2.3 \pm 0.7$	2.73	-2207.45
$H_{2020}$	$18.1 \pm 0.2$	-	-	$5.8 \pm 0.2$	-	$6.6 \pm 0.2$	$-1.0 \pm 0.7$	1.09	-1411.36
$H_{2023}$	$13.3 \pm 0.1$	-	-	$6.3 \pm 0.1$	-	$1.3 \pm 0.1$	$-0.1 \pm 0.1$	1.19	-1454.61

**Table 5.4:** Summary of Model Fits and Parameter MLEs (III). Superscript \* denotes that in the model's original form the dust converged on the nominal stellar surface – and that the model was re-run to enable the dust to converge below the nominal stellar surface. All parameters stated within the present table have been transformed into celestial NE coordinates – angles are defined in the direction East of North. Parameter values stated are the maximum likelihood estimates, with uncertainties given by the posterior distribution widths  $1\sigma$ .  $\ln(B)$  denotes the natural logarithm of the Bayesian Evidence. Larger values (less negative) denote more probable models.

<b>Model</b>	$a$ [mas]	$b$ [mas]	$h$ [mas]	$k$ [mas]	$\alpha$ [°]	$C_D$	$C_B$	$r_b$ [mas]	$\theta_b$ [°]	$\chi_{\text{red}}^2$	$\ln(B)$
$I_{2017}$	$17.0 \pm 0.2$	$24.7 \pm 0.7$	$-0.7 \pm 0.2$	$-3.8 \pm 0.6$	$107.6 \pm 1.6$	$5.2 \pm 0.2$	-	-	-	1.64	-1680.74
$I_{201802}$	$17.7 \pm 1.0$	$19.4 \pm 2.0$	$7.3 \pm 1.2$	$4.3 \pm 5.6$	$128.8 \pm 35.1$	$4.7 \pm 0.5$	-	-	-	0.92	-1321.98
$I_{201805}$	$14.9 \pm 0.5$	$21.9 \pm 0.8$	$3.0 \pm 4.9$	$0.0 \pm 0.4$	$88.5 \pm 2.6$	$6.1 \pm 0.2$	-	-	-	2.72	-2209.93
$I_{2020}$	$15.3 \pm 0.2$	$19.6 \pm 0.5$	$0.9 \pm 0.2$	$-0.4 \pm 0.2$	$76.6 \pm 1.9$	$6.7 \pm 0.2$	-	-	-	0.97	-1355.14
$I_{2023}$	$14.7 \pm 0.5$	$16.7 \pm 0.5$	$1.2 \pm 0.1$	$-0.2 \pm 0.1$	$274.8 \pm 3.3$	$6.6 \pm 0.1$	-	-	-	1.16	-1450.72
$J_{2017}$	$18.4 \pm 0.2$	-	-	-	-	$5.2 \pm 0.1$	$38.8 \pm 7.6$	$87.8 \pm 3.9$	$202.0 \pm 1.1$	0.60	-1172.60
$J_{201802}$	$16.3 \pm 0.3$	-	-	-	-	$5.1 \pm 0.1$	$27.8 \pm 5.7$	$72.7 \pm 3.7$	$166.5 \pm 1.1$	0.46	-1104.45
$J_{201805}$	$16.5 \pm 0.1$	-	-	-	-	$4.8 \pm 0.2$	$16.3 \pm 2.7$	$65.4 \pm 3.4$	$143.3 \pm 1.1$	1.48	-1599.59
$J_{2020}$	$16.7 \pm 0.3$	-	-	-	-	$5.9 \pm 0.1$	$15.1 \pm 4.2$	$73.9 \pm 5.1$	$177.0 \pm 3.4$	0.68	-1212.20
$J_{2023}$	$16.0 \pm 0.4$	-	-	-	-	$6.4 \pm 0.1$	$0.1 \pm 0.0$	$47.9 \pm 5.1$	$194.9 \pm 3.4$	1.48	-1602.36
$K_{2017}$	$17.4 \pm 0.3$	$19.4 \pm 0.2$	-	-	$104.7 \pm 5.6$	$5.1 \pm 0.1$	$40.4 \pm 7.6$	$88.6 \pm 3.8$	$202.6 \pm 5.6$	0.57	-1161.47
$K_{201802}$	$15.2 \pm 0.2$	$17.9 \pm 0.3$	-	-	$91.6 \pm 4.1$	$5.4 \pm 0.1$	$23.4 \pm 5.1$	$74.1 \pm 3.7$	$190.0 \pm 4.1$	0.44	-1091.91
$K_{201805}$	$15.1 \pm 0.3$	$18.9 \pm 1.6$	-	-	$85.8 \pm 8.4$	$5.2 \pm 0.2$	$14.5 \pm 4.1$	$68.5 \pm 4.0$	$184.5 \pm 8.4$	1.37	-1556.04
$K_{2020}$	$15.5 \pm 0.3$	$17.9 \pm 0.3$	-	-	$69.0 \pm 4.5$	$6.2 \pm 0.1$	$14.7 \pm 4.0$	$75.6 \pm 4.9$	$179.9 \pm 4.5$	0.63	-1189.90
$K_{2023}$	$12.9 \pm 0.5$	$15.2 \pm 0.6$	-	-	$271.8 \pm 3.8$	$5.4 \pm 0.2$	$0.4 \pm 0.0$	$7.4 \pm 1.1$	$255.7 \pm 3.4$	1.35	-1546.02
$L_{2017}$	$14.7 \pm 0.3$	-	-	-	-	$5.0 \pm 0.1$	$37.7 \pm 7.3$	$88.2 \pm 4.0$	$158.0 \pm 1.1$	0.59	-1173.14
$L_{201802}$	$13.0 \pm 0.4$	-	-	-	-	$4.9 \pm 0.1$	$25.3 \pm 5.3$	$73.7 \pm 3.8$	$171.2 \pm 1.1$	0.46	-1106.54
$L_{201805}$	$13.1 \pm 0.3$	-	-	-	-	$4.6 \pm 0.1$	$14.6 \pm 2.7$	$66.2 \pm 3.7$	$176.1 \pm 1.1$	1.47	-1601.60
$L_{2020}$	$13.3 \pm 0.3$	-	-	-	-	$5.6 \pm 0.1$	$13.8 \pm 4.1$	$75.1 \pm 5.3$	$183.0 \pm 3.4$	0.68	-1217.81
$L_{2023}$	$16.0 \pm 0.4$	-	-	-	-	$6.4 \pm 0.1$	$0.1 \pm 0.0$	$48.0 \pm 5.0$	$165.1 \pm 3.4$	1.48	-1606.30

**Table 5.5:** Summary of Model Fits and Parameter MLEs (IV). Superscript \* denotes that in the model's original form the dust converged on the nominal stellar surface – and that the model was re-run to enable the dust to converge below the nominal stellar surface. All parameters stated within the present table have been transformed into celestial NE coordinates – angles are defined in the direction East of North. Parameter values stated are the maximum likelihood estimates, with uncertainties given by the posterior distribution widths  $1\sigma$ .  $\ln(B)$  denotes the natural logarithm of the Bayesian Evidence. Larger values (less negative) denote more probable models.

<b>Model</b>	$a$ [mas]	$b$ [mas]	$\alpha$ [°]	$pl$	$C_D$	$C_B$	$r_b$ [mas]	$\theta_b$ [°]	$\chi_{\text{red}}^2$	$\ln(B)$
$M_{2017}$	$14.1 \pm 0.4$	$16.1 \pm 0.5$	$100.4 \pm 8.7$	-	$4.9 \pm 0.1$	$37.4 \pm 7.0$	$89.6 \pm 3.9$	$202.6 \pm 1.1$	0.57	-1163.15
$M_{201802}$	$11.7 \pm 0.4$	$14.6 \pm 0.4$	$92.7 \pm 5.1$	-	$5.2 \pm 0.1$	$22.1 \pm 4.7$	$74.8 \pm 3.7$	$190.0 \pm 1.1$	0.43	-1092.03
$M_{201805}$	$11.7 \pm 0.5$	$15.4 \pm 0.3$	$88.1 \pm 6.2$	-	$5.0 \pm 0.1$	$13.6 \pm 2.8$	$69.6 \pm 4.1$	$184.5 \pm 2.3$	1.37	-1552.89
$M_{2020}$	$11.9 \pm 0.2$	$14.7 \pm 0.3$	$69.2 \pm 4.2$	-	$5.9 \pm 0.1$	$13.2 \pm 3.8$	$76.0 \pm 5.2$	$181.0 \pm 3.4$	0.64	-1194.30
$M_{2023}$	$14.1 \pm 0.4$	$16.1 \pm 0.5$	$272.2 \pm 3.8$	-	$5.4 \pm 0.2$	$0.4 \pm 0.0$	$7.5 \pm 1.0$	$256.3 \pm 4.6$	1.35	-1549.54
$N_{2017}$	$16.7 \pm 0.5$	-	-	$6.4 \pm 1.7$	$5.0 \pm 0.1$	$33.3 \pm 8.5$	$89.3 \pm 4.1$	$202.0 \pm 1.1$	0.59	-1169.23
$N_{201802}$	$14.4 \pm 0.6$	-	-	$4.9 \pm 1.5$	$4.9 \pm 0.1$	$19.6 \pm 6.6$	$74.5 \pm 4.0$	$188.9 \pm 1.1$	0.55	2177.13
$N_{201805}$	$14.1 \pm 0.3$	-	-	$3.9 \pm 0.6$	$4.6 \pm 0.1$	$7.8 \pm 3.1$	$67.1 \pm 3.9$	$184.5 \pm 1.1$	1.47	-1599.11
$N_{2020}$	$15.5 \pm 0.2$	-	-	$5.0 \pm 0.2$	$6.2 \pm 0.1$	$10.4 \pm 3.2$	$77.0 \pm 5.5$	$177.0 \pm 3.4$	0.69	-1217.41
$N_{2023}^*$	$10.1 \pm 1.4$	-	-	$5.2 \pm 0.5$	$4.7 \pm 0.5$	$0.2 \pm 0.1$	$50.0 \pm 5.0$	$194.9 \pm 3.4$	1.46	-1593.04
$O_{2017}$	$14.7 \pm 0.2$	$16.1 \pm 0.3$	$104.0 \pm 4.7$	-	$4.6 \pm 0.1$	$11.3 \pm 2.1$	$95.6 \pm 4.1$	$203.8 \pm 1.1$	0.54	-1140.46
$O_{201802}$	$12.6 \pm 0.3$	$14.2 \pm 0.3$	$93.4 \pm 4.1$	-	$5.1 \pm 0.1$	$6.6 \pm 1.7$	$79.8 \pm 4.7$	$191.1 \pm 2.3$	0.42	-1082.58
$O_{201805}$	$12.8 \pm 0.3$	$14.5 \pm 0.4$	$87.9 \pm 4.4$	-	$5.1 \pm 0.1$	$5.0 \pm 1.9$	$78.6 \pm 8.1$	$186.2 \pm 3.4$	1.29	-1514.30
$O_{2020}$	$13.0 \pm 0.1$	$14.5 \pm 0.2$	$70.2 \pm 4.4$	-	$5.8 \pm 0.1$	$4.6 \pm 3.5$	$85.2 \pm 12.8$	$191.3 \pm 6.9$	0.60	-1280.10
$O_{2023}^*$	$6.2 \pm 0.7$	$12.0 \pm 0.8$	$263.8 \pm 4.3$	-	$4.5 \pm 0.2$	$0.1 \pm 0.0$	$9.5 \pm 1.5$	$250.0 \pm 5.7$	1.41	-1577.46

**Table 5.6:** Summary of Model Fits and Parameter MLEs (V). Superscript \* denotes that in the model's original form the dust converged on the nominal stellar surface – and that the model was re-run to enable the dust to converge below the nominal stellar surface. All parameters stated within the present table have been transformed into celestial NE coordinates – angles are defined in the direction East of North. Parameter values stated are the maximum likelihood estimates, with uncertainties given by the posterior distribution widths  $1\sigma$ .  $\ln(B)$  denotes the natural logarithm of the Bayesian Evidence. Larger values (less negative) denote more probable models.

<b>Model</b>	$a$ [mas]	$b$ [mas]	$\alpha$ [°]	$pl$	$C_D$	$C_B$	$r_b$ [mas]	$\theta_b$ [°]	$h$ [mas]	$k$ [mas]	$\chi_{\text{red}}^2$	$\ln(B)$
$P_{2017}$	$15.2 \pm 0.1$	–	–	–	$4.8 \pm 0.1$	$13.0 \pm 2.2$	$94.5 \pm 3.7$	$202.0 \pm 0.0$	$-0.4 \pm 0.1$	$-0.8 \pm 0.1$	0.55	-1145.06
$P_{201802}$	$13.2 \pm 0.2$	–	–	–	$4.7 \pm 0.1$	$8.0 \pm 2.0$	$75.6 \pm 4.6$	$188.9 \pm 0.0$	$0.3 \pm 0.2$	$0.3 \pm 0.2$	0.46	-1101.25
$P_{201805}$	$13.6 \pm 0.2$	–	–	–	$4.3 \pm 0.2$	$3.4 \pm 0.7$	$62.1 \pm 4.3$	$184.5 \pm 0.0$	$0.6 \pm 0.2$	$-0.5 \pm 0.2$	1.45	-1585.83
$P_{2020}$	$13.9 \pm 0.5$	–	–	–	$4.9 \pm 0.4$	$1.5 \pm 0.7$	$55.2 \pm 9.1$	$178.7 \pm 3.5$	$0.9 \pm 0.3$	$-1.5 \pm 0.6$	0.65	-1194.78
$P_{2023}$	$12.2 \pm 0.2$	–	–	–	$4.6 \pm 0.1$	$0.1 \pm 0.05$	$42.4 \pm 5.9$	$192.7 \pm 0.1$	$-4.2 \pm 0.3$	$1.9 \pm 0.3$	1.10	-1378.34
$Q_{2017}$	$14.8 \pm 0.2$	$16.3 \pm 0.3$	$103.8 \pm 4.1$	–	$4.6 \pm 0.1$	$11.7 \pm 2.1$	$97.0 \pm 4.0$	$203.8 \pm 1.1$	$-0.5 \pm 0.1$	$-0.9 \pm 0.2$	0.48	-1115.69
$Q_{201802}$	$12.7 \pm 0.3$	$14.3 \pm 0.3$	$93.2 \pm 4.2$	–	$5.1 \pm 0.1$	$6.8 \pm 1.8$	$80.2 \pm 4.8$	$190.6 \pm 2.3$	$0.5 \pm 0.3$	$0.1 \pm 0.2$	0.41	-1081.83
$Q_{201805}$	$12.8 \pm 0.2$	$14.7 \pm 0.3$	$88.6 \pm 4.2$	–	$5.1 \pm 0.1$	$5.8 \pm 2.3$	$81.7 \pm 8.4$	$186.2 \pm 3.4$	$0.6 \pm 0.3$	$0.1 \pm 0.2$	1.26	-1499.31
$Q_{2020}$	$14.5 \pm 0.3$	$13.0 \pm 0.3$	$169.8 \pm 6.3$	–	$5.8 \pm 0.1$	$3.2 \pm 1.6$	$77.3 \pm 10.0$	$190.2 \pm 5.7$	$0.3 \pm 0.3$	$-0.6 \pm 0.2$	0.59	-1167.32
$Q_{2023}$	$13.9 \pm 0.3$	$10.8 \pm 0.2$	$256.8 \pm 4.0$	–	$4.9 \pm 0.1$	$0.1 \pm 0.05$	$40.0 \pm 5.9$	$198.4 \pm 5.0$	$-3.9 \pm 0.4$	$2.7 \pm 0.3$	0.99	-1369.33
<b>Model</b>	$a$ [mas]	$b$ [mas]	$\alpha$ [°]	$pl$	$C_D$	$C_B$	$r_b$ [mas]	$\theta_b$ [°]	$E_l$ [°]	$E_a$	$\chi_{\text{red}}^2$	$\log(B)$
$R_{2017}$	$15.0 \pm 0.2$	–	–	–	$4.8 \pm 0.1$	$13.8 \pm 2.5$	$95.1 \pm 3.9$	$202.0 \pm 5.0$	$227.5 \pm 11.3$	$0.6 \pm 0.2$	0.48	-1115.60
$R_{201802}$	$13.3 \pm 0.2$	–	–	–	$4.7 \pm 0.1$	$7.7 \pm 2.1$	$76.4 \pm 5.0$	$188.9 \pm 3.0$	$19.0 \pm 30.8$	$0.3 \pm 0.2$	0.45	-1105.10
$R_{201805}$	$13.9 \pm 0.7$	–	–	–	$3.9 \pm 0.6$	$2.6 \pm 0.9$	$54.6 \pm 9.7$	$184.5 \pm 3.0$	$180.5 \pm 23.4$	$1.4 \pm 2.1$	1.45	-1591.61
$R_{2020}$	$13.6 \pm 0.3$	–	–	–	$5.0 \pm 0.3$	$1.7 \pm 0.5$	$55.6 \pm 7.0$	$178.7 \pm 4.0$	$180.8 \pm 10.2$	$1.8 \pm 1.4$	0.65	-1196.88
$R_{2023}^*$	$8.3 \pm 0.3$	–	–	–	$4.2 \pm 0.1$	$0.01 \pm 0.05$	$32.5 \pm 5.0$	$192.7 \pm 4.1$	$232.1 \pm 3.6$	$3.3 \pm 0.8$	1.00	-1373.23
$S_{2017}$	$14.7 \pm 0.2$	$16.1 \pm 0.3$	$103.6 \pm 4.6$	–	$4.6 \pm 0.1$	$12.2 \pm 2.4$	$98.0 \pm 4.2$	$203.8 \pm 1.1$	$215.8 \pm 10.9$	$0.7 \pm 0.2$	0.48	-1118.14
$S_{201802}$	$14.3 \pm 0.3$	$12.7 \pm 0.3$	$94.3 \pm 4.0$	–	$5.1 \pm 0.1$	$7.2 \pm 1.9$	$80.9 \pm 5.1$	$190.6 \pm 2.9$	$97.9 \pm 27.1$	$0.3 \pm 0.2$	0.41	-1086.90
$S_{201805}$	$14.9 \pm 0.3$	$13.0 \pm 0.3$	$88.6 \pm 4.1$	–	$5.1 \pm 0.1$	$47.0 \pm 13.0$	$133.9 \pm 7.2$	$186.2 \pm 2.3$	$261.1 \pm 18.1$	$0.4 \pm 0.2$	1.21	-1476.30
$S_{2020}$	$14.6 \pm 0.2$	$13.0 \pm 0.1$	$69.0 \pm 4.2$	–	$5.8 \pm 0.1$	$44.2 \pm 19.0$	$139.9 \pm 11.3$	$193.0 \pm 3.4$	$182.3 \pm 18.7$	$0.5 \pm 0.2$	0.57	-1161.44
$S_{2023}^*$	$10.6 \pm 1.0$	$5.5 \pm 0.4$	$247.8 \pm 4.2$	–	$4.2 \pm 0.2$	$0.1 \pm 0.05$	$26.8 \pm 4.5$	$188.1 \pm 3.5$	$200.4 \pm 5.4$	$7.0 \pm 2.9$	0.97	-1366.07

### 5.1.2 Analysis of Rayleigh scattering inference

#### Inadequacy of simple spherical shell models

The thin spherical shell models (here our Model Class A) fit in existing studies (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019) appear to be an insufficient description of  $\mu$  Cephei’s circumstellar dust – correlation plots reveal more clearly that  $\mu$  Cephei’s visibilities are distorted from the typical sinusoidal and cosinusoidal shape of a symmetric shell and that closure phases depart from the zero values of a centro-symmetric model (Figure 5.2, column 1). Despite these limitations, Model Class A still acts as a proxy for the centre of polarised flux. It is interesting to note that our fits locate circumstellar dust between  $1.6-1.8R_*$  which is consistent with the inner dust radii fit to analogous stellar environments (Munoz-Sanchez et al. 2024b; Scicluna et al. 2015; Wittkowski et al. 2012; Ohnaka et al. 2024; Norris et al. 2012; Haubois et al. 2019).

#### Preference for asymmetry

Asymmetric models are strongly preferred by all epochs of data – in pairs of models where one model is symmetric and the other has an additional degree of asymmetric freedom, the asymmetric model is always preferred (indicated by improvements to the  $\chi_{red}^2$  and  $\ln(B)$  values). For examples of this compare  $\ln(B)$  values for model pairs: A/E, C/G, L/M, O/Q, O/S (Tables 5.2-5.5). We were able to simultaneously characterise the asymmetry within the inner circumstellar environment and in extended material – confirming and extending the findings of Norris 2015 and Safonov et al. 2019. We detail our findings on each region below.


#### Inner circumstellar asymmetry

The best fit models to the inner circumstellar environment had two degrees of asymmetry – a) a *geometric asymmetry* which we modelled as an elliptical shell, and b) a *density asymmetry* which we modelled as either an offset of the dust from the star (Model Q), or an angular density enhancement along the limb of the elliptical shell (Figure 5.2, Model S). We found the fits between Models Q and S to be very similar, but that the offsets fit in Model Q consistently caused dust to eclipse the star and so we preferred Model S. The density enhancement of Model S was restricted to  $45^\circ$  to reduce the number of parameters. The best radial function for the dust was an elliptical power law outflow with an exponent of -3 – this was preferred over a thin or thick shell, suggesting that the circumstellar environment does not readily conform to a steady outflow.

#### Extended material asymmetry


We found that asymmetry was also strongly favoured in the extended environment (Models K, M, O, Q, S). We modelled extended material with a simple Gaussian ‘blob’ of fixed 20 mas size. Our best models fitted the blob with a position angle of  $\sim 186-203^\circ$  and with radial position between 80-140 mas (except for 2023 which located this blob at 26 mas – we return to discussion of this below) (Figure 5.2 (column 4, Model S) and Figure 5.3 (b)). In all years but 2023 the ‘blob’ fitted in Model S was co-located with the extended feature identified in wider-field images of  $\mu$  Cephei from 2017 (Safonov et al. 2019), where it was located at a position angle of  $170-220^\circ$  with a centre of radial displacement of  $\sim 120$  mas (Figure 5.3 (c)). It is important to note that our blob is  $\sim 1/5$  the size of the blob in Safonov et al. 2019 and so this difference is accounted for when interpreting our blob as co-located. We were further able to support the persistence of this extended feature into at least 2018 using our own PDI reduction of archival VAMPIRES imaging data obtained on 2018/10/23 (Figure 5.3 (a)).

Our parametric fits to NRM data are therefore able to resolve asymmetries within the inner circumstellar environment at angular resolutions beyond existing imagery with comparable field of view (Safonov et al. 2019) and also beyond that of our own VAMPIRES PDI data that delivers at most a diffraction limited performance (Figure 5.3 (a)). These results demonstrate the extended reach which NRM brings in characterising the geometry of inner-circumstellar environments and how NRM may be used in a complementary fashion to imaging, to probe finer features.



**Figure 5.2:** Images and correlation plots for representative models from Table 5.1 – for dataset 1 (2017). Each column depicts a different model (A, E, O, S) of increasing complexity from left to right. Row 1 depicts the dust density distribution, rows 2 and 3 depict correlation plots for the polarised visibilities, rows 4 and 5 depict correlation plots for the polarised closure phases. Model A (column 1) is the prototypical thin spherical shell model fitted in existing polarimetric interferometry studies and was found to have insufficient complexity to yield good fits to  $\mu$  Cephei. Model E was our first asymmetric morphology – whilst providing some improvement to the fit, the form of asymmetry is not correct. Model O introduced asymmetry into both the inner environment and to extended scales in the form of a Gaussian blob. Model S yielded our best fits – the inner circumstellar environment has a geometric asymmetry and a density enhancement in addition to an extended Gaussian feature. For ease of visualisation the error bars plotted are the intrinsic random errors and do not include the multiplicative systematic factor we derived in Chapter 3 and used in fitting.

We note an interesting behaviour of the Gaussian blob used to model extended material. Both the angular position of the blob and its radial location vary as our models increase in complexity – in some simpler models the blob appears to move by up to  $80^\circ$  and 80 mas between epochs of data (see Model K). This effect becomes less prevalent as complexity is built into the inner circumstellar region. This suggests that when our models for the inner circumstellar environment are inadequately complex, the Gaussian blob moves inwards to compensate, balancing its contribution to fitting wider scale features (a weak contribution to measured signal) with its ability to emulate greater asymmetry in the inner environment (the dominant contribution to measured signal).



**Figure 5.3:** Images and models of  $\mu$  Cephei’s inner circumstellar environment – (a) Our VAMPIRES PDI from 2018/10/23 (polarisation P is plotted), (b) Our best parametric fit for  $\mu$  Cephei from 2018/10/23 (dust density is plotted with a radial stretch  $r^{1.2}$  to make the extended feature clearer), and (c) Image of  $\mu$  Cephei reproduced from Safonov et al. 2019, obtained via speckle interferometry on a 2.3 m telescope (polarisation P is plotted). Our PDI and NRM datasets have an extended feature co-located with the extended feature within Safonov et al. 2019 – at a position angle of  $\sim 170$ - $220^\circ$ , and  $r \sim 120$  mas.

As complexity was built into the inner environment (Models O, Q, S), the radial location of the blob stabilised at  $\sim 80$ - $130$  mas – for all epochs except 2023. We suggest three causes for a closer radial fit of the blob in 2023 – firstly, that our best Model S provides an inadequately complex representation of the inner circumstellar environment in 2023, and as previously described, the blob may move inwards to compensate for this; secondly that the extended feature modelled in 2017-2020 has moved outwards by 2023 and so is not detected as strongly; and thirdly that new circumstellar material has been ejected from the star since 2020 and that we are detecting a new extended feature outside the circumstellar shell but closer to the star than the extended feature observed in 2017-2020. Regardless of the cause, it appears that there is increased asymmetry in the inner circumstellar environment in 2023 as compared to earlier years of our data. We will further support this finding in Chapter 7.

### Proximity of dust to the stellar surface

The location of the dust relative to the stellar surface changes between each of our model classes. From the canonical thin spherical shell Model Class A our estimates for shell radius were between  $1.6$ - $1.8 R_\star$ . These values are much larger than those obtained from our best Models S – most are between  $1.0$ - $1.4 R_\star$ , however in 2023 our elliptical semi-minor axis value is  $\sim 0.5 R_\star$ . We have already noted our suspicion that the circumstellar environment is highly asymmetric in 2023 and thus treat this value with some caution as it could be artificially generated by a too-simple model in an attempt to emulate greater complexity than it can represent in a physical manner. In addition to this, all values for the proximity of dust to the stellar surface must be interpreted with caution as the Rayleigh scattering approximation does not produce *forwards scattering* – which may reduce the apparent diameter of the fitted circumstellar shells (see Section 5.2 for a detailed explanation of this effect).

An additional source of uncertainty lies within what we defined as  $R_\star$  (we used a literature value obtained in the continuum  $R_\star = 10$  mas which is  $2257 R_\odot$  at 1000 pc) (Quirrenbach et al. 1993). Stellar diameter measurements of evolved stars are made significantly more complex by their very low surface gravities

which result in extended photospheres (Arroyo-Torres et al. 2015; Josselin et al. 2007). Red supergiants also have extended molecular layers (Perrin et al. 2005; Wittkowski et al. 2014; Wittkowski et al. 2011). The wavelength dependence of limb darkening, photospheric radius and molecular absorption all make stellar diameter a strong function of wavelength (Quirrenbach et al. 1993). The strong influence of the TiO bands (the closest to our observations are at 623 nm, 672 nm, 712 nm) is thought to cause inflation of  $\mu$  Cephei’s stellar diameter by up to 17% (Quirrenbach et al. 1993). In addition to the resonant ‘bands’ which the TiO molecule imposes, at high circumstellar temperatures the continuum opacity of TiO is a large fraction of the TiO resonant peak opacity (McKemmish et al. 2019). Thus, although our observing wavelength is not in a resonant band, the photospheric radius may also be inflated by TiO. This would imply that the true size of the dust shell, expressed as a fraction of  $R_*$ , is larger.

### Temporal variability

Our model fits provide evidence for temporal evolution of inhomogeneous structure in  $\mu$  Cephei’s inner circumstellar environment. This evolution is a persistent feature of all model classes fitted, from the most simple (Model A) to the most complex (Model S). The timescales of this evolution are difficult to quantify as the variation is most likely not well sampled by our low observing cadence – however we suggest two potential astrophysical causes. The first could be evolution of dust structure in the inner circumstellar environment on account of an inhomogeneous outflow. Reasonable timescales for dust evolution can be derived from wind velocities – estimates vary quite considerably between 19–35 km s<sup>−1</sup> (Montargès et al. 2019; Decin et al. 2023; De Beck et al. 2010; Le Borgne et al. 1989). Adopting even the slowest of these values (19 km s<sup>−1</sup>) and assuming the stellar wind drives dust at the same speed, dust could be displaced radially by  $\sim 4$  AU year<sup>−1</sup> (4 mas at 1000 pc), which would be more than sufficient to account for variations in our fitted asymmetry. Dust variability is also supported by observations of analogous stellar environments – dust within the inner circumstellar environment of AGB stars has been observed to vary on timescales as short as months (Ohnaka et al. 2017b; Khouri et al. 2016).

A second cause for temporal evolution could be the movement of convective cells on the stellar surface. Convection cells generate ‘bright spots’ (also called ‘hot spots’) on the stellar surface (Haubois et al. 2009) which can create regions of enhanced illumination in the dust immediately above the hot spot. Image reconstructions of evolved giant photospheres show convective structures at contrast levels believed plausible for illumination asymmetries in the circumstellar dust (Pilate et al. 2025). Convective cells are expected to comprise a dominant feature of RSGs: high variability of convective structures spanning a range of scales have featured in simulations (Chiavassa et al. 2011) and in observational data (Haubois et al. 2009; Wilson et al. 1997; Young et al. 2000; Anugu et al. 2024; Montargès et al. 2016). Convection is so intense on  $\mu$  Cephei that the tops of up-wellings create ‘plumes’ observed to extend up to  $1.1R_*$ , changing on the scale of years (Ariste et al. 2023). Similar evolutionary timescales have been reported when using tomographic methods (Kravchenko et al. 2019). On Betelgeuse, an analogous red supergiant star, convective cells are reported to occupy as much as 60% of the stellar diameter. Smaller convective structures are seen to vary on a monthly scale (Vlemmings et al. 2024) while large scale convective features are observed to persist for up to 4 years (López Ariste et al. 2018). Observations of Betelgeuse have also been used to demonstrate that photometric variability is correlated with convective time scales and linear polarisation strength, suggesting that convective structures can be responsible for heightened polarisation measurement (Pilate et al. 2024). Thus, it is plausible that illumination changes originating from varying convective cells could also be responsible for evolving asymmetric structure in our models fitted to  $\mu$  Cephei’s circumstellar environment.

### General comments on fitting to circumstellar geometry

For all epochs except dataset 3 (2018) we were able to obtain a  $\chi_{red}^2 < 1$  with our best fit models (Class S). Use of these models was supported by  $\ln(B)$  values which become larger (less negative) as model complexity was increased. Our parametric modelling study added no further complexity beyond Model S for numerical reasons – ensemble methods (like the affine invariant MCMC PIPPI employs) suffer from poor convergence in high dimensions. The number that qualifies as ‘high’ is somewhat ambiguous but occurs at approximately 10 parameters (Carpenter 2017; Huijser et al. 2022). Future implementations of PIPPI should be upgraded to non-ensemble methods such as a HMC code like `numpyro` (NumPyro 2019) which would not place the same restrictions on the number of fitted parameters.

Small grains satisfying the Rayleigh scattering approximation do not exhibit forward scattering and so our measurements of the circumstellar geometry must be interpreted with some caution – an example of a non-physical outcome is the sub-surface value for circumstellar radius obtained in 2023 (Model S). To quantify the forwards scattering present the Mie scattering approximation must be used – requiring multi-wavelength data and full chemical inference over the dust grain material and size. However, even if we adopt these fitted radii as first-order approximations, their values ( $< 2 R_*$ ) combined with our thermodynamic analysis in Chapter 2 would suggest that the dust is most likely an iron-poor silicate or an alternative transparent species like corundum.

Our two epochs obtained on the same evening in 2018 (datasets 2 and 3) exhibit variability in their fits that exceed each fitted parameter’s posterior distribution width, suggestive of un-quantified systematic errors. The 2018 data is complex – in Section 3.7 we diagnosed a probable electronic fault in a wave-plate homing unit and noted in Section 3.6 that AO was configured differently for each dataset. Whilst we do not know the exact cause of the variation in model fits between these epochs, we have identified two sources of systematic error that may contribute, neither of which would be accurately quantified by our multiplicative scaling factor used to estimate the contribution of systematic errors (Section 3.3).

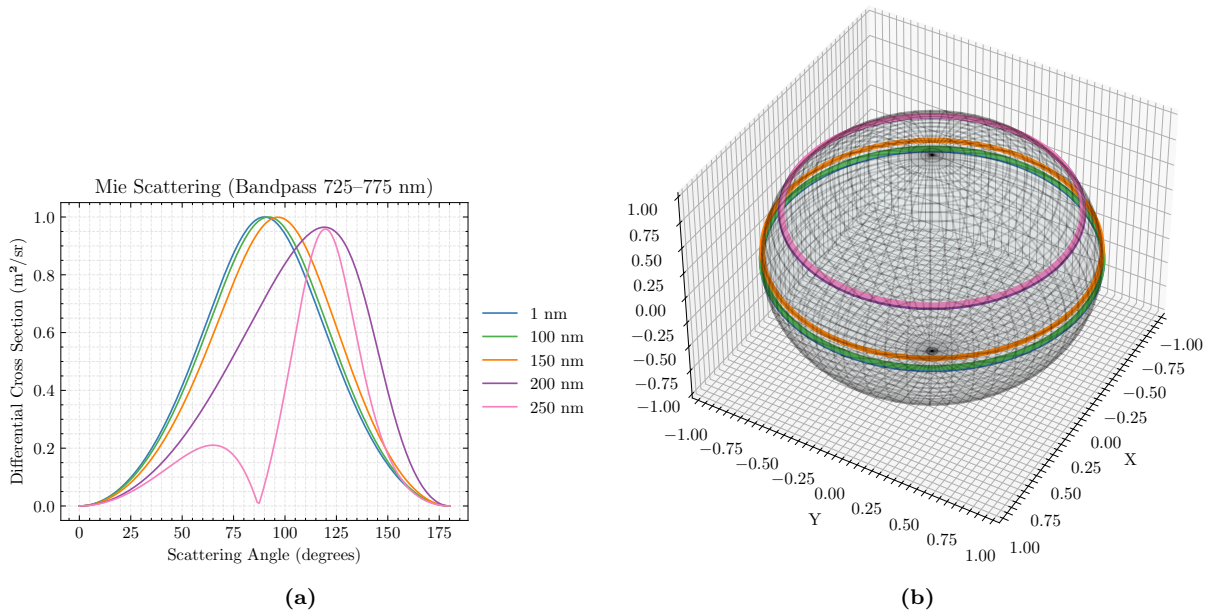
More generally we also note that most of our fitted parameters have very small uncertainty estimates – these must also be interpreted with caution. The width of marginalised posterior distributions reflect the *sensitivity* of the model parameters to our data, which is a fundamentally different thing to how well the model *describes* the data. It is very possible that our data is highly sensitive to the model despite the model itself not adequately representing the data. A good thought experiment for understanding this is to consider fitting real data of a cubic function (with small uncertainties) with a linear function – the posteriors may be very tight but the model itself is *incorrect*. Small posteriors are an indication of a tightly constrained model but do not guarantee that we are able to estimate the true uncertainty in the parameters themselves.

A second contribution to narrow posteriors arises from our errors. With no alternative model we used a multiplicative factor to estimate the contribution of systematic errors (see Section 3.3). However, it is quite likely that the systematic errors we are approximating with this multiplicative factor are largely *polarisation aberrations* – instrumental polarisation effects that occur across the pupil plane from non-uniformity within optics (Ashcraft et al. 2025). A proper quantification of these polarisation aberrations requires ray tracing and is not equivalent to increasing the magnitude of normally distributed noise as polarisation aberrations will be highly spatially correlated. Un-accounted for systematic errors may make our data artificially sensitive to the fitted parameters and narrow our posteriors further. Characterisation of polarisation aberrations within the VAMPIRES instrument is the subject of ongoing work. For all these reasons we do not straightforwardly adopt the posterior distribution widths of our fitted parameters to constitute a physically meaningful measure of how well our model represents the data. We return to these ideas in Chapter 7.

## 5.2 Constraining $\mu$ Cephei’s dust grain size and chemistry with Mie scattering

So far we have demonstrated that computationally convenient Rayleigh scattering models are an effective way to constrain circumstellar geometry using data obtained at a single wavelength. However, use of a Rayleigh scattering approximation has a natural end point which we now explore in greater detail. When the circumstellar grain size is unknown, fitting to circumstellar geometry with a Rayleigh scattering approximation can be misleading as forward scattering is not modelled. The impact which forward scattering has on how circumstellar material is illuminated can be illustrated by considering a spherical dust shell (Figure 5.4 (b)). The polarised scattering function for Rayleigh grains ( $\lesssim 100$  nm) is symmetric, with a degeneracy in the forward and backward directions that causes the strongest polarisation feature to be the shell’s limb (Figure 5.4a (a,b) blue line). However, as grain size is increased and the threshold into the Mie scattering regime is crossed ( $\gtrsim 100$  nm) the scattering function is no longer forwards/backwards symmetric and instead peaks at larger scattering angles (Figure 5.4a (a,b) pink line). From the observer’s perspective this effect causes grains in front of the star to be more strongly polarised than those behind the star – this anisotropy of the scattering function is known as *forwards scattering* (Max Born 1975).

Whilst forwards scattering is a highly useful property of large grains that allows us to probe the ‘depth’ of circumstellar material (Perrin et al. 2009a; Avenhaus et al. 2018a; Wahhaj et al. 2024; Quanz et al. 2011; Perrin et al. 2015), it has the effect of distorting the apparent size of dust features – making them seem smaller than their true size. If a spherically symmetric dust shell of large Mie grains is fit under the assumption of small Rayleigh grains, strong forwards scattering can bias the fitted dust radius towards un-physically small values as the concentration of polarised flux moves forward.



**Figure 5.4:** The absolute polarised differential scattering cross section for dust grains of increasing size at an observing wavelength of  $750 \pm 25$  nm. For small grains satisfying the Rayleigh approximation (radius  $\sim 1$ – $100$  nm) the scattering function is forward/backward symmetric about a peak at a scattering angle of  $\sim 90^\circ$  (on the shell’s limb, bisecting the star in the plane of the sky). For larger Mie-scattering grains ( $\sim 100$  nm+) the scattering phase function becomes anisotropic and the peak scattering angle moves forward to larger scattering angles. Subplot (a) shows this relationship as a function of scattering angle. Subplot (b) projects this onto a sphere (here representing a dust shell) – showing how strong forwards scattering can distort the apparent radius of the circumstellar shell, making it appear smaller when viewed using polarimetry. The colours plotted are consistent between subplots (a) and (b).

Therefore, for grains too large for the Rayleigh approximation it is inescapable that circumstellar geometry, grain size and chemistry are entangled and must be fitted simultaneously. The consequences of modelling large (Mie) grains under the Rayleigh scattering approximation may be significant: derived parameters (such as the dust shell radius) used to place thermodynamic constraint on the dust species may not be valid. For these reasons, Mie scattering should be a necessary ingredient for modelling the dust within inner circumstellar environments – once simpler Rayleigh scattering models have been used to reduce the geometric subspace at a single wavelength. In Chapter 2 we derived a broad and physically motivated set of chemicals and dust grain sizes that may exist in  $\mu$  Cephei’s circumstellar environment. We now deploy the Mie scattering functionality of PIPPI to fit these chemical candidates and grain sizes to our multi-wavelength  $\mu$  Cephei dataset (6) obtained in 2023.


Previous studies of spectro-polarised interferometry data have performed chemical inference under the geometric assumption of a thin spherically symmetric dust shell (Norris et al. 2012; Haubois et al. 2019), which we established in Section 5.1 is an inadequately complex model for  $\mu$  Cephei’s circumstellar dust. It remains unclear how this geometric approximation may have influenced the outcome chemical inference in these studies. To explore this, we conduct our chemical inference on  $\mu$  Cephei under the assumption of three increasingly complex models for the inner circumstellar environment. We test the canonical thin spherical shell as a point of reference against the literature (Model A), an elliptical thin shell as a representative intermediate model (Model E), and an elliptical power law shell with a density enhancement as the best description of  $\mu$  Cephei’s inner circumstellar environment found with Rayleigh

modelling (Model S without extended material). Given our Rayleigh scattering modelling demonstrated that the extended material of Model S contributed a small fraction of the total polarised signal we omit the extended feature from chemical inference here.


The assumption of an optically thin media permits a tremendous number of combinatorial options for grain size and material. As all grains are independent, combinations of different scattering contributors may be linearly averaged to simulate many ‘mixed media’. For now we restrict our attention to the assessment of each dust species defined in Chapter 2 independently, and return to the idea of mixed models after examining the results of preliminary chemical inference. In addition to the novel advance we make in characterising a larger number of candidates for grain chemistry (see Chapter 2), we also implement two improvements in the characterisation of dust grain sizes: firstly by fitting to a much broader and more finely sampled range of dust grain radii, and secondly by fitting to observationally motivated power-law distributions which we anticipate will provide a more physically realistic description of the circumstellar dust than single grain-sizes (see Chapter 2).

### 5.2.1 Results of Mie scattering chemical inference

The result of our chemical inference with single-grain sizes and power-law grain distributions are depicted below in Figures 5.5 and 5.6.



**Figure 5.5:** Chemical and dust grain radius inference for  $\mu$  Cephei in 2023 (dataset 6), for Models A (panel 1), E (panel 2) and S\* (panel 3). Here single-sized grain distributions are assumed. Plotted values are the  $\chi^2_{red}$  values – the optimum value of (chemical, dust grain size) is marked with a red X. Our fits demonstrate that grain size is far better constrained than chemistry – many options for grain chemistry provide equitably good fits to data. There are three regimes of grain sizes that provide good fits,  $\sim 250$ , 750 and above 1200 nm.



**Figure 5.6:** Chemical and dust grain radius inference for  $\mu$  Cephei in 2023 (dataset 6), for Models A (panel 1), E (panel 2) and S\* (panel 3). Here power-law distributions are assumed. The power law distributions of grains are given in terms of a minimum grain size  $a_{min}$ , maximum grain size  $a_{max}$  and a power law exponent  $pl$ . Plotted values are the  $\chi_{red}^2$  values. Our fits demonstrate that grain size is far better constrained than chemistry – many options for grain chemistry provide equitably good fits to data. There are three regimes of grain sizes that provide good fits,  $\sim 250$ ,  $750$  and above  $1200$  nm.

### 5.2.2 Analysis of Mie scattering chemical inference

Several interesting trends emerge from testing a broad range of physically motivated chemical solutions and grain sizes. We performed chemical inference under the assumption of three different dust geometries – a thin spherical shell (Model A), a thin elliptical shell (Model E) and a power law elliptical shell with a density enhancement (Model S\*) – to investigate how approximations to the circumstellar geometry could influence the outcome of chemical inference. From our experiments it appears that whilst more simplified approximations to the circumstellar geometry will naturally result in higher  $\chi_{red}^2$  values, they do not appear to significantly change the relative performance of different chemical and grain sizes *within* each model class. We note that for Model A and E, the best fit solutions are  $250$  nm grains of amorphous Spinel, and that for Model S\* the best fit solution is for  $800$  nm grains of crystalline Forsterite. Whilst these are different solutions, in the paragraphs below we will argue that these should not be interpreted as physically significant differences.

The reason that differing geometric approximations do not significantly change the results of chemical inference here lies in the assumption of an optically thin media. If fit optimally, the best fit geometry for each chemical and grain size will have the same intrinsic amount of approximation in how polarised material is presented in the image plane. This would not necessarily be the case if full radiative transfer code was used as modifying the material and dust grain size could then increase the volume of multiple scattering events which would change how the dust geometry is projected into the image plane (eg. how opaque the dust morphology appears).

From our inference (Figures 5.5 and 5.6) it is evident that grain size and not chemistry is the primary driver of fit quality for each model class. Single-size distributions of small Rayleigh sized dust grains (below  $100$  nm) all yield poor fits to the data. Numerically, the best fits are obtained at  $\sim 250$  nm ( $\chi_{red}^2 < 1$ ), well into the regime where scattering requires Mie theory. The representative grain size fitted with single scattering models has historically been interpreted as the grain size that dominates the scattering cross section (Norris et al. 2012). A radiative environment dominated by scattering from

grains of size 250 nm would facilitate a scattering-driven stellar wind from  $\mu$  Cephei (Höfner 2008).

Whilst numerically the best fits occur at  $\sim 250$  nm, there is periodicity in the fit quality with respect to grain size – generally the fit deteriorates on either side of  $\sim 250$  nm and improves again by  $\sim 700$  nm. This ‘beating’ in the goodness of fit is characteristic of Mie scattering where grains are all approximated to be of a single uniform size and results from the regular pattern of constructive and destructive interference between the incident and reflected components of scattered light from the front and back of a grain (Jenkins 1981). As the grain size becomes larger than the wavelength the strength of this beating is suppressed as an increased number of paths reduces the coherence of re-emitted light. This behaviour is characteristic of the beginning of the *geometric optics* regime (Max Born 1975).

Beating is also seen in fits to the power law distributions – however, the broader range of grain sizes modelled in each distribution make this effect more subtle. The power law grain fits demonstrate agreement with the single grain fits in that power law distributions dominated by small grains provide the worst fit and a minimum grain size of at least  $\sim 250$  nm is required. Also in agreement is that many power law fits containing big grains are indistinguishable from each-other. We had anticipated that use of power law distributions would provide improved fit over single grain distributions on account of providing greater physical realism. However, our results demonstrate that many large-single-grains *and* many large power-law distribution grains provide indistinguishably good fits to our data. As our fits were not improved by broader distributions, this may suggest that the circumstellar dust grains have a narrow distribution – which has been suggested for the silicon-carbon dust of carbon-rich AGBs (Yasuda et al. 2012).

The beating we observe in both the single-grain and power-law grain fits suggests multiple regimes of dust grain sizes which fit data equally well. Figures 5.5-5.6 are plotted on a linear scale to illustrate the similarity of these solutions. From these results we could not confidently rule out grains with radius above  $\sim 1.25 \mu\text{m}$ , which is the upper limit for a grain supporting a scattering driven stellar wind from  $\mu$  Cephei (Höfner 2008) (derived in Chapter 2). Hence, we are unable to use this inference to conclude on whether  $\mu$  Cephei’s circumstellar dust grains can support a scattering driven stellar wind. However, our fits do seem to strongly support that the grains are *not* Rayleigh sized, suggesting a lower limit on the grain size of  $\sim 250$  nm. This grain size is close to the peak of grain sizes observed in the distribution within the ISM ( $\sim 300$  nm) (Srama et al. 2009; Mathis et al. 1977a).

Our fits provide far less constraint on dust chemistry than on dust grain size – whilst numerically there is an optimal solution, many chemical options tested are separated by  $\chi_{red}^2$  differences less than  $\sim 0.01$  and so we interpret these fits as equally good. The similarity of our chemical fits can be understood by examining the optical constant data for the tested chemicals – the real components ( $n$ ) which dominate the single scattering physics applied in PIPPI are all very similar as is their wavelength dependent behaviour. The results of this chemical inference are consistent with existing studies which found amorphous forsterite, enstatite and corundum to be indistinguishable (Haubois et al. 2019; Khouri et al. 2016; Ohnaka et al. 2016; Ohnaka et al. 2017b). It appears that testing our broader range of chemicals under a single scattering approximation did not yield more conclusive constraint of  $\mu$  Cephei’s circumstellar chemistry – all our tested chemical candidates provided approximately equitable fits to the data.


### 5.3 A renewed perspective on single scattering models

The three degenerate regimes of grain size obtained from chemical inference on  $\mu$  Cephei motivate re-analysis of the physical realism of single scattering models like PIPPI (and those implemented in existing spectro-polarimetric interferometry studies (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019)). These types of models are commonly referred to as ‘single scattering’ or ‘optically thin’ models. However, whilst they model the physics of a single scattering process, they are not *constrained* to producing solutions that exhibit optically thin behaviour. The parameter responsible for this lack of constraint is the ‘star-dust-contrast’ parameter common to all single scattering models, which freely modifies the contrast between star and dust to simulate changing the number of dust grains and thereby varying the polarisation amplitude. This parameter makes it possible to indefinitely increase the ‘density’ of the dust whilst continuing to model all grains as optically thin. This permits simple models to add grains beyond the point at which multiple scattering processes would commence and eventually reduce the polarised

visibility amplitudes in a physically realistic system. It is thus possible for single scattering models to enter a subspace of un-physical models in which the optically thin assumption is violated and so the models are self-contradictory.

The relationship between grain size and maximum achievable polarisation amplitude has been briefly mentioned in several direct polarimetric imaging studies on circumstellar environments (Khouri et al. 2016; Ohnaka et al. 2016; Adam et al. 2019). However, as these studies fit to circumstellar material using fully physical (and therefore constrained) radiative transfer code, the consequences of their comments have not been evaluated in the context of simple scattering codes like PIPPI. Specifically, no existing studies have addressed *where* the assumption of an optically thin dust breaks down relative to observationally relevant polarised visibility amplitudes, and how grain size effects where this break down occurs. Simple scattering codes like PIPPI were created to make modelling circumstellar environments computationally tractable – fitting to the circumstellar geometry, chemistry and grain size simultaneously yields an enormous parameter space. However, the state of modern radiative transfer codes (in both computational speed and the level of astrophysical detail) now permit re-analysis of these simple scattering models with these questions of physical realism in mind.

To quantify where the physical realism of simple models will break down, we injected the models fit in Section 5.2.2 into radiative transfer code MCFOST – modelling  $\mu$  Cephei ( $M_\star = 25M_\odot$ ,  $R_\star = 2257R_\odot$ ,  $T_\star = 3750\text{K}$ ) with a representative thin (2 mas) shell at  $1.5 R_\star$  (our Model Class A). For the range of grain sizes fitted in Section 5.2.2 we systematically increased the circumstellar dust mass (ie, added more dust grains) until polarised visibility amplitudes as high as  $\mu$  Cephei’s were obtained ( $V^2 \sim 1.1$ ). At this point we measured a) the polarised images, b) polarised interferometric observables obtained with the g18 mask and c) the optical depth of the model over the observer’s line of sight. Given many of our chemical candidates yielded very similar outcomes in Section 5.2.2, here we restricted our injection testing to a smaller subset of chemicals: amorphous forsterite, corundum, enstatite, spinel, and then as a metal-rich comparison, olivine. The results of these simulations are depicted in Figure 5.7, where the highest polarised visibility amplitude obtainable while the model’s dust has  $\tau \leq 1$  is plotted against grain size. Figure 5.7 demonstrates that for most of the dust grain sizes we have tested in Section 5.2.2, it is not possible to make optically thin dust shells with polarised visibility amplitudes as large as  $\mu$  Cephei’s.




**Figure 5.7:** Maximum obtainable polarised visibility amplitudes (g18 mask) for thin (2 mas) shell radiative transfer models whilst  $\tau \leq 1$  (dust shell is at  $\sim 1.5 R_\star$ ). For each dust grain species ( $\sim \Delta 10$  nm), dust mass is added until the optical depth is  $\sim 1$ , at which point the polarised visibility amplitude is measured. Our results demonstrate that only a small portion of the models in Section 5.2.2 are physically valid solutions – most models would instead require  $\tau > 1$  to have polarised visibility amplitudes as large as  $\mu$  Cephei’s (dashed grey line), violating the internal assumption of a simple scattering model.

At an observing wavelength of 750 nm, grains satisfying the Rayleigh scattering approximation may be as large as  $\sim 100$  nm (Papadakis 1965; Max Born 1975). The first phenomena apparent from our experiment is that for all species except Spinel, very small Rayleigh grains (below  $\sim 50$  nm) and large

grains (above  $\sim 200$  nm) are incapable of producing observationally high polarised visibility amplitudes whilst maintaining their optically thin nature. Rather, only grains of sizes between  $\sim 50$ -150 nm (half of which exceed the upper limit for the Rayleigh scattering approximation) are capable of achieving observationally high polarised visibility amplitudes whilst maintaining an optically thin density. Thus, in the case of  $\mu$  Cephei, a simple scattering code like PIPPI only truly allows us to model a small fraction of the physical solution space required to assess the plausibility of Hofner’s scattering driven stellar wind.

The simulations of Figure 5.7 were built using a thin shell (2 mas) at  $1.5 R_*$ . To confirm that these findings are not simply the result of either an artificially restricted geometry, the chosen dust shell radius, nor the use of a narrow dust grain distribution, we repeated our experiments for thicker shells (10 mas), for smaller shells ( $1.2 R_*$ ) and for broad power law distributions (exponent of -2). Given the similarity of outcomes across all iron-free chemicals tested in Figure 5.7, here we test these variations using amorphous corundum as a representative species. The results of these additional experiments are presented in Figure 5.8 and reveal that neither: increased proximity of dust to the stellar surface, increased thickness of the circumstellar shell nor a broader grain size distribution were able to significantly alter the conclusions of our original experiment.



**Figure 5.8:** Maximum obtainable polarised visibility amplitudes (g18) for various geometric configurations of radiative transfer models of an amorphous corundum circumstellar envelope. For each dust grain species ( $\sim \Delta 10$  nm) dust mass is added until the optical depth is  $\sim 1$  and the polarised visibility amplitude is measured. Changing neither: the proximity of dust to the stellar surface, the thickness of the dust shell nor the grain size distribution were able to dramatically increase the maximum attainable polarised visibility amplitudes of grains larger than  $\sim 200$  nm or below  $\sim 50$  nm.

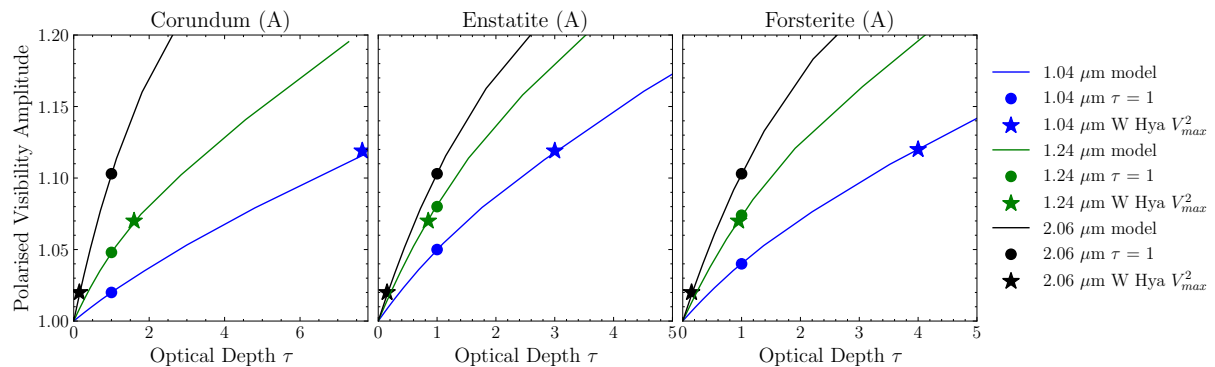
These findings have several interesting consequences for the results of our chemical inference on  $\mu$  Cephei (Section 5.2.2). Our best fit grain sizes were at  $\sim 250$ , 700 and above 1200 nm and yielded indistinguishably good fits to our data. However, Figure 5.7 indicates that most of the PIPPI models used in our inference, including the best fit solutions, *cannot be replicated in radiative transfer code*. If the grain sizes are truly  $\sim 250$ , 700 or above 1200 nm, the circumstellar dust would be of modest optical depth – rendering our assumption that absorption and multiple scattering effects are negligible contributors to wavelength dependence, invalid. In addition, in Section 5.2.2 our chemical inference suggested we could ‘rule out’ Rayleigh grain sizes below  $\sim 100$  nm as they provided a worse fit to our data than larger grains. However, our injection testing here also reveals that the PIPPI models we used to make this assessment also cannot be replicated in radiative transfer code – making our assertion precarious as we have not ruled out Rayleigh grains with a set of physically valid models.

From our re-analysis of the PIPPI models used in our chemical inference of  $\mu$  Cephei (g18) we have revealed that most grain sizes cannot produce observationally high polarised visibility amplitudes whilst maintaining that  $\tau < 1$  – demonstrating that most of the PIPPI models we used in chemical inference

(Section 5.2.2) were physically invalid solutions. The results of our chemical inference are therefore inconclusive (at best). Our findings necessitate the use of fully physical radiative transfer code for chemical inference on  $\mu$  Cephei’s circumstellar material.

The consequences of this analysis and the corresponding limitations of using simple scattering models will be different for every star and observing configuration – ultimately depending on the targets true optical depth and the spatial scales of polarised structure that are probed. Given the state of modern radiative transfer codes, our findings recommend the re-analysis of data within Norris et al. 2012 and Haubois et al. 2019. Though we note that Haubois et al. 2019 confirmed their best fit simple models ( $\sim 300$  nm grains) could be reproduced in radiative transfer code, it is unclear if all tested grain sizes would be robust to injection testing in radiative transfer – as we have demonstrated (Figure 5.7) it is possible for there to remain a portion of grain sizes for which simple scattering models represent physically viable solutions.

We were able to perform a preliminary re-analysis of the model of AGB W Hya fit in Norris et al. 2012 ( $R_\star = 18.7$  mas,  $R_{\text{shell}} = 37.9$  mas), for which large 316 nm grains were fitted using single scattering models and interpreted as evidence for Hofner’s scattering driven stellar wind. In Norris et al. 2012 the chemistry was not constrained from the data itself. Given the proximity of the dust ( $< 2R_\star$ ) to the stellar surface, forsterite, enstatite or corundum were suggested as thermodynamically viable candidates. We simulated the fitted models of each candidate, increasing the dust mass until the polarised visibility amplitudes matched those of W Hya at 1.04, 1.24 and 2.06  $\mu\text{m}$ , which we obtained from Figure 2a of Norris et al. 2012. Figure 5.9 depicts the results of this experiment – for all chemical candidates, a dust shell of 316 nm grains will be optically thick at 1.04  $\mu\text{m}$  when it is dense enough to produce W Hya’s polarised visibility amplitudes. For corundum, such a shell will also be optically thick at 1.24  $\mu\text{m}$ . For enstatite and forsterite, observationally equivalent shells remain optically thin at 1.04 and 1.24  $\mu\text{m}$ .



**Figure 5.9:** Injection testing models from Norris et al. 2012 in radiative transfer code. W Hya is constructed ( $R_\star = 18.7$  mas,  $R_{\text{shell}} = 37.9$  mas) with a shell of 316 nm grains of either: amorphous corundum, enstatite or forsterite. For each chemical candidate (column): solid lines represent the injected radiative transfer models of W Hya, filled in circles denote where  $\tau = 1$  and filled in stars denote W Hya’s polarised visibility amplitude. For all proposed candidates, the fitted dust shell will be optically thick at 1.04  $\mu\text{m}$ .

Our injection testing of the models in Norris et al. 2012 with modern and fully-physical radiative transfer code suggests that at least the 1.04  $\mu\text{m}$  wavelength may have been impacted by the limitations of single scattering models diagnosed within this section. Our findings are consistent with that of Ohnaka et al. 2016 (VLT/SPHERE-ZIMPOL –  $\lambda$  645, 748, 820 nm), where radiative transfer code was used to fit grains between 400 - 500 nm, obtaining a different grain size to Norris et al. 2012 when fully physical and constrained models were deployed. The authors of Ohnaka et al. 2016 noted this point of difference and found when they tested 300 nm grains, their models became optically thick like ours have here. Future analysis of W Hya with spectro-polarimetric interferometry data should re-perform chemical inference in a fully physical radiative transfer code to establish whether the modelling limitations we have explored will influence the results of grain size inference.

## 5.4 Chapter Summary

In this chapter we used the PIPPI package to model the circumstellar material around  $\mu$  Cephei under the assumption that the circumstellar dust is optically thin. We first fit the circumstellar geometry using a Rayleigh scattering approximation, demonstrating that the geometry of  $\mu$  Cephei is substantially more complex than the spherical thin shell models which have been modelled in existing polarimetric interferometry studies (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019).

Our best fit to  $\mu$  Cephei’s circumstellar geometry required an inhomogeneous inner envelope and an extended Gaussian feature. The inner envelope required two forms of asymmetry – both a geometric asymmetry and a density asymmetry. We found that asymmetries on larger spatial scales could be well described by a Gaussian blob at a location consistent with an extended feature reported in a previous study (Safonov et al. 2019). Our modelling also exhibited temporal variability in the asymmetry of the inner circumstellar region, which we suggest may be due to either evolving convective cells or a changing inhomogeneous circumstellar dust distribution. Our modelling demonstrates the powerful reach afforded by VAMPIRES + NRM data into the inner circumstellar environments of evolved stars and provides the most detailed description of  $\mu$  Cephei’s circumstellar geometry to date.

We then used the Mie scattering functionality of PIPPI to infer the chemical and grain size properties of  $\mu$  Cephei’s circumstellar dust, finding that dust grain size is far more constrained than chemistry and that many of our chemical options are indistinguishable. Our analysis found that grain sizes of  $\sim 250$ , 750 and above 1200 nm all provided indistinguishably good fits to our data. Whilst grain sizes of  $\sim 250$  and 750 nm would facilitate the scattering driven stellar wind proposed by Höfner 2008, grains above 1250 nm are too large. The degeneracy in our grain size solutions thus prevented us from using this analysis to ascertain the plausibility of a scattering driven stellar wind from  $\mu$  Cephei.

Finally, we re-analysed the physical realism of single scattering models deployed by PIPPI and by previous studies (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019). By injecting representative models into fully physical radiative transfer code we demonstrated that most of the solutions we fit to  $\mu$  Cephei using PIPPI’s Mie framework are not physically valid solutions. Most truly small Rayleigh grains (below 50 nm) and large Mie grains (above 200 nm) are incapable of producing polarised visibility amplitudes as large as  $\mu$  Cephei’s whilst maintaining an optically thin dust density. As most of the models we used in our chemical inference were self contradictory and violated the assumptions used in their creation, the outcomes of any inference performed on  $\mu$  Cephei using this framework are inconclusive.

We then performed injection testing of the models from Norris et al. 2012 of W Hya in full radiative transfer code, demonstrating that for at least the shortest observing wavelength (1.04  $\mu\text{m}$ ), models of 316 nm grains will be optically thick when dust is dense enough to produce polarised visibility amplitudes as large as W Hya. These findings suggest that re-analysis of this data in fully physical radiative transfer code may be a promising future avenue for research on this star, to confirm the fitted dust size was not affected by the limitations we have diagnosed in simple scattering models.

Simple scattering models like PIPPI provide a fast and effective method for converging on a description of the circumstellar geometry, capitalising on tools like `jax` and the acceleration provided by the current state of GPUs. However, our analysis on the physical realism of these scattering models suggests that they must be compared to radiative transfer for each modelled target – else it is possible that they enter a physically invalid regime which may compromise the results of inference on dust grain size. In the case of our  $\mu$  Cephei data we have demonstrated that use of PIPPI is not possible and that the results for grain size obtained with PIPPI were physically invalid. Our findings necessitate the full use of radiative transfer models for chemical inference on  $\mu$  Cephei and motivate the development of tools to fit radiative transfer models in subsequent Chapters.

## Chapter 6

# The PIRATES Image Reconstruction Algorithm

This chapter consists of the original version of our paper: L Lilley et al. (July 2025b). “PIRATES – a machine-learning framework for polarized, interferometric image reconstruction”. In: *JATIS (accepted)*. The most updated version is available electronically via JATIS.

### Chapter Overview

Within this chapter we present the PIRATES image reconstruction algorithm (Polarimetric Image Reconstruction AI for Tracing Evolved Structures). PIRATES is the first image reconstruction algorithm for optical wavelength polarised interferometry data. Our algorithm is explained and its performance demonstrated in our paper, which comprises the majority of this chapter. In this chapter we also motivate the creation of PIRATES and introduce relevant machine learning concepts.

### Statement of Contribution

The idea to use machine learning to create a polarimetric image reconstruction algorithm was originally my supervisors. We collaboratively built a preliminary framework for Stage 1 (see our paper), after which I was fully responsible for its development, the creation of Stage 2, the design and generation of our custom training data and algorithm training. I wrote our paper with editorial input from my supervisors.

## 6.1 Motivating the creation of PIRATES

In Chapter 5 we demonstrated that the established practice (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019) of using simple scattering codes to model spectro-polarised interferometry data has a significant limitation – simple codes are capable of producing unphysical models that violate their underlying assumption of optically thin dust. Our analysis demonstrated that for  $\mu$  Cephei as observed with VAMPIRES + NRM (g18) *most* of the dust grain sizes tested in chemical inference required moderate optical depth ( $1 < \tau < 10$ ) to produce polarised visibility amplitudes as large as  $\mu$  Cephei’s. Determining the point at which simple scattering codes diverge into an unphysical domain requires rigorous injection testing in radiative transfer codes, as this point is a function of the target and observing apparatus. These findings necessitate use of fully physical radiative transfer code for performing chemical inference on  $\mu$  Cephei. Fitting simultaneously to the geometry, chemistry and grain size of a circumstellar environment in radiative transfer code is a very expensive exercise. Unlike a simple PIPPI model which may be computed in  $\sim 0.3$  seconds, an equivalent radiative transfer model may require  $\sim 10$  minutes even with

parallelised execution on a high performance computer. As such, using an inference framework such as the MCMCs used with PIPPI models would be intractably slow.

Image reconstruction algorithms are complementary to parametric model fitting techniques. It is common for them to be used simultaneously and to serve as inspiration for the design of parametric models which are then used for astrophysical inference (Monnier et al. 2007b; Buscher 2015; Akiyama et al. 2019b). In a large and complex parameter space like that which defines a circumstellar environment, an image reconstruction algorithm would provide a particularly useful tool for narrowing the parameter space. At the time of writing this thesis no such algorithms exist for optical wavelength polarised interferometry data – here we present PIRATES as the first.

One of the central challenges of optical *polarised* interferometric image reconstruction is that it requires the reconstruction of three images (Stokes I, Q and U). These images must display an internally consistent and physically constrained description of the circumstellar environment, as mathematically they are spatial maps of vector components. To establish a physically self-consistent prior, PIRATES is trained on images generated from radiative transfer models of dusty circumstellar environments that are agnostic to specific astrophysical structures (like disks or shells). As demonstrated in our paper (Lilley et al. 2025b), PIRATES is able to reconstruct self-consistent polarised image triplets that also provide optimised fits to polarised (and unpolarised) interferometric observables.

## 6.2 A brief introduction to Machine Learning

Machine Learning (ML) techniques have seen exponential growth in research and commercial markets in the last 10 years. The internet provides a staggeringly large data source of training data – this, coupled with technological advancement and the lowering cost of GPU chips have enabled the rapid development of extremely high performing ML models that have captured the attention of both developers and the general public. Interest in and use of machine learning techniques has been driven by the remarkably high performance of algorithms like AlexNet (Krizhevsky et al. 2012), AlphaFold (Jumper et al. 2021) and GPT-3 (Brown et al. 2020), which are being rapidly integrated into different parts of modern life. The early mathematical principles underpinning some of today’s most powerful ML techniques first formed in the 1940s. In 1943 the first theoretical model of an artificial neural network was proposed (McCulloch et al. 1943), inspired by the functional biology of real neurons in the brain. In 1958 the introduction of the Perceptron model (Rosenblatt 1958) re-framed the artificial neuron as a computational model – becoming the first algorithm to learn from data. However it wasn’t until the concept of ‘back-propagation’ was rediscovered in the 1980s that representational power was enhanced by the ability to train multi-layer networks (Rumelhart et al. 1986).


Machine learning algorithms can be broadly classified as ‘supervised’ or ‘unsupervised’, referring to whether they are taught to *learn* a mapping or to *find* a mapping between a dataset of input and output points. A supervised algorithm that is trained to learn a function  $f(x, \theta)$  consisting of parameters  $\theta$  which maps inputs  $x \in \mathbb{R}$  to outputs  $y \in \mathbb{R}$  based on labelled training data. Internal to the algorithm, this mapping is encoded in an architecture of neural ‘units’ which each have weights  $w$  and biases  $b$ , combined linearly and then parsed as input to a non-linear activation function  $\sigma$  which gives the algorithm the ability of non-linear expression (Equation 6.1). The model’s output prediction  $\hat{y}$  is then compared with the true output  $y$  using a loss function  $\mathcal{L}(y, \hat{y})$ . To train the model, the parameters ( $w, b$ ) are updated by computing the gradient of the loss function with respect to each neural parameter via the chain rule, adjusting each parameter in the direction which reduces the loss (Equations 6.2 and 6.3). The size of each parameter update is tuned with the learning rate  $\eta$  (Deisenroth et al. 2020).

$$\hat{y} = f(x, \theta) = \sigma(w^T x + b) \quad (6.1)$$

$$w \leftarrow w - \eta \frac{\partial \mathcal{L}}{\partial w} \quad (6.2)$$

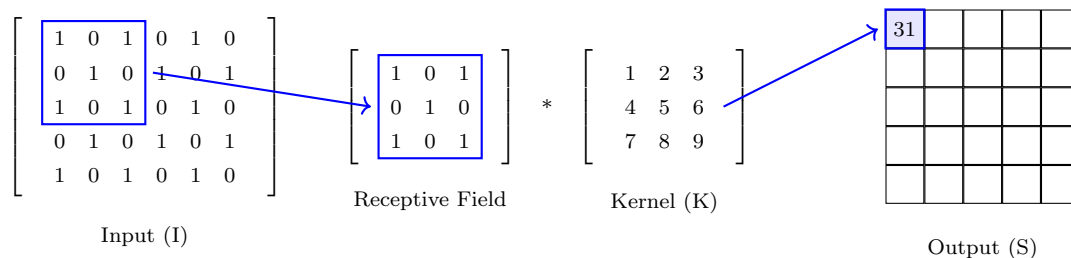
$$b \leftarrow b - \eta \frac{\partial \mathcal{L}}{\partial b} \quad (6.3)$$

This process of computing the gradients and updating the parameters is repeated iteratively across batches of the training examples until the loss function is minimised – representing that a mapping has been learned from the training data (Goodfellow et al. 2016). The internal mechanism for this mapping can take many forms that depend on the task at hand. The simplest form is the perception – the building blocks of fully connected or dense neural networks. Two dimensional architectures of perceptions can be arranged in parallel (width) and in sequence in layers (depth) in a variety of structures to vary the network’s representational power (Figure 6.1).



**Figure 6.1:** A simple fully connected (dense) neural network with 3 input neurons, 4 hidden neurons and 2 output neurons. Functional complexity is added with additional neurons – in either depth (adding more hidden layers) or width (adding more neurons in the same layer). The structure of a dense network requires data to be flattened prior to input – rendering them of limited utility when applied to images which have complex two dimensional structure crucial for information representation.

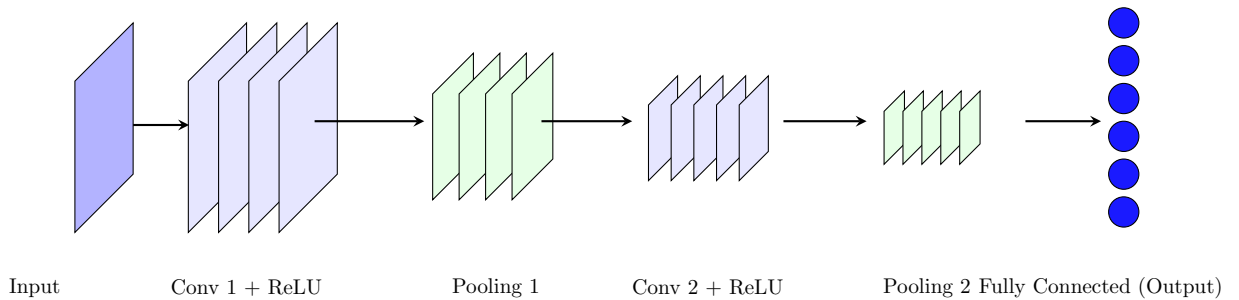
Though powerful, a major limitation of dense architectures is that data must be flattened (made one-dimensional) prior to network input. This reshaping removes information on two-dimensional structure, meaning that dense networks are not an optimal network architecture for processing and understanding images. A solution to this issue is found in the design of Convolutional Neural Networks (CNNs). Unlike the dense architecture, the CNN architecture is designed to preserve spatial information – it permits two-dimensional input to the network and the ‘weights’ and biases are assigned to two-dimensional kernels that perform convolution operations (Figure 6.2). As such, CNNs are architecturally well designed to encode the complex higher order spatial relationships crucial for meaningfully processing images (Wang et al. 2023). The core operation implemented within a CNN is the convolution – from an input image  $I$  and a learned filter (convolutional kernel)  $K$ , the convolution operation produces a feature map  $S$  (Figure 6.2, Equation 6.4). The filter  $K$  slides over the input image  $I$  and at each position the dot product of the overlapping region is computed. This process allows the network to learn patterns from two-dimensional spatial structures like edges or textures.



**Figure 6.2:** An example of the convolution operation at the core of a CNN architecture: a  $3 \times 3$  kernel is applied to an input image, producing a single output value. Each convolutional kernel is shared across the whole image or latent feature space (the two-dimensional ‘hidden layers’ of a CNN), giving the CNN a powerful structurally derived inductive bias. The two-dimensional structure of convolutional kernels permits the interpretation of complex spatial structures.

$$S(i, j) = (I * K)(i, j) = \sum_m \sum_n I(i + m, j + n) \cdot K(m, n) \quad (6.4)$$

Filters are applied across the entire spatial domain of each image, providing CNNs with a powerful structurally derived inductive bias (LeCun et al. 1989). A CNN typically has many kernels which may be added in depth or width to modify the network’s representation power. The application of  $N$  kernels will produce  $N$  feature maps which comprise a ‘latent feature space’ equivalent to the ‘hidden layers’ of a dense network. Convolutional neural networks most commonly take images as inputs and downsize the latent feature space via Pooling layers (which apply aggregate statistics to reduce dimensionality), until the required output dimensionality is obtained (Figure 6.3). However, CNNs can also be built in the reverse order and instead used to *generate* images from a vector – here we exploit this functionality in the architecture of PIRATES, which is a supervised algorithm that learns to *reconstruct* images from a one-dimensional vector of interferometric observables.



**Figure 6.3:** A simple schematic of a typical CNN: two convolutional layers each followed by pooling layers which reduce the dimensionality of the latent space, leading to a fully connected output layer. This structure is typical when the goal is to learn information from input images. In PIRATES we apply the opposite structure – from an input vector of points, images are *reconstructed* by the application of transpose convolutional layers.

Rather than using convolutional layers, *transpose* convolutional layers are used (Equation 6.5), which enable the dimensionality of the input to be increased to that of the three output polarised images (Stokes I, Q and U). Use of this architectural structure encodes the physically realistic scattering relationships embodied by our training data into two dimensional feature maps that are shared between the output images, creating a prior that the images are internally consistent.

$$S(i, j) = (I *^T K)(i, j) = \sum_m \sum_n I(m, n) K(i - m, j - n) \quad (6.5)$$

### 6.3 Paper: PIRATES – a machine learning framework for polarised, interferometric image reconstruction.

We now present our paper (Lilley et al. 2025b), which provides all further details on our algorithm’s design, training and performance. This paper has been accepted to JATIS.

# PIRATES - a machine-learning framework for polarized, interferometric image reconstruction

Lucinda Lilley <sup>\*a,b</sup>, Barnaby Norris<sup>a,b,c</sup>, Peter Tuthill<sup>a,b</sup>, Eckhart Spalding<sup>a</sup>, Miles Lucas<sup>d,e</sup>, Manxuan Zhang<sup>i</sup>, Maxwell Millar-Blanchaer<sup>i</sup>, Christophe Pinte<sup>m</sup>, Michael Bottom<sup>d</sup>, Olivier Guyon<sup>e, f,g,h</sup>, Julien Lozi<sup>e</sup>, Vincent Deo<sup>e,k</sup>, Sébastien Vievard<sup>e,h,n</sup>, Alison P Wong<sup>l,a,b</sup>, Kyohoon Ahn<sup>j, e</sup>, Jaren Ashcraft<sup>f</sup>

<sup>a</sup>Sydney Institute for Astronomy, School of Physics, Physics Road, University of Sydney, NSW 2006, Australia

<sup>b</sup>Sydney Astrophotonic Instrumentation Laboratories, Physics Road, University of Sydney, NSW 2006, Australia

<sup>c</sup>AAO-USyd, School of Physics, University of Sydney, NSW 2006, Australia

<sup>d</sup>Institute for Astronomy, University of Hawai'i, 640 N. Aohoku Pl., Hilo, HI 96720, USA

<sup>e</sup>Subaru Telescope, National Astronomical Observatory of Japan, 650 N. Aohoku Pl., Hilo, HI 96720, USA

<sup>f</sup>College of Optical Sciences, University of Arizona, Tucson, AZ 87521, USA

<sup>g</sup>Steward Observatory, University of Arizona, Tucson, AZ 87521, USA

<sup>h</sup>Astrobiology Center, 2 Chome-21-1, Osawa, Mitaka, Tokyo, 181-8588, Japan

<sup>i</sup>Department of Physics, University of California, Santa Barbara, CA, 93106, USA

<sup>j</sup>Technology Center for Astronomy and Space Science, Korea Astronomy and Space Science Institute, 776 Daedeok-daero, Yuseong-gu, Daejeon 34055, Republic of Korea

<sup>k</sup>Optical Sharpeners SAS, France

<sup>l</sup>Discipline of Business Analytics, University of Sydney Business School, University of Sydney, NSW 2006, Australia

<sup>m</sup>Univ. Grenoble Alpes, CNRS, IPAG, 38000 Grenoble, France

<sup>n</sup>Space Science and Engineering Initiative, College of Engineering, Institute for Astronomy, University of Hawaii, 640 North Aohoku Place, Hilo, HI, 96720, USA

## Abstract.

Optical interferometric image reconstruction is a challenging, ill-posed optimization problem which usually relies on heavy regularization for convergence. Conventional algorithms regularize in the pixel domain, without cognizance of spatial relationships or physical realism, with limited utility when this information is needed to reconstruct images. Here we present PIRATES (Polarimetric Image Reconstruction AI for Tracing Evolved Structures), the first image reconstruction algorithm for optical polarimetric interferometry. PIRATES has a dual structure optimized for parsimonious reconstruction of high fidelity polarized images and accurate reproduction of interferometric observables. The first stage, a convolutional neural network (CNN), learns a physically meaningful prior of self-consistent polarized scattering relationships from radiative transfer images. The second stage, an iterative fitting mechanism, uses the CNN as a prior for subsequent refinement of the images with respect to their polarized interferometric observables. Unlike the pixel-wise adjustments of traditional image reconstruction codes, PIRATES reconstructs images in a latent feature space, imparting a structurally derived implicit regularization. We demonstrate that PIRATES can reconstruct high fidelity polarized images of a broad range of complex circumstellar environments, in a physically meaningful and internally consistent manner, and that latent space regularization can effectively regularize reconstructed images in the presence of realistic noise.

**Keywords:** Image Reconstruction, Polarimetric Image Reconstruction, Polarimetric Interferometry, Interferometry, Machine Learning, VAMPIRES Instrument.

\* Correspondence to: Lucinda Lilley, [lucinda.lilley@sydney.edu.au](mailto:lucinda.lilley@sydney.edu.au)

## 1 Introduction

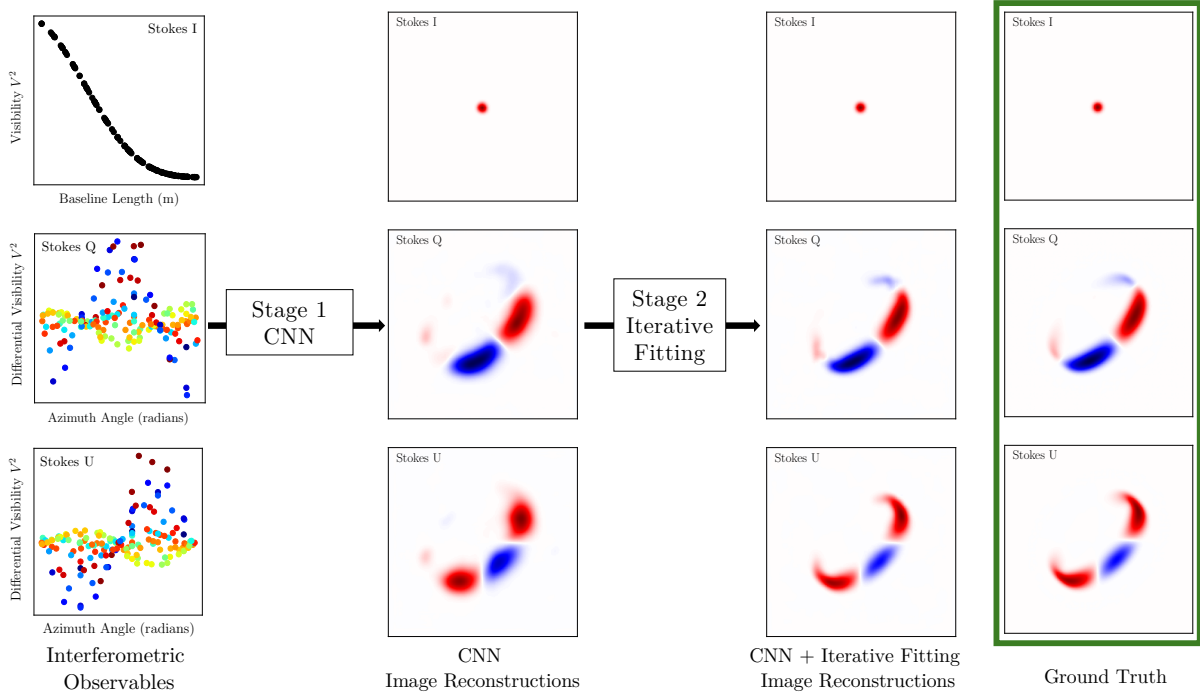
Image reconstruction is an essential technique for deriving physical meaning from interferometric observables and is complementary to forwards modeling that yields plausible descriptions of an astrophysical scene. Reconstructing an accurate image from interferometric observables (visibilities and closure phases) is an inherently ill-posed and non-convex optimization problem.<sup>1</sup> The reconstruction process requires the information contained in interferometric observables to be re-represented in an over-sampled basis (the image), with a significantly higher number of degrees of freedom (pixels) than there are observables - a task with a vast number of degenerate solutions, most of which will be un-physical or dominated by spurious structure or noise. To increase the probability of obtaining physically meaningful solutions, the convergence of an image reconstruction algorithm is often heavily driven by regularization penalties, which place constraints on the statistics of the reconstructed images. Commonly used regularization penalties include; maximum entropy, total variation, L1 and L2 norms.<sup>2-4</sup> These forms of pixel-distribution regularization typically aggregate some function of the pixel values, losing information on spatial connectivity and two dimensional structure. In addition to regularization, a prior on the flux distribution can be specified, for example that flux is expected in the image center.<sup>5</sup> With the assistance of regularization and a prior, image reconstruction algorithms aim to converge on a reconstructed image that is meaningful, contains minimal spurious noise, and most importantly, does not mislead the viewer to infer more complex phenomena than may be mathematically constrained by the original observables.

There are many successful image reconstruction algorithms for un-polarized, optical interferometry,<sup>2-7</sup> however no algorithms currently exist for optical, polarimetric interferometry. The polarization state of light contains critical information otherwise not captured by standard, unpolarized interferometry, and major facilities are now looking towards fully leveraging this information to broaden the observational reach of new and existing instruments.<sup>8,9</sup> Polarized interferometric image reconstruction has an important additional challenge compared to unpolarized interferometric image reconstruction: the images of the different polarization components (e.g. Stokes I, Q and U) are more accurately described as spatial maps of vector components, which at optical wavelengths define the geometry of light-scattering processes. As such, the Stokes images require simultaneous reconstruction, and must mutually depict a physically realistic and internally consistent description of the physics of the scattered scene. Regularizing for these physical constraints from the distribution of flux within and between each Stokes image is a challenging task without a straightforward solution, as each pixel has a complex relationship with its surrounding pixels, both spatially and across each polarization vector component. Polarimetric images of Stokes Q and U also break the assumption of image positivity which is required for the use of many pixel-distribution regularization penalties like maximum entropy or the L1 and L2 norms.<sup>10,11</sup> As such, polarimetric image reconstruction has the simultaneous requirement of a nuanced adaptation of common pixel-distribution regularization penalties, and a way to ensure that reconstructed Stokes image triplets are internally consistent and physically meaningful.

### 1.1 *The PIRATES algorithm*

Here we present PIRATES, a two stage machine learning framework (pre-trained CNN + iterative fitting), that can handle these complex regularization requirements, parsimoniously reconstructing

high fidelity polarized images that accurately match polarized observables (Figure 1).



**Fig 1** Schematic of PIRATES’ dual structure. Stage 1 is a pre-trained CNN which learns the mapping from polarized interferometric observables to the Stokes image triplet (I, Q, U). This mapping between observables and images is learned from ‘probe’ radiative transfer images which are agnostic to specific astrophysical structure. The resulting algorithm then generalizes well to astrophysically meaningful examples like the disk above. Stage 2 is an iterative fitting algorithm - it optimizes the pre-trained CNN images now with respect to their interferometric observables. Changes to the images are made via updates to convolutional kernels in a latent feature space, and not the pixel wise adjustments of conventional unpolarized image reconstruction codes, imparting a structurally derived implicit regularization on reconstructed images.

Firstly, the radiative transfer code MCFOST<sup>12,13</sup> is used to probe a space of physically realistic and relevant polarized images of smooth and diffuse circumstellar dust. A distribution of Gaussian probe ‘volumes’ of dust are randomly generated and injected into the model domain, spanning the physical field of view and spatial scales represented by the observables whilst remaining agnostic to specific astrophysical structures. Similarly, a distribution of dust densities, grain sizes, etc. is used to probe the space of valid scattering solutions. Pre-training the CNN on this data teaches it to learn the mapping between polarized images and polarized interferometric observables (visibilities and closure phases), over a space of self-consistent images which have physically meaningful spatio-polarimetric relationships. Rather than using an explicit term, regularization is encoded in the CNN training data - at inference time, reconstructions are implicitly regularized by the statistics of the training distribution. The architecture of a CNN allows the mapping between observables and images to be learned via non-linear operations of convolutional spatial kernels, and not the pixel-level adjustments performed in conventional image reconstruction codes. Whilst the CNN learns from considerable training data, the enforced generalization of convolutional kernels to the whole of each image gives CNNs a powerful, structurally derived inductive bias.<sup>14</sup> The combination

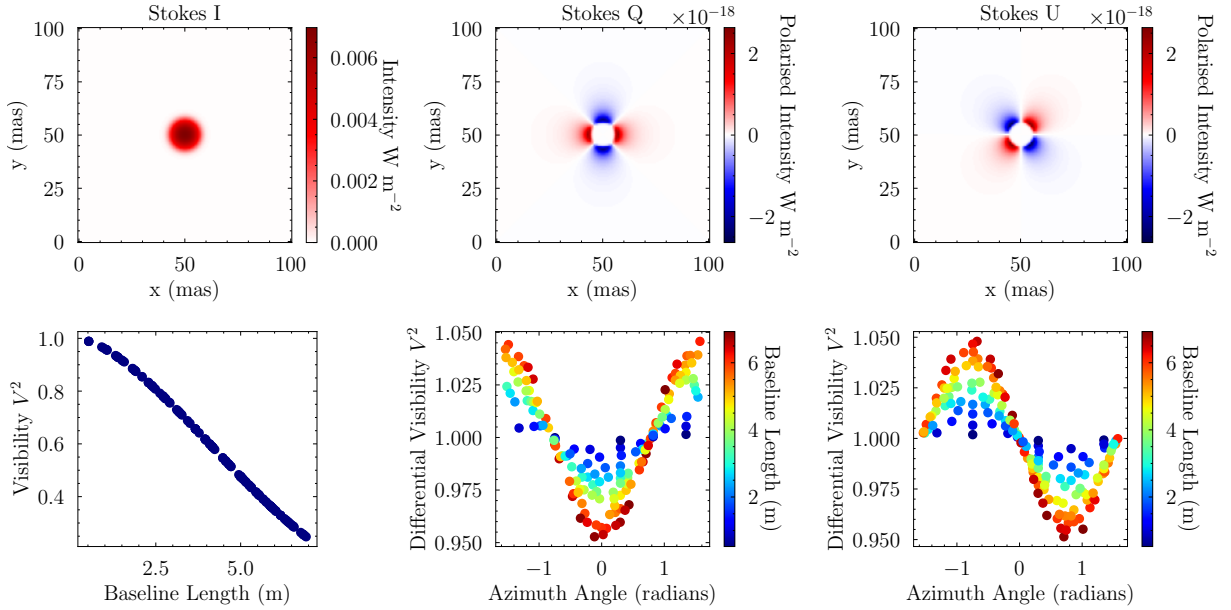
of appropriately regularized pre-training data and the architectural inductive bias of a well built and tuned CNN imparts a powerful implicit regularization on image reconstructions,<sup>15</sup> in this case towards physically consistent Stokes triplets depicting scattered light.

At inference time, we use our pre-trained CNN to reconstruct polarimetric images from observed visibilities and closure phases, for a single test example (stage 1). The CNN-predicted image provides a moderate detail and low noise estimate of the reconstructed images with a single pass, and initializes the iterative fitting in a physically meaningful feature space close to the true minimum (Figure 1, middle column). Iterative fitting (stage 2) then performs additional refinement to the CNN-predicted images by now optimizing them (in the learned feature-space) with respect to the interferometric observables, making physically meaningful updates to the polarized images without introducing spurious structure or noise into the images (Figure 1, right column).

## 1.2 Differential interferometric polarimetry

Here we apply PIRATES to simulated polarized-interferometric observables produced by the VAMPIRES Instrument (Visible Aperture Masking Polarimeter for Resolving Exoplanetary Signatures / Evolved Stars), a differential imaging polarimeter mounted on the SCExAO bench of the SUBARU Telescope.<sup>9,16</sup> VAMPIRES is equipped to perform single telescope interferometry via non-redundant aperture masking (NRM). VAMPIRES + NRM is a prototypical example of polarized interferometry - an instrument built on earlier successes from the now decommissioned NACO + SAMPOL/VLT,<sup>17,18</sup> and SUSI.<sup>19</sup> Operating at optical wavelengths (600 - 800 nm) where polarization signatures induced by dust-scattered starlight are strong,<sup>20</sup> VAMPIRES + NRM performs simultaneous multi-wavelength polarimetric interferometry of linear polarization states Stokes I, Q and U, providing a sub-diffraction limited spatio-spectral description of observationally challenging inner-circumstellar environments that lay host to dusty material,<sup>16,21</sup> such as proto-planetary disks and evolved-star mass-loss shells. Within the present work, we demonstrate our algorithm by simulating observations with the 18 hole ‘g18’ aperture mask, which has 153 baseline vectors.<sup>16</sup> For consistency, the IAU polarimetric coordinate system is adopted throughout this paper.<sup>22</sup>

The reduced data products for polarimetric interferometry are squared visibilities (or visibility ratios) and closure phases for Stokes  $I$ ,  $Q$  and  $U$ . The Stokes formalism represents each point in a wavefront with a Stokes vector,  $\vec{S} = (I, Q, U, V)$ , where  $I$  is intensity, and  $Q$  and  $U$  are the orthogonal linear polarization components.<sup>23</sup> Circular polarization,  $V$  is omitted from the present work (set as  $V = 0$ ), as it is not measured by VAMPIRES. For Stokes  $I$ , visibilities take the conventional format (Figure 2, left panel), however for Stokes  $Q$  and  $U$ , the visibility and closure phase measurements are polarized and differential (Figure 2, middle and right panels). We define the differential visibility as the ratio of orthogonally-polarized raw visibilities, and the differential closure phase as the difference between orthogonally polarized raw closure phases. For VAMPIRES, measurement of Stokes  $Q$  and  $U$  is the result of a triple-differential measurement - linear polarization states H and V (Stokes  $Q$ ), and H45 and V45 (Stokes  $U$ ), are made thrice by modulating three optical components - a ferroelectric liquid crystal (FLC), a half wave plate (HWP) and a polarizing beam-splitter (BS).<sup>9</sup> Triple division of the corresponding visibilities yield very high precision and self-calibrated measurement of Stokes  $Q$  and  $U$ , and remove the conventional need for a PSF calibrator star.<sup>16</sup> For description of the VAMPIRES instrument and data reduction process, please see<sup>16</sup> and.<sup>9</sup>




**Fig 2** Simulated VAMPIRES observables for a thick, spherically symmetric circumstellar envelope: images and visibilities in Stokes I, and polarimetric and self-differential images and visibilities in Stokes Q and U. Stokes I visibilities have the traditional interferometric format, whilst Stokes Q and U visibilities are polarized and self-differential. Polarimetric visibilities of symmetric objects have a characteristic sinusoidal and co-sinusoidal appearance for Stokes Q and Stokes U.

Physical interpretation of polarimetric interferometry data first involves modeling the circumstellar geometry. Historically, this has been achieved via parametric forwards modeling.<sup>24–26</sup> For computational tractability, several significant physical assumptions are typically made - firstly, that the scattering environment is optically thin and homogeneous.<sup>24–26</sup> Secondly, models for scattering allow for additional dust grains to increase polarization amplitudes indefinitely: an approximation that eventually becomes unphysical with the onset of effects from multiple scattering events.<sup>24–27</sup> The limitations of these assumptions are elucidated by recent VAMPIRES datasets for complex evolved stars,<sup>28</sup> for which there are visible and systematic departures from these assumptions. To properly model complex targets, a physically complete treatment of optical depth and inhomogeneity in the dust is required, necessitating full radiative transfer models. Parametric forwards modeling with radiative transfer code is computationally expensive - in contrast to simple parametric scattering models which may run with GPU acceleration in several milli-seconds, a full radiative transfer model may take several minutes. In addition, the design and choice of candidate parametric models becomes counter intuitive when inhomogeneity, optical depth and asymmetry all interact. These limitations and challenges motivate the development of our algorithm, which we design as a tool to reconstruct polarimetric images of challenging and complex circumstellar scenes. PIRATES aims to produce maximally parsimonious and physically self-consistent polarized images which best fit the polarized interferometric observables. The resulting images can also be used to initialize and constrain further parametric model-fitting.

## 2 Materials and Methods

### 2.1 Algorithm Architecture

Here we summarize the structure of PIRATES (Figure 1), the chosen architecture and hyper-parameters; the tuning performed to obtain these values is explained in Section 3. All code within this study is written in Python, neural networks are built in TensorFlow<sup>29</sup> and Keras<sup>30</sup> and run on an NVIDIA GeForce RTX 4090.



**Fig 3** Structure of PIRATES CNN (stage 1). Inputs to the network are vectors of polarized interferometric observables (visibilities and closure phases) ( $1, 2907$ ). These are accepted by a dense layer, and then reshaped into a two dimensional image with many channels ( $4 \times 4 \times 256$ ), which allows for the subsequent application of transpose convolutional layers. Five stride-two transpose convolutional layers then up-sample until the spatial output size ( $128 \times 128$ ) is achieved, with 200 kernels in each layer. Then follow three stride 1 convolutional layers - the first with 200 kernels, and the last two with 8 and 3 to downsize the output channel dimension to 3 (one for each of I, Q and U). The outputs of the network are the polarized image triplets - Stokes I, Q and U ( $128 \times 128 \times 3$ ). Choice of this architecture is discussed in Section 3.1.

Stage 1 of PIRATES is a CNN which takes polarimetric observables as input, and outputs reconstructed polarimetric images (Figure 3). Each training example input consists of a single vector of interferometric observables: visibilities and closure phases for Stokes I, and polarized differential visibilities and closure phases for Stokes Q and U. This vector is then reshaped to have two spatial and one filter dimension, after which it is passed to eight transpose convolutional layers. The first five of these layers have a stride of 2 and up-sample the input until the required spatial dimensions are obtained ( $128, 128$ ). 200 kernels of size ( $4 \times 4$ ) are used in each layer. The last three convolutional layers have a stride of 1 and have 200, 8 and 3 convolutional kernels respectively, to downsize the number of channels to the final output dimensions of ( $128, 128, 3$ ) - corresponding to the image triplet Stokes I, Q and U. The CNN uses the training data (probe images from radiative transfer code) to learn the mapping from observables to images with an image based mean squared error (MSE) as the loss function.

The network is trained on 10,000 training examples, using an automatic learning rate scheduler that commences at  $10^{-4}$ , and training is conducted until there is negligible improvement to the


validation loss. A batch size of 32 was used. Dropout was used as regularization and tuned to 0.3. Leaky ReLU activation functions were used, the negative slope coefficient  $\alpha$  was tuned to 0.25. Training of the CNN took approximately 4 hours on an NVIDIA GeForce RTX 4090.

Stage 2 of PIRATES, an iterative fitting mechanism, is only applied at inference time and to a single observation at a time. First, stage 1 (the pre-trained CNN) is used to reconstruct polarized images from observables. The kernel weights of a copy of the pre-trained CNN are then further updated - wherein the objective function is now replaced with the MSE between the true interferometric observables and those recalculated from the pre-trained CNN images. Then commences stage 2 (iterative fitting) - the pre-trained CNN images are optimized with respect to their fit to the true interferometric observables. Since convolutional kernel weights – rather than the image pixels – are adjusted during iterative fitting, the reconstructions take place within the physics-informed feature space learned by the pre-trained CNN.

At the commencement of stage 2, each copy of the CNN becomes a ‘single use’ item - the pre-trained CNN must be freshly reloaded every time iterative fitting is deployed. Whilst iterative fitting continues to update the CNN weights, we refer to it as ‘fitting’ and not ‘training’, as iterative fitting does not continue to learn a generalizable mapping. Instead, the pre-trained convolutional kernel weights are used as initialization for the iterative fitting, which then optimizes a single image with respect to its observables. When running the iterative fitting, we used an adaptive learning rate scheduler commencing from a learning rate of  $1e-5$ , with a conservative reduction rate of 0.8 and a patience of 5. To ameliorate computational overheads incurred when ending each training epoch, we populated each epoch with  $\sim 75$  instances of the input data, balanced against the requirement to regularly complete epochs to trigger the learning-rate scheduler (a limit which could be removed with custom-written code). Iterative fitting typically takes 5 - 10 minutes and 1500 - 2000 epochs to converge.

## 2.2 Training and Validation ‘Probe’ Datasets and Simulated Astrophysical Data

Our CNN is pre-trained on a custom built dataset of polarimetric ‘probe’ images and resulting polarimetric interferometric observables generated using radiative transfer code MCFOST<sup>12,13</sup> (Figure 4). Choice of training data is crucial for imposing a physically meaningful prior of realistic spatial and polarization distributions. Our training data contains a distribution of scattering representations which feature minimal noise or unwanted artifacts from the radiative-transfer calculations, and have no bias towards any particular astrophysical structure. To achieve this balance, we perform radiative-transfer calculations for an ensemble of randomly generated dust-density volumes, each of which is made up of sets of three-dimensional-Gaussian probe-volumes of randomly chosen location, size and density, and situated around a randomly sized star. Each of these then produces corresponding polarized images and observables (Figure 4). Gaussian volumes spanning a range of sizes probe all spatial scales while avoiding aliasing due to their smoothness, discouraging pixel-level noise or sharp structures in image reconstructions. We convolve all images with a Gaussian kernel of full width half maximum of 1.5 pixels to help suppress pixel-scale stochastic noise arising from the discrete number of photons used within radiative transfer calculations and the edge artifacts from a discrete number of volumetric model cells.



**Fig 4** Five examples of CNN pre-training probe image triplets. Each training example (column) consists of a (log10) Stokes I (first row), Q (second row) and U (third row) image. Images are calculated via a radiative transfer code, using an ensemble of randomly drawn Gaussian probe densities. The CNN learns a mapping between interferometric observables (visibilities and closure phases) and polarized image triplets, that is agnostic to specific astrophysical structure. Moreover, it learns to represent this mapping as a maximally parsimonious, physics-informed feature space.

Input and output data were both normalized prior to training. Inputs of interferometric observables are normalized with respect to their subclasses. Outputs of polarized images are also normalized with respect to their subclasses but require two intermediate modifications. Since the star is orders of magnitude brighter than the surface brightness of the dust, the logarithm of the Stokes I image is used whilst training the network, so that faint dust structures in Stokes I adequately influence loss calculations. In addition, all three images are normalized by the sum of flux in the (natural) Stokes I image, as is conventional in interferometric image reconstruction since the visibilities are independent of absolute flux.


Radiative transfer code was used to make 5500 pairs of polarized images and interferometric observables, which were then subject to augmentation. Images were randomly rotated and flipped, with appropriate Mueller matrices applied to maintain the polarimetric coordinate system, and interferometric observables were resampled from augmented images. A total of 11,000 examples were generated, 10,000 of which were used for training data and 1000 of which were reserved as validation data, which is used to independently evaluate the prediction accuracy of the CNN at the end of each pre-training epoch. Separately, prototypical circumstellar objects like disk and envelopes were injected into radiative transfer code, to create a set of simulated astrophysical data. This dataset was used to evaluate the performance of the iterative fitting, and was not used during training of the CNN (stage 1). More broadly, the simulated astrophysically dataset was used to assess how well PIRATES generalizes to spatially extended and astrophysically meaningful structures.

### 3 Results and Discussion Part I - Architectural Tuning

This section summarizes the tuning of PIRATES’ architecture. For PIRATES’ final results and performance, see Results and Discussion Part II (Section 4). PIRATES hyper-parameters were tuned on the validation data to maximize the accuracy of final iteratively fit images, using the MSE and SSIM (structural similarity index measure) to quantify pixel-wise and structural accuracy.

#### 3.1 Stage 1 - CNN Architecture

Three architectural hyper-parameters were tuned: the number of convolutional layers, the number of convolutional kernels per layer, and the size of the convolutional kernels. We evaluated the impact of the number of layers and kernels simultaneously - conducting a grid search over reasonable values of both parameters. We maintained a ‘rectangular’ network shape - adding the same number of kernels to every layer, with the exception of the final two convolutional layers which downsize the kernel number to the output channel size of 3 (Figure 3). When varying the number of convolutional layers we were constrained by the requirements of adequate receptive field and adequate up-sampling to the final output image size of (128, 128). Up-sampling was achieved by using early convolutional layers with a stride of 2, and layers beyond layer 5 have a stride of 1.

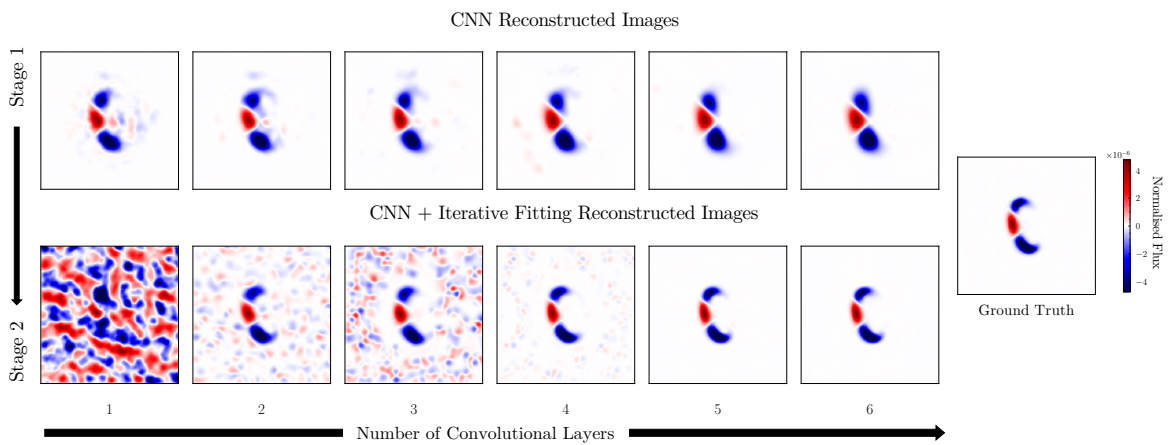


**Fig 5** The influence of changing the number of intermediate convolutional layers on the performance of the a) the pre-trained CNN only (left panel), and b) the entire image reconstruction algorithm - CNN + iterative fitting (right panel), with respect to the validation ‘probe’ data. The influence of increasing layers is evaluated using both the image MSE and SSIM. Adding intermediate convolutional layers improves performance with diminishing returns after  $\sim 6$  layers. In all examples, two final convolutional layers are used to downsize the channel size to 3.

We found that increasing the number of convolutional kernels provided enhanced validation performance with diminishing returns after approximately 200 kernels were added to each layer, and adopted 200 kernels per layer for the remainder of this analysis. We restricted ourselves to even convolutional kernel sizes, to avoid checkerboard noise associated with the combination of odd kernel sizes and stride 2 convolutions,<sup>31</sup> and found that (4 x 4) was the optimal kernel size. The number of convolutional layers had the strongest influence on the accuracy of reconstructed images - see Figure 5. The best final reconstruction accuracy after iterative fitting was found to be

at  $\sim 6$  intermediate layers, with no improvements seen beyond that, suggesting that this number of intermediate layers is sufficient to represent the complexity of this feature space (Figure 5, right).


This trend was also consistently observed with simulated astrophysical data. As convolutional layers are added, the pre-trained CNN images themselves converge to images with modest detail, but more critically, low amounts of noise (Figure 6, top row). As the pre-trained CNN initializes the iterative fitting, this suggests that initialization from a lower resolution but less noisy feature space is more advantageous for the iterative fitting. Using a larger number of layers results in final iteratively fit images that are more accurate and contain minimal noise (Figure 6, bottom row, also see ground truth image on far right). We thus adopt six intermediate convolutional layers - five of stride 2, and one of stride 1, as an optimum architectural configuration for the CNN of PIRATES.



**Fig 6** The influence of increasing the number of intermediate convolutional layers on the images reconstructed at each stage of PIRATES (top row - stage 1, bottom row - stage 2), demonstrated for a simulated astrophysical example of an inclined circumstellar disk. As convolutional layers are added, pre-trained CNN images converge to modest detail and low noise versions of the ground truth, which will subsequently produce the most accurate and low noise final iteratively fit images (compare to ground truth image in bottom right).

### 3.2 Stage 2 - Iterative Fitting Architecture

The iterative fitting is tuned to minimize the residuals between the true observables and those recomputed from the pre-trained CNN images, whilst also ensuring that improvements to the pre-trained CNN images are self-consistent and do not introduce significant spurious noise or structure. The improvement which iterative fitting yields to the observables is used explicitly as the loss function (MSE of observables), however we also monitored the image MSE and the SSIM as metrics to quantify both pixel wise and structural improvements to the images. The iterative fitting process can be tuned by controlling how many CNN layers are iteratively fit - as opposed to having their weights ‘frozen’ to pre-trained values for the duration of the iterative fitting. To assess this, we investigated the impact of freezing layers in both directions (shallow-deep). We found that iterative fitting has the best performance on validation data when all layers are iteratively fit and none are frozen - this optimum is measurable in the improvement to observables, as well as both image based MSE and SSIM scores (Figure 7).



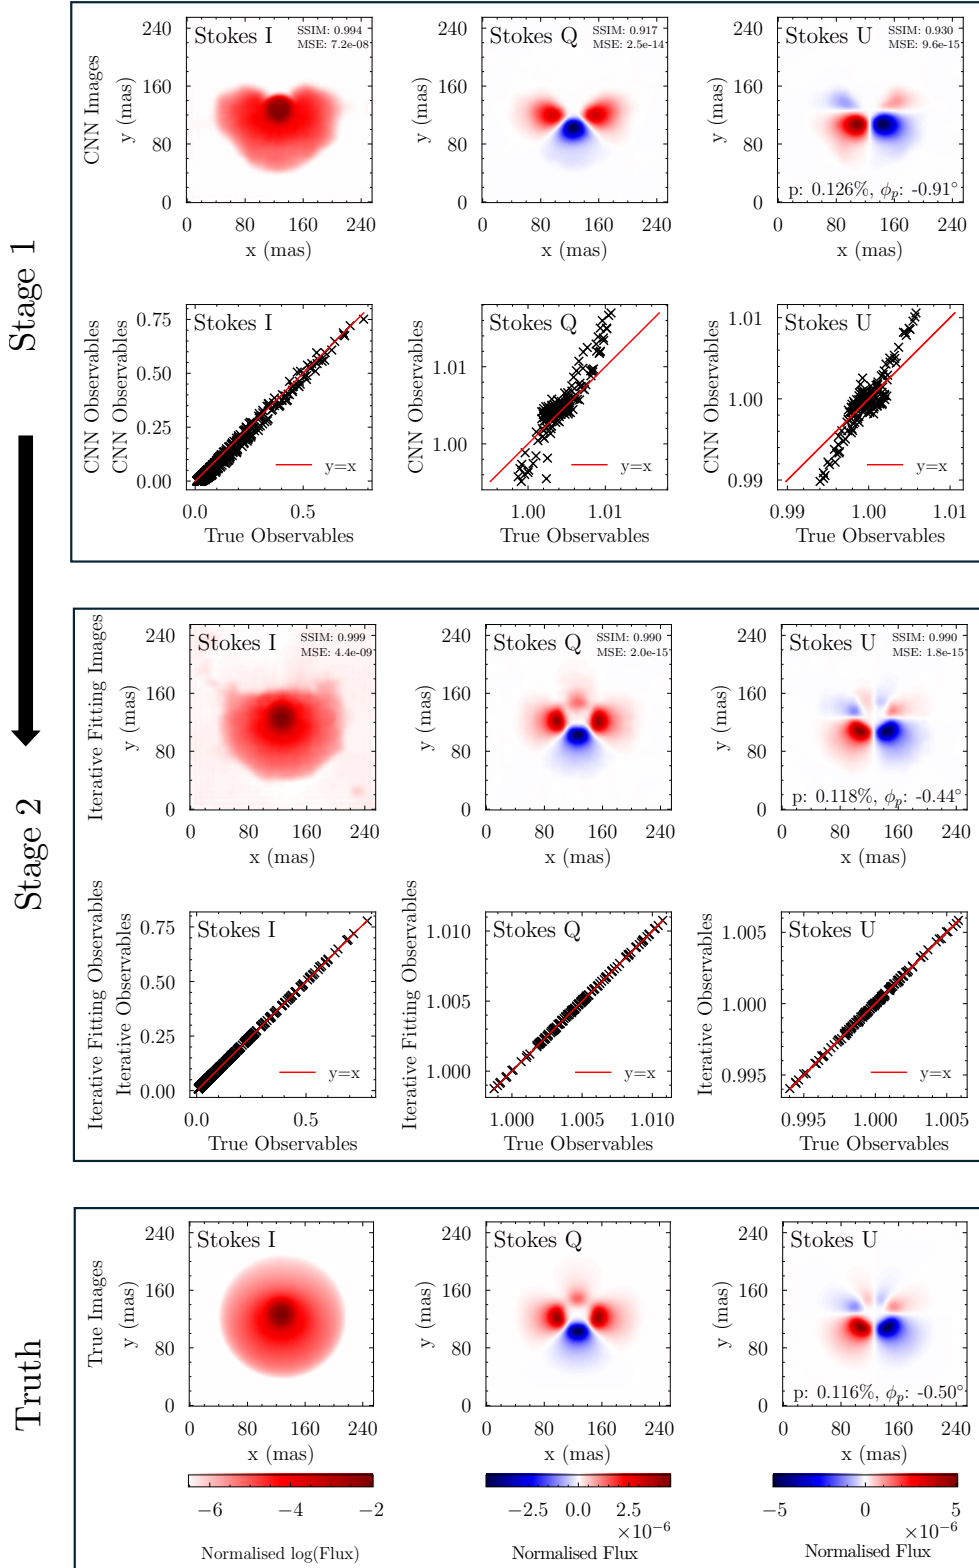
**Fig 7** The influence of freezing convolutional layers during iterative fitting. Here the results are plotted for freezing layers starting from the shallowest layers to the deepest layers, though the results are the same for deep-shallow. The iterative fitting performs optimally with the maximum number of degrees of freedom - where all layers are iteratively fit, and none are frozen to pre-trained values.

#### 4 Results and Discussion Part II - Performance of the two-stage framework on simulated astrophysical examples

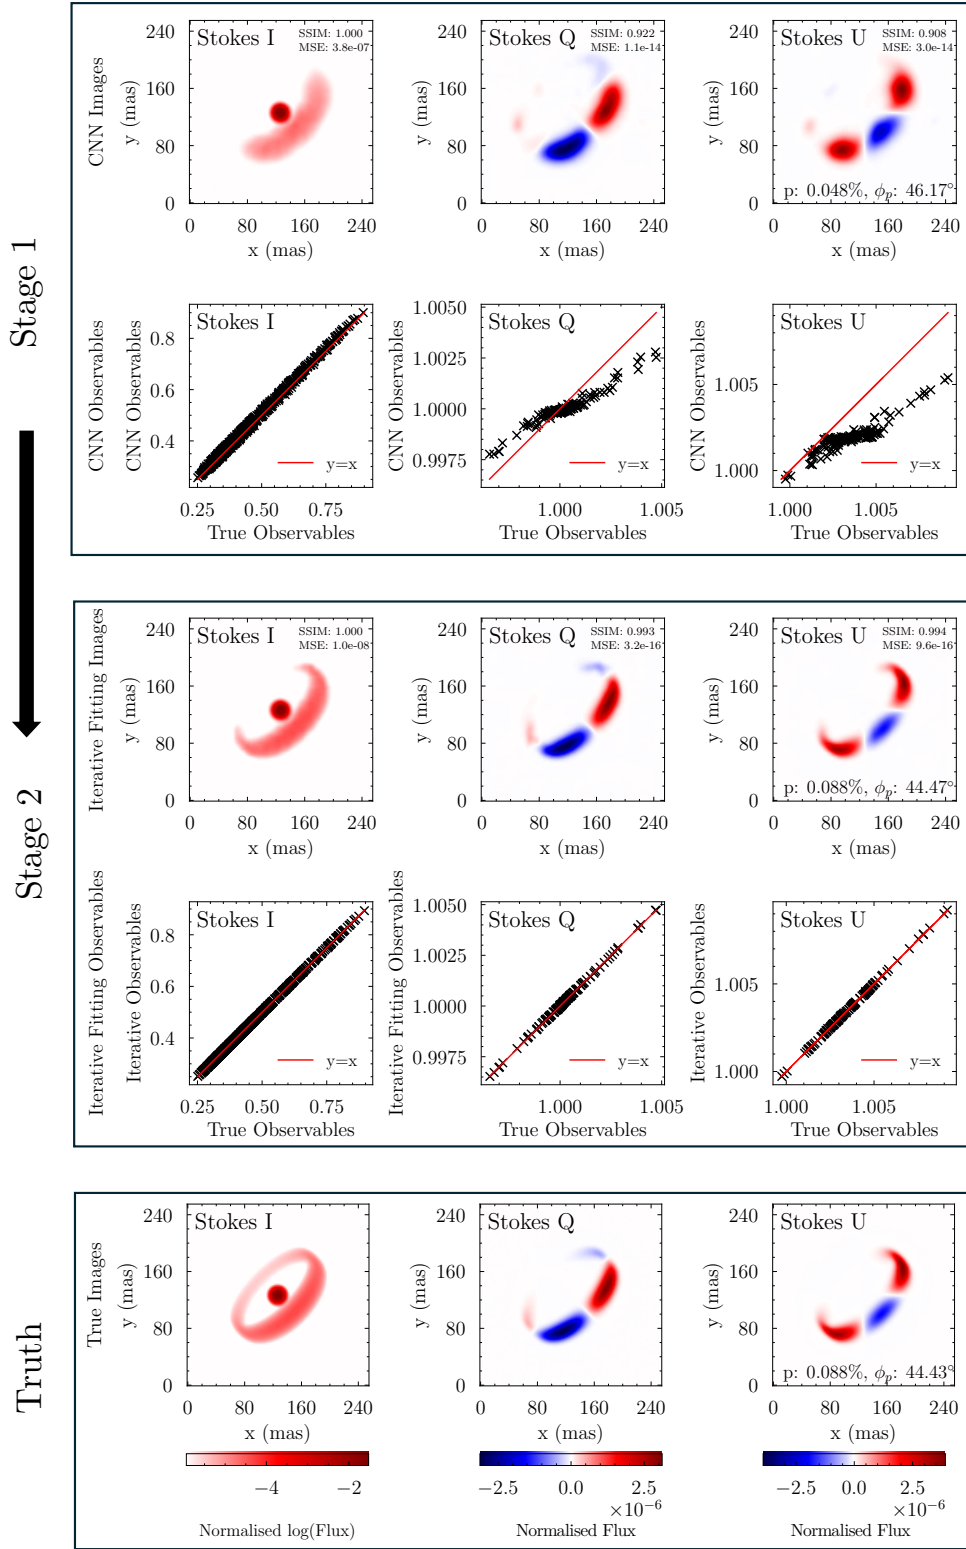
Here we demonstrate the performance of PIRATES using a our simulated astrophysical dataset of disks, spirals and complex circumstellar environments, generated using the MCFOST radiative transfer code. These examples encode the astrophysical complexity we anticipate underlie VAMPIRES NRM (and other polarized-interferometric) datasets - optical depth, non-azimuthal scattering and significant geometric asymmetry. In making these examples, we drew inspiration from real observations of the circumstellar environments of evolved stars like Betelgeuse<sup>32</sup> and  $\mu$  Cephei,<sup>33</sup> debris disks like HR 4796,<sup>34,35</sup> and spiral armed disks like HD 135344B,<sup>35</sup> HD 142527<sup>35</sup> and MWC 758.<sup>35</sup> Figures 8 - 13 demonstrate the image reconstruction performance of PIRATES on a sub-set of the simulated astrophysical dataset, and include the predictions from the pre-trained CNN alone (the ‘initial guess’) (rows 1 and 2) and the final predictions after iterative fitting (rows 3 and 4) - for both the images and interferometric observables.

It can be seen that the predictions from the pre-trained CNN alone give surprisingly accurate results, with broad features of the true image being reproduced (Figures 8 - 13, rows 1). Finer detail is missing, and the appearance of overly connected structure or excess material in the image centers occurs. The observables resampled from the CNN images are relatively poor fits to those sampled from the ground truth images (Figures 8 - 13, rows 2). As iterative fitting proceeds, observables are typically optimized to well below realistic observational levels of uncertainty ( $O(1\%)$ ) (Figures 8 - 13, rows 4); simultaneously the reconstructed images are improved on both a pixel wise and structural level (Figures 8 - 13, rows 3), indicated by improved MSE and SSIM scores. As is visually apparent in Figures 8 - 13, all astrophysical examples see clear improvement in both image accuracy and observables accuracy as a result of the iterative fitting. The final reconstructed images are high fidelity and noise free, accurately resembling the ground truths.

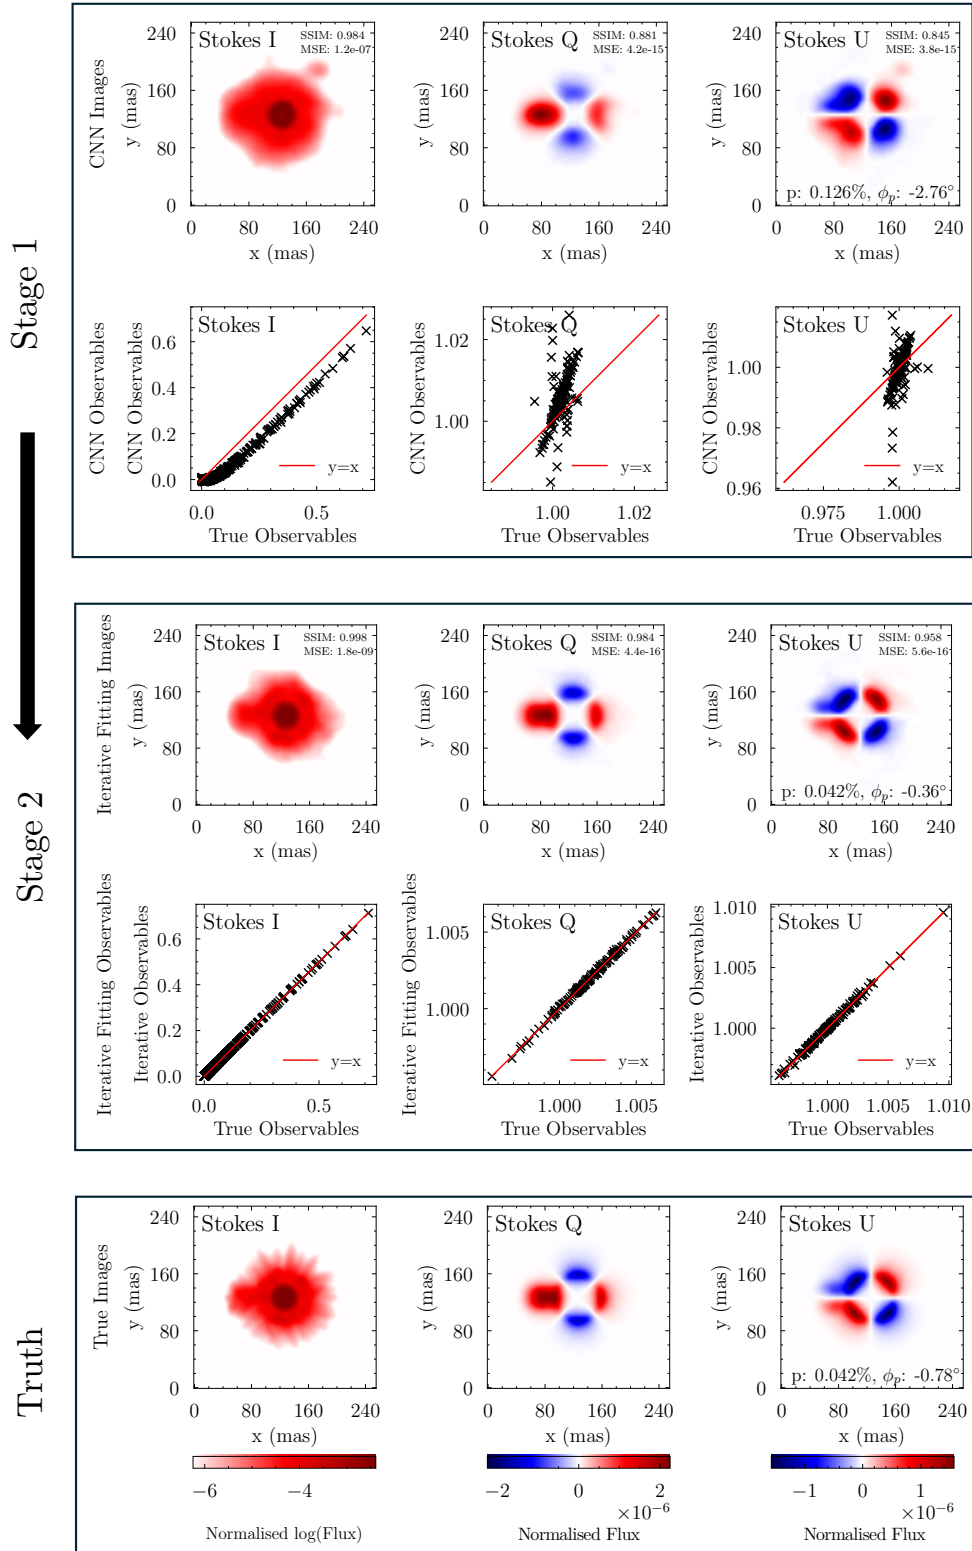
There are two consistent styles of refinement which iterative fitting makes to the reconstructed images. Firstly, iterative fitting enhances the resolution of the CNN pre-trained images to match the information contained in the interferometric observables. This effect is evident across all simulated astrophysical examples - disk profiles are narrowed (Figures 9, 8, 12), the inner cavities of



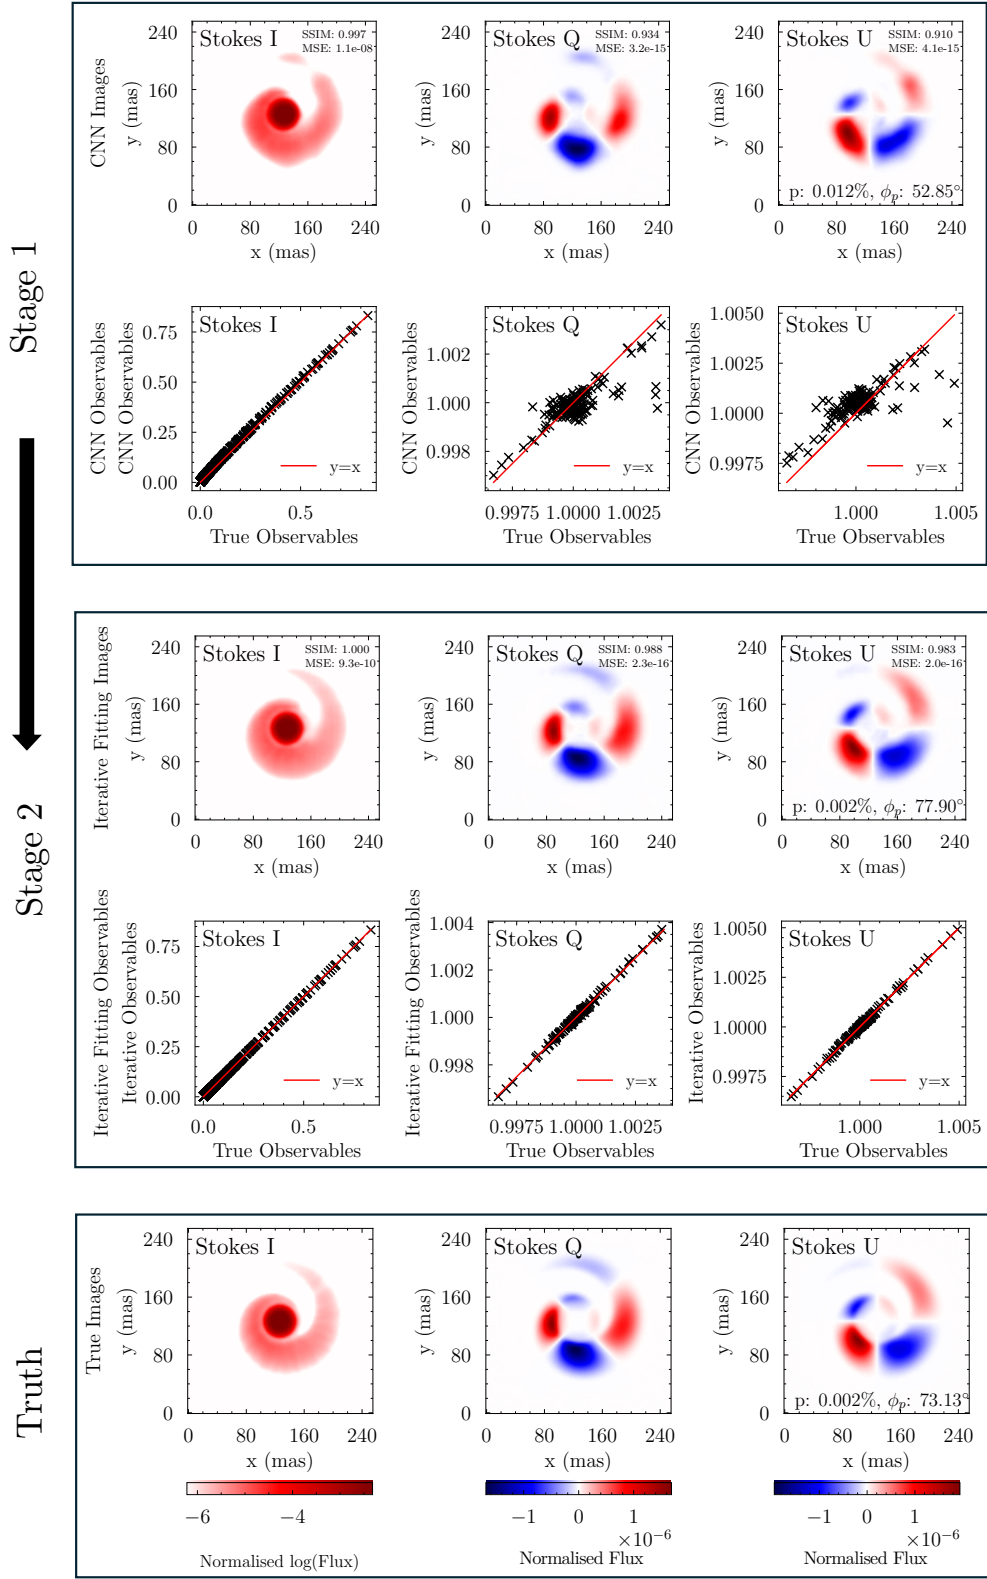
**Fig 8** PIRATES performance on an inclined model circumstellar disk, with a non-azimuthal scattering feature on the top edge of the disk.



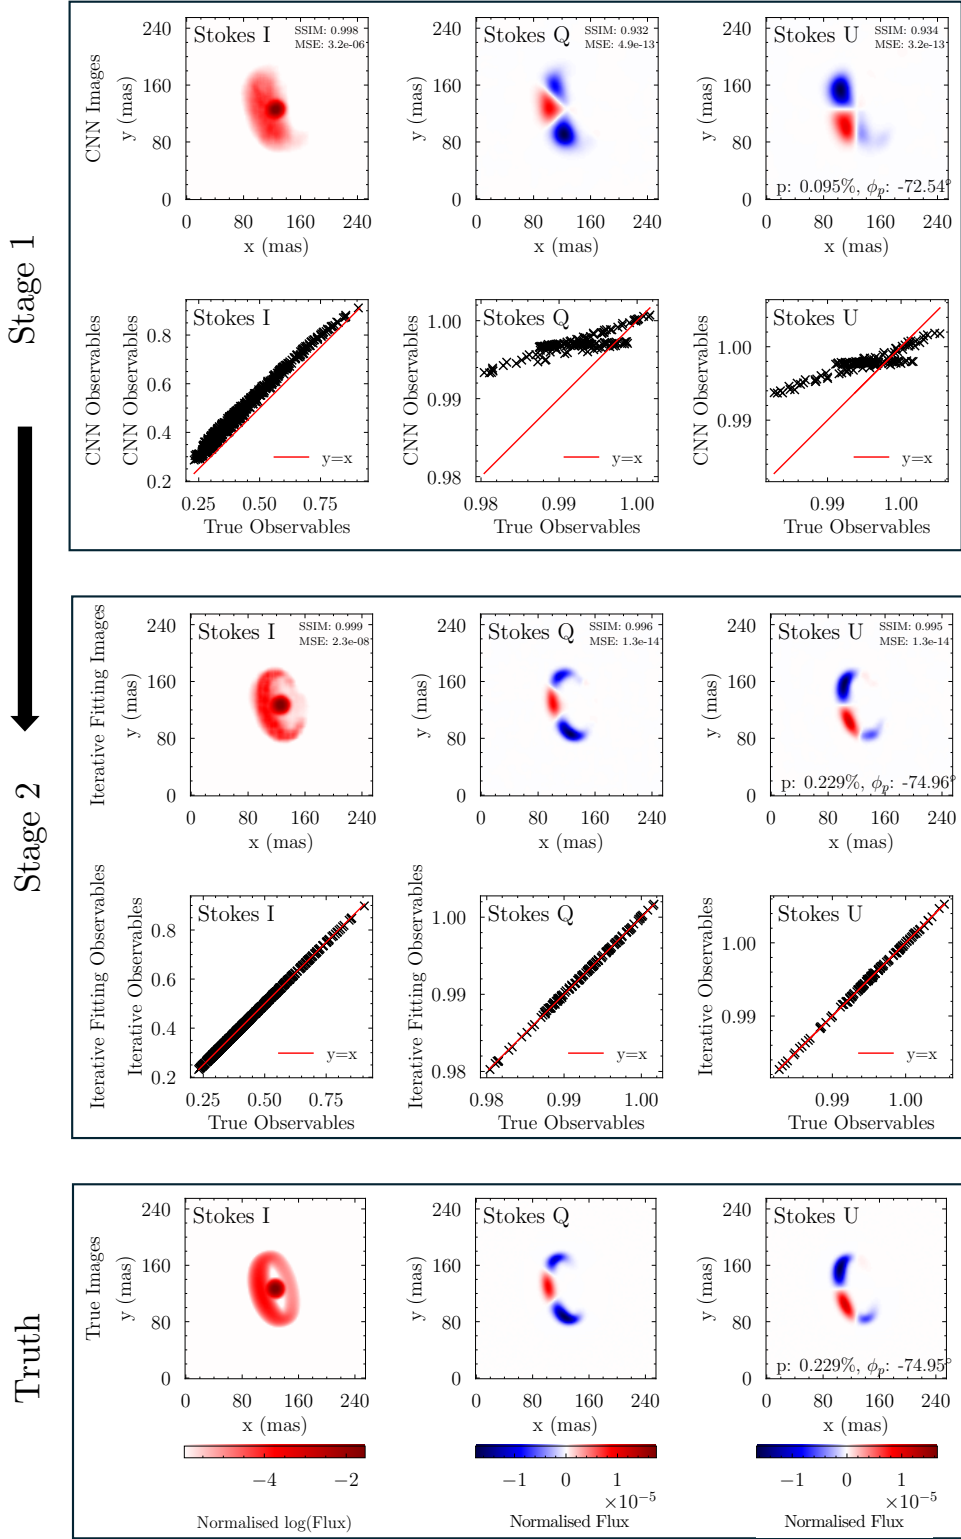
**Fig 9** PIRATES performance on an inclined model circumstellar disk. The effects of large Mie grains are evident - the back of the disk is shadowed and the image is dominated by strong forward scattering within the disk's front.



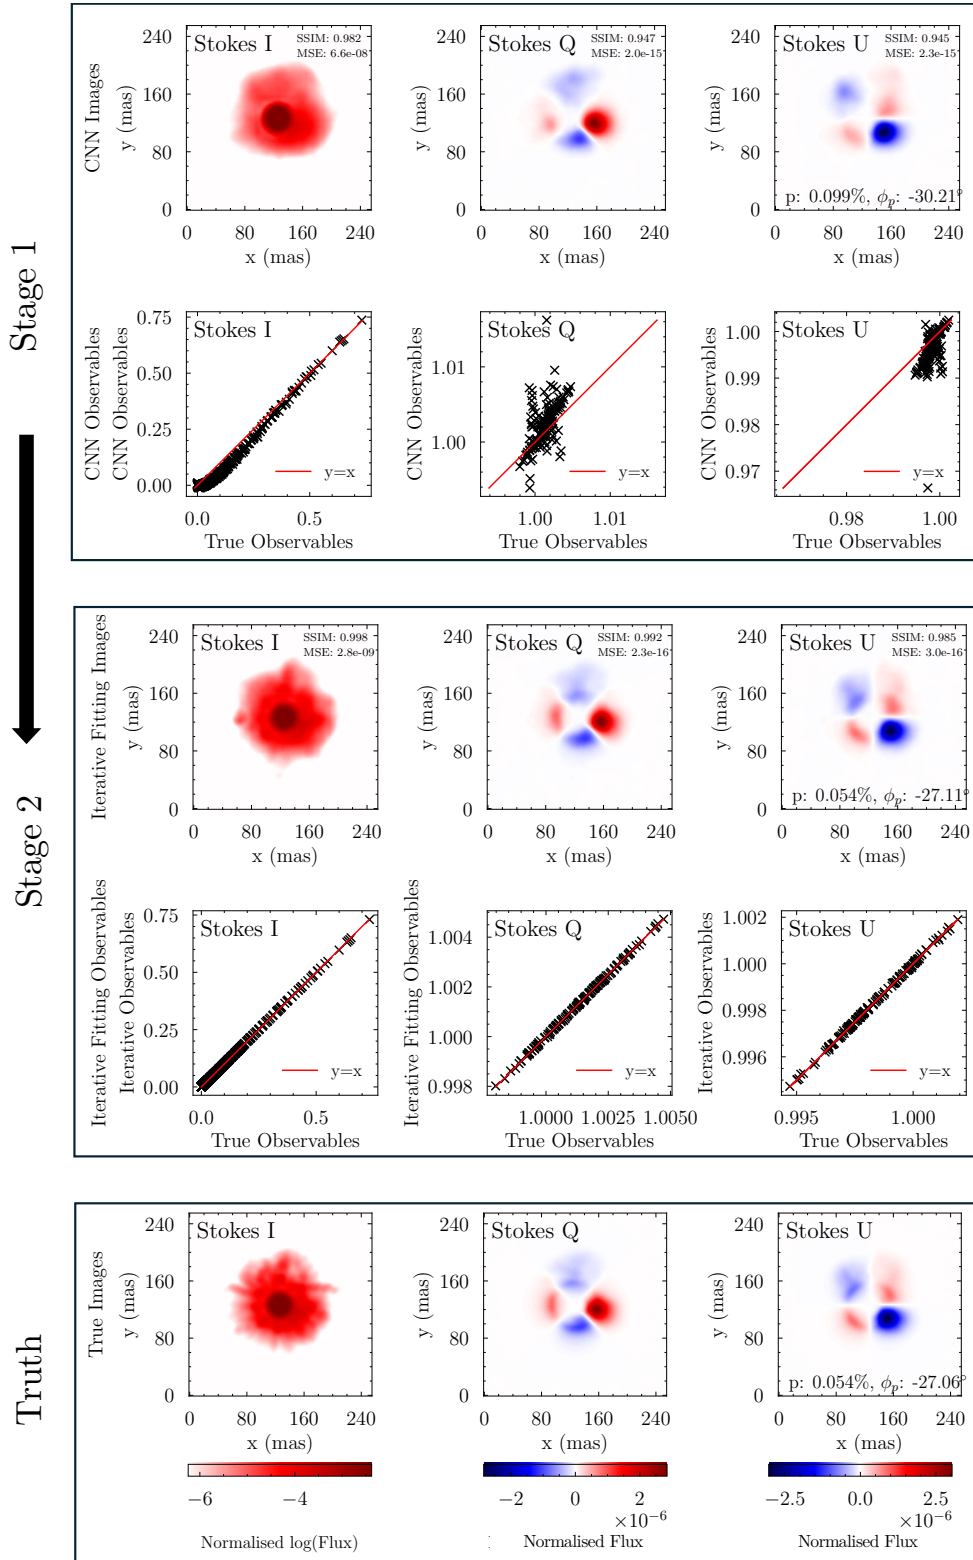
**Fig 10** PIRATES performance on a model clumpy, inhomogeneous circumstellar environment typical of evolved stars.



**Fig 11** PIRATES performance on a model dust spiral.



**Fig 12** PIRATES performance on a model circumstellar disk.



**Fig 13** PIRATES performance on a model dusty circumstellar environment typical of evolved stars.

clumpy circumstellar dust shells are widened (Figure 10, 13) to reveal inner dust radii. The second effect which iterative fitting has on the reconstructed images is that it is capable of reconstructing material that was originally missing from the pre-trained CNN reconstructions, including - fine non-azimuthal lobes of circumstellar disks (Figure 8) and the full length of a spiral tail (Figure 11). Iterative fitting is also capable of removing structure that was incorrectly present within the pre-trained CNN images; a faint circle of dust in the center right of Figure 10, some structure in the center left of Figure 9. These image improvements are quantified by improved MSE and SSIM scores for each image.

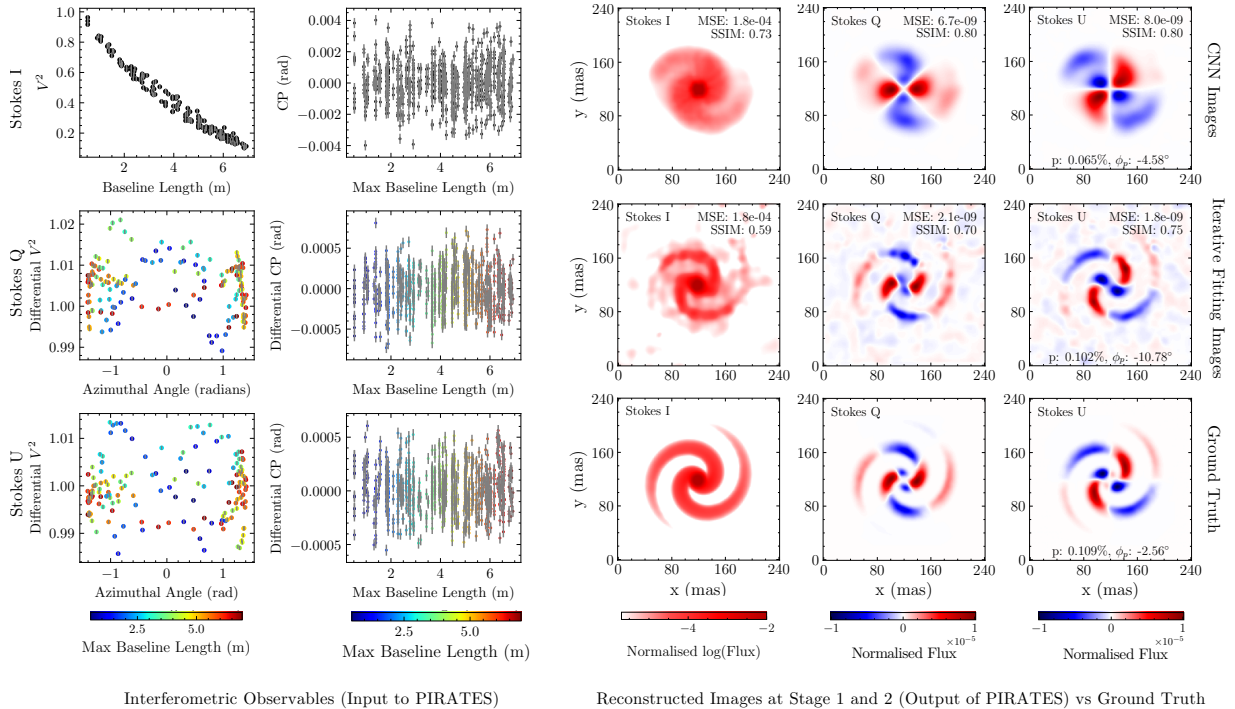
Refinements which the iterative fitting makes to images are mutual and consistent across all three Stokes images - both locally and globally, demonstrated not only by the mutual improvement in polarized structure within the images, but also by measuring the total polarization  $p = \frac{\sqrt{Q^2+U^2}}{I}$  and polarization angle  $p_\phi = \frac{1}{2} \arctan \frac{U}{Q}$ , as annotated in the right column of each example (Figures 8 - 13). Improvements to the images occur without the introduction of spurious noise or structure, as the pre-trained CNN is optimized as a modest detail but low noise initialization and feature-space for the iterative fitting.

## 5 Performance with random noise

We also demonstrate that PIRATES is compatible with the presence of realistic amounts of random error. To investigate the impact of noise, we inject the interferometric observables of our simulated astrophysical examples with error that matches the signal to noise ratio previously seen in on sky NRM datasets (Figure 14 columns 1 and 2). With no adaptation of the PIRATES network, we find that stage 1 (the CNN) still does a good job at providing moderate resolution and low noise initializations for the iterative fitting (Figure 14, row 1, columns 3-5), as it does for noise-free examples (Figures 8 - 13). However, when stage 2 is deployed and images are iteratively fitted with respect to their (now noisy) observables, significant amounts of noise are now introduced into the images, as may be anticipated (Figure 14, row 2, columns 3-5). Thus, to perform optimally on noisy real datasets, PIRATES requires adaptation to reduce the noise and spurious structure that iterative fitting may inject into the reconstructed images as a consequence of noise.

In the presence of noise, the loss function of the iterative fitting is updated to a chi-squared  $\chi^2$  statistic, allowing the iterative fitting to more heavily weight the reconstruction of features represented by higher signal to noise observables. In the spirit of the ‘implicit regularization’ functionality of PIRATES, rather than regularizing the reconstructed images in pixel space like traditional image reconstruction codes, we regularize the CNN’s latent feature space - restricting its functional complexity and encouraging lower complexity solutions which minimize noise. There are many ways to implement latent feature space regularization - here we explore one promising solution, but do not assert this is the only solution.

To provide the additional regularization required to prevent stage 2 (iterative fitting) from introducing noise to reconstructed images, we use an Elastic Weights Consolidation penalty during the iterative fitting stage. Elastic Weights Consolidation (EWC) is useful where retaining historical information benefits future predictive accuracy.<sup>36</sup> It is employed by introducing a continuous penalty for the gradual ‘forgetting’ of network weights which prior information indicates will improve the accuracy of a sequential prediction. In the context of PIRATES, we introduce a penalty during the



**Fig 14** The performance of PIRATES with signal to noise consistent with recent VAMPIRES NRM data, with no algorithmic treatment to constrain the influence of noise. Visibilities and closure phases with injected noise and corresponding error bars are plotted in columns 1 and 2. Stage 1 (CNN) does a good job of a moderate resolution and low noise reconstruction of the ground truth (top row, columns 3-5), however, significant amounts of noise are introduced into the images during the iterative fitting (middle row, columns 3-5). The ground truth images are displayed in the bottom row, columns 3-5.

iterative fitting if iteratively fit weights are modified too much from their pre-trained CNN values. The choice of EWC as regularization is guided by intuition that the CNN’s predictions are already reasonably good, and more crucially, low noise representations of the ground truth. With an EWC penalty, changes to the pre-trained CNN weights are only made where they strongly improve the iteratively fit image’s fit to observables, discouraging non-critical changes to weights that may introduce spurious noise or structure. The elegance of EWC is that when it is applied to an extreme level, it simply converges to the pre-trained CNN prediction. This gives a clear indication where over-regularization has occurred, making it robust to tune.

To create a penalty using EWC, we are first required to quantify the ‘importance’ of each pre-trained CNN parameter,  $\theta_j$ . We then use this ‘importance’ to weight the contributions of each network parameter to the EWC penalty. To quantify the importance of each network parameter, we derive a heuristic for the diagonal of the empirical Fisher Information matrix as is standard practice when defining an EWC penalty for machine learning.<sup>36</sup>

To do so, we first compute an average loss  $\mathcal{L}_b$  over a batch of N input-output pairs  $(x_i, y_i)$  (Equation 1), where  $\ell$  is the pre-trained CNN loss function and  $f_\theta$  denotes the CNN, constructed of network parameters  $\theta_j$  (weights, biases).

$$\mathcal{L}_b = \frac{1}{N} \sum_{i=1}^N \ell(y_i, f_\theta(x_i)) \quad (1)$$

We then use backpropagation to compute the gradient  $g_j^{(b)}$  of  $\mathcal{L}_b$  with respect to each model parameter  $\theta_j$  (Equation 2).

$$g_j^{(b)} = \frac{\partial \mathcal{L}_b}{\partial \theta_j} \quad (2)$$

We use the average squared gradients over  $M$  batches of training data,  $F_j$ , as a heuristic for the Fisher Information of parameter  $j$ , (Equation 3):

$$F_j = \frac{1}{M} \sum_{b=1}^M \left( g_j^{(b)} \right)^2 = \frac{1}{M} \sum_{b=1}^M \left( \frac{\partial \mathcal{L}_b}{\partial \theta_j} \right)^2 \quad (3)$$

$F_j$  is used to weight the squared difference between current  $\theta_j$  (iterative fitting) and previous  $\theta_j^*$  (pre-trained CNN) network parameters (Equation 4), penalizing changes to weights which were more important to the CNN’s original reconstructions.

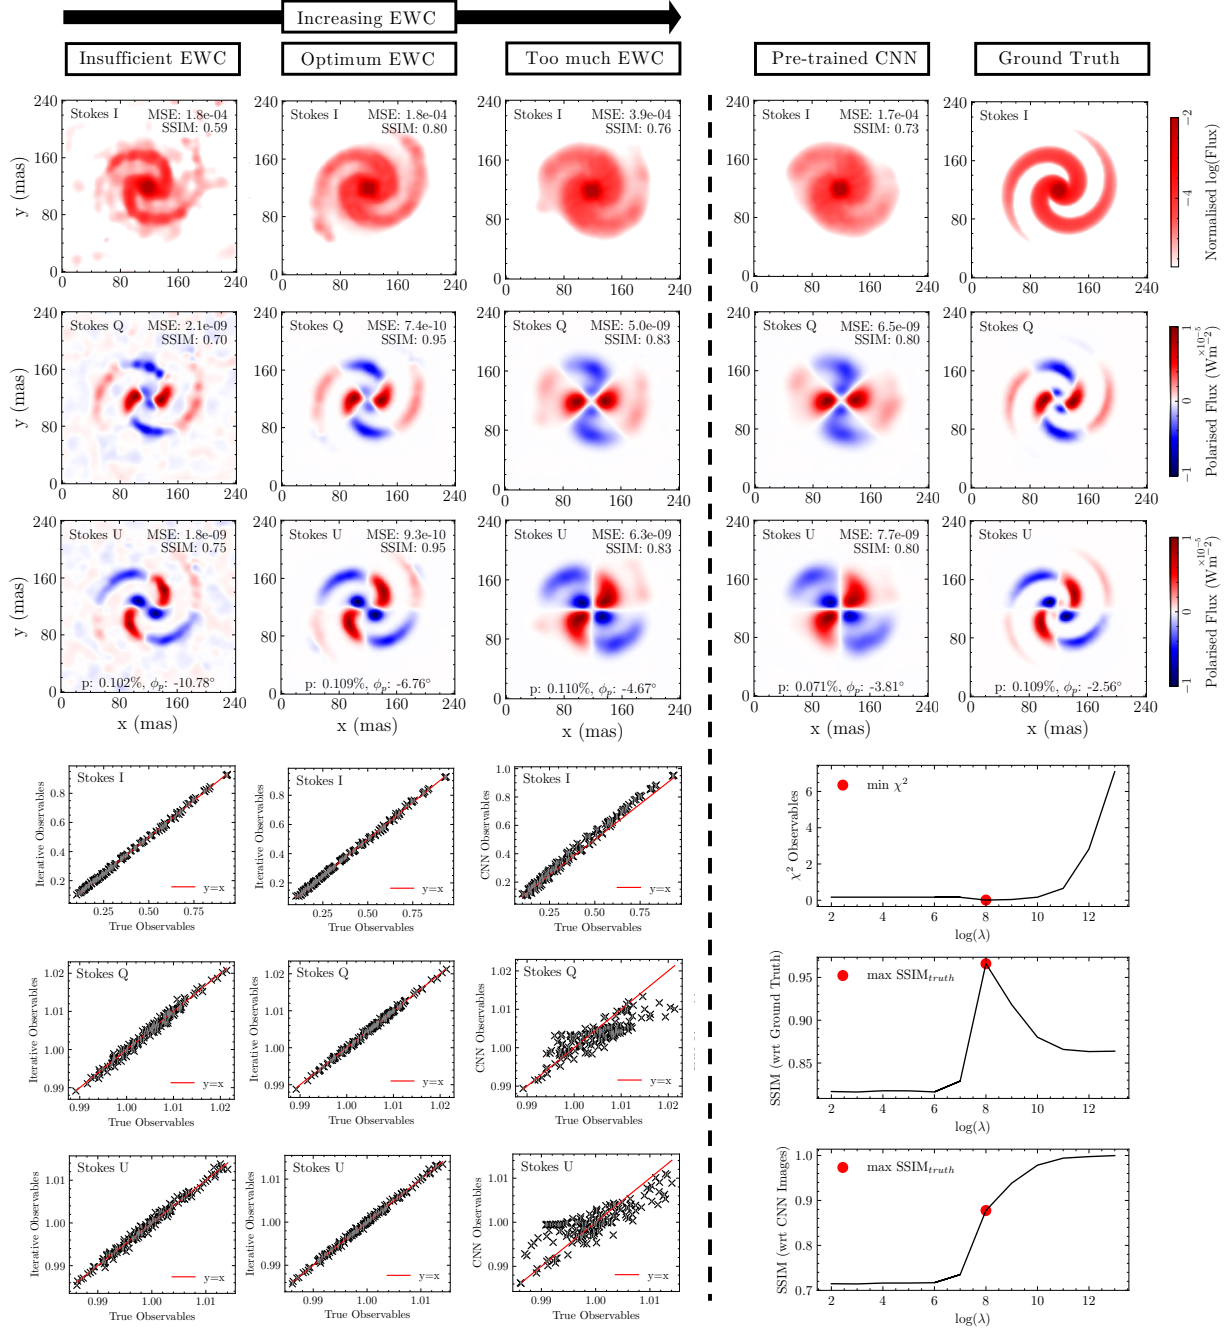
$$EWC = \lambda \sum_i F_j (\theta_j - \theta_j^*)^2 \quad (4)$$

The strength of the EWC penalty is tuned via hyperparameter  $\lambda$ , the final loss function of the iterative fitting,  $L$ , is given by Equation 5.<sup>36</sup>

$$L = \chi^2 + \lambda \sum_j F_j (\theta_j - \theta_j^*)^2 \quad (5)$$

The effect of increasing the strength of EWC regularization via  $\lambda$  is visually striking. The results of iterative fitting with insufficient  $\lambda$  are noisy images (Figure 15, col 1). When  $\lambda$  is of optimal strength, the iterative fitting converges to images with minimal noise (Figure 15, col 2), and the fit to interferometric observables is simultaneously optimized (Figure 15, col 5, row 4). The SSIM scores of iteratively fit images are improved, as are the reproduction of polarization statistics (inset in Stokes U panels). As expected, when EWC regularization is applied to an extreme level, the iteratively fit images do not vary much from the pre-trained CNN images (Figure 15, col 3), nor is there improvement in the fit to observables from the pre-trained CNN initialization.

The influence of the EWC penalty on image reconstruction is easiest to assess by monitoring the SSIM between the pre-trained CNN and iteratively fit images, as a function of  $\lambda$  (Figure 15, col 4-5, row 6). Small values of  $\lambda$  result in noisy iterative fitting images, which are most structurally dis-similar to the CNN images. As  $\lambda$  is increased, the SSIM score increases as the algorithm is encouraged to fit the observables to real structure and not to noise, now maintaining network parameters which were highly informative in the pre-trained CNN. However, eventually  $\lambda$  becomes so large that the iteratively fit images barely deviate from the CNN images (and hence the SSIM converges to  $\sim 1$ ).



**Fig 15** The impact of increasing EWC strength via  $\lambda$  (Equation 5). With low  $\lambda$  (col 1), significant noise is present in reconstructed images. As  $\lambda$  is increased (col 2), the noise present in images is substantially reduced and the SSIM of the reconstructed images is improved relative to the ground truth (col 5). When  $\lambda$  is applied to an extreme level (col 3), reconstructions do not deviate from the CNN pre-trained images (col 4).  $\lambda$  may be tuned by monitoring: the SSIM between iteratively fit images and the pre-trained CNN images (col 4-5, row 6), and the fit to interferometric observables (col 4-5 row 4). We found that the optimal value of  $\lambda$  (where the SSIM between the iteratively fit images and the ground truth is maximized) (col 4-5, row 6) was correlated with an optimized fit to observables (col 4-5, row 4).

To locate an optimum value of  $\lambda$ , we found that a useful heuristic was to monitor the fit to interferometric observables. For our simulated astrophysical examples, with known ground truths, we know a-priori the optimal value of  $\lambda$ , as we can measure the SSIM between the iteratively fit images and the ground truth (Figure 15, col 4-5, row 5). We found that the value of  $\lambda$  where the goodness of fit of interferometric observables was optimized was correlated with the value of  $\lambda$  that maximized the SSIM of the iteratively fit images and their ground truths (Figure 15, col 4-5, row 4 and 5). Initially, increasing  $\lambda$  improves the fit to observables (Figure 15, col 4-5, row 4). As  $\lambda$  is increased, the sub-space of possible image solutions shrinks - as noisy solutions are excluded by regularization. As such, we interpret the improvement which increasing  $\lambda$  yields to the observables as a consequence of more stable and uniquely defined convergence, towards solutions that are free of noise that may perturb the fit to observables. In summary, in real applications with no known ground truth, increasing  $\lambda$  until the goodness-of-fit of the interferometric observables begins to exceed the levels determined by observational error is recommended as a useful heuristic, along with simple inspection of the images.<sup>37,38</sup>

## 6 Conclusions

We have presented PIRATES, the first image reconstruction algorithm for optical polarimetric interferometry, which uniquely performs image reconstruction in a learned, physics-informed feature space rather than in pixel space. PIRATES' dual structure - a pre-trained CNN and an iterative fitting mechanism - allow the algorithm to parsimoniously reconstruct high fidelity polarized images which accurately reproduce polarized, interferometric observables. We have demonstrated that PIRATES accurately reconstructs polarized images of a diverse range of astrophysically realistic radiative transfer models, including disks, spirals and clumpy circumstellar environments.

We found that optimization of PIRATES dual stage structure required a CNN which constructs modest detail but low noise versions of the ground truth, providing a physics-informed feature space to be used as initialization for iterative optimization in the second stage. The iterative fitting stage is capable of improving the reconstructed images consistently across all polarimetric states (Stokes I, Q and U), without introducing noise or spurious structures, simultaneously improving the fit to observables to within realistic observational uncertainty. PIRATES implicitly regularizes reconstructed images via a physically meaningful prior that the CNN learns from pre-training data (produced via radiative transfer models), and the structurally derived inductive bias that forces image updates made by iterative fitting to be statistically consistent across the image domain. In the presence of realistic amounts of noise in interferometric observables, we found that introducing Elastic Weights Consolidation as a latent feature space regularizer was able to prevent the introduction of noise into images during the iterative fitting, and that a single parameter to adjust the strength of regularization was sufficient and robust in tuning the algorithm.

In future work we intend to upgrade PIRATES to streamline flexible polarimetric coordinates and varying (u,v) coverage, by adding an auto-encoder to the beginning of our existing algorithm. Additionally, our treatment of noise only includes random error - strategies for treating systematic sources of error that are common to polarized optical systems are in development. The performance of our algorithm demonstrates the power of dual-structure machine learning algorithms, and suggests they may more broadly be a useful technique for contexts with complex regularization constraints, beyond just those of polarized interferometry.

## 7 Data and Code availability

Code for PIRATES is available on GitHub: [github.com/LucindaLilley/PIRATES](https://github.com/LucindaLilley/PIRATES).

## 8 Acknowledgments

The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community, and are most fortunate to have the opportunity to conduct observations from this mountain.

Based on data collected at Subaru Telescope, which is operated by the National Astronomical Observatory of Japan. The development of SCEXAO is supported by the Japan Society for the Promotion of Science (Grant-in-Aid for Research #23340051, #26220704, #23103002, #19H00703, #19H00695 and #21H04998), the Subaru Telescope, the National Astronomical Observatory of Japan, the Astrobiology Center of the National Institutes of Natural Sciences, Japan, the Mt Cuba Foundation and the Heising-Simons Foundation.

Majority funding for the development of the VAMPIRES instrument came from the Australian Research Council (FT100100953; DP140104065).

LL wishes to acknowledge Adam Taras, Benjamin Pope, Louis Desdoigts and Max Charles for their helpful discussions and advice over figures.

An earlier version of this work was submitted as a Conference Proceedings at SPIE Astronomical Telescopes + Instrumentation, 2024, in Yokohama, Japan.<sup>39</sup>

## 9 Disclosure Statement

The authors declare there are no financial interests, commercial affiliations, or other potential conflicts of interest that have influenced the objectivity of this research or the writing of this paper.

## References

- 1 E. Thiebaut and J.-F. Giovannelli, “Image reconstruction in optical interferometry,” *IEEE Signal Processing Magazine* **27**, pp. 97–109, 1 2010.
- 2 D. Buscher, “Direct Maximum-Entropy Image Reconstruction from the Bispectrum,” *Symposium - International Astronomical Union* **158**, pp. 91–93, 7 1994.
- 3 E. Thiebaut, “MIRA: an effective imaging algorithm for optical interferometry,” *Proceedings of the SPIE* **7013**, p. 70131I, 7 2008.
- 4 F. Baron, J. D. Monnier, and B. Kloppenborg, “A novel image reconstruction software for optical/infrared interferometry,” *The Proceedings of SPIE* **7734**, p. 77342I, 7 2010.
- 5 E. Thiebaut and J. Young, “Principles of image reconstruction in optical interferometry: tutorial,” *Journal of the Optical Society of America A* **34**, p. 904, 5 2017.
- 6 E. Thiebaut and J.-F. Giovannelli, “Image reconstruction in optical interferometry,” *IEEE Signal Processing Magazine* **27**, pp. 97–109, 1 2010.
- 7 F. Baron, “Interferometric image reconstruction: techniques, results, and future direction,” *Proceedings of the SPIE* **9907**, p. 99071C, 8 2016.

- 8 B. R. Setterholm, J. D. Monnier, J.-B. Le Bouquin, N. Anugu, A. Labdon, J. Ennis, K. J. Johnson, S. Kraus, and T. A. ten Brummelaar, “MIRC-X polarinterferometry at CHARA,” in *Optical and Infrared Interferometry and Imaging VII*, A. Mérand, S. Sallum, and P. G. Tuthill, eds., p. 24, SPIE, 12 2020.
- 9 M. Lucas, B. Norris, O. Guyon, M. Bottom, V. Deo, S. Vievard, J. Lozi, K. Ahn, J. Ashcraft, T. Currie, D. Doelman, T. Kudo, L. Leboulleux, L. Lilley, M. Millar-Blanchaer, B. Safonov, P. Tuthill, T. Uyama, A. Walk, and M. Zhang, “Visible-light High-contrast Imaging and Polarimetry with SCExAO/VAMPIRES,” *Publications of the Astronomical Society of the Pacific* **136**, p. 114504, 11 2024.
- 10 J. R. Valenzuela, *Polarimetric Image Reconstruction Algorithms*. PhD thesis, University of Michigan, 2010.
- 11 A. Mus, H. Müller, and A. Lobanov, “Swarm intelligence for full Stokes dynamic imaging reconstruction of interferometric data,” 5 2024.
- 12 C. Pinte, F. Ménard, G. Duchêne, and P. Bastien, “Monte Carlo radiative transfer in protoplanetary disks,” *Astronomy & Astrophysics* **459**, pp. 797–804, 12 2006.
- 13 C. Pinte, T. J. Harries, M. Min, A. M. Watson, C. P. Dullemond, P. Woitke, F. Ménard, and M. C. Durán-Rojas, “Benchmark problems for continuum radiative transfer,” *Astronomy & Astrophysics* **498**, pp. 967–980, 5 2009.
- 14 Z. Wang and L. Wu, “Theoretical Analysis of Inductive Biases in Deep Convolutional Networks,” *arXiv* , 5 2023.
- 15 D. Ulyanov, V. Andrea, and V. Lempitsky, “Deep Image Prior,” *International Journal of Computer Vision* **128**, pp. 1867–1888, 11 2020.
- 16 B. Norris, G. Schworer, P. Tuthill, N. Jovanovic, O. Guyon, P. Stewart, and F. Martinache, “The VAMPIRES instrument: imaging the innermost regions of protoplanetary discs with polarimetric interferometry,” *Monthly Notices of the Royal Astronomical Society* **447**, pp. 2894–2906, 1 2015.
- 17 R. Lenzen, M. Hartung, W. Brandner, G. Finger, N. N. Hubin, F. Lacombe, A.-M. Lagrange, M. D. Lehnert, A. F. M. Moorwood, and D. Mouillet, “NAOS-CONICA first on sky results in a variety of observing modes,” in *SPIE*, M. Iye and A. F. M. Moorwood, eds., p. 944, 3 2003.
- 18 P. Tuthill, S. Lacour, P. Amico, M. Ireland, B. Norris, P. Stewart, T. Evans, A. Kraus, C. Lidman, E. Pompei, and N. Kornweibel, “Sparse aperture masking ({SAM}) at {NAOS}/{CONICA} on the {VLT},” in *Ground-based and Airborne Instrumentation for Astronomy {III}*, I. S. McLean, S. K. Ramsay, and H. Takami, eds., SPIE, 7 2010.
- 19 J. Davis, W. J. Tango, A. J. Booth, T. A. t. Brummelaar, R. A. Minard, and S. M. Owens, “The Sydney University Stellar Interferometer – I. The instrument,” *Monthly Notices of the Royal Astronomical Society* **303**, pp. 773–782, 3 1999.
- 20 F. Snik, *Astronomical Polarimetry: Polarized Views of Stars and Planets*.
- 21 B. R. M. Norris, *Secrets in Stellar Halos : Imaging Against the Glare*. PhD thesis, The University of Sydney, 2015.
- 22 J. P. Hamaker and J. D. Bregman, “Understanding radio polarimetry. III. Interpreting the IAU/IEEE definitions of the Stokes parameters,” *Astronomy and Astrophysics Supplement Series* **117**, pp. 161–165, 5 1996.

- 23 E. Hecht, *OPTICS*, Pearson Education Limited, 5 ed., 2017.
- 24 M. J. Ireland, P. G. Tuthill, J. Davis, and W. Tango, “Dust scattering in the Miras R Car and RR Sco resolved by optical interferometric polarimetry,” *Monthly Notices of the Royal Astronomical Society* **361**(1), pp. 337–344, 2005.
- 25 B. R. Norris, P. G. Tuthill, M. J. Ireland, S. Lacour, A. A. Zijlstra, F. Lykou, T. M. Evans, P. Stewart, and T. R. Bedding, “A close halo of large transparent grains around extreme red giant stars,” *Nature* **484**(7393), pp. 220–222, 2012.
- 26 X. Haubois, B. Norris, P. G. Tuthill, C. Pinte, P. Kervella, J. H. Girard, N. M. Kostogryz, S. V. Berdyugina, G. Perrin, S. Lacour, A. Chiavassa, and S. T. Ridgway, “The inner dust shell of Betelgeuse detected by polarimetric Aperture-Masking interferometry,” *arXiv* (33258), pp. 1–9, 2019.
- 27 A. Doicu, M. I. Mishchenko, and T. Trautmann, “Electromagnetic scattering by discrete random media illuminated by a Gaussian beam I: Derivation of the radiative transfer equation,” *Journal of Quantitative Spectroscopy and Radiative Transfer* **256**, p. 107301, 11 2020.
- 28 L. Lilley, “In Prep,” 2025.
- 29 Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng, *TensorFlow: A System for Large-Scale Machine Learning*, 2016.
- 30 F. Chollet *et al.*, “Keras,” 2015.
- 31 M. Tanaka, “A universal detector of CNN-generated images using properties of checkerboard artifacts in the frequency domain,” 2021.
- 32 E. Cannon, M. Montargès, A. de Koter, A. Matter, J. Sanchez-Bermudez, R. Norris, C. Paladini, L. Decin, H. Sana, J. O. Sundqvist, E. Lagarde, P. Kervella, A. Chiavassa, A. K. Dupree, G. Perrin, P. Scicluna, P. Stee, S. Kraus, W. Danchi, B. Lopez, F. Millour, J. Drevon, P. Cruzalèbes, P. Berio, S. Robbe-Dubois, and A. Rosales-Guzman, “The dusty circumstellar environment of Betelgeuse during the Great Dimming as seen by VLTI/MATISSE,” *Astronomy & Astrophysics* **675**, p. A46, 7 2023.
- 33 J. Borgne, “The polarized dust envelope around the red supergiant mu Cephei,” *Astronomy and Astrophysics* , 1987.
- 34 A.-M. Lagrange, J. Milli, A. Boccaletti, S. Lacour, P. Thebault, G. Chauvin, D. Mouillet, J. C. Augereau, M. Bonnefoy, D. Ehrenreich, and Q. Kral, “An insight in the surroundings of HR 4796,” *Astronomy & Astrophysics* **546**, p. A38, 10 2012.
- 35 S. de Regt, C. Ginski, M. A. Kenworthy, C. Caceres, A. Garufi, T. M. Gledhill, A. S. Hales, N. Huelamo, A. Kospal, M. A. Millar-Blanchaer, S. Perez, and M. R. Schreiber, “Polarimetric differential imaging with VLT/NACO,” *Astronomy & Astrophysics* **684**, p. A73, 4 2024.
- 36 J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, “Overcoming catastrophic forgetting in neural networks,” *Proceedings of the National Academy of Sciences* **114**, pp. 3521–3526, 3 2017.
- 37 P. C. Hansen, “Analysis of Discrete Ill-Posed Problems by Means of the L-Curve,” *SIAM Review* **34**, pp. 561–580, 12 1992.

- 38 S. B. Vandenberghe and Lieven, *Convex Optimization – Boyd and Vandenberghe*, Cambridge University Press, 2004.
- 39 L. Lilley, B. Norris, P. G. Tuthill, E. A. Spalding, M. Lucas, M. Zhang, M. Bottom, M. Millar-Blanchaer, B. S. Safonov, O. Guyon, J. Lozi, V. Deo, S. Vievard, K. Ahn, and J. Ashcraft, “Polarimetric, non-redundant aperture masking with next generation VAMPIRES: new instrumental capabilities, scientific outcomes, and image reconstruction techniques,” in *Optical and Infrared Interferometry and Imaging IX*, S. Sallum, J. Sanchez-Bermudez, and J. Kammerer, eds., p. 35, SPIE, 8 2024.

# Chapter 7

## Application of PIRATES to $\mu$ Cephei

### Chapter overview

This chapter applies the PIRATES algorithm to our observing campaign of  $\mu$  Cephei, reconstructing polarised images of the circumstellar environment over a 7 year period. To our knowledge, these reconstructions provide the highest spatio-temporal resolution description of  $\mu$  Cephei’s inner circumstellar environment to date. The PIRATES image reconstructions demonstrate strong consistency with the outcomes of parametric modelling in Chapter 5, though provide additional information that yields an improved description of  $\mu$  Cephei’s circumstellar geometry.

Using a parametrisation inspired by the PIRATES image reconstructions of  $\mu$  Cephei we fit to the circumstellar geometry, chemistry and grain size in fully physical radiative transfer code via Bayesian Optimisation algorithm `optuna`. The best fit model for  $\mu$  Cephei is an optically thick and inhomogeneous circumstellar envelope with a significant density enhancement of strength  $8.7 \pm 0.3$  at  $267 \pm 26^\circ$  East of North, which we interpret as a ‘clump’ suggestive of an episodic mass loss event. To extended material we fit a Gaussian clump at a radius  $61 \pm 5$  mas at  $200 \pm 5^\circ$  East of North, angularly co-located with an extended feature previously identified by Safonov et al. 2019 in the same region of the South-Western quadrant.

We also place physically meaningful constraint on the circumstellar dust grain size and chemistry – finding that  $\sim 120_{-37}^{+53}$  nm grains of Amorphous Enstatite provide an optimal fit. Grains of this size support the idea that Hofner’s scattering driven stellar wind may explain the high rates of mass loss observed from  $\mu$  Cephei. Our fitted model provides a mass loss rate of  $10^{-5} M_{\odot} \text{yr}^{-1}$ , in healthy agreement with literature values. Our detection of these grains has broader implications for red supergiants, adding to growing evidence that a scattering driven flow may explain the copious contributions of material made by these stars to the interstellar medium.

### Contribution Statement

I re-trained PIRATES for each observational epoch and applied it to the observational data. I wrote the code to fit radiative transfer models using Bayesian Optimisation algorithm `optuna`. I wrote custom code to generate custom density models that can interface with MCFOST. I conducted all analysis.

## 7.1 Managing systematic errors in real data


PIRATES (Chapter 6) has exhibited the potential to be a useful tool for image reconstructions in a number of contexts – in modelling both polarised and unpolarised interferometric data. Here we apply PIRATES to our observing campaign of  $\mu$  Cephei, reconstructing images of the circumstellar environment over a 7 year period. We reconstruct four single-wavelength images at  $\lambda 750 \pm 25$  nm from 2017-2020 (datasets 1-4), and four multi-wavelength images at  $\lambda 610, 670, 720, 760 \pm 25$  nm from 2023 (dataset 6).

As outlined in our paper (Lilley et al. 2025b), PIRATES implements an Elastic Weights Consolidation (EWC) strategy to mitigate the influence with random error in interferometric observables has upon the quality of image reconstructions. However, VAMPIRES Stokes I interferometric observables suffer from significant systematic miscalibration errors, and so a method to manage systematic errors is required when applying PIRATES to our observing campaign of  $\mu$  Cephei.


To manage systematic noise in Stokes I observables we approximate visibilities and closure phases with those of a fully darkened disk (FDD) and then strongly down-weight the contribution of Stokes I observables during the iterative fitting (stage 2). This is an effective method for two reasons. Firstly, Stokes I observables approximated in this way are ‘within the distribution’ of the examples used to pre-train PIRATES, ensuring that the output of the CNN (stage 1) is not a spurious response to out of distribution noise. A physically meaningful output from the CNN (stage 1) is required to provide a reasonable initialisation for the iterative fitting (stage 2). The second reason this method is effective is that the Stokes I observables can then be strongly down-weighted during the iterative fitting by using errors the same size as the FDD signal. This ensures that the iterative fitting stage is as agnostic as possible to the FDD approximation.

When this method is applied to realistic test examples of astrophysical scenes modelled in radiative transfer code we found that the polarised observables maintain their ability to reconstruct high fidelity polarised images, whilst also ‘leading’ the Stokes I image reconstructions to have the same dust distribution – despite the Stokes I observables themselves having only a stellar signal. This behaviour demonstrates the internally consistent weight sharing behaviour of convolutional kernels, and that we can benefit from this inductive bias when missing a modest amount of information (Ulyanov et al. 2020). Figures 7.1 and 7.2 depict two examples of this weight sharing behaviour, providing comparisons between the ground truth images (rows 1), reconstructions with full Stokes I information (rows 2), and reconstructions with FDD approximations to Stokes I (rows 3).

When the Stokes I observables are approximated with a FDD, the Stokes I image reconstructions are of varying fidelity – missing dust cavities (Figure 7.1) and in some cases missing patches of dust (Figure 7.2). However, the overall cost of this approximation to the polarised images is still modest – the reconstruction fidelity and structural accuracy of the polarised images (Stokes Q and U) is still high (see MSE and SSIM values in Figures 7.1 and 7.2). In the test examples of Figures 7.1 and 7.2 we fit Stokes I observables with a FDD, however the systematic errors present within our real VAMPIRES + NRM data are so large that we do not attempt this and instead adopt a literature value for the continuum stellar diameter (21 mas) (Quirrenbach et al. 1993) as we did within parametric modelling in Chapter 5



**Figure 7.1:** PIRATES reconstruction fidelity for a simulated clumpy circumstellar environment – row 1 depicts the true images, row 2 depicts images reconstructed by PIRATES with full Stokes I information, row 3 depicts images reconstructed by PIRATES when Stokes I observables are approximated with a fully darkened disk (FDD). When an approximation to the Stokes I observables is used (bottom row), the Stokes I reconstruction accuracy is lower and dust features can be missing – for example, the dust cavity between the star and dust. However, this comes at a small cost to the quantified accuracy of the polarised images (see MSE and SSIM values), which are able to ‘lead’ the Stokes I reconstructions to produce low order versions of the ground truth.



**Figure 7.2:** PIRATES reconstruction fidelity for a simulated inclined disk – row 1 depicts the true images, row 2 depicts images reconstructed by PIRATES with full Stokes  $I$  information, row 3 depicts images reconstructed by PIRATES when Stokes  $I$  observables are approximated with a fully darkened disk (FDD). When an approximation to the Stokes  $I$  observables is used (bottom row), the Stokes  $I$  reconstruction accuracy is lower and missing some dust on the disk’s closest edge. However, this comes at a small cost to the quantified accuracy of the polarised images (see MSE and SSIM values), which are able to ‘lead’ the Stokes  $I$  reconstructions to produce low order versions of the ground truth.

## 7.2 Reconstructing images of $\mu$ Cephei with PIRATES

For each of our observing epochs we augmented the training dataset with the instrumental Mueller matrix (Chapter 3) and re-trained an instance of PIRATES. In 2023 where data is multi-wavelength, one PIRATES instance is required for each wavelength to account for the varying (u,v) coverage. A future implementation of PIRATES will circumvent the requirement for this retraining with an auto-encoder that can manage both varying (u,v) coverage and varying polarisation coordinates.

### Tuning of the Elastic Weights Consolidation (EWC) factor

Using the procedure outlined Lilley et al. 2025b we optimise the strength of EWC by monitoring the magnitude of the EWC loss and the fit to interferometric observables. EWC scans and the identification of optimal values for each observing epoch are illustrated in Figures A.23 - A.24. All image reconstructions presented within this Chapter were reconstructed with these optimised EWC values.

### Improved fit to observables compared to best parametric models

The PIRATES image reconstructions of  $\mu$  Cephei yielded an improved fit to polarised interferometric observables in all epochs when compared to the best PIPPI models (Chapter 5). Table 7.1 compares the  $\chi^2$  values of each method for each sub-category of observables.

**Table 7.1:** The improvement yielded by PIRATES compared to our fitting with PIPPI –  $\chi^2$  values obtained for each sub-class of observables are displayed.


Year	PIPPI				PIRATES			
	$Q - V^2$	$U - V^2$	$Q - CP$	$U - CP$	$Q - V^2$	$U - V^2$	$Q - CP$	$U - CP$
2017	1.18	2.21	0.26	0.25	1.01	1.86	0.22	0.24
2018-02	1.36	1.95	0.17	0.18	1.29	1.33	0.16	0.16
2018-05	8.84	3.82	0.23	0.24	5.95	4.25	0.19	0.23
2020	3.02	2.04	0.20	0.18	2.40	1.66	0.17	0.16
2023-760	0.66	3.33	0.78	0.86	0.41	3.13	0.30	0.27

### Image Reconstructions


The PIRATES image reconstructions (Stokes I, Q and U) are plotted in Figures 7.3 and 7.4. A logarithmic scale is used for Stokes I (column 1), and a radial stretch ( $r^{1.2}$ ) has been applied to the Stokes Q and U images (columns 2 and 3) to make fainter extended dust structure clearer. All image reconstructions have been transformed back into the Stokes IAU NE celestial coordinate system applying the instrumental Mueller matrix model inverse via least squares. The PIRATES image reconstructions of  $\mu$  Cephei suggest a dominant asymmetric inner circumstellar envelope and fainter extended material concentrated in the South-Western quadrant.

### Dust both in the inner environment and in extended regions

Our image reconstructions reveal that for all epochs of data the scene is dominated by polarised flux from the inner circumstellar environment. However, there is also consistent extended material present in the South-Western quadrants of all epochs, as also seen in our parametric modelling and detected an PDI study (see our Chapter 5 and Safonov et al. 2019). In the PIRATES image reconstructions this extended feature has polarisation strength of at most 6% of the signal from the inner circumstellar environment. There is also a fainter feature in the North/North-Eastern quadrants, which was also tentatively identified by Safonov et al. 2019. In our PIRATES images this feature has at most 2% of the polarised signal of the inner circumstellar environment, and is significantly varied in its reconstruction between epochs. We suspect that this demonstrates this feature is limited by the intrinsic signal to noise of our data. This suggests that PIRATES is able to reconstruct images with internal contrast of up to  $\sim 1/50$ .



**Figure 7.3:** PIRATES image reconstructions for archival VAMPIRES data of  $\mu$  Cephei at  $\lambda 750 \pm 25$  nm from 2017-2020. All epochs of data reveal an asymmetric circumstellar envelope and consistent extended material in the South-Western quadrant. Stokes I (column 1) is plotted on a logarithmic scale, Stokes Q and U (columns 2 and 3) are plotted with a radial stretch ( $r^{1.2}$ ) to accentuate fainter dust on extended scales.




**Figure 7.4:** PIRATES image reconstructions for our multi-band VAMPIRES data of  $\mu$  Cephei from 2023, obtained at  $\lambda$  610, 670, 720, 760  $\pm$  25 nm. All wavelengths reveal an asymmetric circumstellar envelope and consistent extended material in the South-Western quadrant. Stokes I (column 1) is plotted on a logarithmic scale, Stokes Q and U (columns 2 and 3) are plotted with a radial stretch ( $r^{1.2}$ ) to accentuate fainter dust on extended scales.

## Asymmetry in inner environment

All image reconstructions reveal asymmetry in both the inner circumstellar environment and in extended material. The asymmetry in the inner environment is most clearly visualised in images of polarisation  $P = \sqrt{Q^2 + U^2}$  (Figure 7.5). In the inner circumstellar region, image reconstructions for datasets 1 - 4 (2017-2020) exhibit both a mild geometric asymmetry and a brightness asymmetry, whilst image reconstructions of dataset 6 (2023) exhibit geometric symmetry and a much stronger brightness asymmetry. Both of these forms of asymmetry were identified and fitted in earlier parametric modelling with PIPPI. Pleasingly, the brightness asymmetries reconstructed by PIRATES are co-incident with the brightness asymmetries fitted parametrically with PIPPI in Chapter 5 (Figure 7.5 – compare the top row of parametric models to the bottom row of image reconstructions where the brightness asymmetries are generally aligned).

It was noted in Chapter 5 that we suspected the density asymmetry fitted in Model S (which was restricted to an angular size of  $45^\circ$ ) was insufficiently small and that this led to an unphysically small semi-minor axis fitted to the ellipse – that the geometric asymmetry was attempting to emulate greater complexity than it could represent in a physically meaningful manner. We also suspected that the 2023 data was more asymmetric than in earlier epochs. These suspicions are confirmed by the PIRATES image reconstructions for 2023 which have the strongest brightness asymmetry of all years, of size that exceeds  $45^\circ$  with no artificially small semi-minor axis.



**Figure 7.5:** The consistency of asymmetry fitted with PIPPI and reconstructed by PIRATES. Here the middle 70 mas of each image is depicted to highlight the inhomogeneity of the inner-most circumstellar environment. Row 1 depicts the best PIPPI models (Model S from Chapter 5), and row 2 depicts the PIRATES image reconstructions. All images are of Polarisation ( $P = \sqrt{Q^2 + U^2}$ ) at  $\lambda 750 \pm 25$  nm or  $\lambda 760 \pm 25$  nm (see headings). There is general alignment of asymmetric features between the PIPPI fits and PIRATES images.

We note that the orientation of the ellipses fitted parametrically are consistently *not* aligned with the small geometric asymmetries indicated within the image reconstructions. However, interestingly, the parametric ellipses *are* aligned along the axis of the image reconstructions which have the most brightness asymmetry. This suggests that the density asymmetry of Model S may have been inadequately small for *all* epochs of data – and that parameters defining the geometric asymmetry were thus fitted to compensate for this, rather than fitting to the true alignment of the circumstellar shell.

## Temporal Evolution

The PIRATES image reconstructions exhibit temporal evolution of the dust in both the inner circumstellar environment and in extended regions between 2017-2023. This evolution was also suggested by our parametric modelling (Chapter 5), affirming that this is a systematic feature of the data. We interpret this variability differently depending on the geometric region and strength of the polarised material. The

inner-most circumstellar environment (the circumstellar shell) strongly dominates the signal and reconstructed images. We interpret the temporal variability of this material with confidence as there is clear evolution to both the dust shape and the changing orientation and size of a strong brightness asymmetry. We also note that the image reconstructions from the two 2018 datasets (2 and 3) are consistent, which suggests that we can interpret broader variability of this region across other years as significant.

However, there is also significant variability in the presentation of fainter extended material. Firstly, we note that in Chapter 5 we found that a low order representation of this extended material (a Gaussian blob fitted in the South-Western quadrant) was an essential feature of geometric modelling of  $\mu$  Cephei’s circumstellar environment. If we interpret the extended material in our PIRATES image reconstructions to first order, our image reconstructions would also support that this is a good low-order representation of the extended material in this quadrant. This extended feature is faint (6% the polarisation of the inner region) though persists in our image reconstructions through all epochs of data with a small amount of variability to its centre of mass.

When the reconstructed Stokes I images are stretched on a logarithmic scale, significantly more variability is revealed in the appearance of the extended dust. This includes a faint variable feature (2% of the maximum polarisation) in the North or North-Eastern quadrant, which was also tentatively suggested by imagery from Safonov et al. 2019. We first confirmed that variations in the location and morphology of this faint extended material were not caused by stochastic variability in the iterative fitting process – we obtained indistinguishable images by running the iterative fitting multiple times. True variations in the dust morphology between our observing epochs are entirely plausible – analogous evolved stellar environments are observed to exhibit variability in the inner-circumstellar dust distribution on timescales as short as months to a year (Khoury et al. 2016; Ohnaka et al. 2017b). However, we would expect that our two reconstructions from the same evening in 2018 (datasets 2 and 3) should demonstrate consistency in the faintest extended material, which does not appear to be the case. As discussed, there are likely to be unquantified systematic variations between our 2018 datasets which may explain this variation (Chapter 3).


We interpret these results broadly to suggest that variations in the location and detailed morphology of faint extended material in our images reconstructions are likely the result of: variations in observing conditions, varying or uncharacterised systematic errors in our data, differing values required for the EWC, or the result of these dust features having very low signal to noise in our datasets. It seems most reasonable to interpret the extended material to first order – supporting the presence of an extended feature in the South-Western quadrant from 2017 through to 2023. In our original observing plan we had hoped to observe temporal evolution of this extended feature – expecting that it may travel radially outward with the stellar wind. When comparing the image reconstructions across epochs, there is indeed some suggestion that the extended feature exists at a larger radial distance in 2023 than in 2017. In the circumstellar environment of red supergiant Antares, similar ‘clumps’ were observed move outwards radially by  $0.4''$  over 12 years (Ohnaka 2014). However, to explore this idea with sufficient confidence, we await polarimetric calibration models for our 2024 and 2025 data in which we would required analysis of to confirm this motion is persistent.

## Polarisation Statistics


$\mu$  Cephei exhibits variable total polarisation in strength and direction (Larsson-Leander 1964; Coyne et al. 1968; Polyakova 2003). One of the limitations of reconstructing images from interferometric observables is that there is an unknown normalisation factor that renders comparison to literature values of  $P$  and  $\phi_P$  impossible on a pixel scale without photometric calibration of Stokes I flux.<sup>1</sup> However, it is interesting to compare the net *fractional* variability of  $P = \sqrt{Q^2 + U^2}/I$  and  $\phi_P = 1/2 \arctan(U/Q)$  values computed from reconstructed images to historically reported variability. Our reconstructed images exhibit variability in  $P$  of up to  $\sim 0.6\%$ , and  $\phi_P$  of up to  $\sim 25^\circ$  (Figure 7.7), which is fractionally consistent with historical data (V band, 550 nm) that shows  $P$  and  $\phi_P$  can vary by 2% and  $160^\circ$  within as little as a  $\sim 2 - 4$  year period (Figure 7.6, reproduced from Polyakova 2003). Throughout this thesis we have noted subtle geometric variation in models and images fit or reconstructed from the two 2018 datasets

<sup>1</sup>This is a planned part of future work

obtained on the same evening (datasets 2 and 3) – however, we note here that their net polarisation statistics are very similar.




**Figure 7.6:**  $\mu$  Cephei’s light curve (V band, 550 nm), degree of polarisation strength  $P\%$  and angle  $\theta^\circ$  vary significantly over  $\sim 20$  years. The smooth curve in row 1 is superimposed, with a period of  $850 \pm 150$  days. Figure and caption reproduced from Polyakova 2003.



**Figure 7.7:** Variation of degree of polarisation strength  $P\%$  and angle  $\phi_P^\circ$  derived from the PIRATES image reconstructions of  $\mu$  Cephei from 2017-2023. All years depict data obtained at  $\lambda 750 \pm 25$  nm, except for 2023 which is obtained at  $\lambda 760 \pm 25$  nm. Whilst the pixel-wise values cannot be compared, the summed polarisation statistics demonstrate fractional variability that falls well within the typical scales reflected historically (Figure 7.6).

Whilst we do not have  $P$  and  $\phi_P$  values covering the time periods from which we have reconstructed images of  $\mu$  Cephei, it is useful to confirm the fractional polarisation values appear consistent with historically reported trends. From our PIRATES image reconstructions we constructed vector maps of the pixel-wise polarisation vectors, confirming that the scattering symmetry is predominantly azimuthal in both the inner circumstellar environment and in extended material (Figure 7.8).



**Figure 7.8:** Polarisation vector fields of PIRATES images reconstructed from each observing epoch (dates, dataset codes and observing wavelengths are in the top right of each panel). The polarisation field has been multiplied by a  $r^{0.1}$  stretch function to make all polarisation vectors comparable on the same axes – polarisation vector lengths are thus not comparable at different points in (RA, DEC), but polarisation vector angles are comparable. The polarisation vector angles demonstrate that the circumstellar structure is dominated by azimuthal scattering, with no persistent radial features evident.

### 7.3 Parametrisation of $\mu$ Cephei’s circumstellar geometry

We used the PIRATES reconstructed images of  $\mu$  Cephei as inspiration for a parametric model of the circumstellar environment. We then fit to this parametrisation in radiative transfer code via Bayesian Optimisation algorithm `optuna`, solving for the geometry, chemistry and dust grain size simultaneously within a fully physical framework that should account for the limitations of PIPPI addressed in Chapter 5. Bayesian Optimisation algorithms are probabilistic estimators that balance the need to fully explore a probability space with the task of finding optimal parameter values (Brochu et al. 2010). They are designed for use on high dimensional and computationally expensive models like the fully physical and multi-wavelength radiative transfer models we fit here.


Due to the computational expense of radiative transfer models (which may take  $\sim 10$  minutes per multi-wavelength model) we first approximate the circumstellar geometry with the most dominant component of the signal – the inner-most circumstellar shell, which we model as a thick spherically symmetric shell with a density enhancement – we refer to this as Model 1. We fit  $\mu$  Cephei with Model 1 for each of the candidates for dust chemistry derived in Chapter 2. Then, once the best dust chemistry is established, we fit a second model combining this enhanced circumstellar shell with the extended dust feature we identified in Chapter 5 and our PIRATES image reconstructions – we refer to this as Model 2.

When fitting multi-wavelength data, the final  $\chi_{red}^2$  score depends on both the geometric approximation and the appropriateness of the multi-wavelength behaviour set by chemistry and grain size. In Chapter 5 we fit to a geometric approximation equivalent to Model 1 and obtained a  $\chi_{red}^2 < 1$  during chemical inference. Despite limitations with the physical realism of PIPPI models, the ability to obtain this good fit suggests that we should be able to obtain a similarly satisfactory fit to  $\mu$  Cephei here (with either Model 1 or 2), and are not limited by either geometric approximations nor systematic errors. A schematic defining all parameters for both models is provided in Figure 7.9.

Instead of hypothesis testing different representative grain sizes like we did with PIPPI (Chapter 5), here we fit continuously to a power law distribution and maintain the assumption of spherical Mie grains. With fully physical radiative transfer code, multiple scattering effects may have a significant influence on the fitted grain size, as was suggested by our injection testing of radiative transfer models in Chapter 5. The image reconstructions from PIRATES (Figure 7.4) allow us to use reasonably narrow priors for some geometric parameters – eg. the density enhancement can be localised to within  $\sim 90^\circ$ , the extended material to within  $\sim 30$  mas and our shell radius to within  $\sim 10$  mas. For parameters we cannot constrain from the images (eg. the dust grain size), wide priors that cover the ranges that we derived in Chapter 2 are necessary to assess the plausibility of grains supporting a scattering driven stellar wind. To initialise each Bayesian Optimisation we randomly drew 50 points from an approximate region of high probability (found manually). We then ran each Bayesian Optimisation algorithm for 700 evaluations which we assessed was sufficient to adequately span the parameter space and to observe stable numerical convergence (the loss function had plateaued for  $\sim 200$  iterations). Initial conditions and priors are summarised in Table 7.2. To ensure that the whole prior for each parameter is fully explored, some parameters (like the dust mass and grain size) were explored on a logarithmic scale, whilst others could be satisfactorily explored linearly (also noted in Table 7.2).

### 7.4 Chemical Inference with Radiative Transfer


All fitted models, their parameters estimates and uncertainties are summarised in Table 7.3, where uncertainties are the  $1\sigma$  values obtained from using Kernel Density Estimation (KDE) on the likelihood-weighted samples. Estimated posteriors for each chemical are provided in Appendix A.4.2, Figures A.25 - A.33. The dust density distribution of our best fit model (Model 2 – Amorphous Enstatite) is depicted in Figure 7.10. Images made from the best fit model are depicted in Figure 7.11. Correlation plots of the best model and real observables are depicted in Figure 7.12.



**Figure 7.9:** An overview of the parametrization of the dust models fit to  $\mu$  Cephei in radiative transfer code. Shell (grey) has radius  $R$ , thickness  $\Delta R$ , dust mass  $M_{\text{dust}}$ , enhancement ( $E_\alpha$ ,  $E_\theta$ ,  $E_\phi$ ,  $E_\sigma$ ), and a Gaussian blob (blue) ( $B_r$ ,  $B_\phi$ ,  $B_\sigma$ ,  $B_\theta$ ,  $B_{\text{mass}}$ ). All angles denoted with subscripts  $\theta$  are angles into the plane of the sky, measured from the  $z$  axis (where  $\theta = 0^\circ$ ) towards the  $x$  axis (where  $\theta = 90^\circ$ ).

**Table 7.2:** Bayesian Optimisation Parameters – names, meanings, units, initial conditions, priors and the scale with which the prior space was explored. Model 1 is a thick circumstellar spherical shell with a density enhancement – fitted parameters are those above the table’s midline. Model 2 is a thick circumstellar spherical shell with a density enhancement and an extended Gaussian feature – fitted parameters are all those defined in the table. Angles that represent angular coordinates in the image plane ( $E_\phi$ ,  $B_\phi$ ) are defined in celestial NE coordinates, in degrees East of North.

Parameter	Meaning	Unit	Initial	Prior	Linear/Log
$a_{min}$	Minimum dust grain size	$\mu\text{m}$	[0.05, 0.15]	[0.001, 1]	Log
$a_{width}$	Dust grain distribution width	$\mu\text{m}$	[0.001, 0.05]	[0.001, 1.0]	Log
$a^{pl}$	Power law exponent pl	-	[-2,-4]	[-2, -8]	Linear
R	Shell radius	mas	[13, 18]	[10, 20]	Linear
$\Delta R$	Shell thickness	mas	[0.1 ,5]	[1, 40]	Log
$M_{dust}$	Dust mass	$M_\odot$	[2e-9, 6e-8]	[1e-11 1e-6]	Log
$E_\sigma$	Enhancement width $\sigma$	$^\circ$	[0.001, 0.05]	[20, 150]	Linear
$E_\theta$	Enhancement location $\theta$	$^\circ$	[70, 110]	[45, 145]	Linear
$E_\phi$	Enhancement location $\phi$	$^\circ$	[290, 330]	[200, 300]	Linear
$E_a$	Enhancement amplitude	-	[3,5]	[1, 10]	Linear
$B_{mass}$	Blob dust mass	$M_\odot$	[1e-10, 4e-8]	[1e-11, 1e-8]	Log
$B_r$	Blob radial distance	mas	[50, 70]	[45, 75]	Linear
$B_\phi$	Blob location $\phi$	$^\circ$	[190, 210]	[155, 265]	Linear
$B_\theta$	Blob location $\theta$	$^\circ$	[80, 100]	[45, 135]	Linear
$B_\sigma$	Blob size $\sigma$	$^\circ$	[25, 35]	[5, 50]	Linear



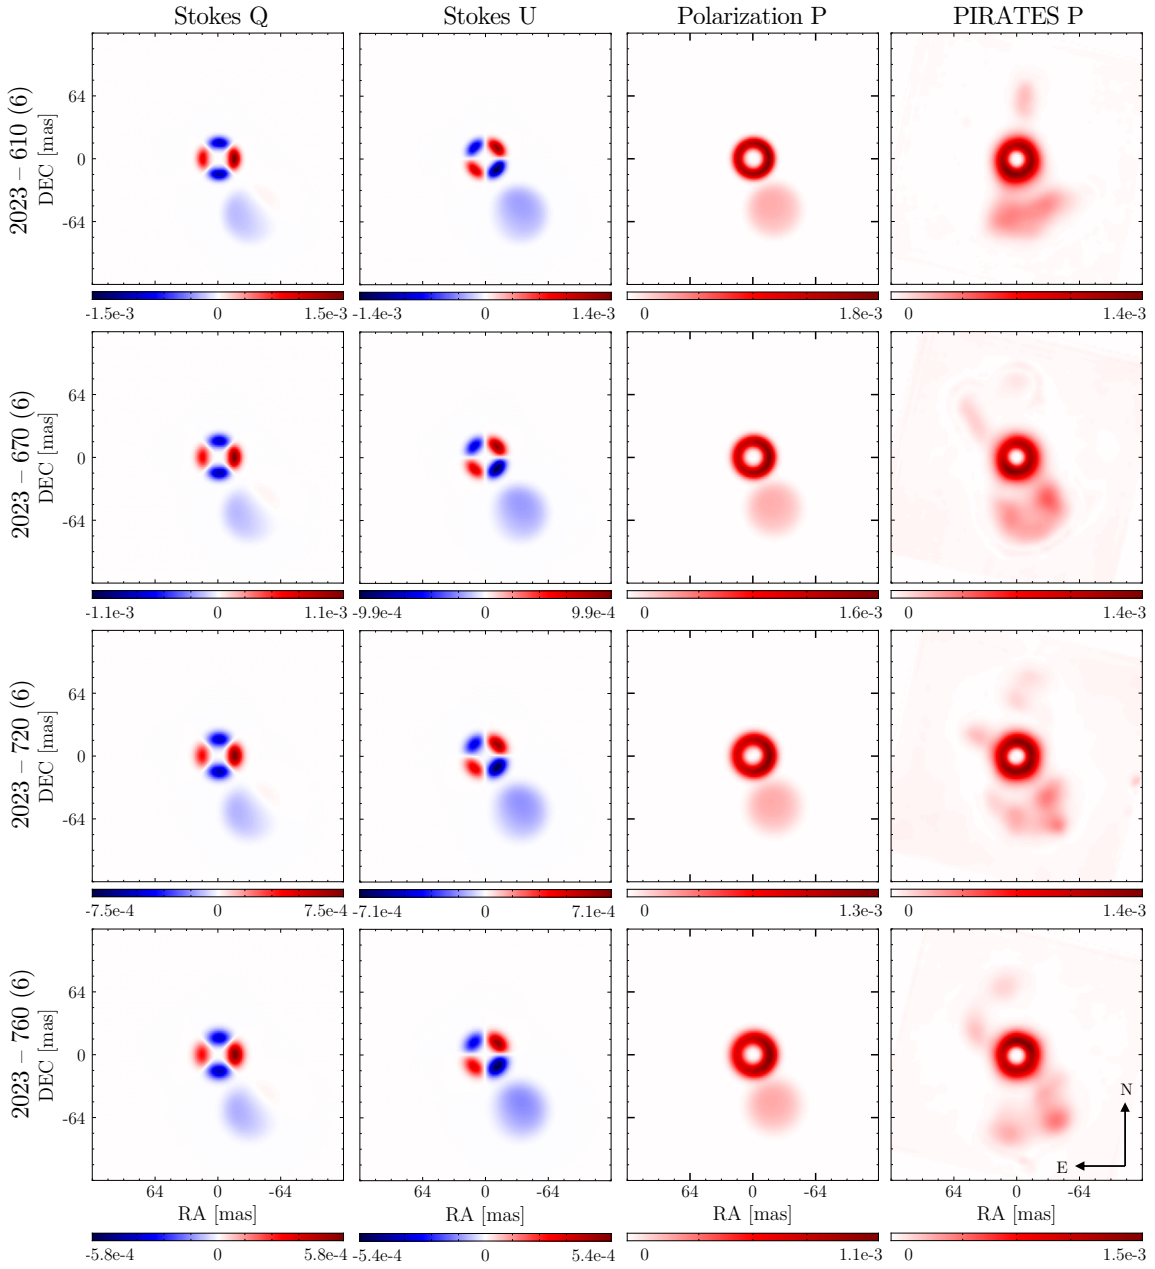
**Figure 7.10:** The dust density distribution of the best model fit to  $\mu$  Cephei in 2023 – Model (2) consisting of Amorphous Enstatite. Model was fit in radiative transfer code MCFOST via Bayesian Optimisation algorithm `optuna`. Image is formed by summing over the line of sight and has been radially scaled by  $r^{1.2}$  to make the extended feature in the South Western quadrant visible.

**Table 7.3:** Bayesian Optimisation Results for all fits to Model 1 and 2. All chemical options derived in Chapter 2 are fitted to Model 1, and the best fit chemical from Model 1 is then fit to Model 2. Model fits are ordered in descending  $\chi_{red}^2$ . Parameter uncertainties are obtained via Kernel Density Estimation with a kernel size set by Scott’s rule – values provided are  $1\sigma$ . Angles that represent angular coordinates in the sky plane ( $E_\phi$ ,  $B_\phi$ ) are defined in celestial NE coordinates in degrees East of North.

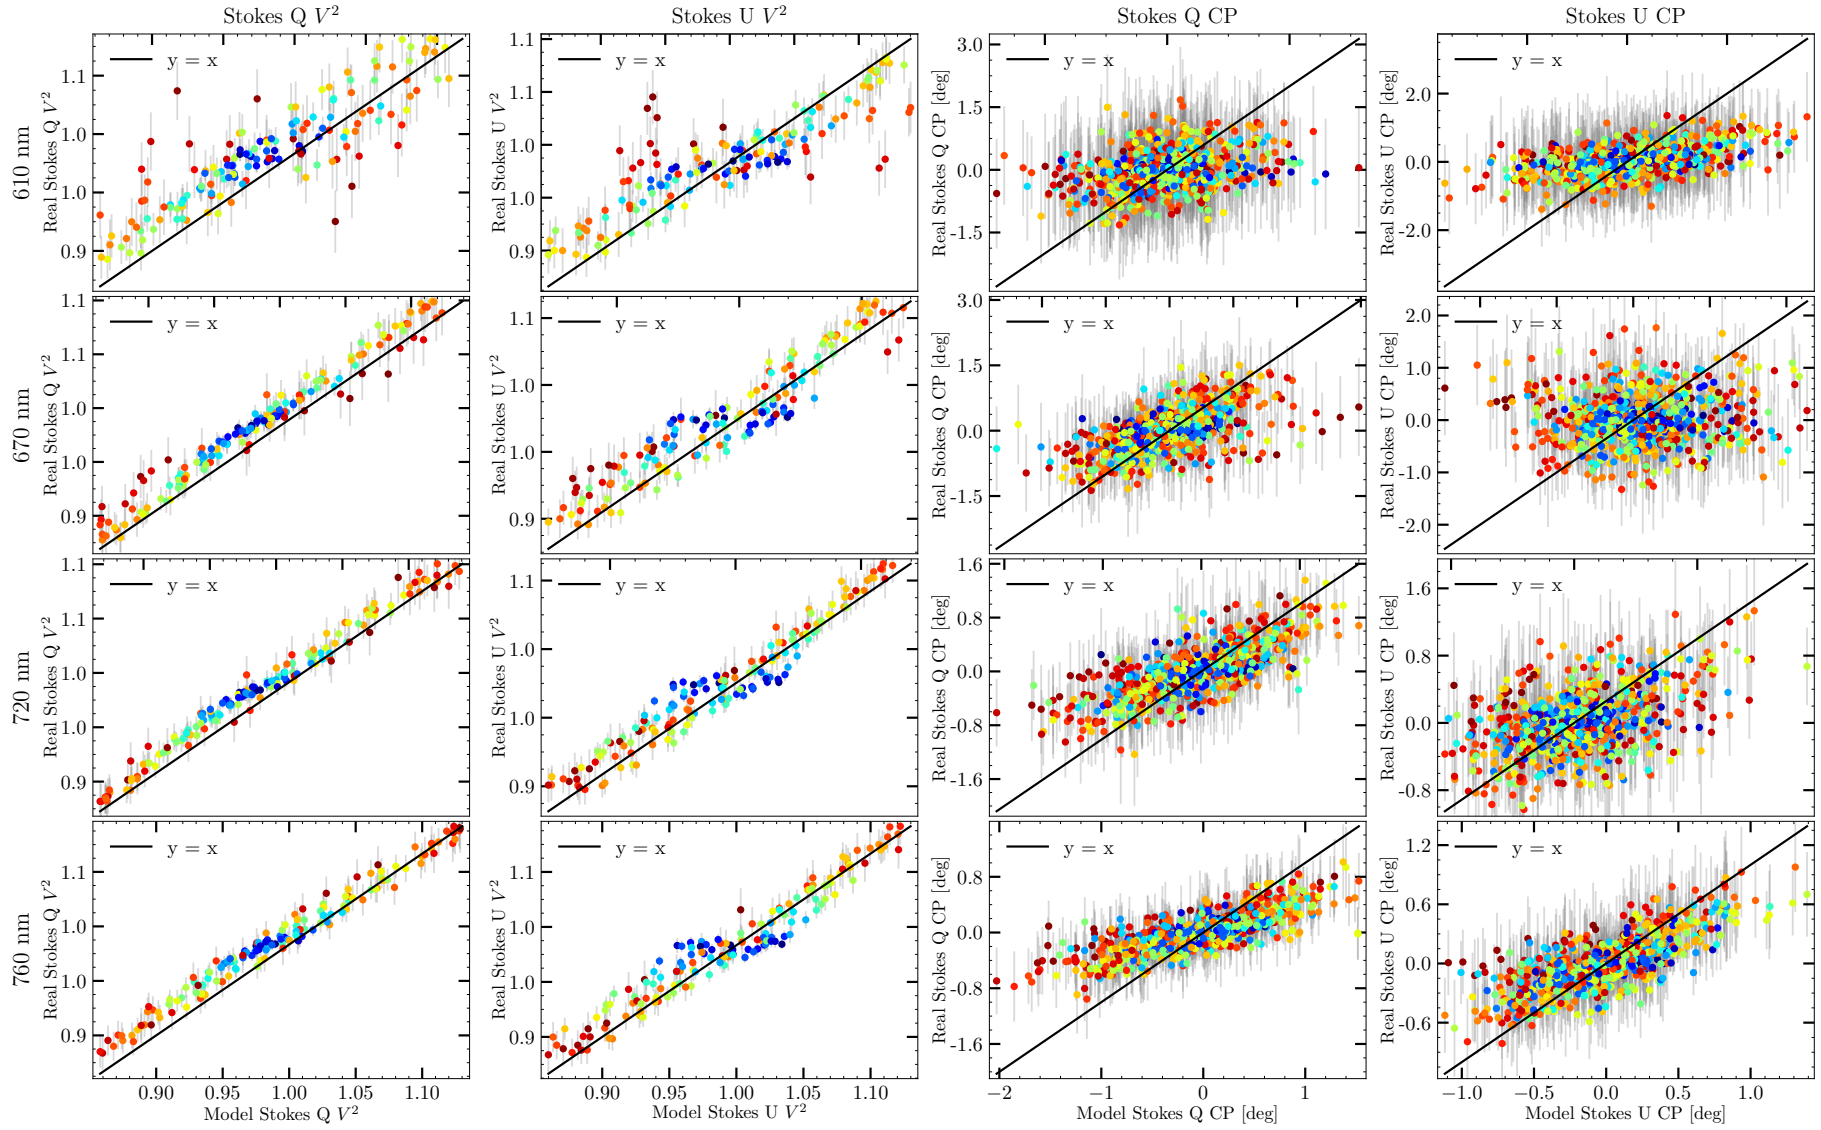
Chemical	Model	$a_{min}$ [nm]	$a_{width}$ [nm]	pl	R [mas]	$\Delta R$ [mas]	$M_{dust}$ [ $M_\odot$ ]	$E_\sigma$ [°]
Forsterite (C)	1	$62_{-22}^{+34}$	$2_{-1}^{+2}$	$-3.9 \pm 0.5$	$16 \pm 1$	$1.2_{-0.6}^{+0.8}$	$5.4_{-2.0}^{+3.3} \times 10^{-8}$	$108 \pm 10$
Enstatite (C)	1	$51_{-23}^{+40}$	$8_{-4}^{+10}$	$-2.2 \pm 0.2$	$17 \pm 1$	$0.2_{-0.1}^{+0.1}$	$1.2_{-0.4}^{+0.5} \times 10^{-8}$	$59 \pm 11$
Silica (A)	1	$15_{-6}^{+11}$	$340_{-174}^{+356}$	$-4.8 \pm 0.2$	$17 \pm 2$	$0.2_{-1.4}^{+0.1}$	$6.9_{-2}^{+3} \times 10^{-8}$	$66 \pm 17$
Pyroxene <sub>95</sub> (C)	1	$109_{-52}^{+100}$	$201_{-113}^{+258}$	$-5.3 \pm 0.3$	$17 \pm 1$	$0.5_{-0.2}^{+0.3}$	$0.7_{-2.8}^{+5.6} \times 10^{-8}$	$103 \pm 19$
Forsterite (A)	1	$10_{-5}^{+8}$	$197_{-76}^{+122}$	$-3.6 \pm 0.1$	$17 \pm 1$	$5.6_{-0.8}^{+0.9}$	$2.0_{-0.3}^{+0.3} \times 10^{-8}$	$80 \pm 14$
Corundum (A)	1	$107_{-39}^{+62}$	$69_{-36}^{+74}$	$-4.5 \pm 0.2$	$17 \pm 1$	$0.1_{-0.1}^{+0.1}$	$1.0_{-0.4}^{+0.6} \times 10^{-8}$	$126 \pm 8$
Spinel (A)	1	$138_{-46}^{+69}$	$127_{-72}^{+164}$	$-5.2 \pm 0.3$	$17 \pm 1$	$7.3_{-1.5}^{+1.8}$	$1.1_{-0.5}^{+1} \times 10^{-8}$	$105 \pm 9$
Enstatite (A)	1	$123_{-44}^{+69}$	$2_{-1}^{+1}$	$-3.2 \pm 0.3$	$17 \pm 1$	$0.8_{-0.4}^{+0.8}$	$1.1_{-0.4}^{+0.6} \times 10^{-8}$	$45 \pm 6$
Enstatite (A)	2	$120_{-37}^{+53}$	$2_{-1}^{+1}$	$-2.2 \pm 0.2$	$16 \pm 1$	$5.6_{-2.4}^{+1.7}$	$1.6_{-0.8}^{+1.6} \times 10^{-8}$	$43 \pm 8$

$E_\theta$ [°]	$E_\phi$ [°]	$E_a$	$B_{mass}$ [ $M_\odot$ ]	$B_r$ [mas]	$B_\phi$ [°]	$B_\theta$ [°]	$B_\sigma$ [°]	$\chi_{red}^2$
$84 \pm 10$	$259 \pm 16$	$4.0 \pm 0.6$	–	–	–	–	–	4.6
$137 \pm 11$	$262 \pm 13$	$4.9 \pm 0.5$	–	–	–	–	–	3.9
$117 \pm 10$	$276 \pm 29$	$1.9 \pm 0.4$	–	–	–	–	–	3.1
$115 \pm 21$	$266 \pm 21$	$5.0 \pm 1.0$	–	–	–	–	–	3.1
$119 \pm 18$	$285 \pm 24$	$1.8 \pm 0.4$	–	–	–	–	–	2.8
$92 \pm 13$	$272 \pm 16$	$5.5 \pm 0.4$	–	–	–	–	–	2.8
$96 \pm 18$	$237 \pm 25$	$7.0 \pm 0.4$	–	–	–	–	–	2.5
$150 \pm 15$	$230 \pm 11$	$6.6 \pm 0.4$	–	–	–	–	–	2.3
$156 \pm 10$	$267 \pm 26$	$8.7 \pm 0.3$	$7.3_{-4.0}^{+8.9} \times 10^{-11}$	$61 \pm 5$	$200 \pm 5$	$99 \pm 7$	$27 \pm 2$	2.0



**Figure 7.11:** Images of the best model fit to  $\mu$  Cephei in 2023 (dataset 6) using radiative transfer code MCFOST – Model 2 consists of a spherical thick shell with a density enhancement and an extended Gaussian blob. Amorphous Enstatite was adopted as it was found to yield the best chemistry for the inner circumstellar environment when all options were tested with Model 1. The first three columns depict images of the best model – Stokes Q (column 1), Stokes U (column 2) and polarisation (column 3). Columns 1-3 have been blurred by an Airy function of FWHM corresponding to our longest baseline length. Column 4 depicts the PIRATES reconstructed images of  $\mu$  Cephei. All images have a radial stretch of  $r^{1.2}$ , and all Stokes images were sum normalised by Stokes I and are dimensionless. Comparison of columns 3 and 4 demonstrates that both the inner circumstellar cavity and extended material are well reproduced by our best model. Each row depicts a different observational wavelength –  $\lambda$  610, 670, 720,  $760 \pm 25$  nm (denoted as row labels).




**Figure 7.12:** Observables of the best model fit to  $\mu$  Cephei in 2023 (dataset 6) using radiative transfer code MCFOST (Model 2 as depicted and described in Figure 7.11 – Amorphous Enstatite). Each row depicts the fit to observables at a different observational wavelength (610, 670, 720, 760 nm). Columns 1 and 2 depict the correlation of polarised visibilities, columns 3 and 4 depict the correlation of polarised closure phases. The error bars plotted are the combination of random (obtained by bootstrapping) and systematic (estimated by comparison with a calibrator).

## 7.5 Analysis of Radiative Transfer models

### Chemical Differentiation in systems with $\tau > 1$

We fitted all chemical options derived in Chapter 2 with Model 1 and found that the best fit dust chemistry was Amorphous Enstatite. In Chapter 5 we found that with optimised grain size all chemical options provided very similar fits to the data with  $\Delta\chi_{red}^2 \sim 0.01$  between chemicals. In contrast, here we found that the differences between each optimised chemical species are much larger with  $\Delta\chi_{red}^2 \sim 0.2 - 2.3$ , which implies that we have a greater ability to differentiate between chemicals when fitting using radiative transfer code. The most crucial improvement made by using radiative transfer code is the inclusion of multiple scattering processes. Photons emitted from the star no longer have a single polarising interaction with a single dust grain. Instead, they navigate stochastically through the dust and may have multiple interactions that re-scatter and de-polarise the emerging light in an anisotropic fashion that depends on grain size (Max Born 1975). In Chapter 5 it was noted that when modelling only a single scattering process we anticipated the equivalence of chemical fits on account of the similarity of optical constant data – in particular the real component ( $n$ ) which dominates scattered signal when  $\tau < 1$ . However, here our fitted radiative transfer models depart significantly from  $\tau < 1$  (Figure 7.13).




**Figure 7.13:** Optical Depth of best fit model – Amorphous Enstatite (Model 2) as a function of wavelength. Most of the fitted shell has an optical depth above 1, the peak of density and optical depth are observed in the region where the density enhancement is maximised.

When  $\tau > 1$ , absorption processes become increasingly dominant as each additional photon-dust interaction increases the probability that the photon will be absorbed. The absorption cross section is a strong function of the imaginary component  $k$  of the dust refractive index. Though the values of  $k$  are small for all our tested chemicals, they differ from each-other by several orders of magnitude (Figures A.1 - A.3). To illustrate how rapidly these differences compound to distinguish chemical species, we used `miepython` to compute the absorption and scattering cross sections for 100 nm grains of varying  $k$  (using a range of  $10^{-2}$  -  $10^{-6}$  that spans the variation in our candidate chemicals from Chapter 2). By plotting the probability of photon survival after  $N$  interactions, the differing behaviour of each species is clear after only a single interaction (Figure 7.14). Thus when  $\tau > 1$ , we would expect from a theoretical standpoint to have greater chemical differentiation on account of both the real and imaginary components of the refractive index now being modelled in a physically complete manner. We previously noted that existing studies have been unable to constrain the circumstellar chemistry from optical and near-IR data alone (Hauboiss et al. 2019; Khouri et al. 2016; Ohnaka et al. 2016; Ohnaka et al. 2017a). However, we note an important point of difference – these studies found that the dust was best fitted as optically thin when using radiative transfer code. In contrast, the circumstellar dust shell we have fitted to  $\mu$  Cephei has a minimum optical depth of  $\sim 1$  which becomes as large as  $\tau = 3.6$  at 610 nm – the greatest optical depth occurs within the fitted density enhancement (Figure 7.13).

Thus, our results do not contradict existing studies which struggled to differentiate between chemicals. Instead, they suggest that greater optical depth (which we can now model in radiative transfer code) provides greater chemical differentiation on account of absorption processes further distinguishing each chemical. We proceed to analyse our best fit solution (Amorphous Enstatite) as the numerical optimum

– however we note that our subsequent arguments remain true for many of the better fitting models and withstand variations to the fitted parameters. We return to interpreting our spread of fitted models in later discussion.



**Figure 7.14:** Survival probability of photons scattering through dust media with fixed real refractive index  $n$  and varying imaginary refractive index values  $k$  (given in the legend).  $P_{survival} = (1 - P_{abs})^N$ , where  $P_{abs}$  is computed from Mie scattering theory as  $C_{abs}/(C_{abs} + C_{scat})$ . Though the values of  $k$  of our tested chemicals are all small ( $10^{-2}$  -  $10^{-6}$ ), they are distinguished from one another after a small number of photon interactions – suggesting that theoretically we have greater chemical differentiation when  $\tau > 1$ .

The enhanced optical depth of our best model also confirms the suspicions raised in Chapter 5: that a fully physical solution for  $\mu$  Cephei’s circumstellar dust requires  $\tau > 1$ , meaning it cannot be accurately modelled with simple scattering models like PIPPI and those of Ireland et al. 2005, Norris et al. 2012 and Haubois et al. 2019. In Chapter 2 we also hypothesised that if circumstellar conditions permit, crystalline structure may form and be rapidly re-amorphised by cosmic rays – which would lead to a dominant amorphous signature in wide field spectroscopy and a crystalline signature in the inner circumstellar environment. The greater chemical differentiation yielded by modelling optical depth here has revealed that  $\mu$  Cephei’s inner circumstellar environment is best described by an amorphous silicate and not a crystalline alternative (which produced some of the worst fits to data), which does not support this hypothesis.

### Clumps of dust with $\tau > 1$

Our best model of  $\mu$  Cephei’s inner circumstellar environment features a significant density enhancement of amplitude  $8.7 \pm 0.3$  at  $267 \pm 26^\circ$  East of North and at  $156 \pm 10^\circ$  into the plane of the sky (see model description in Figure 7.9). In Chapter 1 we reviewed the growing observational evidence suggesting that ‘clumps’ (regions of enhanced density) are common in circumstellar environments, often interpreted as evidence for episodic mass loss events. Such ‘clumps’ are typically modelled with a density enhancement like ours (eg. AGB W Hydrae in Ohnaka et al. 2016 with contrast  $4 \pm 1$ ), and have been demonstrated to be optically thick (eg. RSG VY Canis Majoris and YSG IRC +10420 in Shenoy et al. 2013 and Shenoy et al. 2015). Our best model’s density enhancement is more intense than that fit for W Hydrae. If these clumps are interpreted as signatures of episodic mass loss events then is reasonable that we might observe more intense clumps around  $\mu$  Cephei given W Hydrae’s lower mass loss rate (Khouri et al. 2015b).

### Extended material – clump in the South Western quadrant

Model 2 allowed us to fit the extended material around  $\mu$  Cephei – the model located a Gaussian blob at a radial distance of  $61 \pm 5$  mas at an angular location of  $\phi$ :  $200 \pm 5^\circ$  East of North and  $\theta$ :  $99 \pm 7^\circ$  (see model description in Figure 7.9). We note the agreement of the angular position of this blob ( $\phi$ ) with that which we fitted parametrically with PIPPI, and with the extended material identified by Safonov

et al. 2019. As we predicted in Chapter 5, with the representation freedom provided by PIRATES, a more complex inner circumstellar environment is reconstructed for the 2023 dataset than was permitted with PIPPI’s Model S. This then leaves this extended clump free to exist at a greater radial distance, no longer needing to compensate for a lack of asymmetric representation in the inner environment. Here the blob is located at more than twice the radial distance found in Chapter 5. This extended clump is different to the enhancement fitted in the inner circumstellar environment – it is optically thin ( $\tau \sim 0.15$ ) and has a dust mass of  $7.3_{-4.0}^{8.9} \times 10^{-11} M_{\odot}$ .

## Consistency with PIRATES image reconstructions

Images of our best fit model are in excellent agreement with the PIRATES reconstructed images (Figure 7.11). The density enhancement fitted within the circumstellar shell is consistent with the brightness asymmetry presented in the reconstructed images (both in contrast and in location, though we note less agreement is exhibited in the angular location at 760 nm than at other wavelengths – we return to this in later discussion). The Gaussian clump fitted to extended material appears to approximate the extended material’s centre of mass, co-located with the angular location and radial distance suggested by the PIRATES image reconstructions.

## Proximity of dust to the stellar surface

Our best model fits a thin shell ( $2_{-1}^{+1}$  mas) with a radius of  $1.6 R_{\star}$  ( $16 \pm 1$  mas) which is consistent with the range of dust shell radii fitted to analogous stellar environments (Munoz-Sanchez et al. 2024b; Scicluna et al. 2015; Wittkowski et al. 2012; Ohnaka et al. 2024; Norris et al. 2012; Haubois et al. 2019). Our fitted dust shell is co-located with  $\mu$  Cephei’s disconnected molecular layers at  $\sim 2 R_{\star}$ , coined the ‘Mol-Sphere’ (Tsuji 2000). The co-location or mixing of circumstellar gas and dust is an understood phenomena in O-type evolved stars (Kervella et al. 2018; Gobrecht et al. 2016; Karovicova et al. 2013). Like our parametric modelling in Chapter 5, the thin nature of the circumstellar shell fitted confirms that  $\mu$  Cephei’s circumstellar dust does not appear to readily conform to a steady outflow.

## Grain size inference


The grain distribution fit to our best model is very narrow if only the MLE values are considered (120–122 nm). Our fits to all chemical options via Model 1 reveal that this is a persistent feature of all fits – most of the better fit distributions are less than  $\sim 200$  nm wide, and for both Models 1 and 2 fit (independently) to Amorphous Enstatite, the grain distribution widths are very narrow. We will return to discuss in great detail how these narrow distributions should be interpreted – first we assess the MLE solutions for their physical consequences, accepting their narrow parametrisation as the numerical optimum.

In Chapter 5 we used PIPPI to find that multiple dust grain sizes provided very similar fits to our data – grains of radius  $\sim 250$ , 750 and above 1200 nm. Whilst grains of 250 and 750 nm would support Hofner’s scattering driven stellar wind from  $\mu$  Cephei, grains of above 1200 nm are too large and hence we were unable to conclusively comment on the plausibility of a scattering driven wind from this inference. In contrast, the posterior for our best fitted model fitted in radiative transfer here demonstrates that we can rule out very large grains above 1200 nm, and instead supports a distribution of grains around  $\sim 120$  nm. There are several comments to be made about these points of difference.


Firstly, to understand why it is possible to rule out very large (1200 nm) grains when using radiative transfer code, we constructed a set of models containing a thin shell of Amorphous Enstatite around  $\mu$  Cephei ( $R_{\text{dust}} = 1.2 R_{\star}$ ) at representative grain sizes (2 nm–2  $\mu\text{m}$ ) and simulated observing at  $\lambda 750 \pm 25$  nm. For each of these models we increased dust shell mass until polarised visibility amplitudes were as large as  $\mu$  Cephei’s ( $\sim 1.1$ ). Our results demonstrate why radiative transfer permits us to rule out large grains of 750 nm and above 1200 nm – when models with these large grains have sufficient dust density to generate observationally high polarised visibility amplitudes, they are *very* optically thick with  $\tau \sim 10$  (Figure 7.15).

Clarity on why these large grain models with high optical depth provide poor fits to data can be obtained by examining the polarised images of these models and comparing them to our PIRATES image

reconstructions. Multiple scattering induced by high optical depth causes the ‘shell’ structure of these models to appear in the polarised images as filled in ‘discs’ (Figure 7.16). In Figure 7.16 we plot several representative models from Figure 7.15 that convey how this structural trend relates to grain size – each column depicts a representative grain size, the first row is the Stokes I image and the second row is the polarisation  $P = \sqrt{Q^2 + U^2}$  image made from each model.



**Figure 7.15:** Simulated thin shells of Amorphous Enstatite around  $\mu$  Cephei ( $R_{\text{shell}} = 1.2 R_\odot$ ), observed at  $\lambda 750 \text{ nm}$ . For each grain size, dust mass is increased until polarised visibility amplitudes match  $\mu$  Cephei’s ( $\sim 1.1$ ). Plotted are the optical depths of these observationally comparable models (left) and the required dust mass (right). These simulations demonstrate that most large grains (radii  $> 500 \text{ nm}$ ) capable of matching  $\mu$  Cephei’s polarised visibility amplitudes have densities that correspond to very high optical depths ( $\tau \sim 10$ ).



**Figure 7.16:** The cost of obtaining  $\mu$  Cephei’s polarised visibility amplitudes – images are of our models from Figure 7.15 where dust mass has been increased until the polarised visibility amplitude of the model matches  $\mu$  Cephei’s. This point occurs in a different physical regime for each grain size, which the present images depict. For small grain sizes (2 nm), obtaining  $\mu$  Cephei’s polarised visibility amplitudes requires modest optical depth, reducing the clarity of the shell’s cavity. For intermediate grain sizes ( $\sim 100 \text{ nm}$ ), high polarisation is obtainable whilst the shell’s cavity is maintained. However, as dust grain size is increased up to and beyond 500 nm, the prominence of the cavity is reduced and the morphology becomes heavily embedded.

For small grains ( $\sim 2 \text{ nm}$ ) the central cavity of the dust shell is partially obscured by the intermediate optical depth required. For intermediate sized grains ( $\sim 100 \text{ nm}$ ), high polarised visibility amplitudes can be obtained whilst retaining the shell’s cavity, however as grains grow larger ( $\sim 250 \text{ nm}$  and above) the cavity becomes opaque. By  $\sim 500 \text{ nm}$  optical depth is so high ( $\tau = 10$ ) that multiple scattering dominates

the images and the cavity has ‘closed’ – the star is heavily embedded in dust. Such a morphology contradicts the structure of our PIRATES image reconstructions (Figures 7.3 and 7.4) – which feature a strong central cavity at all wavelengths, indicating that the dust cannot be so optically thick as to remove this feature.

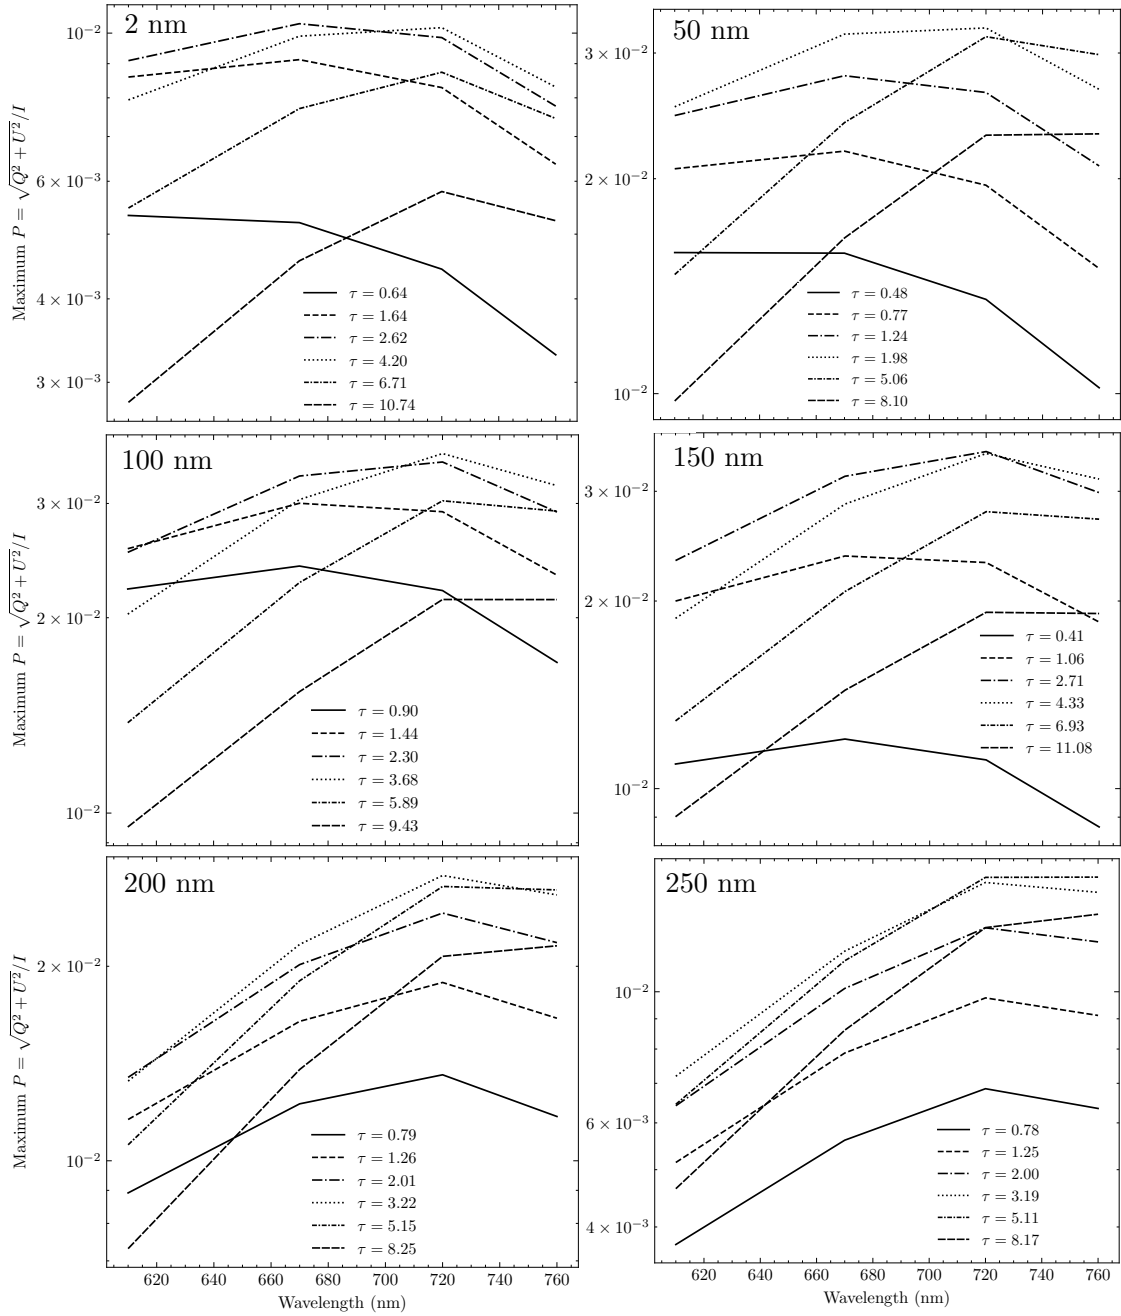
Thus with fully physical radiative transfer code we have so far obtained far more constraint on dust grain size than we did using simple scattering models like PIPPI. The major point of difference is inclusion of optical depth which provides full physical constraint of the modelled dust systems. Whilst our previous use of PIPPI left ambiguity over  $\mu$  Cephei’s true dust grain size and the consequences for Hofner’s scattering driven stellar wind, with radiative transfer code we have broken this degeneracy and can now rule out grain sizes of  $\sim 500$  nm and above.

The second significant point of comparison to be made is that the grain sizes we have fitted in the present radiative transfer framework (a narrow distribution around  $120_{-37}^{+53}$  nm) are significantly smaller than even the smallest grain size indicated by our modelling with PIPPI in Chapter 5 ( $\sim 250$  nm). This suggests that the wavelength dependence of a system with modest optical depth and  $\sim 120$  nm grains was best emulated by our (un-physical) simple scattering models with  $\sim 250$  nm grains. To understand why we obtained different grain sizes using each modelling framework, it is helpful to consider the two key changes that dust exhibits when its optical depth is increased beyond  $\tau = 1$ . Firstly, the increase in multiple scattering effects causes a net *de-polarisation* – photons navigate stochastically through the dust and subsequent interactions reduce the coherence of emerging polarisation as a result of randomising the original polarisation vector.

The second important change is that increased optical depth does not simply rescale the wavelength dependence of polarisation but *reshapes it*, typically moving the scattering angle at which maximum polarisation is observed forwards (Shrestha et al. 2020; Lin et al. 2022; Yang et al. 2017). It is this behaviour that explains why we obtained different grain sizes with each modelling technique, and why those obtained with PIPPI may be larger – increased optical depth and larger grain sizes may both move the scattering angle of maximum polarisation *forwards*. As PIPPI could not model optical depth, it instead fit a larger grain size and compensated for the change in polarised flux with the star-dust-contrast parameter (see Chapter 5 for a discussion of this free parameter).

We confirmed that optical depth re-shapes the wavelength dependence of polarisation in our own simulated circumstellar shells by tracking how the wavelength dependence of  $P = \sqrt{Q^2 + U^2}/I$  varies as a function of observing wavelength and optical depth. To do so we use the same representative Amorphous Enstatite shell at  $1.2R_*$ . Our results are plotted in Figure 7.17 for several representative grain sizes, revealing that it requires only modest optical depth ( $\sim \tau = 2 - 3$ ) to change the wavelength dependence of polarisation significantly from optically thin behaviour. Our models clearly show that we should expect to obtain different values for fitted grain size if we fit an optically thicker system using simple scattering models restricted to optically thin wavelength dependence – in other words, that simple models may bias grain size inference if they are not applied to targets that are *strictly* optically thin.

Given that there is growing empirical evidence (including our own analysis) that enhanced density ‘clumps’ cause the inner circumstellar environment to have  $\tau > 1$  (Shenoy et al. 2013; Shenoy et al. 2015; Ohnaka et al. 2016), our results reinforce the necessity of using fully physical radiative transfer code for inference on grain size. The significance of the discrepancy between these two modelling practices will depend on the astronomical target and observing apparatus – in our case for  $\mu$  Cephei we found a discrepancy of  $\sim 130$  nm between methods, which for some targets would be large enough to change whether inference supports the detection of circumstellar grains facilitating a scattering driven stellar wind (Höfner 2008). It is also satisfying to visually compare the similarity of wavelength dependence of: the grains we fitted with PIPPI (250 nm, optically thin,  $\tau = 0.78$  Figure 7.17) to those we have fitted in radiative transfer ( $\sim 100$  nm, using  $\tau = 3.68$  Figure 7.17), which have similar wavelength dependent polarisation.



**Figure 7.17:** The effect of increased optical depth on the wavelength dependence of polarisation, each panel displays this relationship for a different representative grain size. A thin circumstellar shell of Amorphous Enstatite was simulated around  $\mu$  Cephei at  $1.2 R_*$ . Dust mass was increased to span a range of optical depths  $\tau$  as indicated within each legend. The polarisation  $P = \sqrt{Q^2 + U^2} / I$  of the shell is measured at each observational wavelength  $\lambda 610, 670, 720, 760 \pm 25$  nm. As optical depth is increased above  $\sim \tau = 2$ , the wavelength dependence changes significantly, and may be quite different to that of the optically thin case. Optical depth  $\tau$  is a wavelength dependent property – we recorded the value of  $\tau$  at 760 nm in the figure legends.

## Support for Hofner’s scattering driven stellar wind hypothesis

In the case of our best fit model, the fitted grains have minimum size of  $120_{-37}^{+53}$  nm and a grain distribution width of  $2_{-1}^{+1}$  nm. If we treat only the MLE values, this dust grain distribution is very narrow. In this section we will explore both the physical and numerical origins for the sizes of these fitted grain distributions.

The general expectation of these inner circumstellar environments is that they may host larger grains than the broader distribution of the interstellar medium (Asano et al. 2013; Srama et al. 2009; Mathis et al. 1977a), and that their grain distributions may be narrow (though we note the theoretical research on this has mostly been conducted for carbon-rich AGBs) (Groenewegen 2006; Winters et al. 1997; Yasuda et al. 2012). In the case of silicate dust, dust grains are anticipated to have a narrow distribution with radii of  $\sim 500$  nm due to the interaction of grain growth and wind acceleration (Höfner 2008).

These ideas have guided existing studies to fit an ‘effective’ or ‘averaged’ grain radius as a reasonable approximation for a narrow grain distribution – representative values between 0.02-0.7  $\mu\text{m}$  are common (Norris et al. 2012; Ohnaka et al. 2016; Khouri et al. 2020; Scicluna et al. 2015; Adam et al. 2019). Norris et al. 2012 argued that grains smaller than the effective radius may exist but will contribute a significantly smaller amount to the polarised signal on account of a smaller polarised scattering cross section. An alternative treatment of dust grain size has been to assume a Mathis et al. 1977b ISM distribution consisting of small (Rayleigh) grains, a power law exponent of -3.5 and a maximum grain size that is fitted freely (Khouri et al. 2015a). In our present radiative transfer modelling we attempted to advance the level of constraint placed on the circumstellar grain size distribution by fitting freely to a power law distribution with a minimum size  $a_{\text{min}}$ , width  $a_{\text{width}}$  and power law exponent  $pl$ .


Considering our best fit solution, the grains fitted to  $\mu$  Cephei ( $\sim 120_{37}^{+53}$ ) are significantly smaller than the  $\sim 500$  nm predicted by Höfner 2008, and we demonstrated in Figures 7.15 and 7.16 that a representative shell of 500 nm grains that generates polarised visibility amplitudes comparable to  $\mu$  Cephei’s will be too optically thick to be consistent with our PIRATES image reconstructions of  $\mu$  Cephei (Figures 7.3 and 7.4).

Our analysis in Chapter 5 found that the provision of power law distributions did not significantly improve the fit to data – although the differences between fits obtained with single grains and power law distributions were small, the single grain size distributions provided a slight improvement. These models were later shown to be physically unrealistic, however our present fit’s preference for a very narrow grain distribution seems to be consistent with the wavelength dependence of results presented in Chapter 5. Our finding is also consistent with research on analogous RSG Betelgeuse – where spectro-polarimetric interferometry was used to constrain the grain distribution to around  $\sim 300$  nm with distribution width of  $\sim 20$  nm (Haubois et al. 2019). Thus it may be possible that the grain size distributions around O-rich RSGs are as narrow as those found in C-rich AGBs (Groenewegen 2006; Winters et al. 1997; Yasuda et al. 2012).


To comment on the idea raised in Norris et al. 2012 – that small grains may be present but do not contribute significantly to the signal – it is useful to be reminded of the relationship between grain size and polarisation. To demonstrate this we simulated 3D thin shells of Amorphous Enstatite using `miepython` and for each grain size computed the total contribution to scattering and polarised scattering cross sections. Whilst the scattered cross section  $\sigma_{\text{sca}}$  (which is the product of both the geometric cross section and the total scattering cross section) rapidly rises with grain size, the *polarised* cross section  $\sigma_{\text{pol}}$  is also a function of the scattering phase function (Figure 7.18). As grain size increases above the Rayleigh scattering threshold ( $\sim 100$  nm at 750 nm), the scattering function becomes increasingly anisotropic and dominated by strong (in direction anisotropy) forwards scattering that is comparatively weaker in contribution to polarised signal. The change in strength of polarised scattering can be visualised by plotting the scattering cross section  $\sigma_{\text{sca}}$  and polarised scattering cross section  $\sigma_{\text{pol}}$  against grain size (Figure 7.19).

The resulting function has a ‘flat top’ – indicating that in the presence of a uniform (in number) grain distribution, the emerging polarised signal would be dominated by grains between  $\sim 40$ -120 nm. In the case of Norris et al. 2012 and their fit of 300 nm grains to W Hya, from Figure 7.19 we can infer that a

distribution of grains around 300 nm *must* be narrow. If smaller grains between 50-120 nm were present in any significant fraction, they would rapidly down-weight the effective grain size. Given that we do not expect or have evidence to support the idea of a multi-modal grain distribution (clustered around perhaps  $\sim 10$  nm and  $\sim 300$  nm), this reasoning combined with the result of Norris et al. 2012 seem to support that a narrow grain size distribution is a feature of these inner circumstellar environments, and that in terms of the grain sizes we can measure with spectro-polarimetric interferometry, Norris et al. 2012's 300 nm grains may actually be a *lower limit* for the representative polarising grain size.




**Figure 7.18:** The relationship between grain size and scattering anisotropy – for Amorphous Enstatite grains observed at  $760 \pm 25$  nm. Each panel depicts the polarised scattering phase function as a function of observation angle. Here the source of light (incident) is travelling leftward from the  $180^\circ$  labels. The size of the lobes indicate relative polarisation strength and the orientation of the lobes inscribe the anisotropy of the polarised scattering function for that grain size. As grains increase in size towards the upper threshold for the Rayleigh scattering approximation ( $\sim 100$  nm at  $760$  nm), polarisation strength increases. However for grains larger than this threshold, the scattering function becomes increasingly anisotropic and dominated by forwards scattering. Strong forwards scattering contributes weaker polarisation than smaller grains – and so a decrease in polarised scattering is observed.



**Figure 7.19:** The relationship between grain size and polarised scattering cross section – for simple thin shells of single sized grains. Amorphous Enstatite was simulated at  $\lambda 760 \pm 25$  nm using  $n = 1.6$  and  $k = 0.0001$  (derived from optical constant data). `miepython` was used to compute the scattering cross sections and they were averaged over the scattering angles that constitute a 3D spherically symmetric shell.

With these concepts in mind we now examine the fitted grain distributions of all of our chemical options (Table 7.3), regardless of fit quality, as we found they yielded interesting insight on the appropriateness of the power-law distribution as a description of the grain number density in these inner circumstellar environments. Between chemicals there is considerable variability in the dust distribution widths fitted. Whilst we would anticipate variation in the best fit dust grain size distributions for each chemical, it was unexpected that some chemicals (eg. best fit Amorphous Enstatite Models 1 and 2) have widths of 2 nm, whilst others (eg. Amorphous Silica or Pyroxene Mg95) have widths of 200-340 nm. An explanation for

this variability may be found by scaling the fitted number density distributions by the polarised scattering cross section function depicted in Figure 7.19, which provides a first order estimate of how each grain size contributes to the measured signal. Figure 7.20 depicts the grain distributions fitted to each chemical and model (x axis), where the space between the minimum and maximum grain size (y axis) is filled with a gradient that is the fitted number density distribution multiplied by the polarised scattering cross section. This scaling demonstrates that whilst there is numeric variability in the distribution widths, typically  $\sim 80\%$  of the polarised signal is obtained from  $\sim 20 - 40$  nm of each distribution. It seems that *all* our distributions are narrow when weighted by their contribution to the polarised signal.




**Figure 7.20:** Dust grain distributions fitted in radiative transfer code to each chemical candidate and model with Bayesian Optimisation package `optuna`. For each model, the gradient between the minimum and maximum grain sizes (blue) is given by the fitted number density scaled by the polarised scattering cross section, which provides a first order indication of which of the fitted grains contribute most to the polarised signal. The  $\chi^2_{red}$  values for each model are plotted in forest green. The minimum grain size required to facilitate a scattering driven stellar wind is indicated by the dashed grey line (80 nm). The upper limit for the size of grains supporting a scattering driven stellar wind is 1250 nm and was eliminated within our inference.

These results (Figure 7.20) suggest that whilst our minimum grain sizes appear to be well constrained, the power law widths and exponents are not well constrained. We hypothesise both numerical and physical reasons for this. From a numerical perspective it is possible that this lack of constraint is due to the intrinsic signal to noise of our data, and in that we have attempted to fit 3 grain size distribution parameters with 4 wavelengths of data – the complexity of this prescription may have been overly ambitious. We may find that with improved signal to noise in our data (namely, the appropriate characterisation of systematic errors), the upper limits of these power-law distributions (which contribute weakly to signal) are better constrained. Alternatively, there may be too much covariance between the power law exponent and maximum grain size, which results in the fitted solution collapsing to a representative grain size  $a_{min}$ .

From a physical perspective there are several reasons that our fitted power-laws may not be well constrained. Firstly, a power law may not be well constrained given that all grain sizes do not contribute equally to the polarised scattered flux – as indicated by Figure 7.19. Both the number of grains and their relative contribution to the scattered signal will affect the fitted number density. Less polarising grains (here, those on the upper limits of our fitted grain distributions) will be harder to constrain. The second physical reason a power-law distribution may not be well constrained by our data is that it is *not the appropriate number density function*. Characterisations of number density distributions of circumstellar grains are limited – we could only find an explicit number density distribution for the silicon-carbon

grains of Carbon-rich AGBs (Yasuda et al. 2012). We attempted to reproduce this distribution and found it could not be reproduced with a power law, instead requiring a very narrow (at most  $\sigma = 0.02 \mu\text{m}$ ) distribution such as a log-normal with a cut off at  $\sim 1 \mu\text{m}$ . Figure 7.21 depicts the distribution of Yasuda et al. 2012 (right column) alongside our reproduction of this distribution with respect to grain number density (left column, black) and scaled grain number density (middle column, black). We also overlaid a representative power law distribution on this figure (blue lines, left and middle columns), which we manually aligned to best match the Yasuda et al. 2012 distribution (minimum grain size 30 nm and power law exponent of -2).



**Figure 7.21:** Our reproduction of the number density distribution for silicon-carbon dust from Yasuda et al. 2012. Column 3 depicts the distribution from Yasuda et al. 2012, column 1 depicts our reproduction of this distribution with respect to number density (black) and column 2 depicts our reproduced distribution on scaled axes (black). To approximate the distribution from Yasuda et al. 2012 we required a log-normal distribution ( $\mu = 0.03 \mu\text{m}$ ,  $\sigma = 0.02 \mu\text{m}$ ) with a cut off at  $\sim 1 \mu\text{m}$ . Overlaid onto columns 1 and 2 is a representative power law distribution that we manually aligned to best match the Yasuda et al. 2012 distribution (minimum grain size of 30 nm and power law exponent of -2).

Though carbon dust has different properties to silicon dust (namely enhanced absorption which will result in grain formation at much larger values of  $R_*$ ), it is interesting to consider that if the dynamics of grain formation are similar, silicon grains could also be described by a similar more symmetric distribution than a power-law. If this is the case, a power law is not the appropriate number density function, we would expect the maximum grain size to be poorly defined, and for the power law distribution’s minimum grain size to be fitted to the most strongly contributing grain size effectively collapsing the parametrisation to a ‘representative’ grain size.

The above reasoning – for both numeric and physical origins of our fitted solutions – suggest that it is reasonable to treat the outcome of our model fits in Table 7.3 as best providing representative grain sizes – equivalent to models fit in existing studies (Ireland et al. 2005; Norris et al. 2012; Hauboiss et al. 2019). Our results suggest that assessing the plausibility of narrow distributions more symmetric than the power law would be an interesting direction of future work. The radiative transfer code used within this analysis (MCFOST) (Pinte et al. 2006; Pinte et al. 2009) at present only treats power law distributions of grain sizes. Future package development could include flexible number density functionality to allow the testing of these ideas.

The MLE for the minimum grain size fitted to our best model (Enstatite (A) (2)) is  $120_{-53}^{+37}$  nm – large enough to support a scattering driven stellar wind from  $\mu$  Cephei. Uncertainty in this grain size is asymmetric on account of the logarithmic scaling used when fitting the parameter via Bayesian Optimisation. This value and its uncertainties support at least a  $1\sigma$  detection of grains supporting Hofner’s scattering driven stellar wind in the circumstellar environment of  $\mu$  Cephei.

## Mass Loss Estimates

From our best fit model we can estimate the mass loss rate. Given that the fitted grain size is large enough to support Hofner’s scattering driven stellar wind, we assume that all the dust in our model is lost from the star at the speed of a typical stellar wind with velocity  $19 - 35 \text{ km s}^{-1}$  (Montargès et al. 2019; Decin et al. 2023; De Beck et al. 2010; Le Borgne et al. 1989). We adopt an average of these values ( $27 \text{ km s}^{-1}$ )

and assume a gas/dust ratio of 200 as is typical for  $\mu$  Cephei and analogous red supergiants (Montargès et al. 2019; Mauron et al. 2011). Combining these assumptions with the geometric parameters fitted, we derive a total mass loss rate of  $10^{-5} M_{\odot} \text{ yr}^{-1}$ , which is in agreement with the upper range of literature values between  $10^{-7} - 10^{-5} M_{\odot} \text{ yr}^{-1}$  (Shenoy 2016; Gehrz 1971; Montargès et al. 2019).

### A missing piece of the puzzle?

We can make several statements from the analysis presented within this thesis:

1. When fitting to a single wavelength of our 2023 data using PIPPI ( $\lambda$  760 nm) (Chapter 5), and when examining our PIRATES image reconstructions (Figure 7.3 - 7.4), we were able to obtain fits to our observables that were  $\chi_{red}^2 < 1$ . This tells us that our geometric description of  $\mu$  Cephei’s environment is not a limitation to our modelling.
2. When fitting to multi-wavelength data from 2023 using PIPPI, we were also able to produce chemical fits with  $\chi_{red}^2 \sim 1$  – however, we subsequently demonstrated these models were un-physical when injected into radiative transfer code, an effect we identified to be caused by a lack of optical depth modelling.
3. When using fully physical radiative transfer, and an equivalent (or more detailed) geometry equivalent to (2) – *we were unable to obtain a  $\chi_{red}^2 < 1$ .*

These findings suggest that when fully physical radiative transfer code is used, and the wavelength dependence of our tested dust species is constrained to that of physically valid solutions – *that there remains a missing element in our model for wavelength dependence.*

We hypothesise several possible causes for this apparent limitation to our final models of  $\mu$  Cephei. Firstly, it is possible that whilst our Bayesian Optimisations converged, our estimation of the underlying posteriors could be improved. Variations to the stochastic convergence of the trials may cause variations to the  $\chi_{red}^2$  but are unlikely to cause the larger discrepancy of  $\Delta\chi_{red}^2 \sim 1$  between our best model in the present Chapter (which had  $\chi_{red}^2 = 2$ ) and our best PIPPI chemical solution (which had  $\chi_{red}^2 = 1$ ). Under the assumption that  $\chi_{red}^2 = 2$  does (approximately) represent the best obtainable model using radiative transfer code, we can hypothesise several origins of a ‘missing’ component of our present model’s wavelength dependence.

Firstly, we assumed that the chemical composition of our extended circumstellar feature (SW quadrant) was the same as that of the inner-circumstellar shell. The extended material is modelled as a clump located at  $\sim 61 \pm 5 \text{ mas}$  ( $\sim 6R_{\star}$ ). As simulated in Chapter 2, at these radii many of the iron-richer chemicals tested will now be cool enough to condense into solids. The presence of higher iron content at larger radii could contribute to the wavelength dependence of polarisation. However, given how faint this clump is compared to the inner circumstellar environment, it seems this would be unlikely to yield sufficient improvement to obtain a  $\chi_{red}^2 < 1$ .

The second option is that what we have modelled as a ‘density enhancement’ is actually an illumination enhancement – we outlined this as a possibility in Chapter 5. Large and persistent convective cells may cause heightened illumination in the circumstellar dust in the region above the cell. This would modify the wavelength dependence of our illumination feature – an illuminated region will have very different wavelength dependence to a density enhanced region with  $\tau \sim 3.5$ . We were not able to test this possibility in the present work as we do not have Stokes I observables from our VAMPIRES data and so could not confirm the presence of a hot-spot on the surface of the star.

The third possible cause, which we consider most likely, is that we have not modelled the effects of *limb polarisation* – polarisation from scattering in the extended photosphere of a star, originally predicted by (Harrington 1969). Since the prediction of limb polarisation, the injection of atmospheric models into radiative transfer has confirmed the significant contribution which the photosphere makes to total polarisation measurements (Kostogryz et al. 2015; Kostogryz et al. 2016; Wittkowski et al. 2016; Ohnaka et al. 2016; Gustafsson et al. 2008). Two studies have examined the contribution of atmospheric scattering to RSG Betelgeuse (Haubois et al. 2019; López Ariste et al. 2018), demonstrating that limb polarisation may contribute up to  $Q/I = 12\%$  at the limb at 550 nm (Haubois et al. 2019). At observing wavelengths

closer to our own (750 nm), values provided are between  $Q/I = 1 - 2\%$  for a surface gravity of  $\log g = 0.5$  or 0 (Haubois et al. 2019; López Ariste et al. 2018).

$\mu$  Cephei’s surface gravity has estimates between  $\log g = -0.36$  and 0 cgs (Kipper 2010; Josselin et al. 2007). We can estimate from these models of Betelgeuse that the *lower limit* of polarisation from the limb may be  $Q/I = 1 - 2\%$ . As modelled in Haubois et al. 2019 for Betelgeuse, we also expect that the molecular atmosphere around  $\mu$  Cephei will lower the total polarisation as a function of increased absorption which will dilute the coherence of polarisation in the escaping light (Turner 2014; Magalhães 1981; Coyne et al. 1979).  $\mu$  Cephei’s nominal stellar diameter at  $\lambda 750\text{nm}$  (21 mas) and the resolving power of VAMPIRES + NRM (longest baseline length has a resolution of 20 mas) should mean that limb polarisation is resolvable. At the present time no limb polarisation models for  $\mu$  Cephei exist, though we are currently involved in their development as a result of the findings presented within this Chapter. The integration of limb polarisation models with our dust modelling framework is a part of our ongoing work. <sup>2</sup>

## 7.6 Chapter Summary

This Chapter discussed the application of PIRATES to data from our observing campaign of  $\mu$  Cephei, providing 7 years of high spatio-temporal coverage of  $\mu$  Cephei’s inner circumstellar environment. Our image reconstructions provide the most detailed description of  $\mu$  Cephei’s inner circumstellar environment to date, and demonstrate the prevalence of asymmetry on all scales – adding to the growing body of observational evidence on such features in circumstellar environments. Our image reconstructions indicated that both a significant density asymmetry in the inner circumstellar environment and an extended clump are required to model  $\mu$  Cephei’s circumstellar dust.

By fitting to the parametrisation suggested by our PIRATES reconstructions in fully physical radiative transfer code we found that Amorphous Enstatite provided the best fit to our data by a slim margin. Our best fit model consisted of a circumstellar shell at  $1.6 R_*$ , with a large density enhancement of contrast  $\sim 8.7 \pm 0.3$  at  $\sim 267 \pm 26^\circ$  East of North, brighter than equivalent enhancements fitted to AGB stars. We interpreted this density enhancement in a similar fashion to previous studies – as evidence of a large clump of dust in the inner circumstellar environment. Our fitted model consists of dust with  $\tau \sim 1 - 3.5$ , which supports a growing body of evidence that circumstellar environments contain optically thick features that cannot be modelled without treatment of multiple scattering. Our model also included an extended clump in the South-Western quadrant, at a radial distance of  $61 \pm 5$  and an angular location of  $200 \pm 5^\circ$  East of North, co-located with material previously detected by Safonov et al. 2019 and within our own PDI reduction (see Chapter 5). The extended clump is much less dense than the inner material, with  $\tau \sim 0.15$  and a total mass of  $7.3_{-4.0}^{+8.9} \times 10^{-11}$ .

We found that use of radiative transfer code for chemical inference was able to resolve the degeneracy in grain size solutions obtained from simple scattering codes like PIPPI. Through radiative transfer analysis we demonstrated that the cause of better constraint was the modelling of optical depth – which not only changes the appearance of a model in the image plane, but changes the wavelength dependence of the polarisation signal. Critically, we demonstrated that use of simple scattering models can bias the outcome of grain size inference if simple models are fitted to targets that are not strictly optically thin – in this scenario, simple models may attempt to emulate optical depth by increasing the fitted grain size. Using radiative transfer code,  $\mu$  Cephei was fit with grains of size  $120_{-37}^{+53}$  nm, occupying a narrow distribution which we evaluated may have physical or numeric origins, both of which suggested it was appropriate to treat this grain size as a representative fit to the data. Amorphous Enstatite grains with radius of  $\sim 120$  nm are of appropriate size to support a Hofner scattering driven stellar wind from  $\mu$  Cephei. This analysis adds to a small but growing body of evidence that scattering is a critical element in the launch of winds from O-rich evolved stars – and that this mechanism may apply not only to the AGB stars for which it was originally construed, but to RSG stars as well. Our estimated mass loss rate is within the higher range but in agreement with values for  $\mu$  Cephei from the literature –  $10^{-5} M_\odot \text{ yr}^{-1}$ .

<sup>2</sup>Private communication with J.P Harrington, N Kostogryz, X Haubois. We were unable to obtain models for the limb polarisation of  $\mu$  Cephei during this candidature, but its proper treatment is a part of ongoing work and will be included in our subsequent paper.

However, we found that we were unable to obtain a  $\chi_{red}^2 \sim 1$  for our chemical inference in fully physical radiative transfer code (despite being able to do so with the simpler unphysical models of Chapter 5). Given that our analysis in the present chapter has tested the same pool of candidates for grain chemistry but with full physical constraint, this suggests that there is a component of the wavelength dependence of  $\mu$  Cephei missing from our modelling. Whilst there are no existing models at this time, we hypothesise that this missing element is limb polarisation, and suggest this to be a promising avenue for future studies of  $\mu$  Cephei's circumstellar environment.

# Chapter 8

## Conclusions and Future Work

### 8.1 Summary

This thesis aimed to advance the status of optical-wavelength spectro-polarimetric interferometry as a method for probing the inner circumstellar environments of evolved stars. Its primary scientific focus was RSG  $\mu$  Cephei, for which we aimed to constrain the circumstellar geometry, chemistry and dust grain size with a core objective of ascertaining the plausibility of Hofner’s scattering-driven stellar wind from  $\mu$  Cephei (Höfner 2008).

We have achieved this aim by advancing three core initiatives: (1) the development of novel instrumentation – VAMPIRES + NRM, (2) the creation of a suite of modelling and image reconstruction tools for analysing optical spectro-polarised interferometric observables – which have addressed the limitations of existing techniques – and (3) demonstrating the viability of these methods via their application to  $\mu$  Cephei. Subsequent constraint of the circumstellar material has yielded novel science outcomes underpinning the mechanisms that contribute to a scattering driven stellar wind.

The foundation of this research was the instrumentation work conducted in Chapter 3 – within which we have created facility level reduction software for the NRM mode of VAMPIRES. We upgraded the existing treatment of VAMPIRES + NRM data by integrating empirically fitted Muller matrix instrument models into software and modelling frameworks, permitting the disentanglement of instrumental polarisation and non-azimuthal science signal. Our engineering work – in particular the correction of a computing issue which improved observing efficiency by  $\sim 30\%$  and our archival data analysis – have improved observing with and interpretation of VAMPIRES + NRM data, which will provide a solid foundation for planned future science.

We developed two modelling tools for VAMPIRES + NRM, that may be applied more generally to all optical spectro-polarised interferometric observables. Firstly, we created PIPPI (Chapter 3), which is a differentiable and GPU accelerated version of the simple scattering models which constitute the existing treatment of this data type (Ireland et al. 2005; Norris et al. 2012; Haubois et al. 2019). The acceleration provided by PIPPI’s design permitted the testing of broader and more finely sampled circumstellar chemistry and grain size than has historically been computationally tractable. Secondly, we created PIRATES (Chapter 6) – the first image reconstruction algorithm for optical wavelength polarimetric interferometry data (Lilley et al. 2025b). PIRATES was designed to be a complementary technique to parametric model fitting and is trained on polarised images and observables generated by fully physical radiative transfer code. In our paper (Lilley et al. 2025b) we demonstrated that PIRATES may be applied to a broad range of astrophysically relevant targets, and can reconstruct high fidelity images whilst simultaneously optimising fits to the polarised interferometric observables.

The application of these two codes (PIPPI and PIRATES) to  $\mu$  Cephei and the comparison of their outcomes yielded several valuable insights which are helpful in guiding future modelling practices for this data type. Our results demonstrated that whilst simple models (PIPPI) are a very effective way to constrain circumstellar geometry, they are limited and potentially misleading if used to constrain the chemistry and grain size of circumstellar dust that is not *strictly* optically thin at all observational wavelengths. In Chapter 5 we applied PIPPI to fit  $\mu$  Cephei’s grain size – finding three degenerate grain

size regimes could equally well explain the data ( $\sim 250, 750, \text{above } 1200 \text{ nm}$ ). We then proceeded to re-analyse the simple scattering model approach by injecting the fitted PIPPI models into radiative transfer code. Doing so revealed that most of the fitted models were not physically valid solutions – most of the fitted dust grain sizes require an optical depth greater than 1 to produce polarised visibility amplitudes as large as  $\mu$  Cephei’s.

Subsequently, informed by the PIRATES image reconstructions of  $\mu$  Cephei, in Chapter 7 we fit to the geometry, chemistry and grain size of  $\mu$  Cephei’s circumstellar dust using fully physical radiative transfer code. We obtained a narrow distribution of best fit grain sizes around  $120_{-37}^{+53} \text{ nm}$  – a value significantly below even the smallest grain size fit with PIPPI (250 nm). To understand why use of radiative transfer code gave us a) a non-degenerate solution for grain size and b) a different solution for grain size to PIPPI, we analysed the radiative transfer models used to constrain  $\mu$  Cephei’s dust. Doing so demonstrated that the proper modelling of optical depth is responsible for explaining both points of difference. Increased optical depth changes both the appearance of the dust geometry in the image plane and the wavelength dependence of polarisation. These results demonstrate that if targets which are not strictly optically thin are fit with simple models like PIPPI, the results of grain size inference will be biased. Without strong prior information to support that a target has optically thin dust at all observational wavelengths, fully physical radiative transfer code should be used for chemical inference.

Throughout application of our developed modelling techniques we constrained several elements of  $\mu$  Cephei’s circumstellar geometry, providing the highest resolution spectro-spatial coverage of the inner circumstellar environment to date. We used PIRATES to reconstruct images of  $\mu$  Cephei over a 7-year period, consistently observing a strong polarised circumstellar shell and a persistent extended dust feature. The circumstellar shell is best fit at  $1.6 R_*$ , with a large density enhancement of contrast  $\sim 8.3 \pm 0.3$  at  $\sim 267 \pm 26^\circ$  East of North, brighter than equivalent enhancements fitted to AGB stars. We interpreted this density enhancement in a similar fashion to previous studies – as evidence of a large clump in the inner circumstellar environment. Our fitted model consists of dust with  $\tau \sim 1 - 3.5$ , which supports a growing body of evidence that circumstellar environments contain optically thick features that cannot be modelled without treatment of multiple scattering. We observed temporal evolution of the inhomogeneity within the inner circumstellar environments, on scales that could be explained by typical stellar wind speeds. To the extended environment we fitted a Gaussian clump in the South-Western quadrant, at a radial distance of  $61 \pm 5$  and an angular location of  $200 \pm 5^\circ$  East of North, angularly co-located with material previously detected by Safonov et al. 2019 and within our own PDI reduction (see Chapter 5). The extended clump is much less dense than the inner material, with  $\tau \sim 0.15$  and a total mass of  $7.3_{-4.0}^{+8.9} \times 10^{-11}$ .

Our radiative transfer fitting yielded that the best fit chemistry was Amorphous Enstatite, with grains narrowly distributed about  $\sim 120_{-37}^{+53} \text{ nm}$ . These grains are of appropriate size and transparency to facilitate Hofner’s scattering driven stellar wind, adding to a small but growing body of evidence that this may explain the deficit remaining from our theoretical understanding of RSG mass loss (Scicluna et al. 2015; Haubois et al. 2019).

## 8.2 Suggestions for Future Work

This research has revealed great scope for future work employing optical-wavelength spectro-polarised interferometry data on  $\mu$  Cephei and on other evolved targets – here we outline what we think are the most important avenues to advance instrumentation, modelling and the astrophysics of evolved dusty stars.

Firstly, we suggest that the next direction of instrumental characterisation is to model polarisation aberrations – which we estimated in this work with a multiplicative scaling factor. We suspect the biggest improvement to results (in particular to physical model constraint) will be yielded by understanding the effect that these highly spatially correlated errors have on VAMPIRES + NRM data.

Secondly, from the perspective of refining our best-fit model of  $\mu$  Cephei – as we outlined at the end of

Chapter 7 – we suggest that limb polarisation may play an important role in  $\mu$  Cephei’s circumstellar environment and in modelling the wavelength dependence of our data. Although models for  $\mu$  Cephei’s limb polarisation do not yet exist, our dust framework is now ready to integrate them for a more physically complete description of low surface gravity stars like  $\mu$  Cephei. We are currently involved in the development and integration of these models and will include them in a forthcoming publication (Lilley et al. 2025a). There are many exciting potential astronomical targets which our work can now be applied to, directly expanding the scope and relevance of our findings to wider classes of dust-producing stellar systems. Specific targets likely to yield excellent data with available observational resources include evolved stars Omicron Ceti and R Aquarii, each has been previously observed with VAMPIRES + NRM but (unfortunately) only in the faintest part of their variability cycles (Wong 2023; Charles 2022). More generally, we suggest that alternative targets with equivalently low surface gravities should be probed for limb polarisation, and that these efforts may lead to the spatially resolved detection of limb polarisation around an evolved star. Such an effort might be best targeted to dust-free objects so as to avoid ambiguity over the origin of polarised signals.

The research within this PhD focused on singular evolved stars, not addressing the role which obscured and un-detected binary companions play in the stellar mass loss process (Vermeulen et al. 2025; Smith 2014; Decin et al. 2020). The instrumentation discussed and modelling techniques developed would nicely lend themselves to research on this subject matter, and would be an interesting avenue of further research.

At the heart of our analysis is the outcome that fully physical and constrained modelling is required to capture the effects of optical depth – which plays a dominant role in  $\mu$  Cephei’s circumstellar environment. Our preliminary re-analysis of the models in Norris et al. 2012 suggest that data should be reanalysed now that detailed radiative transfer code and modelling techniques (like those described here-in) exist to facilitate modelling optical depth. Our analysis suggests that tools like PIPPI are best used for modelling the circumstellar *geometry*, but should not be used for modelling the circumstellar *chemistry or grain size* without a strong prior that the circumstellar dust is optically thin. In cases where this prior does not exist, chemistry and grain size inference should be performed in radiative transfer code.

Within the present work we fitted radiative transfer models using a Bayesian Optimisation algorithm, as it provided an efficient but thorough way to sample the parameter spaces of interest given the computational expense of radiative transfer models. Fitting continuously to parameters like the circumstellar grain size is ideal so long as numerical convergence is guaranteed. A stronger method (which is not computationally tractable at the present time) is to fix the grain distribution and then constrain all other parameters under the assumption of fixed grains. This was possible to do with PIPPI given GPU acceleration but is not currently computationally tractable to employ with radiative transfer codes like MCFOST. However, the computational tractability of such an approach may be on the horizon given the rapid and exciting development of faster radiative transfer codes like RAD-Jax which obtain a speed up of  $10^4$  compared to existing methods (Levis et al. 2025). Although RAD-Jax does not yet support scattering, its implementation would make more rigorous hypothesis testing of grain size in full radiative transfer code computationally tractable.

When fitting to circumstellar grain size using a power law distribution we found that parameters for the distribution width and exponent were not well constrained, and that it was most appropriate to interpret the minimum fitted grain size as a representative grain size for the dust. Whilst we hypothesised numerical and physical origins for this lack of constraint, our results may also indicate that narrow grain distributions could provide better characterisation of the grains within inner circumstellar environments than power law distributions. Adding flexible grain distribution functionality to MCFOST would permit the investigation of this possibility.

One idea that we pursued during our use of PIPPI was the viability of performing ‘volume reconstruction’ of circumstellar environments using the three-dimensional information yielded by polarimetry. This idea stemmed from the need for physical consistency between polarised images, and the difficulty of regularising images that are not positive definite (which Stokes Q and U are not). Implemented in `jax`, this algorithm computed the gradients of every 3D voxel with respect to the fit to observables and a volume-based regularisation term. To make this computationally tractable we were required to assume the scene was optically thin and multiply the volume by the pre-computed scattering grids we described

in Chapter 4. Whilst we found this technique worked, it suffered from the same limitations we identified in PIPPI in Chapter 5. However, if packages like RAD-Jax have scattering functionality in the near future, reconstructing circumstellar volumes in fully physical radiative transfer will be computationally tractable and would be a very interesting and promising direction of future modelling development.

# Chapter A

## Appendix Material

This chapter contains supplementary Appendix material for the following chapters:

For Chapter 2:

- Optical constant data collated and used in all chemical modelling (Figures [A.1](#) - [A.3](#)).

For Chapter 3:

- Equations for reducing the triple differential visibilities and closure phases (Equations [A.1](#) - [A.8](#)).

For Chapter 5:

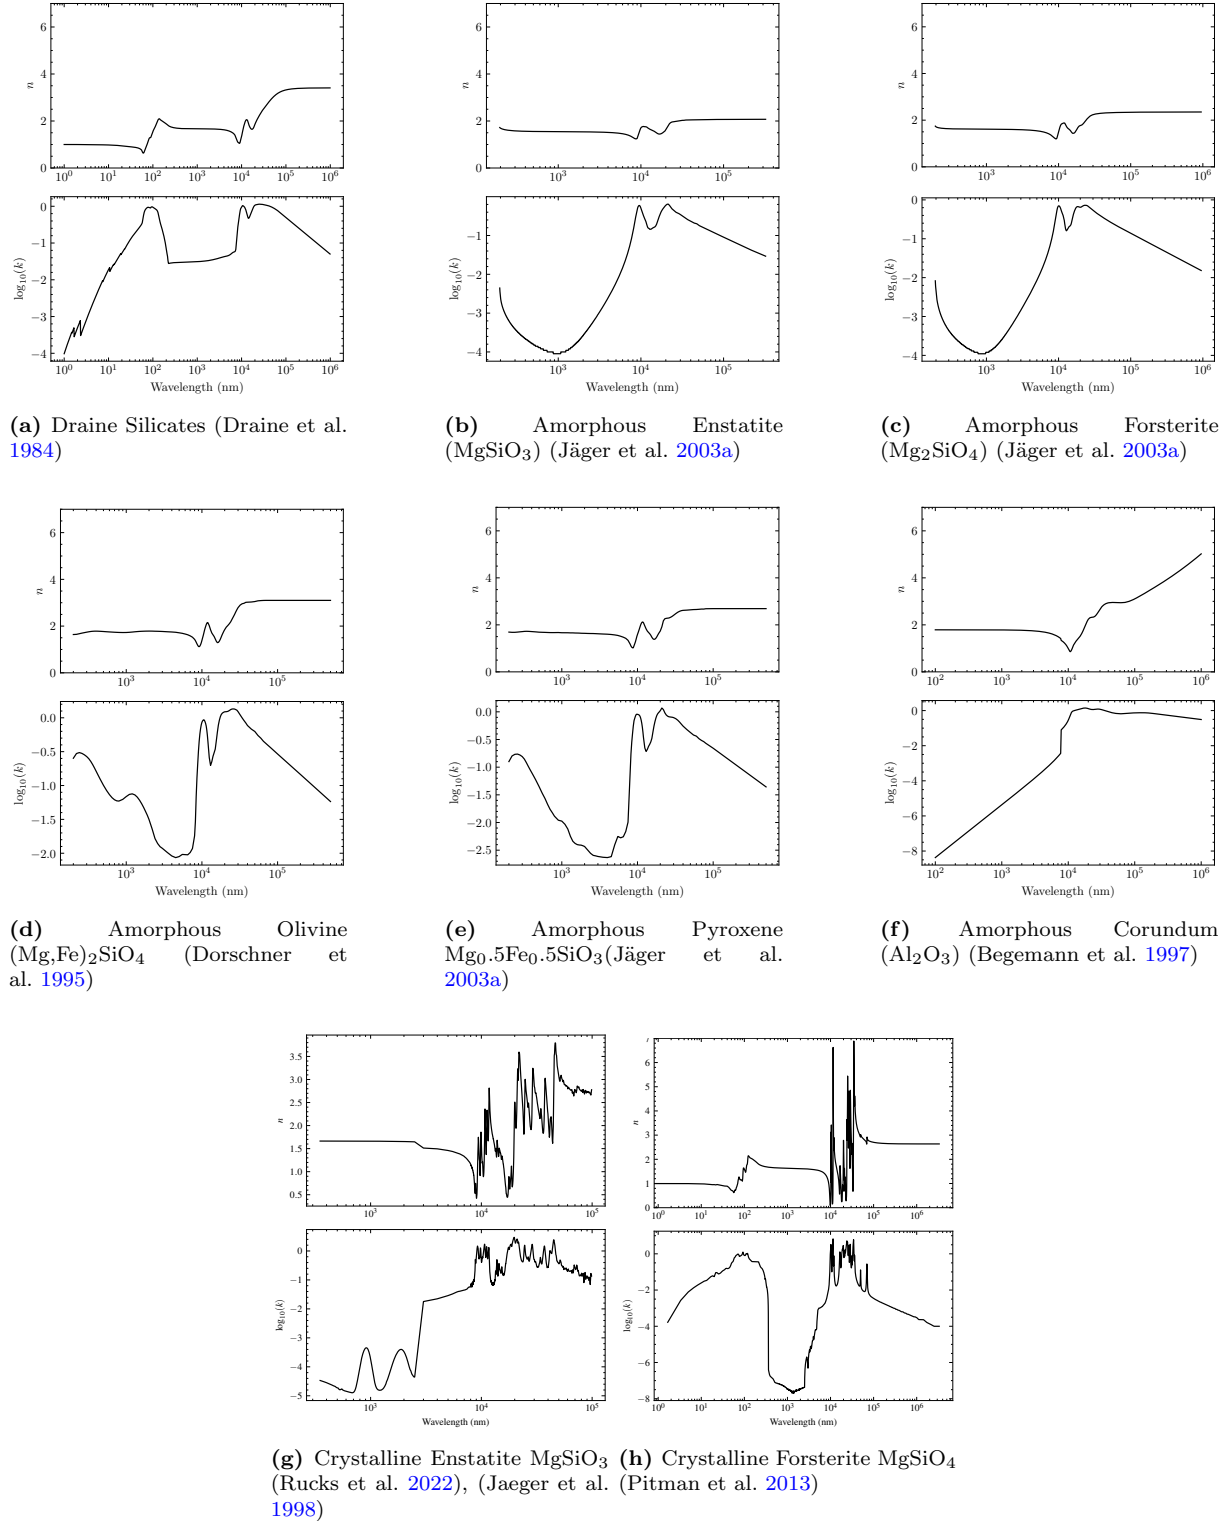
- MCMC Hyper-parameters (Table [A.1](#)).
- MCMC priors, initial conditions (Tables [A.2](#) and [A.3](#)).
- Corner plots for fitted models (Figures [A.4](#) - [??](#))

For Chapter 7:



- EWC scans for application of PIRATES to real  $\mu$  Cephei data (Figures [A.23](#) - [A.24](#)).
- Bayesian Optimisation Posteriors for each fitted chemical (Figures [A.25](#) - [A.33](#))

## A.1 Chapter 2


### A.1.1 Complex Refractive Indices $m(\lambda)$




**Figure A.1:** Optical constants for various dust species used in radiative transfer models. Data obtained from Jena University Optical Constants Database. Where species have a crystalline form, the contributions from all three optic axes are ‘mixed’.

(a) Amorphous Silica  $\text{SiO}_2$  (Franta et al. 2016)


(b) Crystalline Spinel (Zeidler et al. 2013), (Chernova et al. 2017)




(c) Olivine Mg/Fe 95/5 (Jäger et al. 1994)




(d) Olivine Mg/Fe 80/20 (Jäger et al. 1994)



(e) Olivine Mg/Fe 70/30 (Jäger et al. 1994)




(f) Olivine Mg/Fe 60/40 (Jäger et al. 1994)



(g) Pyroxene Mg/Fe 95/5 (Jäger et al. 2003a) (h) Pyroxene Mg/Fe 80/20 (Jäger et al. 2003a)

**Figure A.2:** Optical constants for various dust species used in radiative transfer models. Data obtained from Jena University Optical Constants Database. Where species have a crystalline form, the contributions from all three optic axes are ‘mixed’.



(a) Pyroxene Mg/Fe 60/40 (Jäger et al. 2003a)

(b) Pyroxene Mg/Fe 70/30 (Jäger et al. 2003a)

(c) Crystalline Corundum (Querry 1985)

**Figure A.3:** Optical constants for various dust species used in radiative transfer models. Data obtained from Jena University Optical Constants Database. Where species have a crystalline form, the contributions from all three optic axes are ‘mixed’.

## A.2 Chapter 3

HN denotes the HWP state (0,22.5,45,67.5) FN denotes the FLC state (A,B) and CN denotes the camera state (1,2).

$$\begin{aligned}
 Q_{\text{vis},1} &= \frac{H0\_FA\_C1,\text{vis}}{H0\_FA\_C2,\text{vis}} & Q_{\text{vis},2} &= \frac{H0\_FB\_C1,\text{vis}}{H0\_FB\_C2,\text{vis}} & Q_{\text{vis},3} &= \sqrt{\frac{Q_{\text{vis},1}}{Q_{\text{vis},2}}} \\
 Q_{\text{vis},4} &= \frac{H45\_FA\_C1,\text{vis}}{H45\_FA\_C2,\text{vis}} & Q_{\text{vis},5} &= \frac{H45\_FB\_C1,\text{vis}}{H45\_FB\_C2,\text{vis}} & Q_{\text{vis},6} &= \sqrt{\frac{Q_{\text{vis},4}}{Q_{\text{vis},5}}}
 \end{aligned} \tag{A.1}$$

$$Q_{\text{final, vis}} = \sqrt{\frac{Q_{\text{vis},3}}{Q_{\text{vis},6}}} \tag{A.2}$$

$$\begin{aligned}
 U_{\text{vis},1} &= \frac{H22.5\_FA\_C1,\text{vis}}{H22.5\_FA\_C2,\text{vis}} & U_{\text{vis},2} &= \frac{H22.5\_FB\_C1,\text{vis}}{H22.5\_FB\_C2,\text{vis}} & U_{\text{vis},3} &= \sqrt{\frac{U_{\text{vis},1}}{U_{\text{vis},2}}} \\
 U_{\text{vis},4} &= \frac{H67.5\_FA\_C1,\text{vis}}{H67.5\_FA\_C2,\text{vis}} & U_{\text{vis},5} &= \frac{H67.5\_FB\_C1,\text{vis}}{H67.5\_FB\_C2,\text{vis}} & U_{\text{vis},6} &= \sqrt{\frac{U_{\text{vis},4}}{U_{\text{vis},5}}}
 \end{aligned} \tag{A.3}$$

$$U_{\text{final, vis}} = \sqrt{\frac{U_{\text{vis},3}}{U_{\text{vis},6}}} \tag{A.4}$$

$$\begin{aligned}
 Q_{\text{CP},1} &= H0\_FA\_C1,\text{CP} - H0\_FA\_C2,\text{CP} & Q_{\text{CP},4} &= H45\_FA\_C1,\text{CP} - H45\_FA\_C2,\text{CP} \\
 Q_{\text{CP},2} &= H0\_FB\_C1,\text{CP} - H0\_FB\_C2,\text{CP} & Q_{\text{CP},5} &= H45\_FB\_C1,\text{CP} - H45\_FB\_C2,\text{CP} \\
 Q_{\text{CP},3} &= \frac{Q_{\text{CP},1} - Q_{\text{CP},2}}{2} & Q_{\text{CP},6} &= \frac{Q_{\text{CP},4} - Q_{\text{CP},5}}{2}
 \end{aligned} \tag{A.5}$$

$$Q_{\text{final, CP}} = \frac{Q_{\text{CP},3} - Q_{\text{CP},6}}{2} \tag{A.6}$$

$$\begin{aligned}
 U_{\text{CP},1} &= H22.5\_FA\_C1,\text{CP} - H22.5\_FA\_C2,\text{CP} & U_{\text{CP},4} &= H67.5\_FA\_C1,\text{CP} - H67.5\_FA\_C2,\text{CP} \\
 U_{\text{CP},2} &= H22.5\_FB\_C1,\text{CP} - H22.5\_FB\_C2,\text{CP} & U_{\text{CP},5} &= H67.5\_FB\_C1,\text{CP} - H67.5\_FB\_C2,\text{CP} \\
 U_{\text{CP},3} &= \frac{U_{\text{CP},1} - U_{\text{CP},2}}{2} & U_{\text{CP},6} &= \frac{U_{\text{CP},4} - U_{\text{CP},5}}{2}
 \end{aligned} \tag{A.7}$$

$$U_{\text{final, CP}} = \frac{U_{\text{CP},3} - U_{\text{CP},6}}{2} \tag{A.8}$$

## A.3 Chapter 5

### A.3.1 MCMC Hyper-parameters, Priors, Initial Conditions

**Table A.1:** MCMC steps, walkers and burn in discard

Model	Steps	Walkers	Burn in
A	5000	6	2500
B	5000	9	2000
C	5000	9	2000
D	8000	12	3000
E	8000	12	3000
F	9000	15	3000
G	9000	15	3000
H	8000	12	3000
I	13000	18	3000
J	9000	15	3000
K	11000	21	3000
L	7000	8	2000
M	12000	24	3000
N	13000	18	3000
O	12000	24	3000
P	12000	24	3000
Q	15000	30	5000
R	12000	24	3000
S	15000	30	5000

**Table A.2:** Model Priors and Initial Conditions – Rayleigh Scattering with PIPPI (I)

Model	$p_0/[p_{min}, p_{max}]$	a	$C_D$	$S_T$	$pl$	$a_2$	$S_C$	$\alpha$	b	h	k
A	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12,200]	4 [0.01,10]	- -	- -	- -	- -	- -	- -	- -	- -
B	$p_0$ [ $p_{min}, p_{max}$ ]	15 [1,200]	4 [0.01,10]	15 [1,150]	- -	- -	- -	- -	- -	- -	- -
C	$p_0$ [ $p_{min}, p_{max}$ ]	15 [1,200]	4 [0.01,10]	- -	3 [1, 10]	- -	- -	- -	- -	- -	- -
D	$p_0$ [ $p_{min}, p_{max}$ ]	15 [1,200]	4 [0.01,10]	- -	- -	15 [12, 200]	2 [0.01, 50]	- -	- -	- -	- -
E	$p_0$ [ $p_{min}, p_{max}$ ]	15 [1,200]	4 [0.01,10]	- -	- -	- -	0 [-25, 25]	15 [12, 200]	- -	- -	- -
F	$p_0$ [ $p_{min}, p_{max}$ ]	15 [1,200]	4 [0.01,10]	3 [1,6]	- -	- -	0 [-25, 25]	15 [12, 200]	- -	- -	- -
G	$p_0$ [ $p_{min}, p_{max}$ ]	15 [1,200]	4 [0.01,10]	- -	- -	- -	0 [-25, 25]	15 [12, 200]	- -	- -	- -
H	$p_0$ [ $p_{min}, p_{max}$ ]	15 [1,200]	4 [0.01,10]	- -	- -	- -	- -	- -	- -	0 [-25, 25]	0 [-25, 25]
I	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12,200]	4 [0.01,10]	- -	- -	- -	- -	0 [-25, 25]	15 [12, 200]	0 [-25, 25]	0 [-25, 25]

**Table A.3:** Model Priors and Initial Conditions – Rayleigh Scattering with PIPPI (II)

Model	$p_0/[p_{min}, p_{max}]$	a	$C_D$	$\alpha$	b	$\phi_b$	$r_b$	$C_B$	h	k	$E_a$	$E_l$
J	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	- -	- -	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	- -	- -
K	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	0 [-25, 25]	15 [12, 200]	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	- -	- -
L	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	- -	- -	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	- -	- -
M	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	0 [-25, 25]	15 [12, 200]	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	- -	- -
N	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	- -	- -	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	- -	- -
O	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	0 [-25, 25]	15 [12, 200]	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	- -	- -
P	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	- -	- -	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	0 [-9,-9]	0 [-9,-9]	- -	- -
Q	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	0 [-25, 25]	15 [12, 200]	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	0 [-9,-9]	0 [-9,-9]	- -	- -
R	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	- -	- -	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	2 [0.01, 15]	270 [90, 360]
S	$p_0$ [ $p_{min}, p_{max}$ ]	15 [12, 200]	4 [0.01, 10]	0 [-25, 25]	15 [12, 200]	$2/3\pi$ [ $\pi, 2\pi$ ]	120 [0, 500]	0.5 [0.01, 500]	- -	- -	2 [0.01, 15]	270 [90, 360]

This page is intentionally left blank.

## A.3.2 Corner Plots

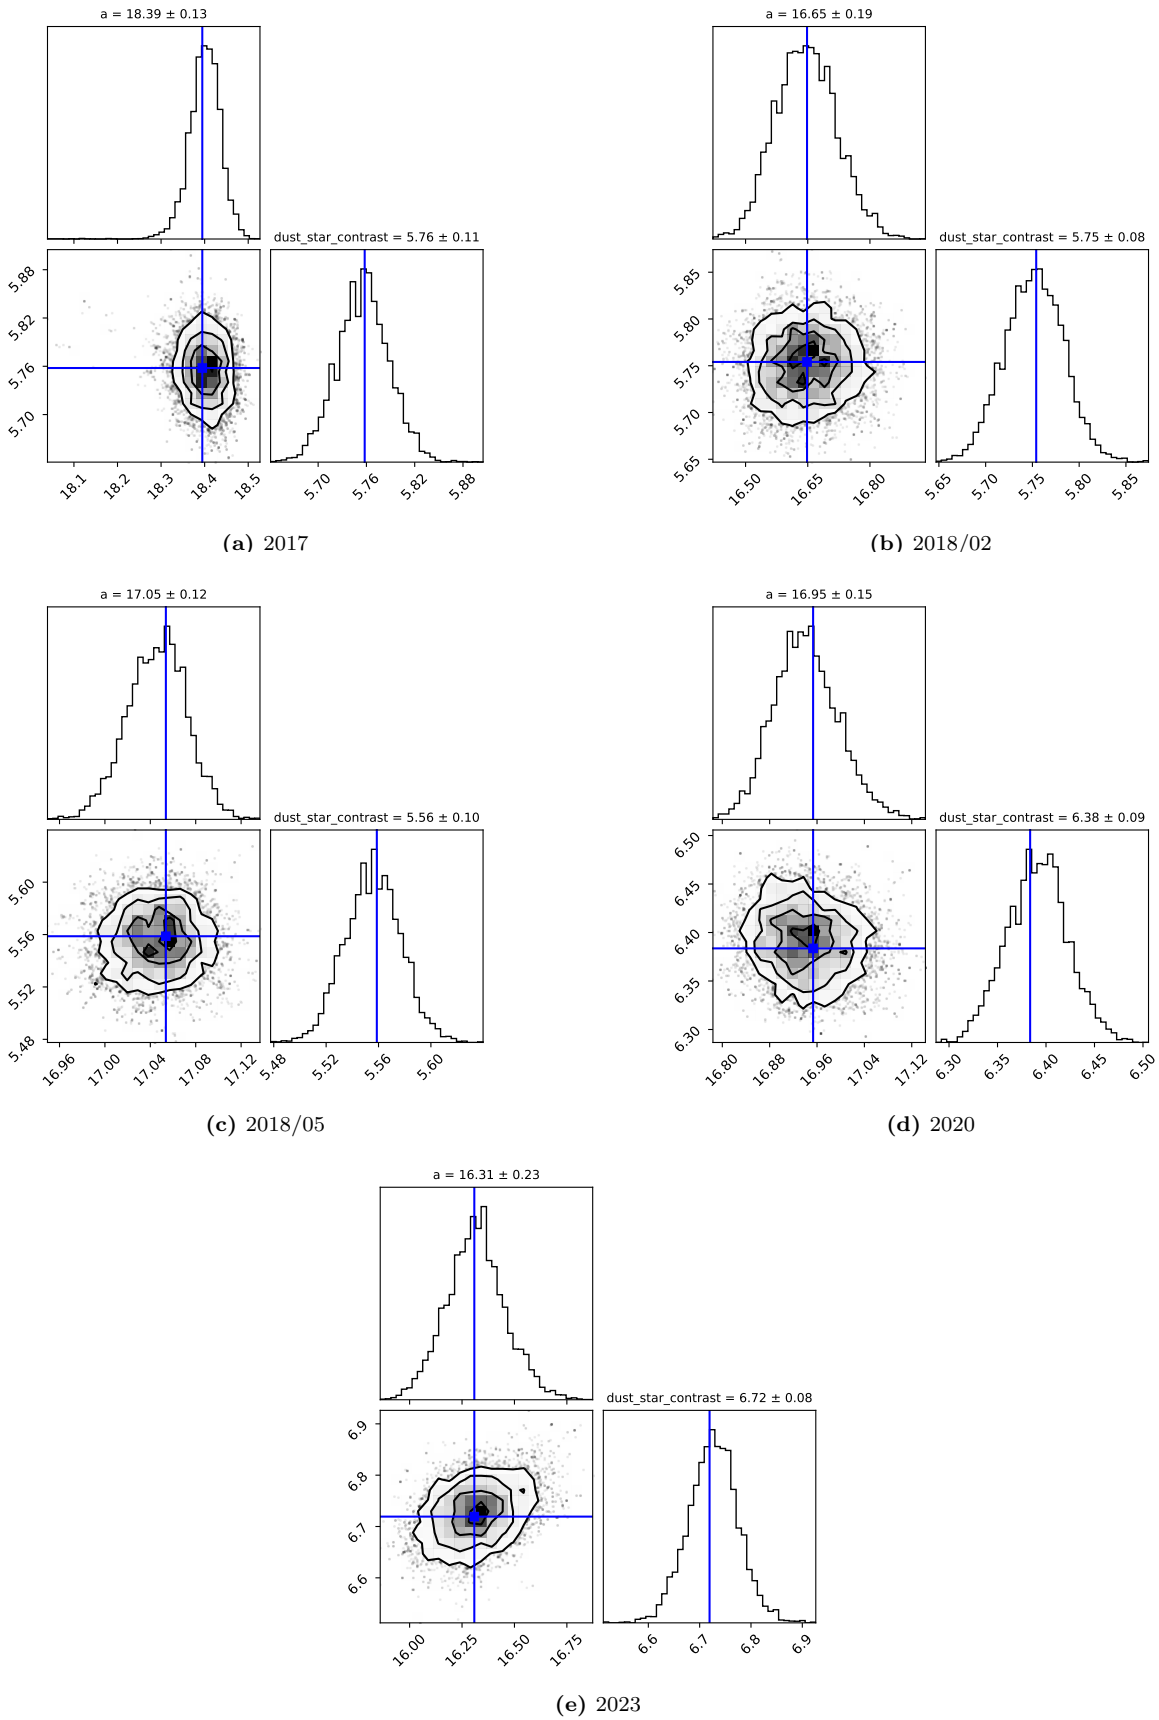


Figure A.4: Corner plots for Model Class A — thin spherical shell.

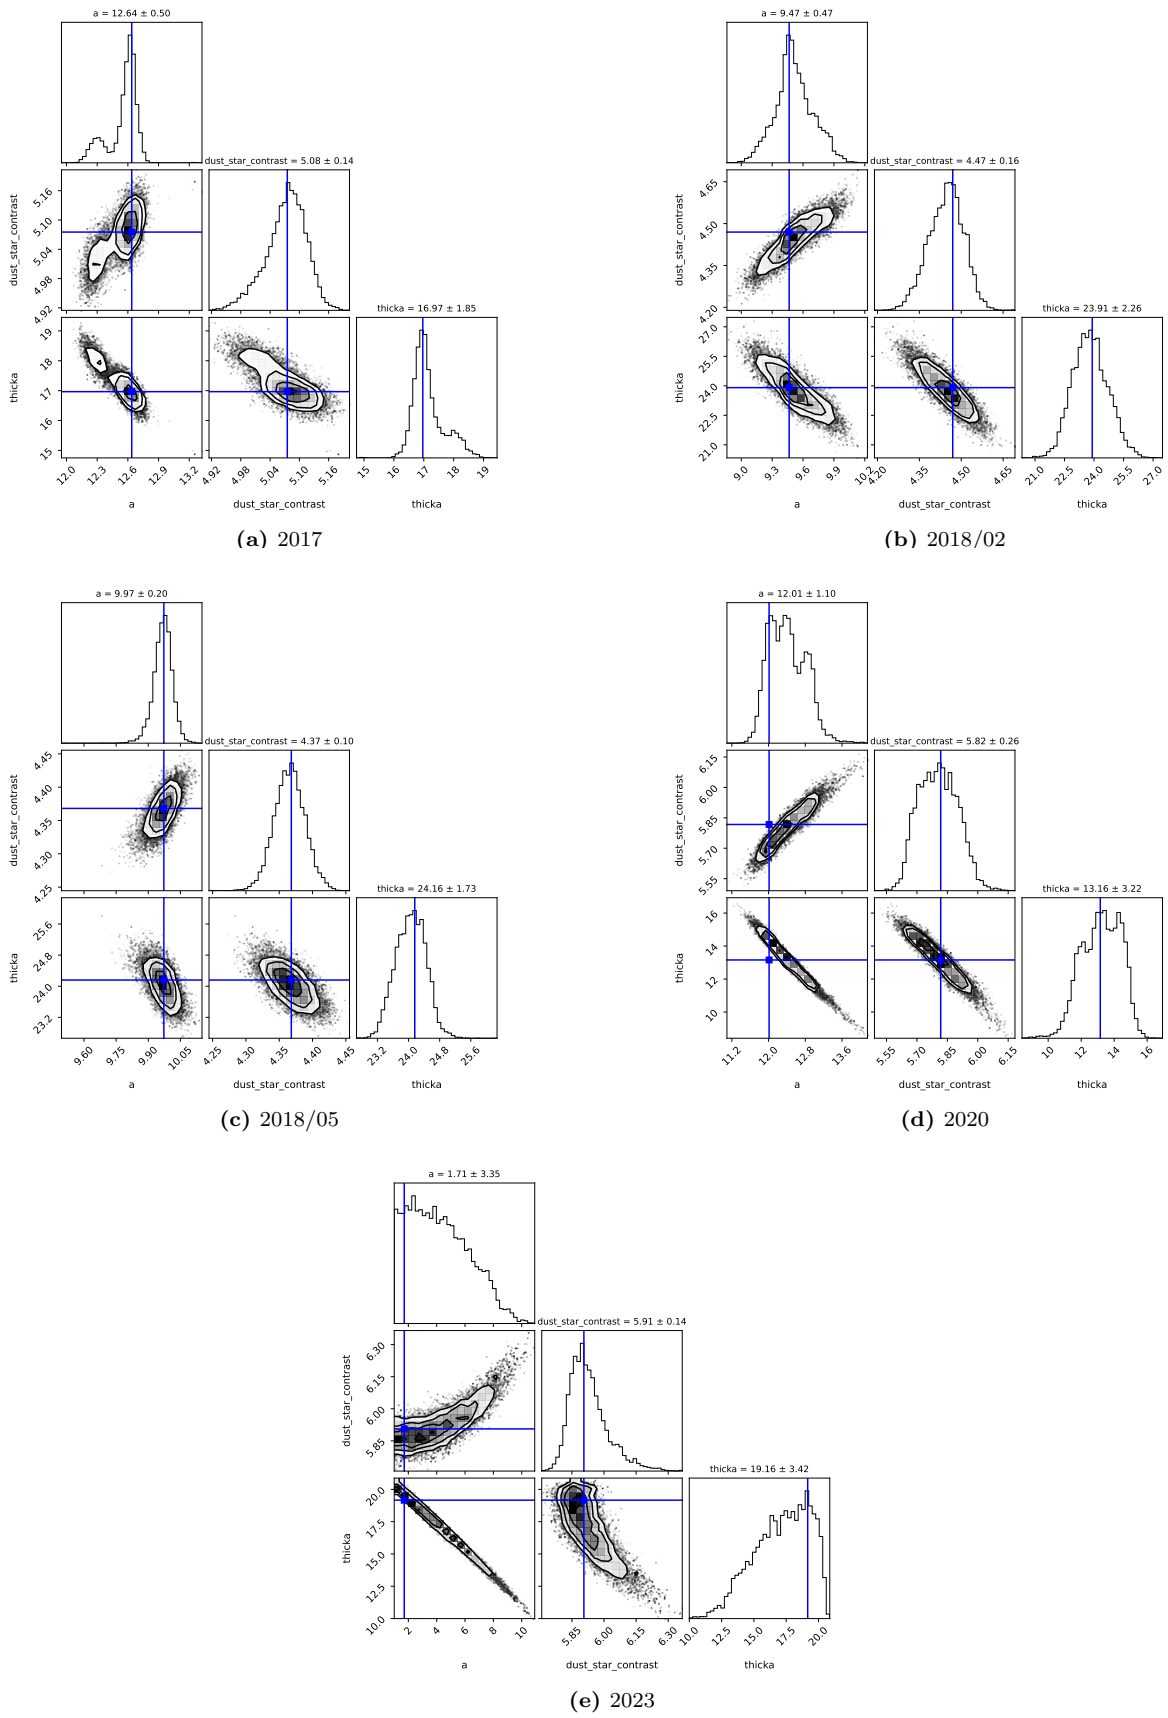


Figure A.5: Corner plots for Model Class B — thick spherical shell.

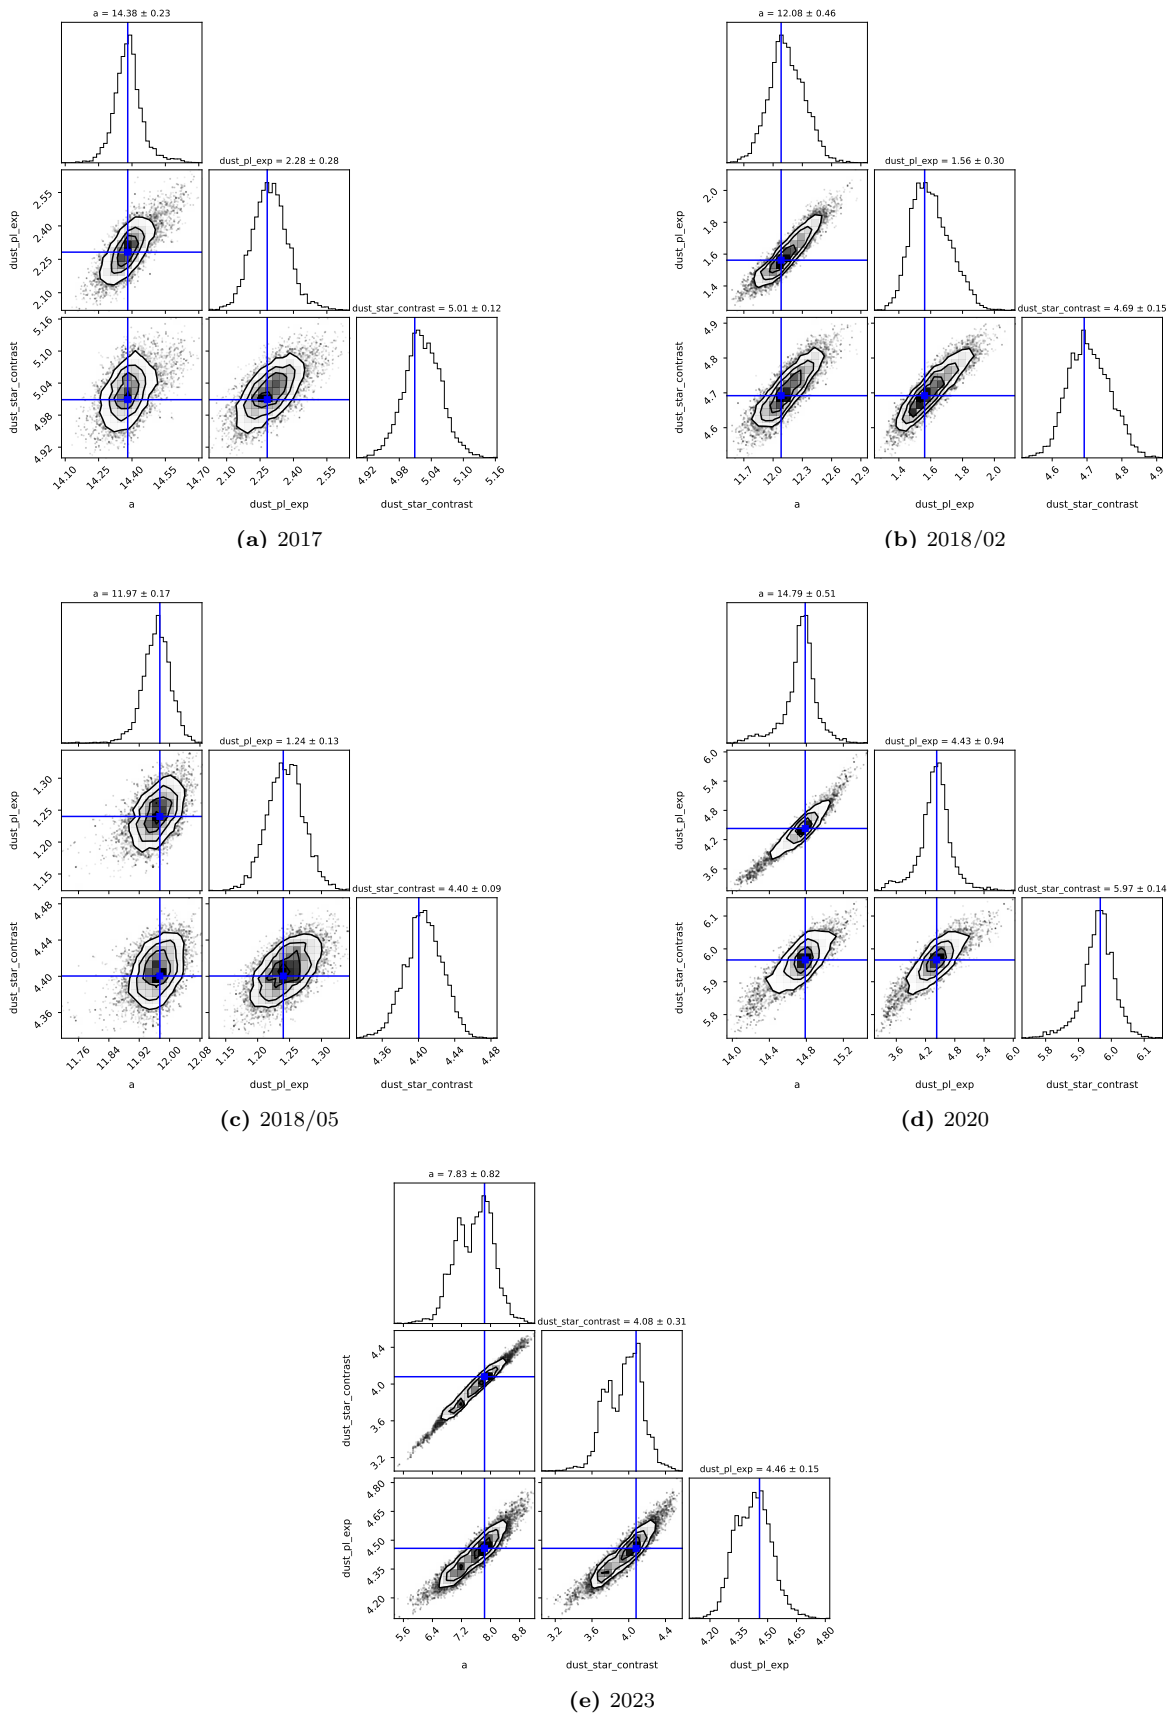
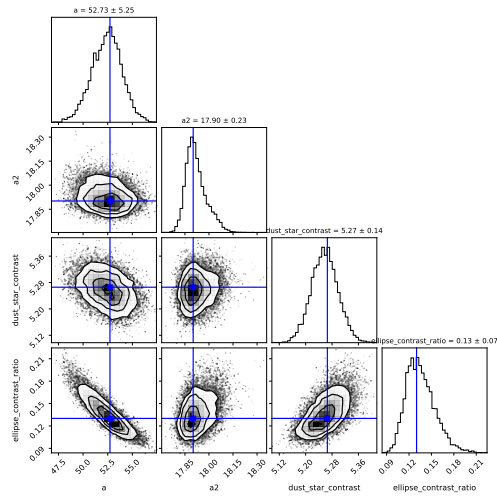
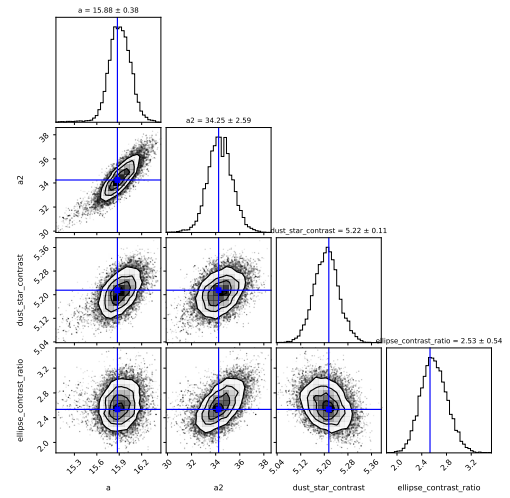


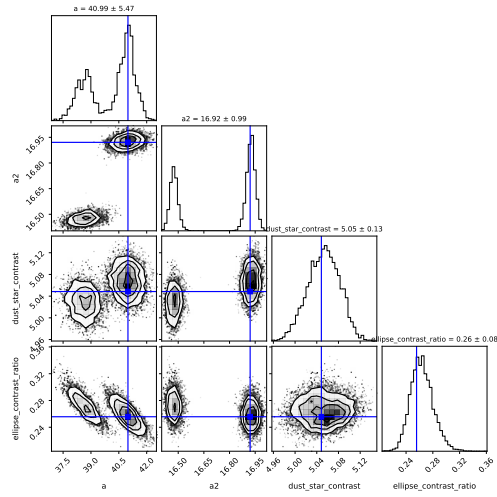
Figure A.6: Corner plots for Model Class C — power-law spherical shell.



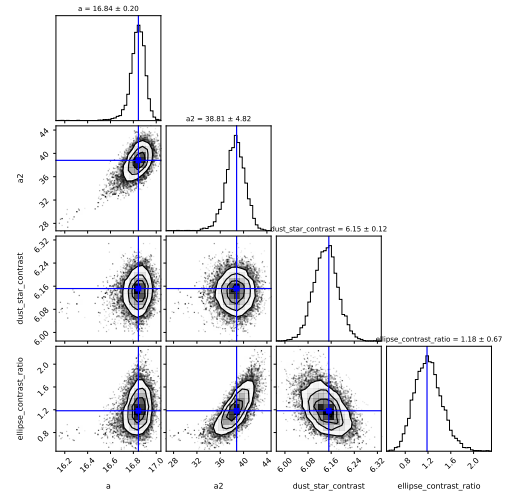
(a) 2017



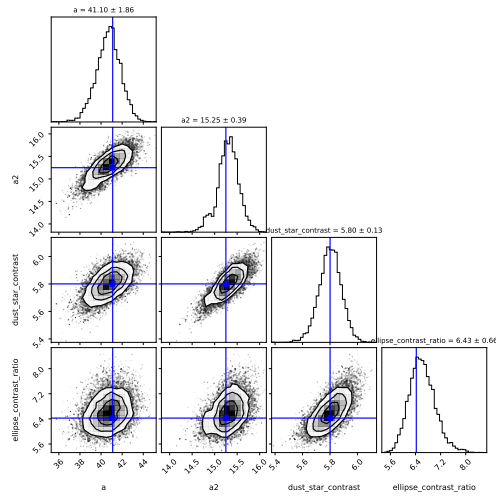
(b) 2018/02



(c) 2018/05



(d) 2020



(e) 2023

Figure A.7: Corner Plots for Model Class D — Two thin spherical shells.

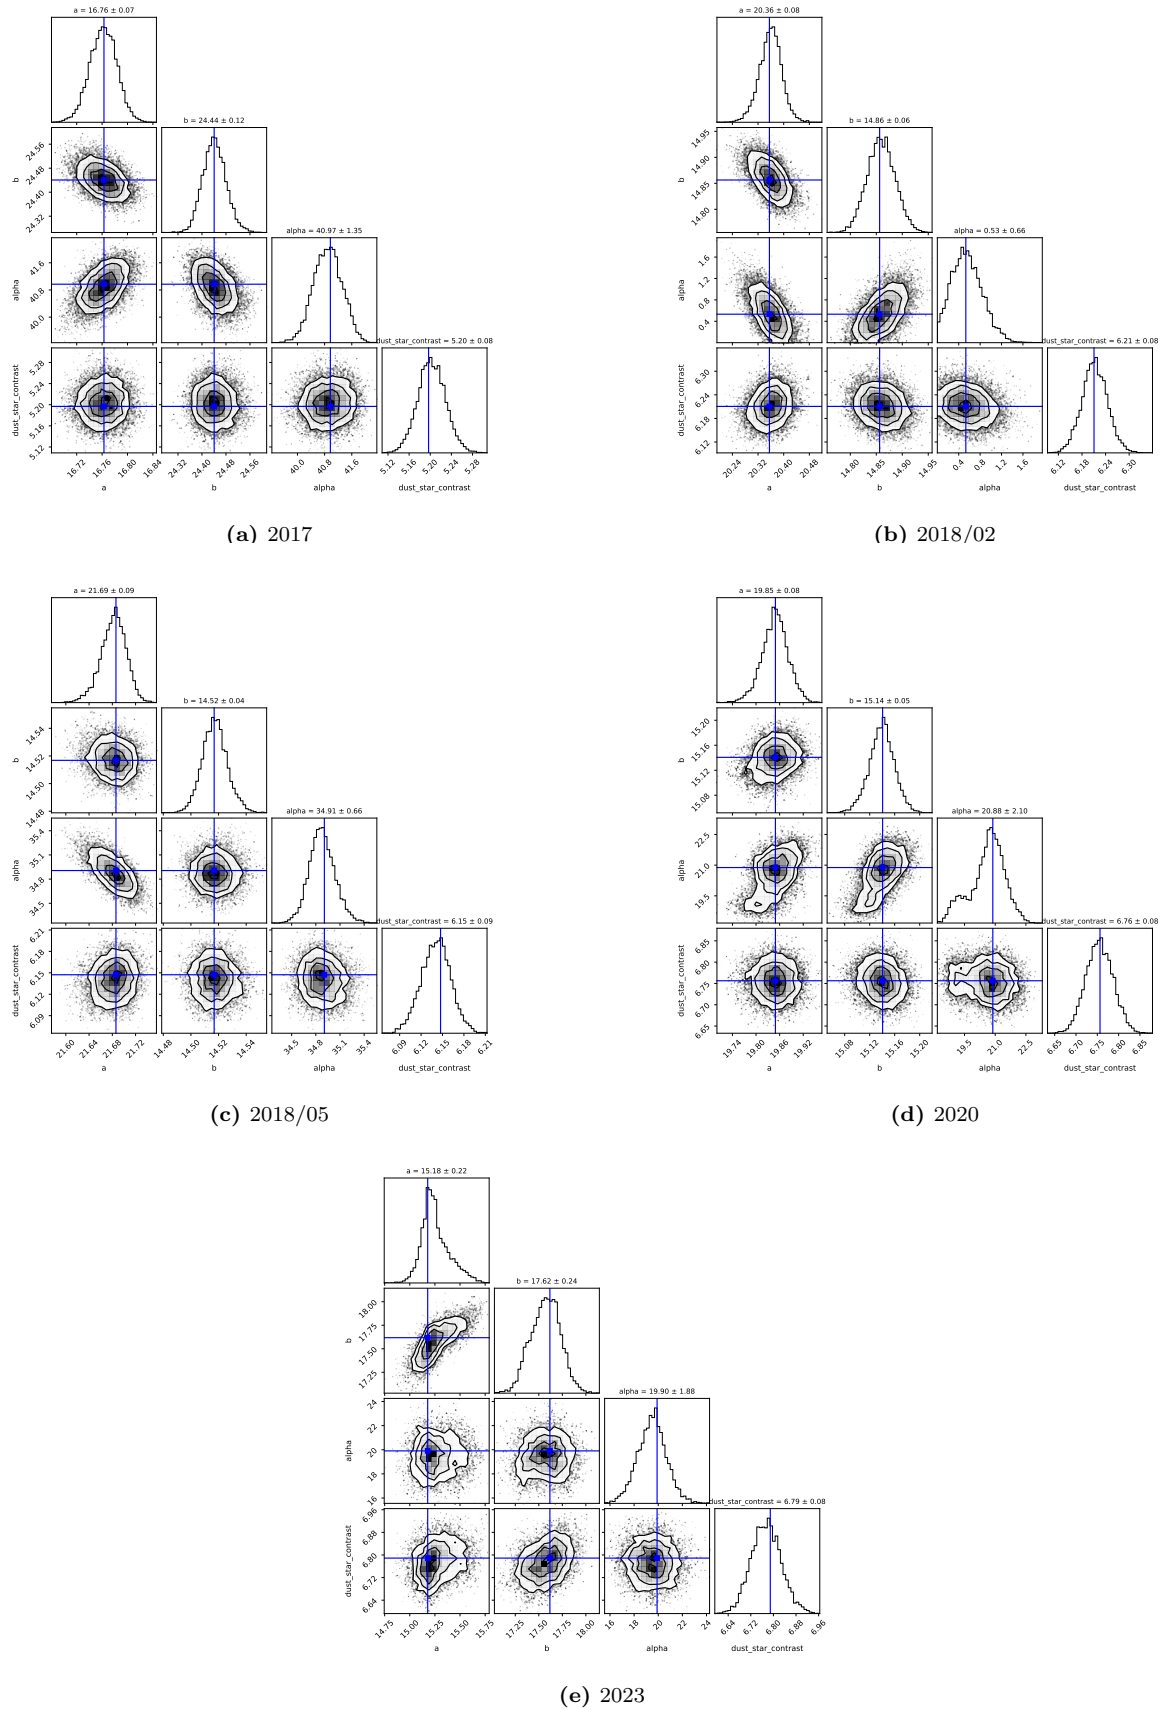
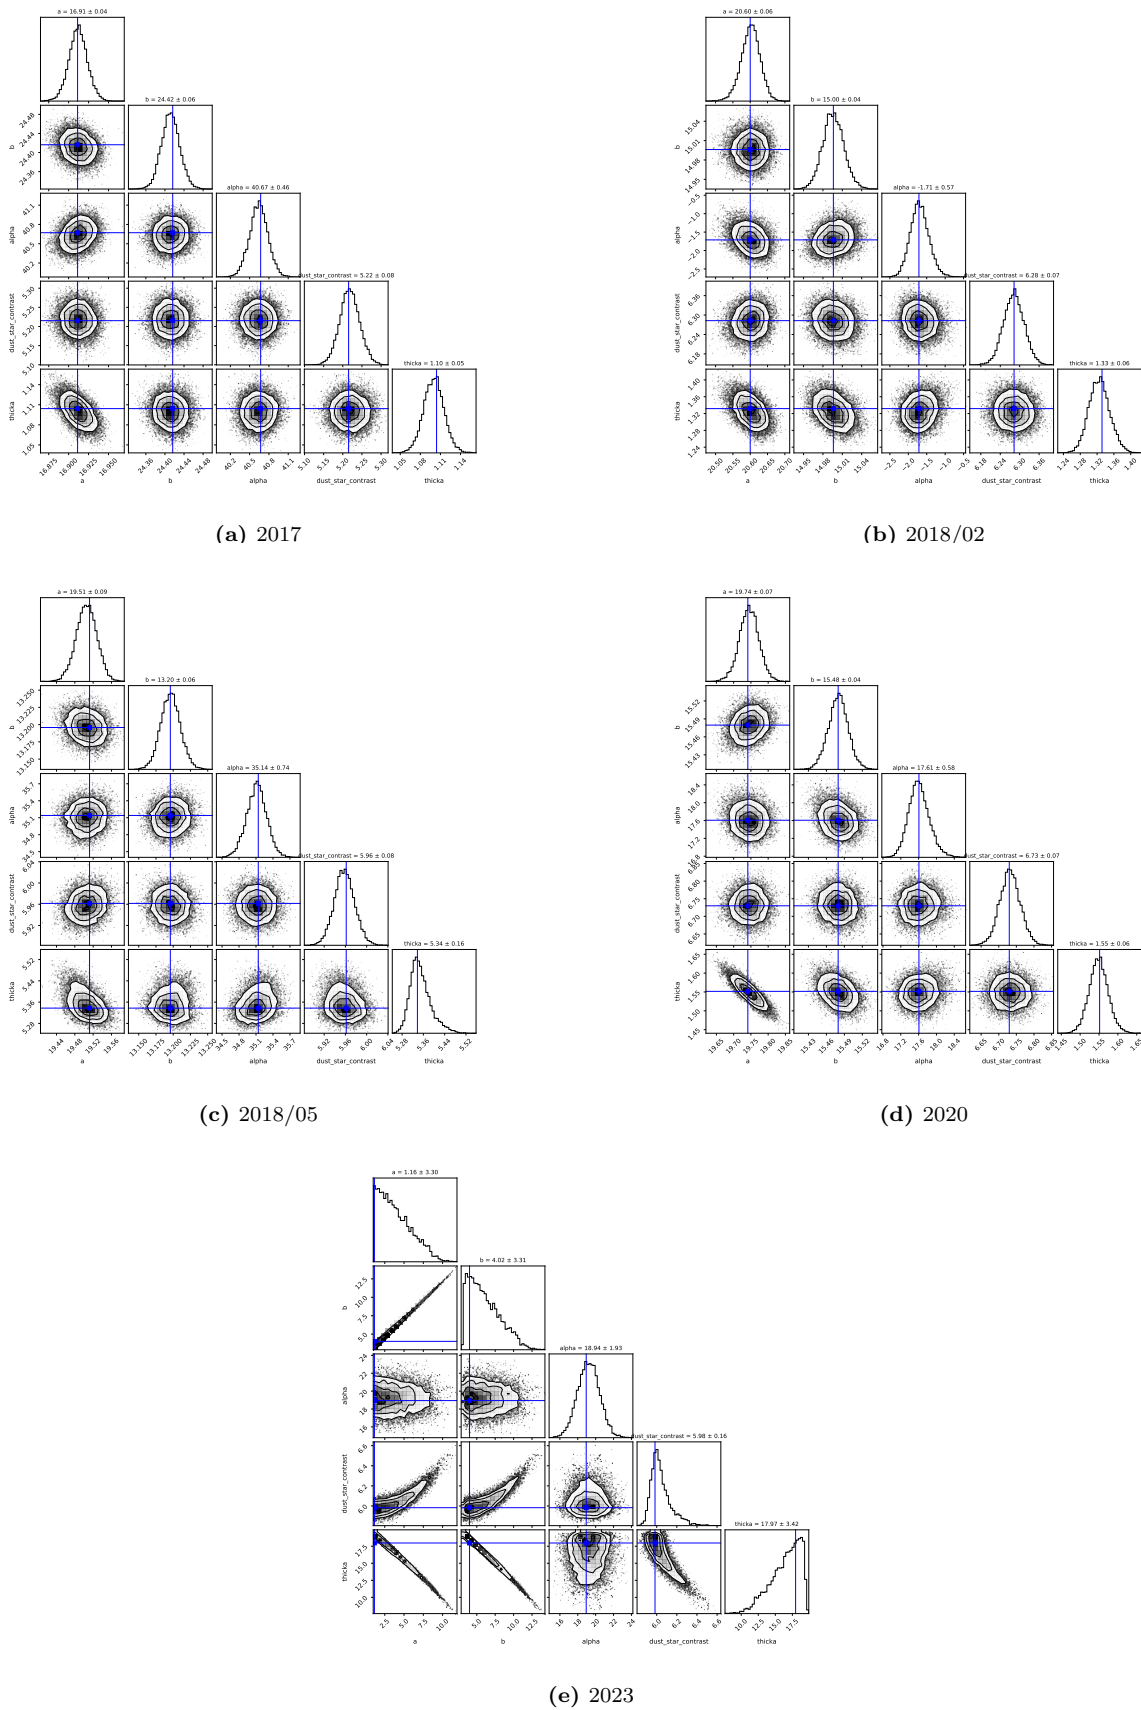


Figure A.8: Corner Plots for Model Class E — Thin elliptical shell.



**Figure A.9:** Corner Plots for Model Class F — Thick elliptical shell. Elliptical thickness constrained to 1–6.

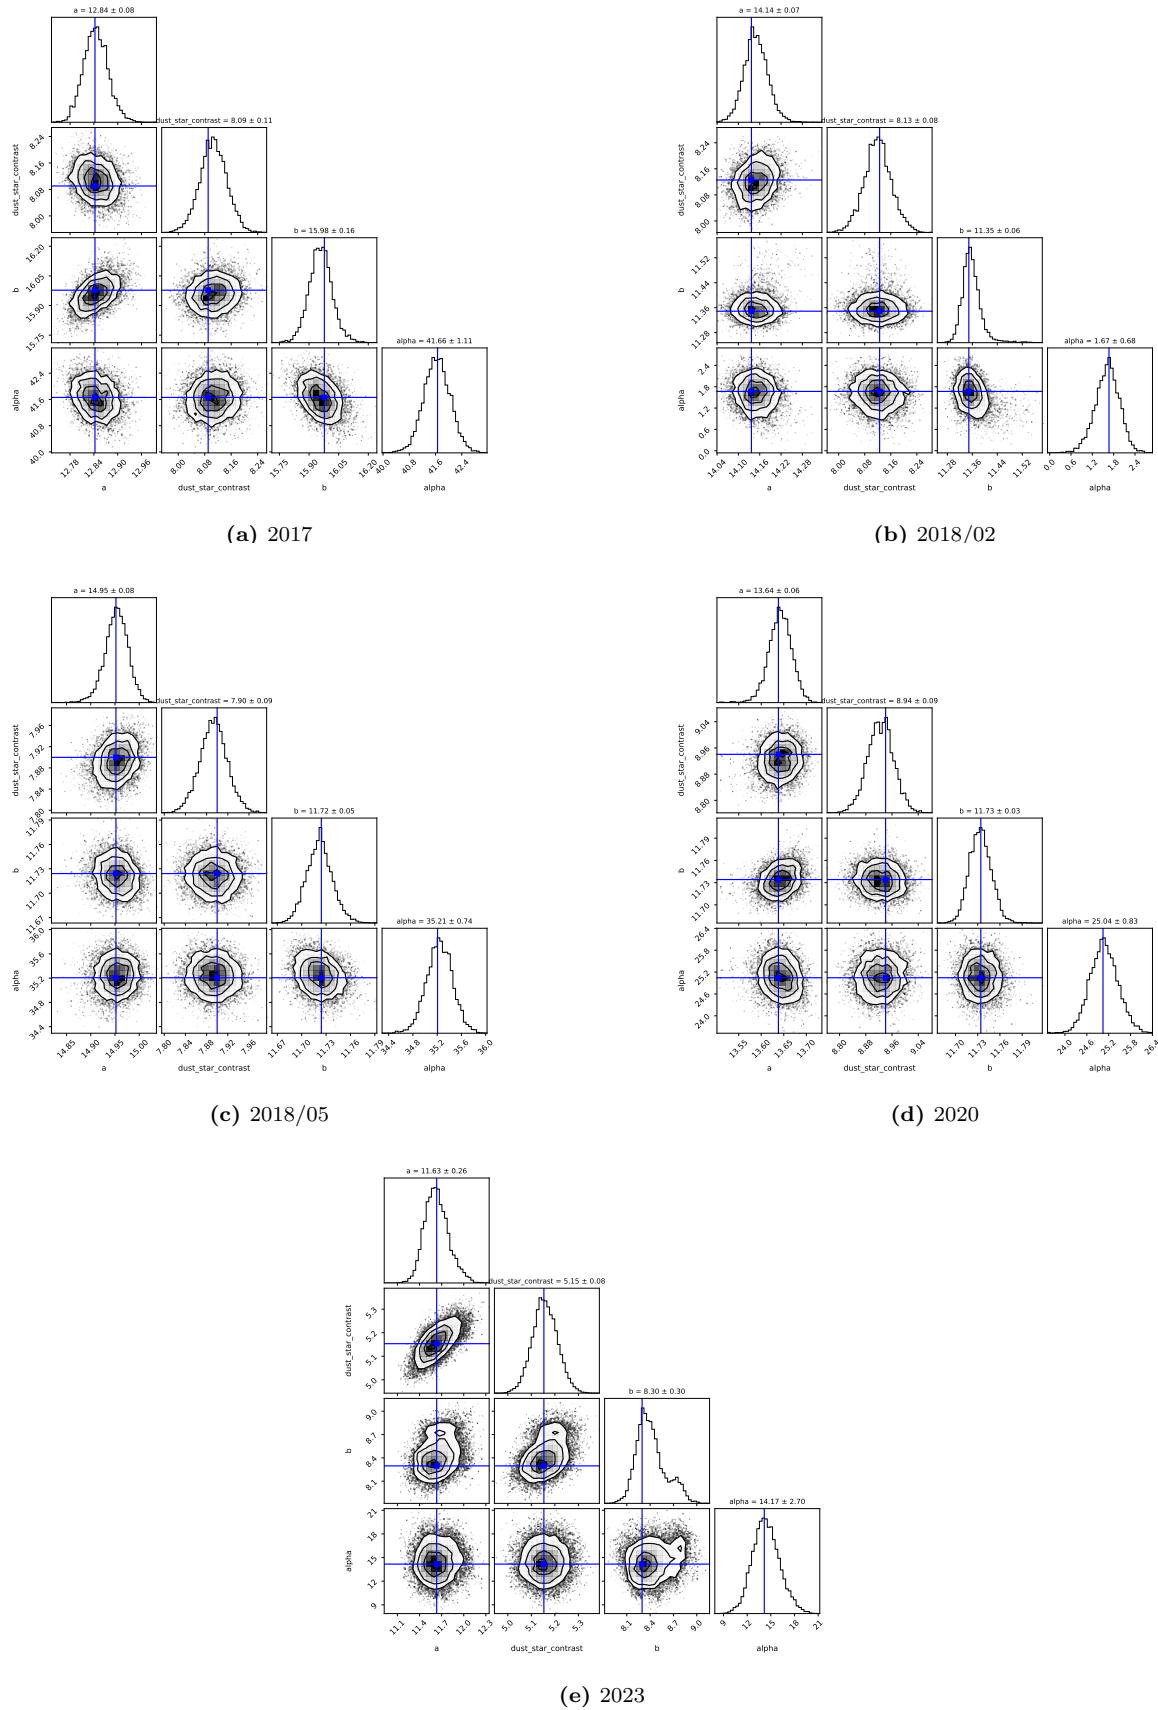
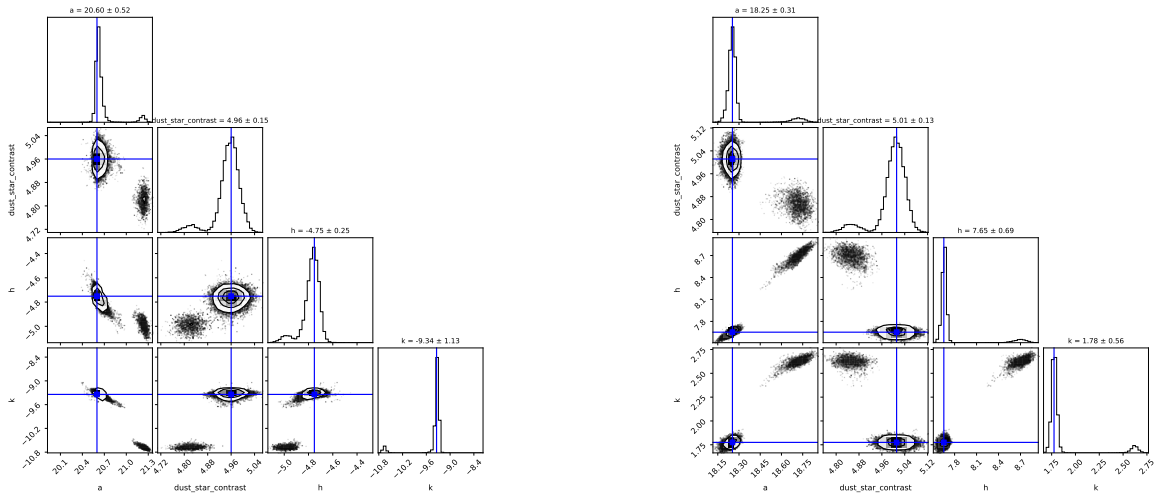
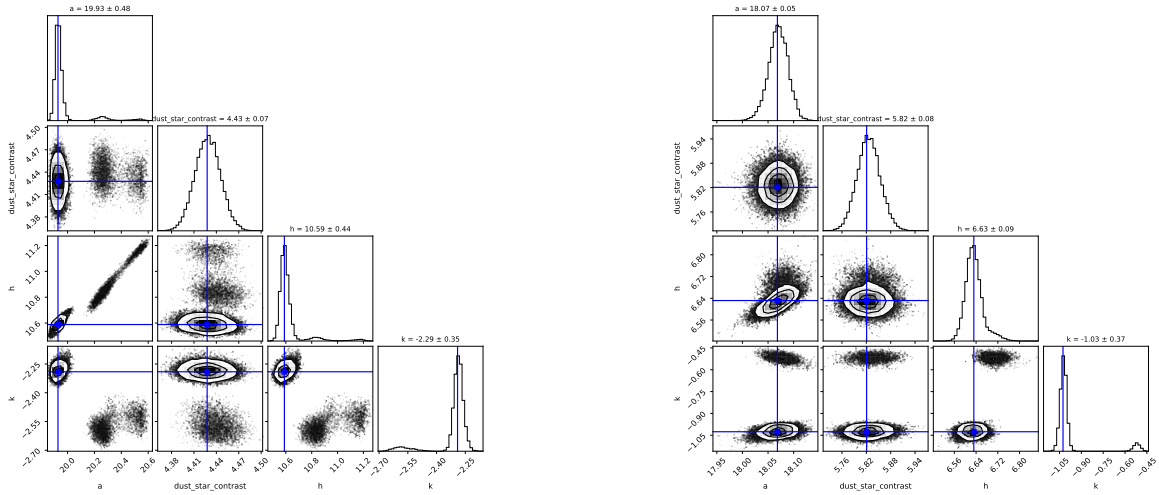


Figure A.10: Corner Plots for Model Class G — Power-law elliptical shell (exponent 3).



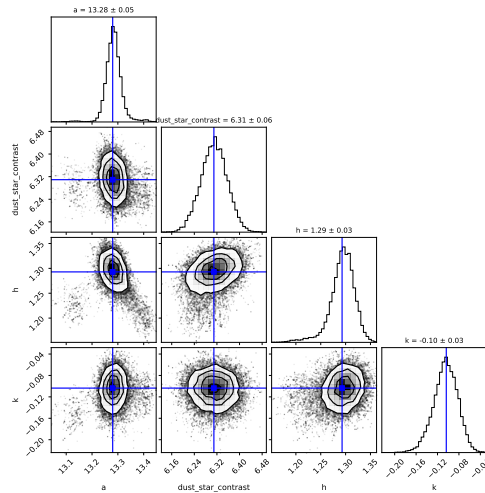
(a) 2017

(b) 2018/02



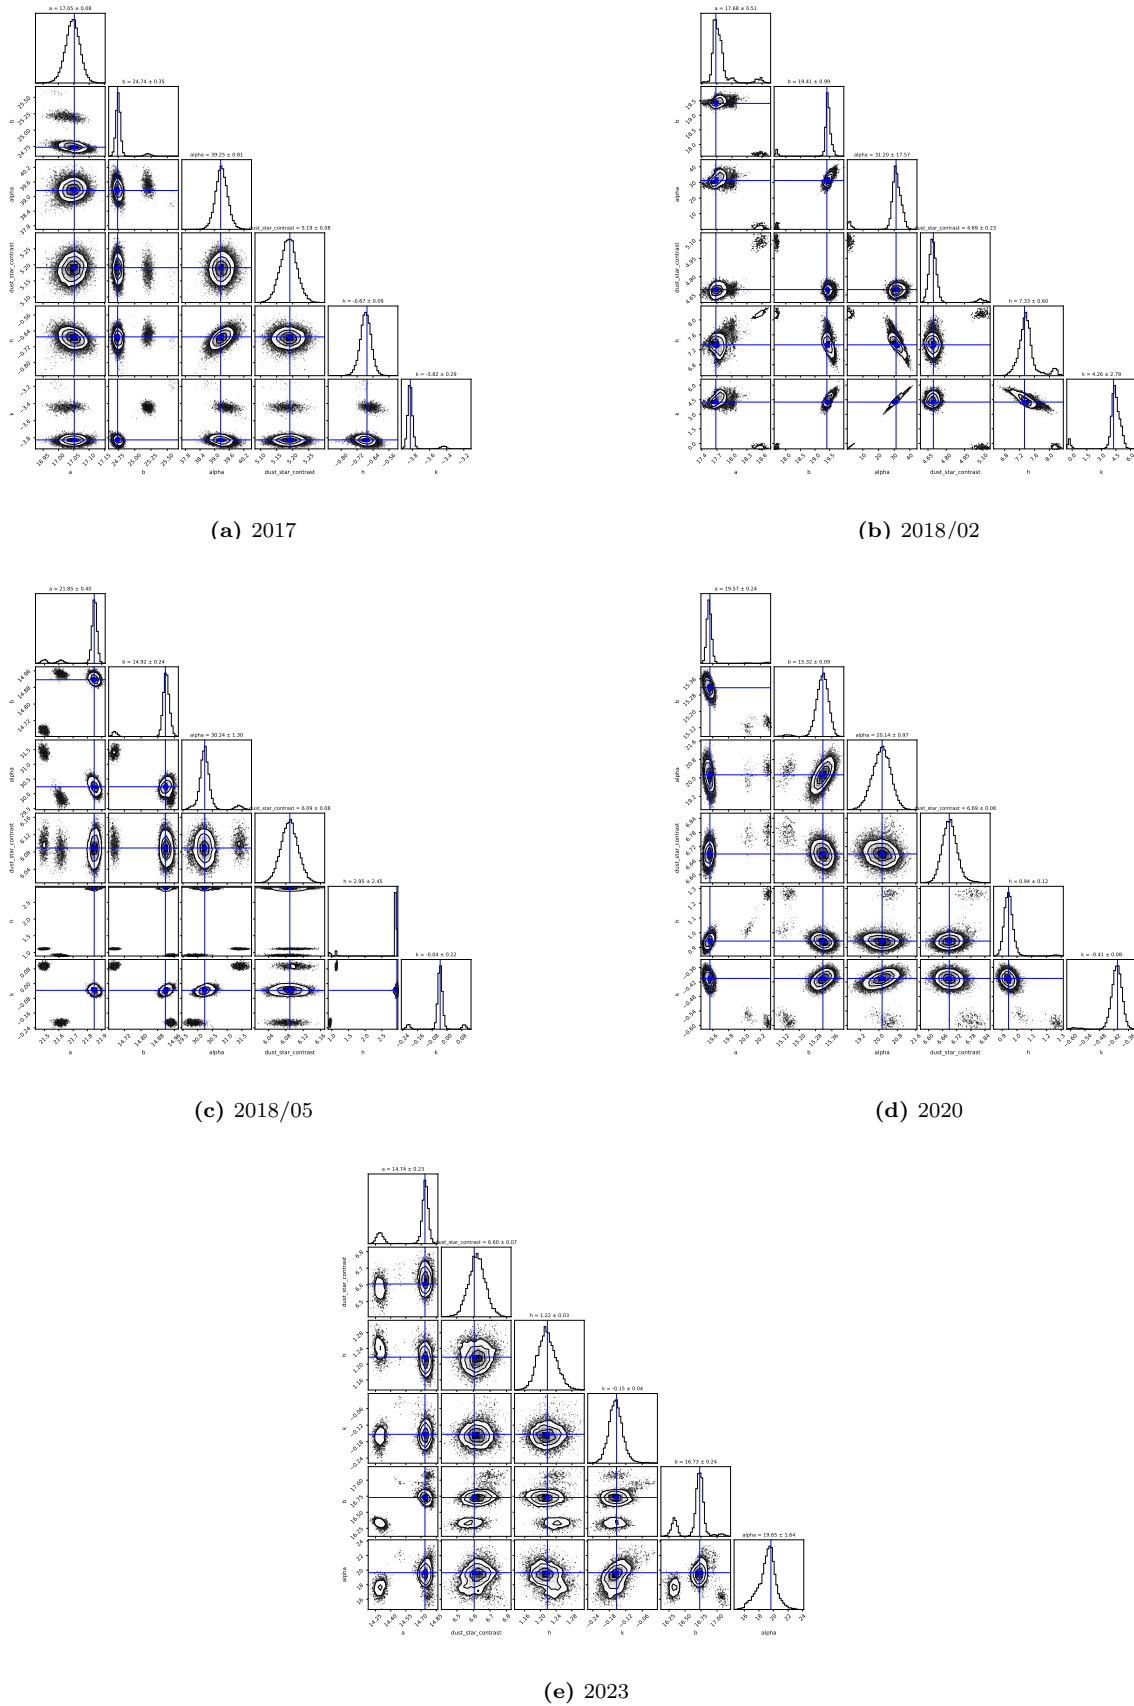
(c) 2018/05

(d) 2020

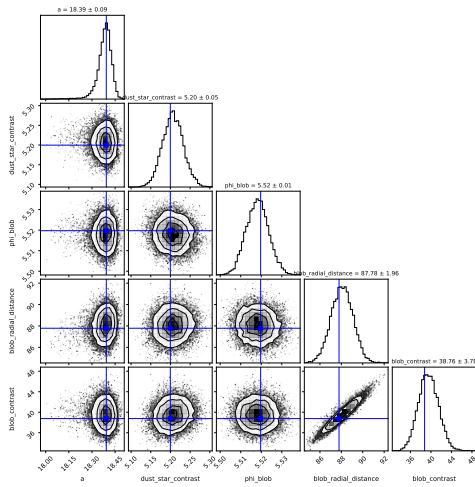


(e) 2023

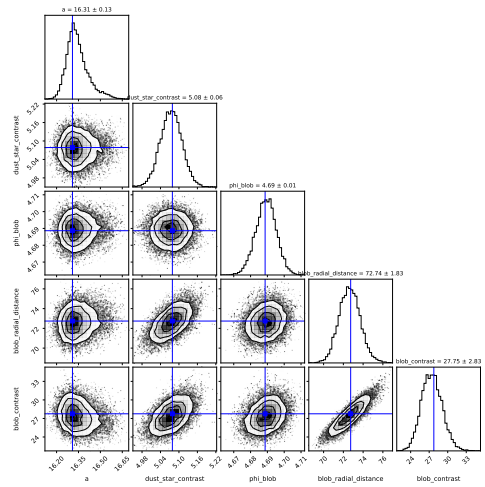
**Figure A.11:** Corner Plots for Model Class H — Thin offset spherical shell. Multi-modal posteriors but clear MLEs.



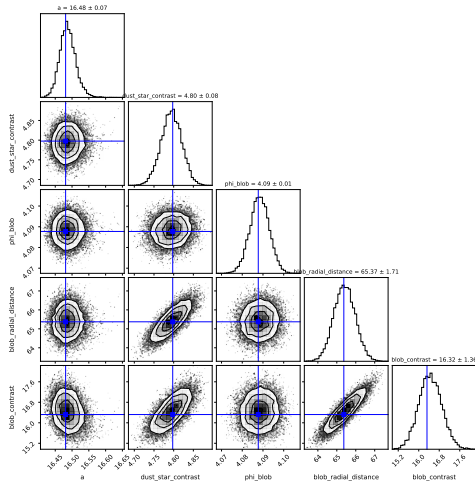
**Figure A.12:** Corner Plots for Model Class I — Thin offset elliptical shell. 2018 epochs are notably more multimodal/noisy.



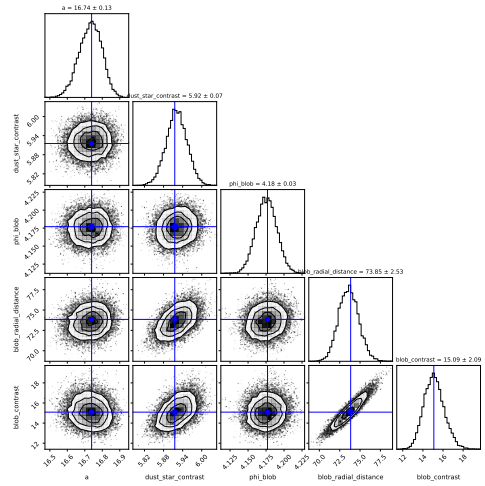
(a) 2017



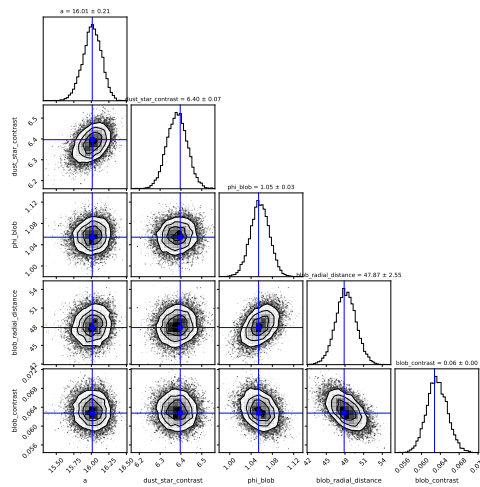
(b) 2018/02



(c) 2018/05

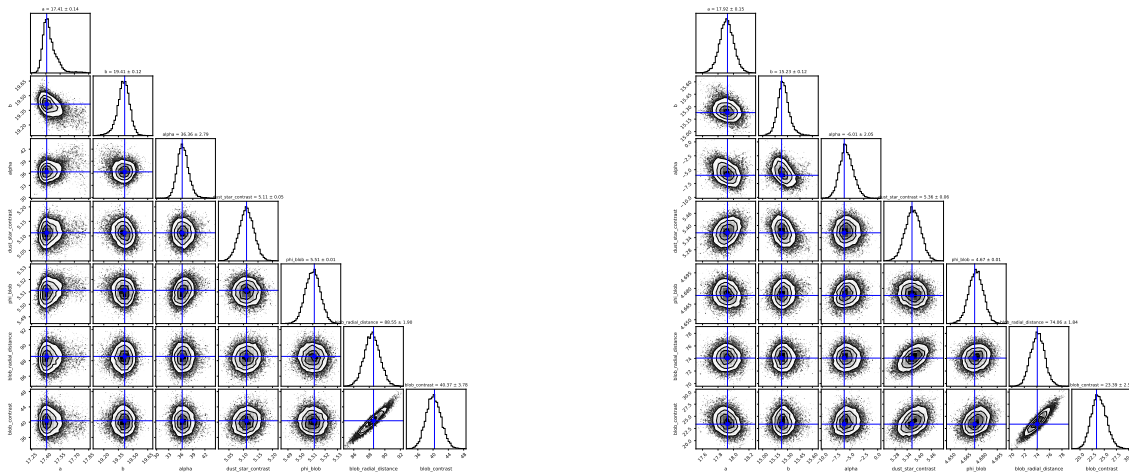


(d) 2020



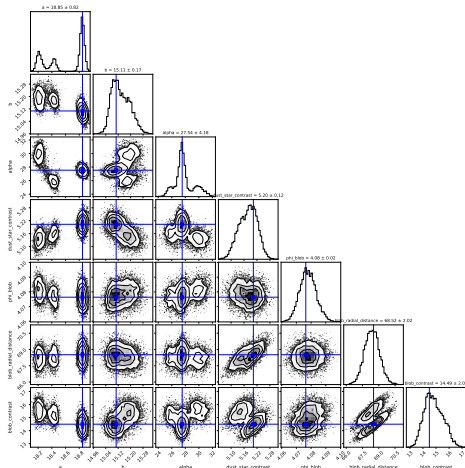
(e) 2023

**Figure A.13:** Corner Plots for Model Class J — Thin spherical shell with blob.

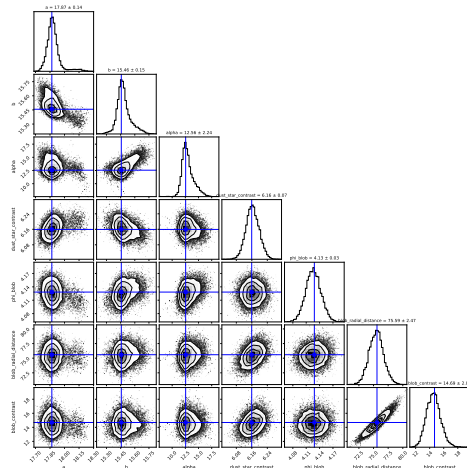


(a) 2017

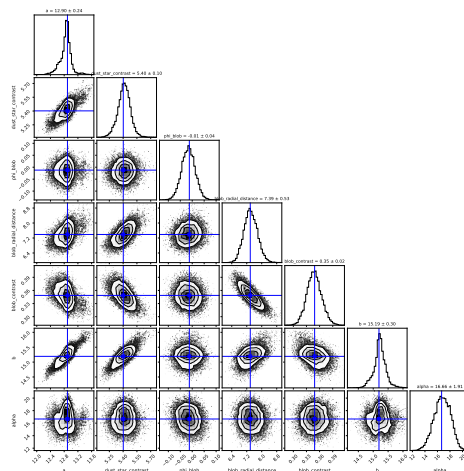
(b) 2018/02



(c) 2018/05

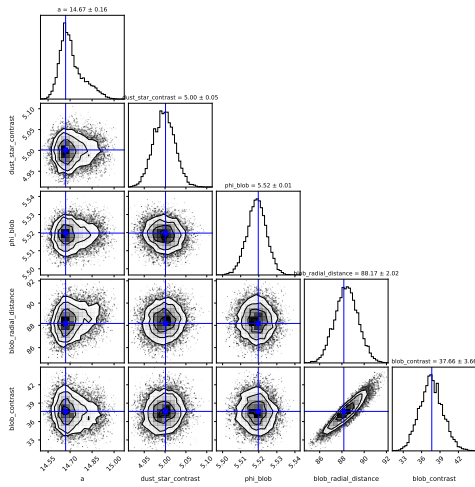


(d) 2020

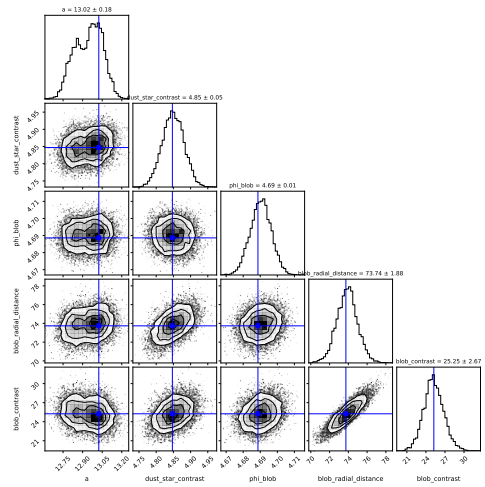


(e) 2023

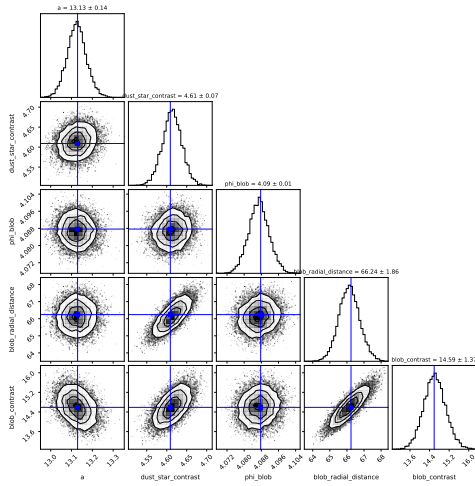
Figure A.14: Corner Plots for Model Class K — Thin elliptical shell with blob.



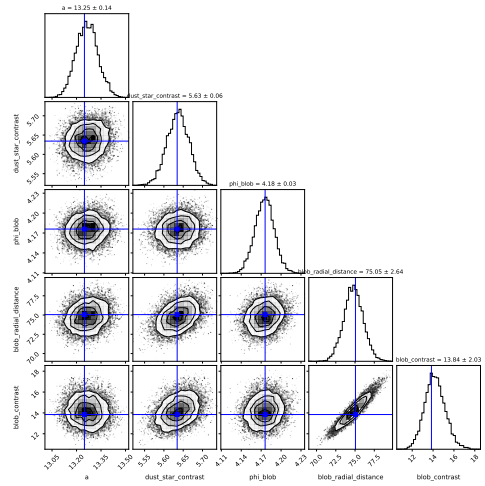
(a) 2017




(b) 2018/02



(c) 2018/05

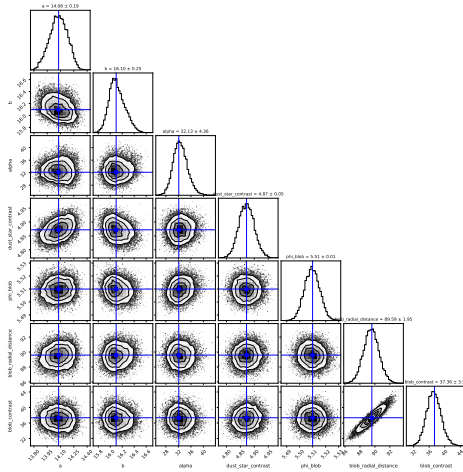


(d) 2020

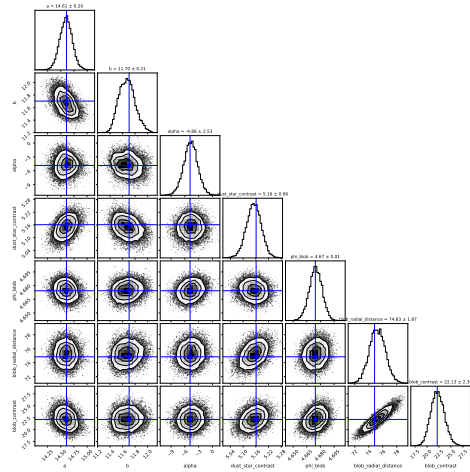


(e) 2023

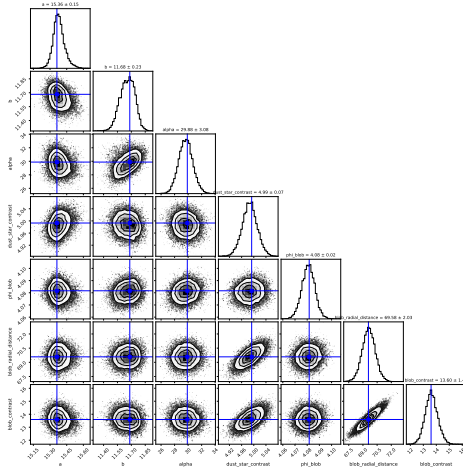
Figure A.15: Corner Plots for Model Class L — Thick spherical shell with blob.




(a) 2017



(b) 2018/02



(c) 2018/05



(d) 2020

**Figure A.16:** Corner Plots for Model Class M — Thick elliptical shell with blob.

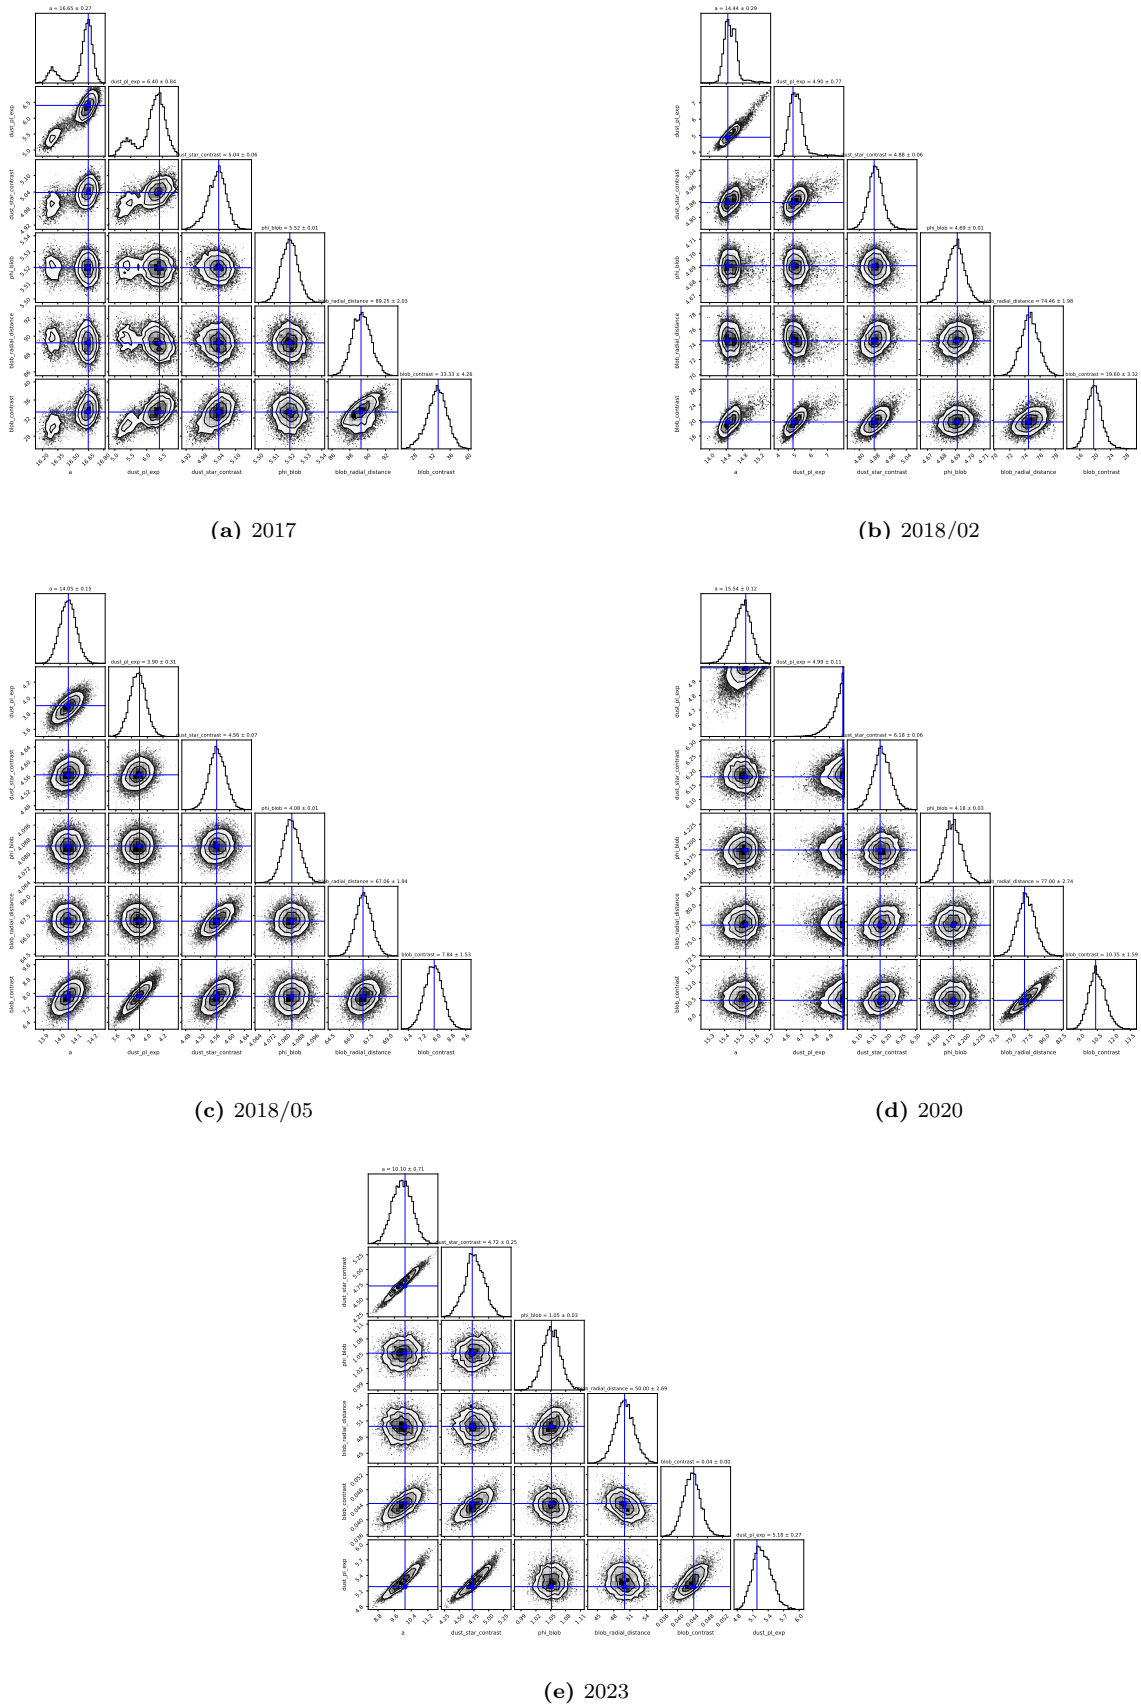
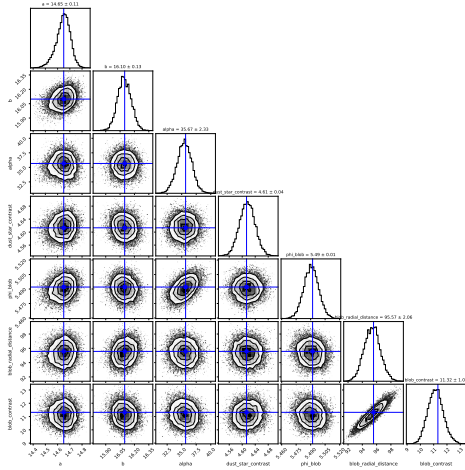
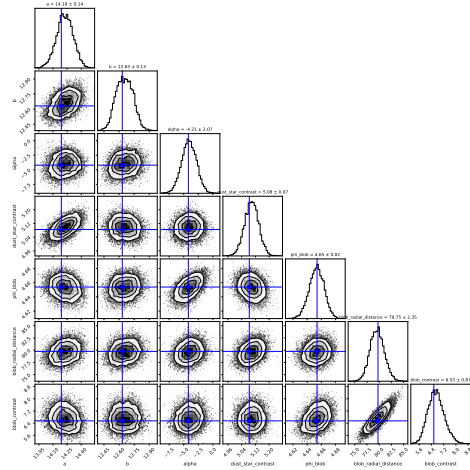


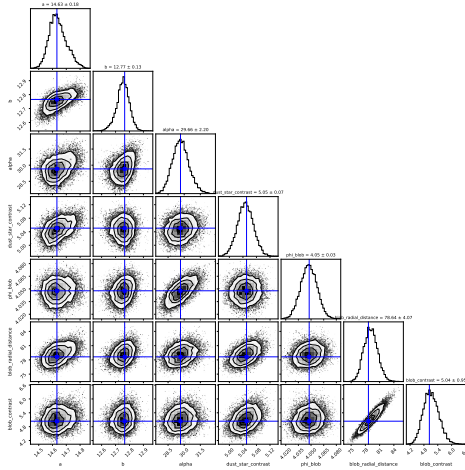
Figure A.17: Corner Plots for Model Class N — Power-law spherical shell with blob.



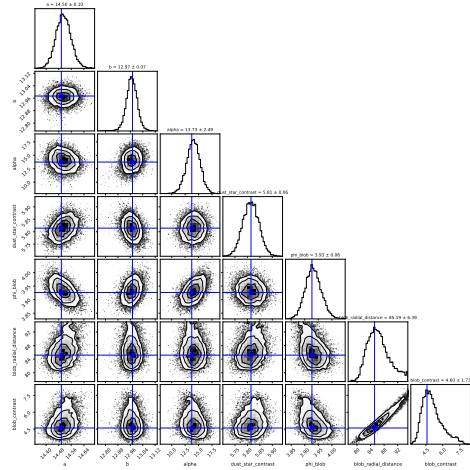
(a) 2017



(b) 2018/02




(c) 2018/05

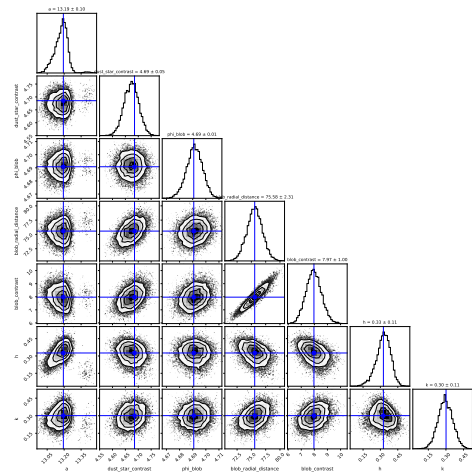


(d) 2020

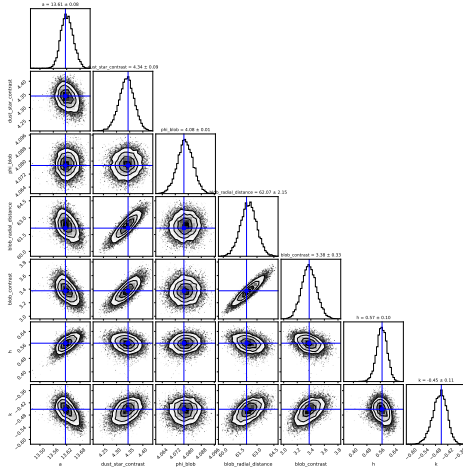
**Figure A.18:** Corner Plots for Model Class O — Power-law elliptical shell with blob.



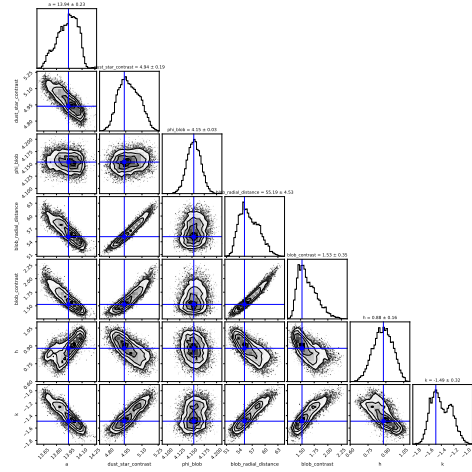
(a) 2017



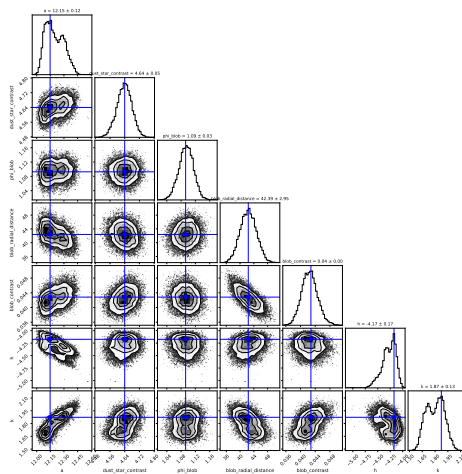
(b) 2018/02



(c) 2018/05




(d) 2020




(e) 2023


Figure A.19: Corner Plots for Model Class P — Offset power-law spherical shell + blob.




(a) 2017




(b) 2018/02



(c) 2018/05




(d) 2020

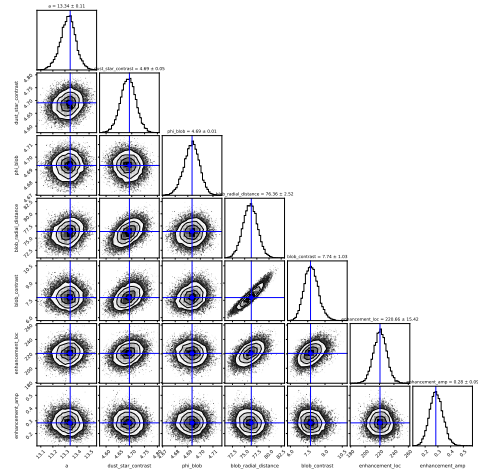


(e) 2023

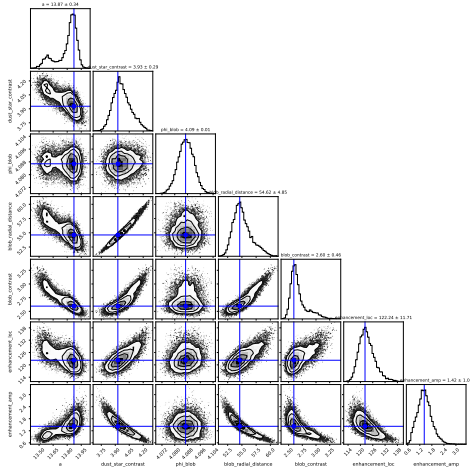
**Figure A.20:** Corner Plots for Model Class Q — Offset power-law elliptical shell + blob.



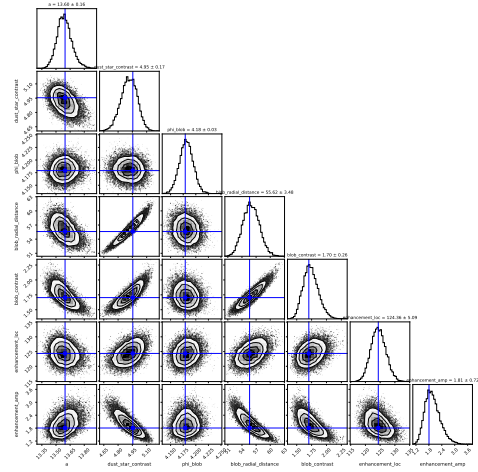
(a) 2017



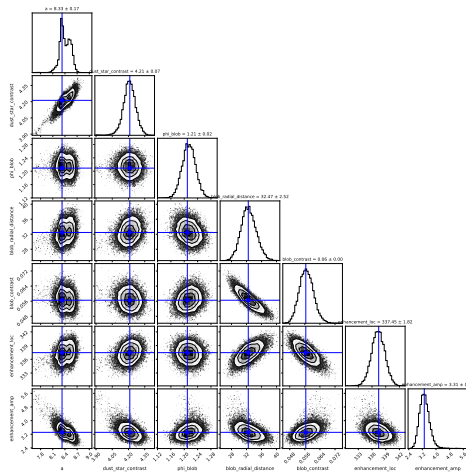
(b) 2018/02



(c) 2018/05




(d) 2020




(e) 2023


Figure A.21: Corner Plots for Model Class R — Enhanced power-law spherical shell + blob .




(a) 2017




(b) 2018/02



(c) 2018/05



(d) 2020



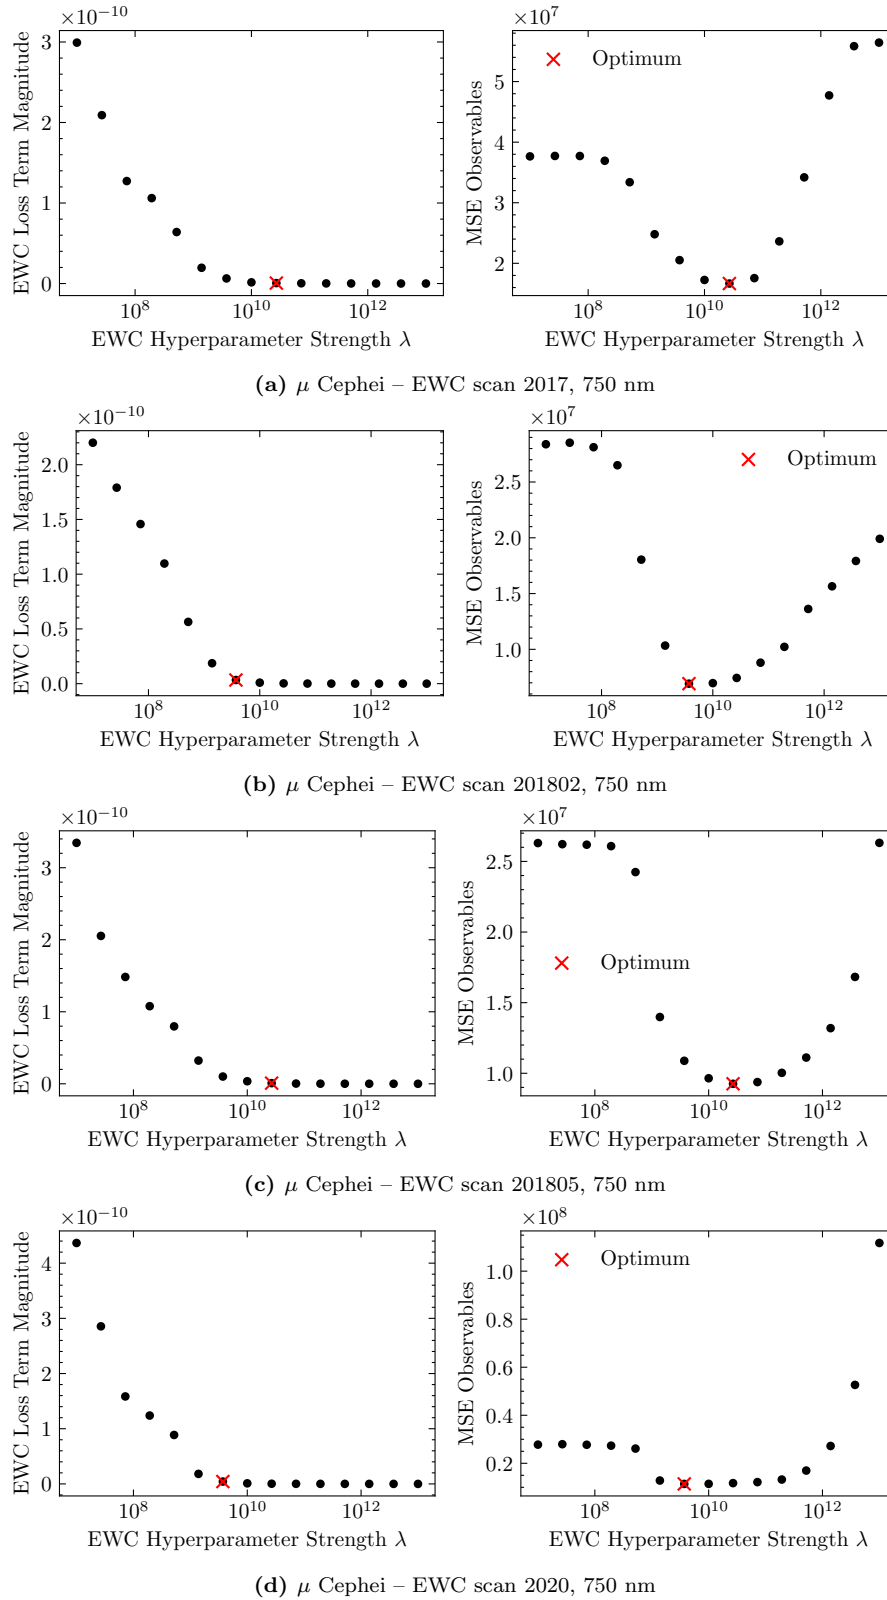
(e) 2023

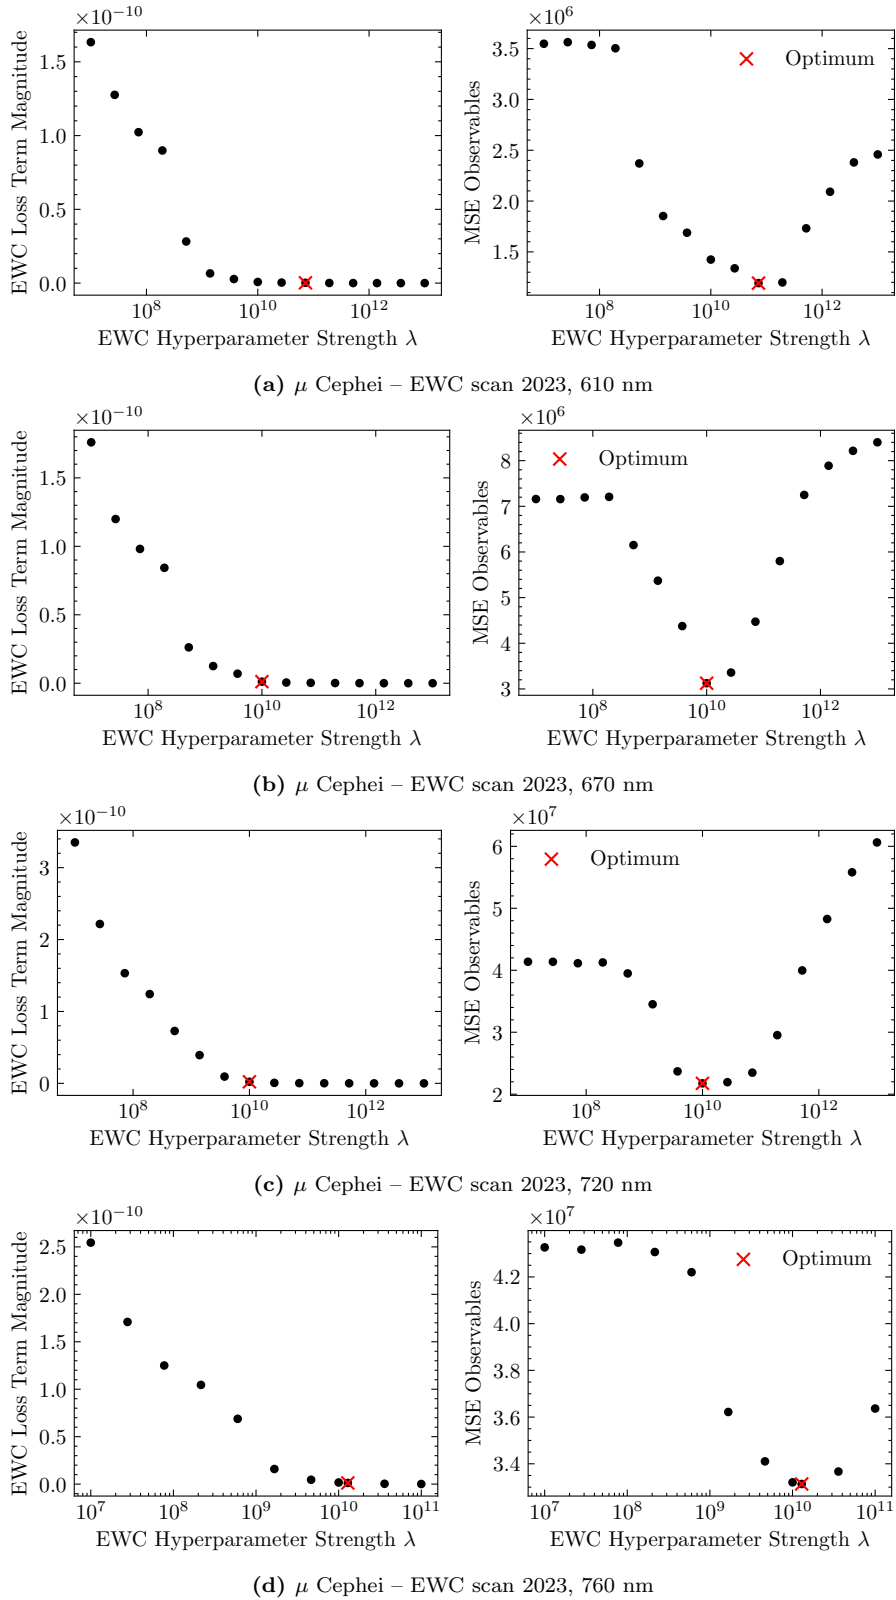
**Figure A.22:** Corner Plots for Model Class S — Enhanced power-law elliptical shell + blob.

This page is intentionally left blank.

## A.4 Chapter 8

## A.4.1 EWC Scans

Figure A.23:  $\mu$  Cephei EWC scans for archival datasets 1-4.



**Figure A.24:**  $\mu$  Cephei EWC scans for multi-band dataset 6.

## A.4.2 Bayesian Optimisation Posteriors

Shaded regions indicate the mode used to estimate the  $1\sigma$  uncertainties of the maximum likelihood estimates for each parameter.

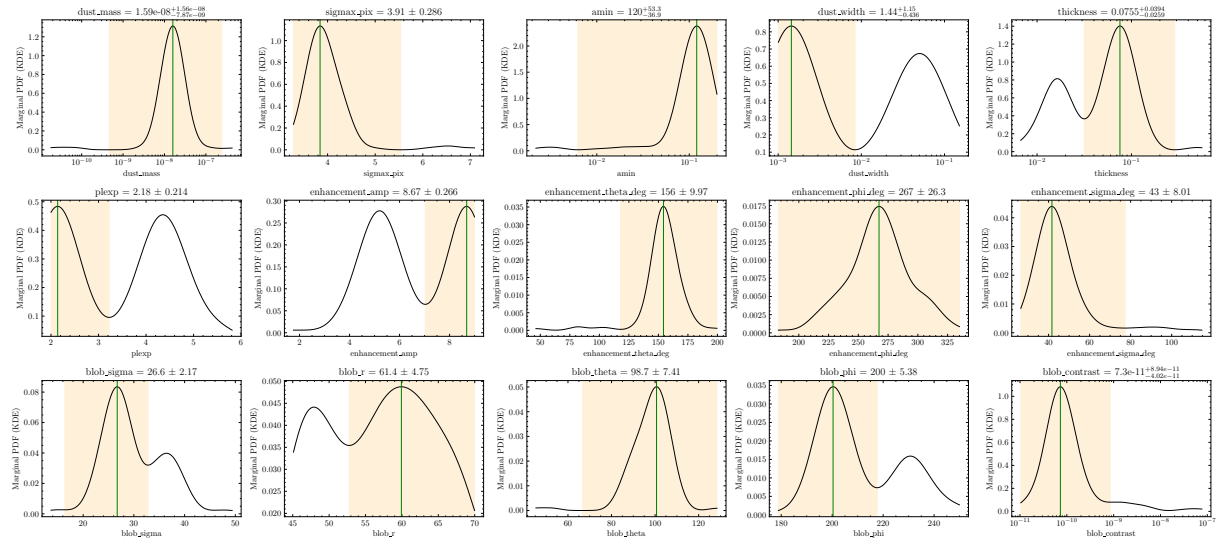


Figure A.25: Bayesian Optimisation Posteriors: Amorphous Enstatite – Model 2.

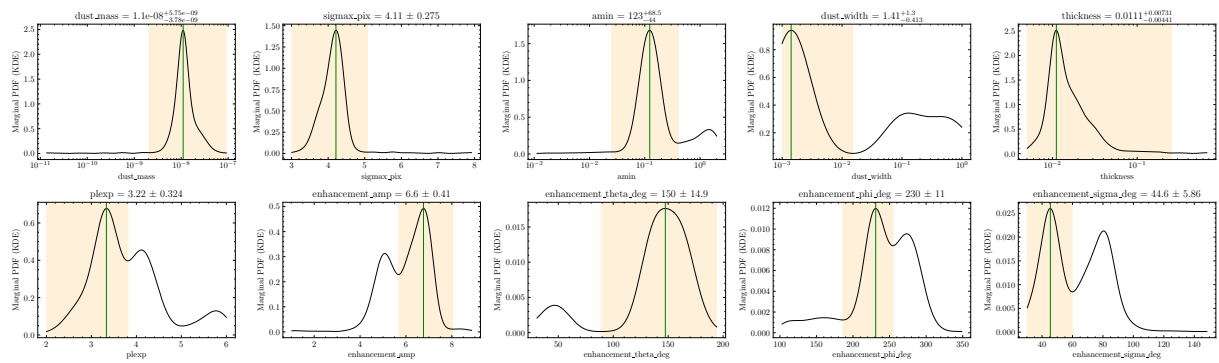


Figure A.26: Bayesian Optimisation Posteriors: Amorphous Enstatite – Model 1.

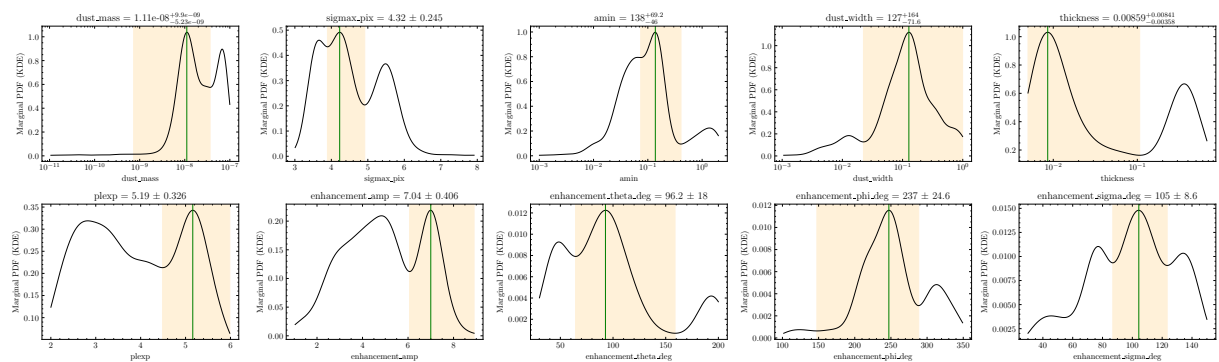


Figure A.27: Bayesian Optimisation Posteriors: Amorphous Spinel – Model 1.

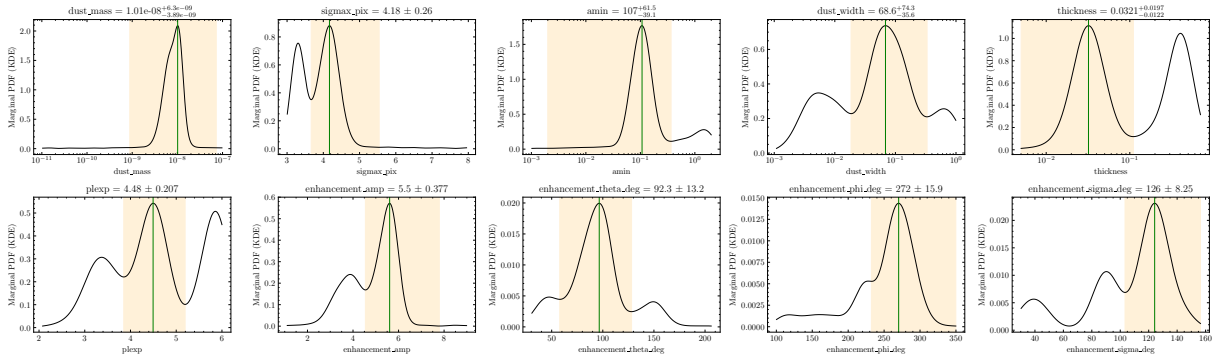


Figure A.28: Bayesian Optimisation Posteriors: Amorphous Corundum – Model 1.

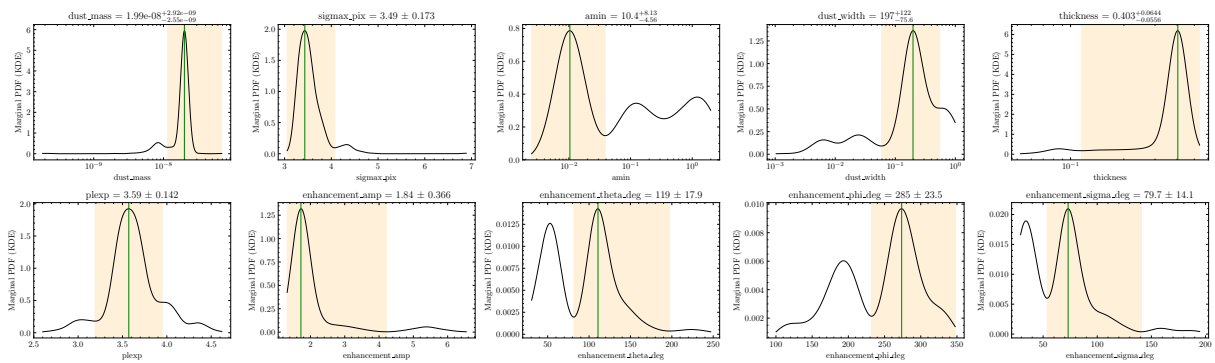


Figure A.29: Bayesian Optimisation Posteriors: Amorphous Forsterite – Model 1.

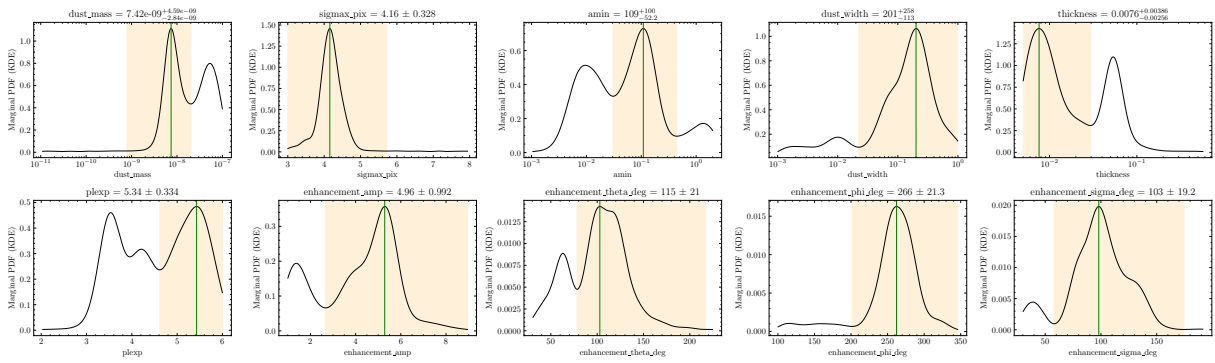


Figure A.30: Bayesian Optimisation Posteriors: Crystalline Pyroxene 95– Model 1.

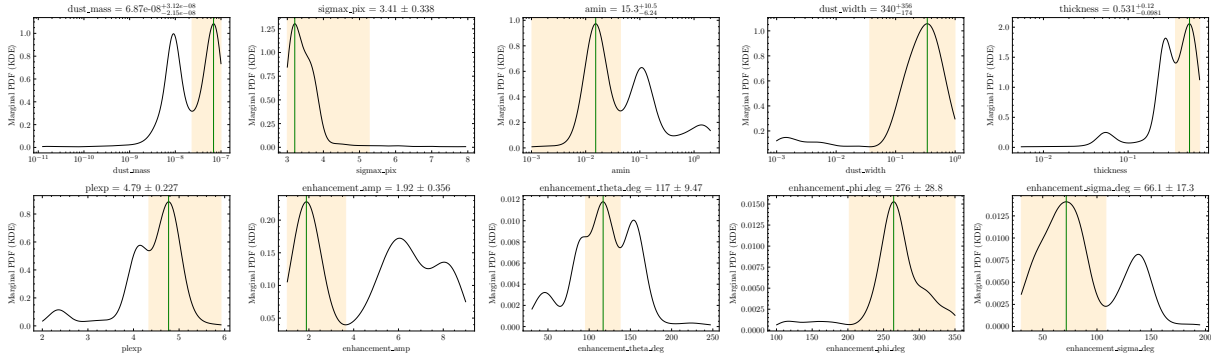


Figure A.31: Bayesian Optimisation Posteriors: Amorphous Silica– Model 1.

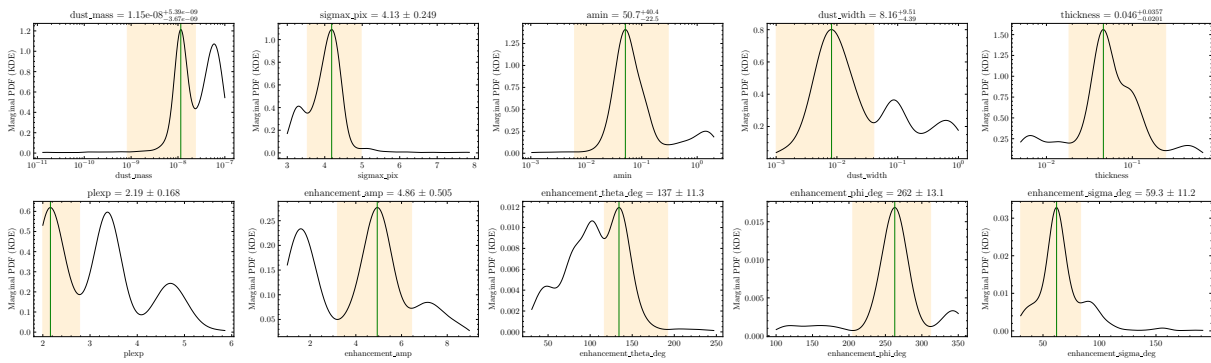


Figure A.32: Bayesian Optimisation Posteriors: Crystalline Enstatite – Model 1.

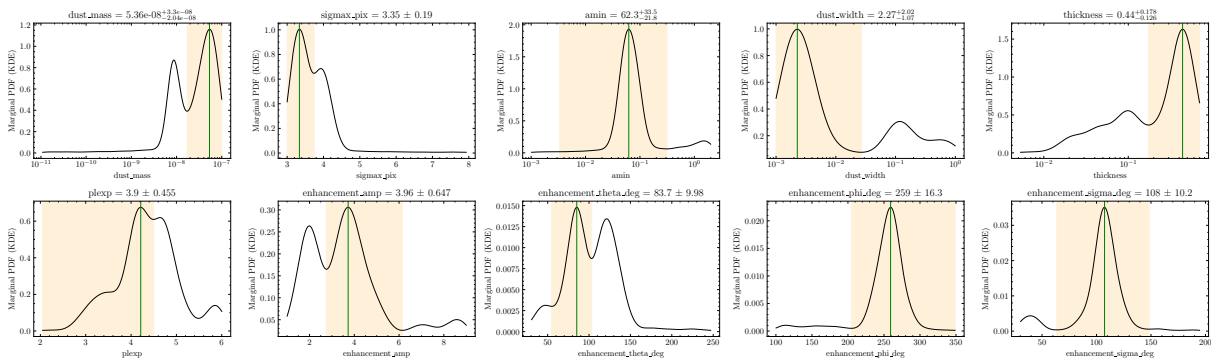


Figure A.33: Bayesian Optimisation Posteriors: Crystalline Forsterite – Model 1.

# Bibliography

- Adam, C. et al. (Aug. 2019). “Exploring the innermost dust formation region of the oxygen-rich AGB star IK Tauri with VLT/SPHERE-ZIMPOL and VLTI/AMBER”. In: *Astronomy & Astrophysics* 628, A132. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/201834999>.
- Ahn, Kyohoon et al. (Sept. 2021). “SCEXAO, a testbed for developing high-contrast imaging technologies for ELTs”. In: URL: <http://arxiv.org/abs/2109.13353>.
- Akiyama, Kazunori et al. (Apr. 2019a). “First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole”. In: *The Astrophysical Journal Letters* 875.1, p. L1. ISSN: 2041-8205.
- Akiyama, Kazunori et al. (Apr. 2019b). “First M87 Event Horizon Telescope Results. VI. The Shadow and Mass of the Central Black Hole”. In: *The Astrophysical Journal Letters* 875.1, p. L6. ISSN: 2041-8205.
- Alharbi, S.S. (2018). “Mass Loss in Evolved Stars”. In: URL: <http://arxiv.org/abs/1412.0803>.
- Aller, L H et al. (Jan. 1951). “The Spectra of R Andromedae (maximum) and {MU} Cephei Between 7400 and 8800 A.” In: *The Astrophysical Journal* 113, p. 72.
- Ando, Tomohiro (2010). *Bayesian Model Selection and Statistical Modeling*. CRC Press. ISBN: 1439836159. URL: [https://books.google.com.au/books?id=PJzLUEt-VoYC&dq=bayes+information+criterion+textbook&lr=&source=gbs\\_navlinks\\_s](https://books.google.com.au/books?id=PJzLUEt-VoYC&dq=bayes+information+criterion+textbook&lr=&source=gbs_navlinks_s).
- Andrych, Kateryna et al. (July 2023). “Second-generation protoplanetary discs around evolved binaries: a high-resolution polarimetric view with SPHERE/IRDIS”. In: *Monthly Notices of the Royal Astronomical Society* 524.3, pp. 4168–4195. ISSN: 0035-8711.
- Andrych, Kateryna et al. (Nov. 2024). “Multiwavelength high-resolution polarimetric imaging of second-generation disc around post-AGB binary IRAS08544–4431 with SPHERE”. In: *Monthly Notices of the Royal Astronomical Society* 535.2, pp. 1763–1777. ISSN: 0035-8711.
- Anugu, Narsireddy et al. (Oct. 2024). “CHARA Near-infrared Imaging of the Yellow Hypergiant Star  $\rho$  Cassiopeiae: Convection Cells and Circumstellar Envelope”. In: *The Astrophysical Journal* 974.1, p. 113. ISSN: 0004-637X. URL: <https://iopscience.iop.org/article/10.3847/1538-4357/ad6b2b>.
- Ariste, A. López (2023). “The height of convective plumes in the red supergiant  $\mu$  Cep”. In: *Astronomy and Astrophysics*.
- Ariste, A. López et al. (Jan. 2023). “The height of convective plumes in the red supergiant  $\mu$  Cep”. In: URL: <http://arxiv.org/abs/2301.01326><http://dx.doi.org/10.1051/0004-6361/202244285>.
- Aristotle (1981). *Analytica Priora et Posteriora*. Ed. by E Ross et al. Oxford University Press. ISBN: ISBN 9780198145622.
- Arroyo-Torres, B. et al. (Mar. 2015). “What causes the large extensions of red supergiant atmospheres?” In: *Astronomy & Astrophysics* 575, A50. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201425212>.
- Asano, Ryosuke S. et al. (June 2013). “What determines the grain size distribution in galaxies?” In: *Monthly Notices of the Royal Astronomical Society* 432.1, pp. 637–652. ISSN: 1365-2966.
- Ashcraft, Jaren N. et al. (Mar. 2025). “Polarization aberrations in next-generation Giant Segmented Mirror Telescopes (GSMs)”. In: *Astronomy & Astrophysics* 695, A28. ISSN: 0004-6361.
- Asplund, Martin et al. (Sept. 2009). “The Chemical Composition of the Sun”. In: *Annual Review of Astronomy and Astrophysics* 47.1, pp. 481–522. ISSN: 0066-4146.
- Avenhaus, Henning et al. (Aug. 2018a). “Disks around T Tauri Stars with SPHERE (DARTTS-S). I. SPHERE/IRDIS Polarimetric Imaging of Eight Prominent T Tauri Disks\*.” In: *The Astrophysical Journal* 863.1, p. 44. ISSN: 0004-637X.

- Avenhaus, Henning et al. (Mar. 2018b). “Disks ARound TTauri Stars with Sphere (DARTTS-S) I: Sphere / IRDIS Polarimetric Imaging of 8 prominent TTauri Disks”. In: URL: <http://arxiv.org/abs/1803.10882><http://dx.doi.org/10.3847/1538-4357/aab846>.
- Baron, F (2020). *OITOOOLS: the All-in-One Tool Package for Optical Interferometry*.
- Beasor, Emma R et al. (Mar. 2020). “A new mass-loss rate prescription for red supergiants”. In: *Monthly Notices of the Royal Astronomical Society* 492.4, pp. 5994–6006. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article/492/4/5994/5727330>.
- Begemann, B. et al. (Feb. 1997). “Aluminum Oxide and the Opacity of Oxygen-rich Circumstellar Dust in the 12–17 Micron Range”. In: *The Astrophysical Journal* 476.1, pp. 199–208. ISSN: 0004-637X.
- Bickel, Peter J. et al. (Mar. 2015). *Mathematical Statistics*. Chapman and Hall/CRC. ISBN: 9780429174780. URL: <https://www.taylorfrancis.com/books/9781498723817>.
- Bickel, William S. et al. (May 1985). “Stokes vectors, Mueller matrices, and polarized scattered light”. In: *American Journal of Physics* 53.5, pp. 468–478. ISSN: 0002-9505. URL: <https://pubs.aip.org/ajp/article/53/5/468/1052590/Stokes-vectors-Mueller-matrices-and-polarized>.
- Bjorkman, Karen S. (Apr. 2000). “Polarimetry as a Diagnostic of Circumstellar Envelopes”. In: *International Astronomical Union Colloquium* 175, pp. 384–395. ISSN: 0252-9211.
- Bladh, S. et al. (Oct. 2012). “Exploring wind-driving dust species in cool luminous giants”. In: *Astronomy & Astrophysics* 546, A76. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201219138>.
- Bladh, S. et al. (May 2013). “Exploring wind-driving dust species in cool luminous giants”. In: *Astronomy & Astrophysics* 553, A20. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201220590>.
- Bohren, Craig F. et al. (Apr. 1998). *Absorption and Scattering of Light by Small Particles*. Wiley. ISBN: 9780471293408.
- Bradbury, J et al (2018). *JAX: composable transformations of Python+NumPy programs*.
- Brochu, Eric et al. (Dec. 2010). “A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning”. In: URL: <http://arxiv.org/abs/1012.2599>.
- Brown, T et al. (2020). “Language models are few-shot learners”. In: *NIPS’20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, pp. 1877–1901.
- Bruch, Rachel J. et al. (May 2021). “A Large Fraction of Hydrogen-rich Supernova Progenitors Experience Elevated Mass Loss Shortly Prior to Explosion”. In: *The Astrophysical Journal* 912.1, p. 46. ISSN: 0004-637X. URL: <https://iopscience.iop.org/article/10.3847/1538-4357/abef05>.
- Buscher, D (July 2015). “Model fitting and image reconstruction”. In: *Practical Optical Interferometry*. Cambridge University Press, pp. 218–247.
- Cannon, E et al. (Jan. 2021). “The inner circumstellar dust of the red supergiant Antares as seen with VLT/SPHERE/ZIMPOL”. In: *Monthly Notices of the Royal Astronomical Society* 502.1, pp. 369–382. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article/502/1/369/6074268>.
- Canovas, H. et al. (July 2011). “Data-reduction techniques for high-contrast imaging polarimetry”. In: *Astronomy & Astrophysics* 531, A102. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201116918>.
- Carpenter, B (2017). *Ensemble Methods are Doomed to Fail in High Dimensions*. URL: <https://statmodeling.stat.columbia.edu/2017/03/15/ensemble-methods-doomed-fail-high-dimensions/>.
- Carroll, B.W. et al. (2007). *An Introduction to Modern Astrophysics*. Second. San Francisco: Pearson Education.
- Chael, Andrew A. et al. (Apr. 2018). “Interferometric Imaging Directly with Closure Phases and Closure Amplitudes”. In: *The Astrophysical Journal* 857.1, p. 23. ISSN: 0004-637X.
- Charles, M. (2022). “Taking Mira to the Diffraction Limit”. In.
- Cheetham, Anthony C. et al. (Aug. 2016). “Sparse aperture masking with SPHERE”. In: ed. by Fabien Malbet et al., 99072T. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2231983>.
- Chernova, E. et al. (Nov. 2017). “Optical NIR-VIS-VUV constants of advanced substrates for thin-film devices”. In: *Optical Materials Express* 7.11, p. 3844. ISSN: 2159-3930.

- Chiavassa, A. et al. (Nov. 2011). “Radiative hydrodynamics simulations of red supergiant stars”. In: *Astronomy & Astrophysics* 535, A22. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201117463>.
- Chiavassa, A. et al. (May 2022). “Probing Red Supergiant dynamics through photo-center displacements measured by Gaia”. In: URL: <http://arxiv.org/abs/2205.05156%20http://dx.doi.org/10.1051/0004-6361/202243568>.
- Claudi, R et al. (2025). “High-Contrast Imaging: Hide and Seek with Exoplanets”. In: *Arxiv*. URL: <https://arxiv.org/abs/2501.07976>.
- Coyne, G. V. et al. (Feb. 1968). “Wavelength dependence of polarization. XI. MU Cephei.” In: *The Astronomical Journal* 73, p. 20. ISSN: 00046256. URL: [http://adsabs.harvard.edu/cgi-bin/bib\\_query?1968AJ....73...20C](http://adsabs.harvard.edu/cgi-bin/bib_query?1968AJ....73...20C).
- Coyne, G. V. et al. (Aug. 1979). “Wavelength dependence of polarization. XXXVI - Changes in polarization across TiO bands in cool stars”. In: *The Astronomical Journal* 84, p. 1200. ISSN: 00046256. URL: [http://adsabs.harvard.edu/cgi-bin/bib\\_query?1979AJ....84.1200C](http://adsabs.harvard.edu/cgi-bin/bib_query?1979AJ....84.1200C).
- Davies, Ben et al. (Mar. 2020). “The ‘red supergiant problem’: the upper luminosity boundary of Type II supernova progenitors”. In: *Monthly Notices of the Royal Astronomical Society* 493.1, pp. 468–476. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article/493/1/468/5709941>.
- Davis, J. et al. (Mar. 1999). “The Sydney University Stellar Interferometer – I. The instrument”. In: *Monthly Notices of the Royal Astronomical Society* 303.4, pp. 773–782. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article/303/4/773/1062978>.
- De Beck, E. et al. (Nov. 2010). “Probing the mass-loss history of AGB and red supergiant stars from CO rotational line profiles”. In: *Astronomy & Astrophysics* 523, A18. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/200913771>.
- Decin, L. et al. (Sept. 2020). “(Sub)stellar companions shape the winds of evolved stars”. In: *Science* 369.6510, pp. 1497–1500. ISSN: 0036-8075.
- Decin, Leen et al. (Mar. 2023). “ALMA detection of CO rotational line emission in red supergiant stars of the massive young star cluster RSGC1 – Determination of a new mass-loss rate prescription for red supergiants”. In: URL: <http://arxiv.org/abs/2303.09385>.
- Deisenroth, M et al. (2020). *Mathematics for Machine Learning*. Vol. 1. Cambridge University Press.
- Dell’Agli, F. et al. (June 2014). “On the alumina dust production in the winds of O-rich asymptotic giant branch stars”. In: *Monthly Notices of the Royal Astronomical Society* 441.2, pp. 1115–1125. ISSN: 1365-2966.
- Dorfi, A et al. (1998). “AGB Stars and Mass Loss”. In: *Astronomische Gesellschaft*.
- Dorschner, J et al. (1995). “Steps toward interstellar silicate mineralogy. II. Study of Mg-Fe-silicate glasses of variable composition.” In: *Astronomy and Astrophysics* 300, pp. 503–520. URL: <https://adsabs.harvard.edu/full/1995A%26A...300..503D>.
- Draine, B. T. et al. (Oct. 1984). “Optical properties of interstellar graphite and silicate grains”. In: *The Astrophysical Journal* 285, p. 89. ISSN: 0004-637X.
- Dupree, Andrea K. et al. (Sept. 2022). “The Great Dimming of Betelgeuse: A Surface Mass Ejection and Its Consequences”. In: *The Astrophysical Journal* 936.1, p. 18. ISSN: 0004-637X. URL: <https://iopscience.iop.org/article/10.3847/1538-4357/ac7853>.
- Dykes, E et al. (2024). “SCEXAO/CHARIS Near-infrared Scattered-light Imaging and Integral Field Spectropolarimetry of the AB Aurigae Protoplanetary System”. In: *The Astrophysical Journal* 977.2.
- Efron, B. (Jan. 1979). “Bootstrap Methods: Another Look at the Jackknife”. In: *The Annals of Statistics* 7.1. ISSN: 0090-5364.
- Eisenhauer, Frank et al. (Aug. 2023). “Advances in Optical/Infrared Interferometry”. In: *Annual Review of Astronomy and Astrophysics* 61.1, pp. 237–285. ISSN: 0066-4146. URL: <https://www.annualreviews.org/doi/10.1146/annurev-astro-121622-045019>.
- Esposito, Thomas M. et al. (July 2020). “Debris Disk Results from the Gemini Planet Imager Exoplanet Survey’s Polarimetric Imaging Campaign”. In: *The Astronomical Journal* 160.1, p. 24. ISSN: 0004-6256. URL: <https://iopscience.iop.org/article/10.3847/1538-3881/ab9199>.
- Fabian, D et al. (2000). “Steps toward interstellar silicate mineralogy. V. Thermal Evolution of Amorphous Magnesium Silicates and Silica”. In: *Astronomy and Astrophysics* 364, pp. 282–292. URL: <https://ui.adsabs.harvard.edu/abs/2000A%26A...364..282F/abstract>.

- Foreman-Mackey, Daniel et al. (Mar. 2013). “emcee : The MCMC Hammer”. In: *Publications of the Astronomical Society of the Pacific* 125.925, pp. 306–312. ISSN: 00046280. URL: <http://iopscience.iop.org/article/10.1086/670067>.
- Fraknoi, A. et al. (2016). *Astronomy*. ISBN: 978-1-938168-28-4.
- Franta, Daniel et al. (Apr. 2016). “Optical characterization of SiO<sub>2</sub> thin films using universal dispersion model over wide spectral range”. In: ed. by Christophe Gorecki et al., p. 989014.
- Gail, H et al. (1999). “Mineral formation in stellar winds. I. Condensation sequence of silicate and iron grains in stationary oxygen rich outflows”. In: *Astronomy and Astrophysics* 347, pp. 594–616.
- Gail, H. -P et al. (1984). “Formation of crystalline and amorphous carbon grains”. In: *Astronomy and Astrophysics* 132.1, pp. 163–167.
- Gail, Hans-Peter et al. (Dec. 2020). “Non-stoichiometric amorphous magnesium-iron silicates in circumstellar dust shells”. In: *Astronomy & Astrophysics* 644, A139. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/202038090>.
- Gardner, Jonathan P. et al. (June 2023). “The James Webb Space Telescope Mission”. In: *Publications of the Astronomical Society of the Pacific* 135.1048, p. 068001. ISSN: 0004-6280. URL: <https://iopscience.iop.org/article/10.1088/1538-3873/acd1b5>.
- Garufi, A. et al. (Jan. 2020). “Disks Around T Tauri Stars with SPHERE (DARTTS-S)”. In: *Astronomy & Astrophysics* 633, A82. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/201936946>.
- Gehrz, R (1988). “Sources of Stardust in the Galaxy”. In: *Interstellar Dust: Proceedings of the 135th Symposium of the International Astronomical Union*. Ed. by L Allamandola et al. California: Kluwer Academic Publishers, p. 445.
- (1971). “Mass Loss from M Stars”. In: *The Astrophysical Journal*. URL: <https://adsabs.harvard.edu/full/1971ApJ...165..285G>.
- Girard, Julien H V (2010). “Status and new operation modes of the versatile VLT/NaCO”. In: *SPIE*. URL: <https://www.eso.org/sci/libraries/SPIE2010/7736-95.pdf>.
- Gledhill, T. M. et al. (Apr. 2001). “Axisymmetry in protoplanetary nebulae: using imaging polarimetry to investigate envelope structure”. In: *Monthly Notices of the Royal Astronomical Society* 322.2, pp. 321–342. ISSN: 0035-8711.
- Gledhill, T.M (1991). “Optical polarization in the disc around Beta Pictoris”. In: *MNRAS*.
- Glindemann, Andreas (Nov. 2012). “Optical interferometry”. In: *Astrometry for Astrophysics*. Cambridge University Press, pp. 154–174.
- Glindemann, Andreas et al. (July 2000). “The VLT Interferometer: a unique instrument for high-resolution astronomy”. In: ed. by Pierre J. Lena et al., p. 2.
- Gobrecht, D. et al. (Jan. 2016). “Dust formation in the oxygen-rich AGB star IK Tauri”. In: *Astronomy & Astrophysics* 585, A6. ISSN: 0004-6361.
- Golay, Marcel J. E. (Feb. 1971). “Point Arrays Having Compact, Nonredundant Autocorrelations”. In: *Journal of the Optical Society of America* 61.2, p. 272. ISSN: 0030-3941.
- González Delgado, D. et al. (Mar. 2003). “Imaging polarimetry of stellar light scattered in detached shells around the carbon stars R Scl and U Ant”. In: *Astronomy & Astrophysics* 399.3, pp. 1021–1035. ISSN: 0004-6361.
- Goodfellow, I et al. (2016). *Deep Learning (Adaptive Computation and Machine Learning series)*. Vol. 1. The MIT Press.
- Goodman, Jonathan et al. (Jan. 2010). “Ensemble samplers with affine invariance”. In: *Communications in Applied Mathematics and Computational Science* 5.1, pp. 65–80. ISSN: 2157-5452. URL: <http://msp.org/camcos/2010/5-1/p04.xhtml>.
- Greenbaum, Alexandra Z. et al. (Nov. 2014). “An image-plane algorithm for JWST’s non-redundant aperture mask data”. In: URL: <http://arxiv.org/abs/1411.3446>  
<http://dx.doi.org/10.1088/0004-637X/798/2/68>.
- Groenewegen, M. A. T. (Mar. 2006). “The mid- and far-infrared colours of AGB and post-AGB stars”. In: *Astronomy & Astrophysics* 448.1, pp. 181–187. ISSN: 0004-6361.
- Gustafsson, B. et al. (Aug. 2008). “A grid of MARCS model atmospheres for late-type stars”. In: *Astronomy & Astrophysics* 486.3, pp. 951–970. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361:200809724>.

- Guyon, Olivier (Sept. 2018). “Extreme Adaptive Optics”. In: *Annual Review of Astronomy and Astrophysics* 56.1, pp. 315–355. ISSN: 0066-4146. URL: <https://www.annualreviews.org/doi/10.1146/annurev-astro-081817-052000>.
- Hamacher, Duane W. et al. (2019). “Indigenous use of stellar scintillation to predict weather and seasonal change”. In: *Proceedings of the Royal Society of Victoria* 131.1, p. 24. ISSN: 0035-9211.
- Hamaker, J. P. et al. (May 1996). “Understanding radio polarimetry. III. Interpreting the IAU/IEEE definitions of the Stokes parameters”. In: *Astronomy and Astrophysics Supplement Series* 117.1, pp. 161–165. ISSN: 0365-0138. URL: <http://aas.aanda.org/10.1051/aas:1996147>.
- Hampson, Karen M. et al. (Oct. 2021). “Adaptive optics for high-resolution imaging”. In: *Nature Reviews Methods Primers* 1.1, p. 68. ISSN: 2662-8449. URL: <https://www.nature.com/articles/s43586-021-00066-7>.
- Hans, G et al. (Dec. 2013). “Theory of Circumstellar Dust Shells”. In: *Physics and Chemistry of Circumstellar Dust Shells*. Cambridge University Press, pp. 63–94.
- Hans-Peter, Gail (1941). *Physics and Chemistry of circumstellar dust shells*. 1st ed. Cammbridge. URL: [https://sydney.primo.exlibrisgroup.com/discovery/fulldisplay?docid=alma991031633740405106&context=L&vid=61USYD\\_INST:sydney&lang=en&search\\_scope=MyInst\\_and\\_CI&adaptor=Local%20Search%20Engine&tab=Everything&query=any,contains,Physics%20and%20Chemistry%20of%20Circumstellar%20Dust%20Shells&mode=basic](https://sydney.primo.exlibrisgroup.com/discovery/fulldisplay?docid=alma991031633740405106&context=L&vid=61USYD_INST:sydney&lang=en&search_scope=MyInst_and_CI&adaptor=Local%20Search%20Engine&tab=Everything&query=any,contains,Physics%20and%20Chemistry%20of%20Circumstellar%20Dust%20Shells&mode=basic).
- Harald Mutschke (2020). *Database of Optical Constants for Cosmic Dust*.
- Harrington, J. P. (1969). “The Intrinsic Polarization of Mira Variables”. In: *Astrophysical Letters*, Vol. 3, p.165.
- Hart, G. J. Joost ‘t et al. (Aug. 2021). “Full characterization of the instrumental polarization effects of the spectropolarimetric mode of SCEXAO-CHARIS”. In: URL: <http://arxiv.org/abs/2108.04833>.
- Hartquist, E et al. (2006). *Diffuse matter from Star Forming Regions to Active Galaxies*. Vol. 1, pp. 83–102.
- Hashimoto, J. et al. (Oct. 2012). “POLARIMETRIC IMAGING OF LARGE CAVITY STRUCTURES IN THE PRE-TRANSITIONAL PROTOPLANETARY DISK AROUND PDS 70: OBSERVATIONS OF THE DISK\*”. In: *The Astrophysical Journal Letters* 758.1, p. L19. ISSN: 2041-8205.
- Haubois, X. et al. (Dec. 2009). “Imaging the spotty surface of Betelgeuse in the H band”. In: *Astronomy & Astrophysics* 508.2, pp. 923–932. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/200912927>.
- Haubois, X. et al. (2019). “The inner dust shell of Betelgeuse detected by polarimetric Aperture-Masking interferometry”. In: *arXiv* 33258, pp. 1–9. ISSN: 23318422.
- Heidelberg-Jena Database (2024). *Heidelberg-Jena Database of Optical Constants (HEJDOC)*. URL: <https://www2.mpia-hd.mpg.de/HJPD0C/create.php>.
- Heras, A. M. et al. (Jan. 1997). “ISO-SWS Spectral Characterization of Stars”. In: *Astrophysics and Space Science* 255.1-2, pp. 251–253. ISSN: 0004-640X.
- Hofmann, K H (2005). “Interferometric Observations of the Mira Star O Ceti with the VLT/VINCI Instrument in the Near Infrared”. In.
- Hofmann, K.H et al. (1998). “Limb-darkening and radii of non-Mira M giant models”. In: *Astronomy and Astrophysics*.
- Höfner, S. (Nov. 2008). “Winds of M-type AGB stars driven by micron-sized grains”. In: *Astronomy & Astrophysics* 491.2, pp. L1–L4. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361:200810641>.
- Höfner, Susanne et al. (Nov. 2018). “Mass loss of stars on the asymptotic giant branch”. In: *The Astronomy and Astrophysics Review* 26.1, p. 1. ISSN: 0935-4956. URL: <http://link.springer.com/10.1007/s00159-017-0106-5>.
- Hogg, David W. et al. (May 2018). “Data Analysis Recipes: Using Markov Chain Monte Carlo\*”. In: *The Astrophysical Journal Supplement Series* 236.1, p. 11. ISSN: 0067-0049. URL: <https://iopscience.iop.org/article/10.3847/1538-4365/aab76e>.
- Holstein, R.G. van (2021). “High-contrast imaging polarimetry of exoplanets and circumstellar disk”. PhD thesis. Leiden University.
- Honda, M et al. (2022). “Subaru/IRCS L-band spectro-polarimetry of the HD142527 disk scattered light”. In: *Publications of the Astronomical Society of Japan* 74.4, pp. 851–856.
- Hough, James (June 2006). “Polarimetry: a powerful diagnostic tool in astronomy”. In: *Astronomy and Geophysics* 47.3, pp. 31–3. ISSN: 1366-8781.

- Huijser, David et al. (Mar. 2022). “Properties of the affine-invariant ensemble sampler’s ‘stretch move’ in high dimensions”. In: *Australian & New Zealand Journal of Statistics* 64.1, pp. 1–26. ISSN: 1369-1473. URL: <https://onlinelibrary.wiley.com/doi/10.1111/anzs.12358>.
- Humphreys, Roberta M. et al. (Mar. 2022). “Episodic Gaseous Outflows and Mass Loss from Red Supergiants”. In: *The Astronomical Journal* 163.3, p. 103. ISSN: 0004-6256. URL: <https://iopscience.iop.org/article/10.3847/1538-3881/ac46ff>.
- Hunziker, S et al. (2021). “HD142527: Quantitative disk polarimetry with SPHERE”. In: *Astronomy and Astrophysics* 648.
- Hutchison, Mark A. et al. (Dec. 2016). “On the maximum grain size entrained by photoevaporative winds”. In: *Monthly Notices of the Royal Astronomical Society* 463.3, pp. 2725–2734. ISSN: 0035-8711.
- Ireland, M. J. (Aug. 2013). “Phase errors in diffraction-limited imaging: contrast limits for sparse aperture masking”. In: *Monthly Notices of the Royal Astronomical Society* 433.2, pp. 1718–1728. ISSN: 1365-2966. URL: <http://academic.oup.com/mnras/article/433/2/1718/1751534/Phase-errors-in-diffractionlimited-imaging>.
- Ireland, M. J. et al. (2005). “Dust scattering in the Miras R Car and RR Sco resolved by optical interferometric polarimetry”. In: *Monthly Notices of the Royal Astronomical Society* 361.1, pp. 337–344. ISSN: 00358711.
- Jaeger, C. et al. (1998). “Steps toward interstellar silicate mineralogy. IV. The crystalline revolution”. In: *Astronomy and Astrophysics* 339, pp. 904–916.
- Jager, C et al. (1994). “Steps toward interstellar silicate mineralogy I. Laboratory results of a silicate glass of mean cosmic composition”. In: *Astronomy and Astrophysics* 292, pp. 641–655.
- Jäger, C. et al. (Sept. 2003a). “Steps toward interstellar silicate mineralogy”. In: *Astronomy & Astrophysics* 408.1, pp. 193–204. ISSN: 0004-6361. URL: <https://articles.adsabs.harvard.edu/pdf/1994A%26A...292..641J>.
- (Sept. 2003b). “Steps toward interstellar silicate mineralogy: VII. Spectral properties and crystallization behaviour of magnesium silicates produced by the sol-gel method”. In: *Astronomy & Astrophysics* 408.1, pp. 193–204. ISSN: 0004-6361.
- Jäger, C. et al. (Apr. 2003c). “Structural processing of enstatite by ion bombardment”. In: *Astronomy & Astrophysics* 401.1, pp. 57–65. ISSN: 0004-6361.
- Janet Akyuz Mattei (1997). “Introducing Mira Variables”. In: *JAAVSO* 25.
- Jeffers, S. V. et al. (Mar. 2012). “Direct imaging of a massive dust cloud around R Coronae Borealis”. In: *Astronomy & Astrophysics* 539, A56. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201117138>.
- Jeffreys, Harold (1998). *The Theory of Probability*. 1st ed. Vol. 1. Oxford Classic Texts in the Physical Sciences. ISBN: 9780198503682.
- Jenkins, F. (1981). *Fundamentals of Optics*. ISBN: 0-07-032330-5.
- Jennings, Jeff et al. (July 2020). “frankenstein: protoplanetary disc brightness profile reconstruction at sub-beam resolution with a rapid Gaussian process”. In: *Monthly Notices of the Royal Astronomical Society* 495.3, pp. 3209–3232. ISSN: 0035-8711.
- JMMC (n.d.). *OITools JMMC*.
- Josselin, E. et al. (July 2007). “Atmospheric dynamics and the mass loss process in red supergiant stars”. In: *Astronomy & Astrophysics* 469.2, pp. 671–680. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361:20066353>.
- Jumper, John et al. (Aug. 2021). “Highly accurate protein structure prediction with AlphaFold”. In: *Nature* 596.7873, pp. 583–589. ISSN: 0028-0836.
- Kama, M. et al. (Aug. 2009). “The Inner Rim Structures of Protoplanetary Discs”. In: URL: <http://arxiv.org/abs/0908.1692><http://dx.doi.org/10.1051/0004-6361/200912068>.
- Karovicova, I. et al. (Dec. 2013). “New insights into the dust formation of oxygen-rich AGB stars”. In: *Astronomy & Astrophysics* 560, A75. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201322376>.
- Kawaguchi, N. (1983). “Coherence loss and delay observation error in very-long-baseline interferometry”. In: *Radio Research Laboratories* 30, pp. 59–87. URL: <https://ui.adsabs.harvard.edu/abs/1983RaRLJ...30..59K/abstract>.
- Kelley, D et al. (2005). *Exploring Ancient Skies*. Vol. 1. Springer.
- Kerrison, Emily F. et al. (Dec. 2024). “From terrestrial weather to space weather through the history of scintillation”. In.

- Kervella, Pierre et al. (Jan. 2018). “The close circumstellar environment of Betelgeuse”. In: *Astronomy & Astrophysics* 609, A67. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/201731761>.
- Khouri, T. et al. (May 2015a). “Dusty wind of W Hydrae”. In: *Astronomy & Astrophysics* 577, A114. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201425092>.
- (May 2015b). “Dusty wind of WHydrae”. In: *Astronomy & Astrophysics* 577, A114. ISSN: 0004-6361.
- Khouri, T. et al. (July 2016). “Study of the inner dust envelope and stellar photosphere of the AGB star R Doradus using SPHERE/ZIMPOL”. In: *Astronomy & Astrophysics* 591, A70. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201628435>.
- Khouri, T. et al. (Mar. 2020). “Inner dusty envelope of the AGB stars W Hydrae, SW Virginis, and R Crateris using SPHERE/ZIMPOL”. In: *Astronomy & Astrophysics* 635, A200. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/201834618>.
- Khouri, T. et al. (May 2024). “An empirical view of the extended atmosphere and inner envelope of the asymptotic giant branch star R Doradus”. In: *Astronomy & Astrophysics* 685, A11. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/202348382>.
- Kipper, Tõnu (Jan. 2010). “On the Optical Spectrum of Mu Cephei”. In: *Open Astronomy* 19.3-4. ISSN: 2543-6376. URL: <https://www.degruyter.com/document/doi/10.1515/astro-2017-0418/html>.
- Kobayashi, Hiroshi et al. (Oct. 2011). “Sublimation temperature of circumstellar dust particles and its importance for dust ring formation”. In: *Earth, Planets and Space* 63.10, pp. 1067–1075. ISSN: 1343-8832. URL: <http://link.springer.com/10.5047/eps.2011.03.012>.
- Kolmogorov, A. (1941). “Energy dissipation in locally isotropic turbulence”. In: URL: [https://www.researchgate.net/publication/270592304\\_Energy\\_dissipation\\_in\\_locally\\_isotropic\\_turbulence](https://www.researchgate.net/publication/270592304_Energy_dissipation_in_locally_isotropic_turbulence).
- Kolokolova, L et al. (2015). *Polarimetry of Stars and Planetary Systems*. Vol. 1. United Kingdom: Cambridge University Press. URL: [https://assets.cambridge.org/97811070/43909/frontmatter/9781107043909\\_frontmatter.pdf?utm\\_source](https://assets.cambridge.org/97811070/43909/frontmatter/9781107043909_frontmatter.pdf?utm_source).
- Kostogryz, N. M. et al. (Mar. 2015). “Center-to-limb polarization in continuum spectra of F, G, K stars”. In: *Astronomy & Astrophysics* 575, A89. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201424844>.
- Kostogryz, N. M. et al. (Feb. 2016). “Center-to-limb variation of intensity and polarization in continuum spectra of FGK stars for spherical atmospheres”. In: *Astronomy & Astrophysics* 586, A87. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201527598>.
- Kouveliotou, C et al. (2014). “Enduring Quests-Daring Visions (NASA Astrophysics in the Next Three Decades)”. In: *ArXiv*.
- Kravchenko, K. et al. (Dec. 2019). “Tomography of cool giant and supergiant star atmospheres”. In: *Astronomy & Astrophysics* 632, A28. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/201935809>.
- Krisciunas, Kevin L. (1986). “Photometry of Alpha Orionis, CE Tauri, and Mu Cephei”. In: *The Journal of the American Association of Variable Star Observers* 15.1, pp. 15–21.
- Krizhevsky, A et al. (2012). “ImageNet Classification with Deep Convolutional Neural Networks”. In: *Advances in Neural Information Processing Systems 25 (NIPS 2012)*.
- Labeyrie, A. et al. (2006a). *Introduction to Optical Stellar Interferometry*.
- Labeyrie, A. et al. (2006b). *Introduction to Optical Stellar Interferometry*. ISBN: 9780521828727.
- Lacour, S. et al. (Aug. 2011). “Sparse aperture masking at the VLT”. In: *Astronomy & Astrophysics* 532, A72. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201116712>.
- Lacour, Sylvestre, Tuthill, Peter (2023). “Sparse Aperture Masking on Paranal”. In: URL: <https://www.eso.org/sci/publications/messenger/archive/no.146-dec11/messenger-no146-18-23.pdf>.
- Larsson-Leander, G. (1964). “Photoelectric observations of mu Cephei.” In: *Arkiv for Astron* 3, pp. 285–297.
- Lauer, Tod R. et al. (Mar. 1991). “The core of the nearby S0 galaxy NGC 7457 imaged with the HST planetary camera”. In: *The Astrophysical Journal* 369, p. L41. ISSN: 0004-637X. URL: <http://adsabs.harvard.edu/doi/10.1086/185954>.
- Le Borgne, J. et al. (1989). “The polarized dust envelope around the red supergiant  $\mu$  Cephei”. In: *Astronomy and astrophysics (Berlin. Print)* 210.1, pp. 198–210. ISSN: 0004-6361.
- LeCun, Y. et al. (Dec. 1989). “Backpropagation Applied to Handwritten Zip Code Recognition”. In: *Neural Computation* 1.4, pp. 541–551. ISSN: 0899-7667.

- Leech, K et al. (2003). *ISO Handbook Volume VI: The Short-Wavelength Spectrometer (SWS)*.
- Lenzen, Rainer et al. (Mar. 2003). “NAOS-CONICA first on sky results in a variety of observing modes”. In: *SPIE*. Ed. by Masanori Iye et al., p. 944. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.460044>.
- LeVan, P (1989). “Ten Micron Observations of Bright Circumstellar Shells - Spectral Properties and a Search for Extended Emission”. In: *Publications of the Astronomical Society of the Pacific*.
- Levis, A et al. (Sept. 2025). “Revealing Fine Structure in Protoplanetary Disks with Physics Constrained Neural Fields”. In: *ArXiv*.
- Lilley, L et al. (Dec. 2025a). “Beneath the shadows: characterising the circumstellar dust of RSG mu Cephei with VAMPIRES”. In: *A&A (In Prep)*.
- Lilley, L et al. (July 2025b). “PIRATES – a machine-learning framework for polarized, interferometric image reconstruction”. In: *JATIS (accepted)*.
- Lilley, Lucinda et al. (Aug. 2024). “Polarimetric, non-redundant aperture masking with next generation VAMPIRES: new instrumental capabilities, scientific outcomes, and image reconstruction techniques”. In: *Optical and Infrared Interferometry and Imaging IX*. Ed. by Stephanie Sallum et al. SPIE, p. 35. ISBN: 9781510675131. URL: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/13095/3018210/Polarimetric-non-redundant-aperture-masking-with-next-generation-VAMPIRES/10.1117/12.3018210.full>.
- Lin, Zhe-Yu Daniel et al. (Apr. 2022). “Thermal emission and scattering by aligned grains: Plane-parallel model and application to multiwavelength polarization of the HL Tau disc”. In: *Monthly Notices of the Royal Astronomical Society* 512.3, pp. 3922–3947. ISSN: 0035-8711.
- Lodders, Katharina (July 2003). “Solar System Abundances and Condensation Temperatures of the Elements”. In: *The Astrophysical Journal* 591.2, pp. 1220–1247. ISSN: 0004-637X.
- López Ariste, A. et al. (Dec. 2018). “Convective cells in Betelgeuse: imaging through spectropolarimetry”. In: *Astronomy & Astrophysics* 620, A199. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/201834178>.
- Lucas, Miles et al. (Nov. 2024). “Visible-light High-contrast Imaging and Polarimetry with SCEXAO VAMPIRES”. In: *Publications of the Astronomical Society of the Pacific* 136.11, p. 114504. ISSN: 0004-6280.
- Magalhães, Antônio Mário (1981). “Linear Polarization Changes across TiO Bands in Cool Variables: V CVn”. In: pp. 231–236. URL: [http://link.springer.com/10.1007/978-94-009-8492-9\\_23](http://link.springer.com/10.1007/978-94-009-8492-9_23).
- Martí-Vidal, I. et al. (Mar. 2014). “UVMULTIFIT: A versatile tool for fitting astronomical radio interferometric data”. In: *Astronomy & Astrophysics* 563, A136. ISSN: 0004-6361.
- Mathis, J. S. et al. (Oct. 1977a). “The size distribution of interstellar grains”. In: *The Astrophysical Journal* 217, p. 425. ISSN: 0004-637X.
- (Oct. 1977b). “The size distribution of interstellar grains”. In: *The Astrophysical Journal* 217, p. 425. ISSN: 0004-637X. URL: <http://adsabs.harvard.edu/doi/10.1086/155591>.
- Mauron N.; Querci, F. (1990). “Resonant scattering of NaI and KI in the circumstellar envelopes of alpha Orionis and MU Cephei : observations”. In: *Astronomy and Astrophysics* 86, p. 513.
- Mauron N., M Cailloux et al (1986). “First image of the circumstellar envelope of the red supergiant Mu Cep in the Na I 5890 A line.” In: *Astronomy and Astrophysics* 165.
- (1986). “First image of the circumstellar envelope of the red supergiant Mu Cep in the Na I 5890 A line.” In: *Astronomy and Astrophysics* 165.
- Mauron, N. et al. (Feb. 2011). “The mass-loss rates of red supergiants and the de Jager prescription”. In: *Astronomy & Astrophysics* 526, A156. ISSN: 0004-6361.
- Max Born, Emil Wolf (1975). *Principles of Optics*. Pergamon Press.
- McCulloch, Warren S. et al. (Dec. 1943). “A logical calculus of the ideas immanent in nervous activity”. In: *The Bulletin of Mathematical Biophysics* 5.4, pp. 115–133. ISSN: 0007-4985.
- McKemmish, Laura K et al. (Sept. 2019). “ExoMol molecular line lists – XXXIII. The spectrum of Titanium Oxide”. In: *Monthly Notices of the Royal Astronomical Society* 488.2, pp. 2836–2854. ISSN: 0035-8711.
- Mehner, Andrea (Aug. 2021a). “Observations of outflows of massive stars”. In: *Proceedings of the International Astronomical Union* 17.S370, pp. 37–55. ISSN: 1743-9213. URL: [https://www.cambridge.org/core/product/identifier/S1743921323000108/type/journal\\_article](https://www.cambridge.org/core/product/identifier/S1743921323000108/type/journal_article).

- Mehner, Andrea (Aug. 2021b). “Observations of outflows of massive stars”. In: *Proceedings of the International Astronomical Union* 17.S370, pp. 37–55. ISSN: 1743-9213. URL: [https://www.cambridge.org/core/product/identifiaer/S1743921323000108/type/journal\\_article](https://www.cambridge.org/core/product/identifiaer/S1743921323000108/type/journal_article).
- Mérand, Antoine (Aug. 2022). “Flexible spectro-interferometric modeling of OIFITS data with PMOIRED”. In: *Optical and Infrared Interferometry and Imaging VIII*. Ed. by Antoine Mérand et al. SPIE, p. 61. ISBN: 9781510653474.
- Meynet, G. et al. (Mar. 2015). “Impact of mass-loss on the evolution and pre-supernova properties of red supergiants”. In: *Astronomy & Astrophysics* 575, A60. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201424671>.
- Millour, F. (Sept. 2014). “Interferometry concepts”. In: *EAS Publications Series* 69-70. Ed. by F. Millour et al., pp. 17–52. ISSN: 1633-4760. URL: <http://www.eas-journal.org/10.1051/eas/1569003>.
- Millour, Florentin (Apr. 2008). “All you ever wanted to know about optical long baseline stellar interferometry, but were too shy to ask your adviser”. In: URL: <http://arxiv.org/abs/0804.2368%20http://dx.doi.org/10.1016/j.newar.2008.04.012>.
- Molster, F. (2000). “Crystalline silicates in circumstellar dust shells”. PhD thesis. Amsterdam, Netherlands: Sterrenkundig Instituut Anton Pannekoek.
- Molster, F. J. et al. (Jan. 2002a). “Crystalline silicate dust around evolved stars”. In: *Astronomy & Astrophysics* 382.1, pp. 184–221. ISSN: 0004-6361.
- Molster, F. J. et al. (Jan. 2002b). “Crystalline silicate dust around evolved stars III - A correlations study of crystalline silicate features”. In: *Astronomy & Astrophysics* 382.1, pp. 241–255. ISSN: 0004-6361.
- Molster, F.J et al. (1999). “The composition and nature of the dust shell surrounding the binary AFGL 4106???” In: *Astronomy and Astrophysics* 350, pp. 163–180.
- Monnier, J (2002). “Astrophysics with Closure Phases”. In: URL: <https://www.jmmc.fr/mirrors/obsvlti/book/Monnier.pdf>.
- Monnier, J. D. et al. (Feb. 2007a). “The Keck Aperture-masking Experiment: Near-Infrared Sizes of Dusty Wolf-Rayet Stars”. In: *The Astrophysical Journal* 655.2, pp. 1033–1045. ISSN: 0004-637X.
- Monnier, J.D (2003). *Astrophysics with Closure Phases*. Vol. 6. EAS Publications Series, Observing with the VLTI. URL: <https://www.jmmc.fr/mirrors/obsvlti/book/Monnier.pdf>.
- Monnier, John D. et al. (July 2007b). “Imaging the Surface of Altair”. In: *Science* 317.5836, pp. 342–345. ISSN: 0036-8075.
- Montargès, M et al. (May 2019). “NOEMA maps the CO J = 2 1 environment of the red supergiant  $\mu$  Cep”. In: *Monthly Notices of the Royal Astronomical Society* 485.2, pp. 2417–2430. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article/485/2/2417/5309994>.
- Montargès, M. et al. (Apr. 2016). “The close circumstellar environment of Betelgeuse”. In: *Astronomy & Astrophysics* 588, A130. ISSN: 0004-6361.
- Munoz-Sanchez, G. et al. (Oct. 2024a). “Episodic mass loss in the very luminous red supergiant [W60] B90 in the Large Magellanic Cloud”. In: *Astronomy & Astrophysics* 690, A99. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/202450737>.
- Munoz-Sanchez, G. et al. (Nov. 2024b). “The dramatic transition of the extreme Red Supergiant WOH G64 to a Yellow Hypergiant”. In: URL: <http://arxiv.org/abs/2411.19329>.
- NAOJ (2024). *Reflectivity of the Subaru Telescope’s Mirrors*. URL: <https://www.naoj.org/Observing/Telescope/Parameters/Reflectivity/>.
- National Academies of Sciences (2023). *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. Washington, DC: The National Academies Press.
- Norris, B. et al. (Jan. 2015a). “The VAMPIRES instrument: imaging the innermost regions of protoplanetary discs with polarimetric interferometry”. In: *Monthly Notices of the Royal Astronomical Society* 447.3, pp. 2894–2906.
- Norris, Barnaby R M (2015). “Secrets in Stellar Halos : Imaging Against the Glare”. PhD thesis. The University of Sydney.
- Norris, Barnaby R.M. et al. (2012). “A close halo of large transparent grains around extreme red giant stars”. In: *Nature* 484.7393, pp. 220–222. ISSN: 00280836.
- Norris, Ray P. et al. (2015b). “Australian Aboriginal Astronomy - An Overview”. In: *Handbook of Archaeoastronomy and Ethnoastronomy*. New York, NY: Springer New York, pp. 2215–2222.
- NumPyro (2019). *Markov Chain Monte Carlo (MCMC)*. URL: <https://num.pyro.ai/en/stable/mcmc.html>.

- O’Gorman, E. et al. (Jan. 2015). “ALMA observations of anisotropic dust mass loss in the inner circumstellar environment of the red supergiant VY Canis Majoris”. In: *Astronomy & Astrophysics* 573, p. L1. ISSN: 0004-6361.
- Ohnaka, K. (Aug. 2014). “Imaging the outward motions of clumpy dust clouds around the red supergiant Antares with VLT/VISIR”. In: *Astronomy & Astrophysics* 568, A17. ISSN: 0004-6361.
- Ohnaka, K. et al. (May 2016). “Clumpy dust clouds and extended atmosphere of the AGB star W Hydrae revealed with VLT/SPHERE-ZIMPOL and VLTI/AMBER”. In: *Astronomy & Astrophysics* 589, A91. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201628229>.
- (Jan. 2017a). “Clumpy dust clouds and extended atmosphere of the AGB star W Hydrae revealed with VLT/SPHERE-ZIMPOL and VLTI/AMBER”. In: *Astronomy & Astrophysics* 597, A20. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201629761>.
- (Aug. 2017b). “Vigorous atmospheric motion in the red supergiant star Antares”. In: *Nature* 548.7667, pp. 310–312. ISSN: 0028-0836. URL: <https://www.nature.com/articles/nature23445>.
- Ohnaka, K. et al. (Dec. 2024). “Imaging the innermost circumstellar environment of the red supergiant WOH G64 in the Large Magellanic Cloud”. In: URL: <http://arxiv.org/abs/2412.01921%20http://dx.doi.org/10.1051/0004-6361/202451820>.
- Papadakis, Emmanuel P. (Apr. 1965). “Revised Grain-Scattering Formulas and Tables”. In: *The Journal of the Acoustical Society of America* 37.4, pp. 703–710. ISSN: 0001-4966. URL: <https://pubs.aip.org/jasa/article/37/4/703/746897/Revised-Grain-Scattering-Formulas-and-Tables>.
- Perrin, G. et al. (June 2005). “Study of molecular layers in the atmosphere of the supergiant star  $\mu$  Cep by interferometry in the K band”. In: *Astronomy & Astrophysics* 436.1, pp. 317–324. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361:20042313>.
- Perrin, Marshall D. et al. (Dec. 2009a). “THE CASE OF AB AURIGAE’S DISK IN POLARIZED LIGHT: IS THERE TRULY A GAP?” In: *The Astrophysical Journal* 707.2, pp. L132–L136. ISSN: 0004-637X.
- (Dec. 2009b). “THE CASE OF AB AURIGAE’S DISK IN POLARIZED LIGHT: IS THERE TRULY A GAP?” In: *The Astrophysical Journal* 707.2, pp. L132–L136. ISSN: 0004-637X.
- Perrin, Marshall D. et al. (Jan. 2015). “POLARIMETRY WITH THE GEMINI PLANET IMAGER: METHODS, PERFORMANCE AT FIRST LIGHT, AND THE CIRCUMSTELLAR RING AROUND HR 4796A”. In: *The Astrophysical Journal* 799.2, p. 182. ISSN: 1538-4357.
- Pesce, Dominic W. (Apr. 2021). “A D-term Modeling Code (DMC) for Simultaneous Calibration and Full-Stokes Imaging of Very Long Baseline Interferometric Data”. In: *The Astronomical Journal* 161.4, p. 178. ISSN: 0004-6256.
- Pilate, Q et al. (2025). “First images of Antares photosphere from spectropolarimetry”. In: *Astronomy & Astrophysics*.
- Pilate, Q. et al. (Nov. 2024). “The variability of Betelgeuse explained by surface convection”. In: *Astronomy & Astrophysics* 691, A297. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/202450987>.
- Pinte, C. et al. (Dec. 2006). “Monte Carlo radiative transfer in protoplanetary disks”. In: *Astronomy & Astrophysics* 459.3, pp. 797–804. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361:20053275>.
- Pinte, C. et al. (May 2009). “Benchmark problems for continuum radiative transfer”. In: *Astronomy & Astrophysics* 498.3, pp. 967–980. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/200811555>.
- Pitman, Karly M. et al. (Mar. 2013). “Revisiting astronomical crystalline forsterite in the UV to near-IR”. In: *Earth, Planets and Space* 65.3, pp. 129–138. ISSN: 1343-8832. URL: <http://link.springer.com/10.5047/eps.2012.05.009>.
- Planquart, L. et al. (May 2024). “An impressionist view of V Hydrae. When MATISSE paints Asymmetric Giant Blobs”. In: URL: <http://arxiv.org/abs/2405.07821>.
- Polyakova, T. A. (Jan. 2003). “Variations in the Brightness and Polarization of  $\mu$  CEP”. In: *Astrophysics* 46.1, pp. 18–23. ISSN: 0571-7256.
- Prahl, S (2025). *miepython: A Python library for Mie scattering calculations*. URL: <https://github.com/scottprahl/miepython>.
- Quanz, Sascha P. et al. (Sept. 2011). “VERY LARGE TELESCOPE/NACO POLARIMETRIC DIFFERENTIAL IMAGING OF HD100546—DISK STRUCTURE AND DUST GRAIN PROPERTIES BETWEEN 10 AND 140 AU”. In: *The Astrophysical Journal* 738.1, p. 23. ISSN: 0004-637X.
- Querry, M.R (1985). *Optical Constants*. Tech. rep. CRDC-CR-85034.

- Quirrenbach, A. et al. (Mar. 1993). “Angular diameter measurements of cool giant stars in strong TiO bands and in the continuum”. In: *The Astrophysical Journal* 406, p. 215. ISSN: 0004-637X. URL: <http://adsabs.harvard.edu/doi/10.1086/172432>.
- Rauer, H. et al. (Nov. 2014). “The PLATO 2.0 mission”. In: *Experimental Astronomy* 38.1-2, pp. 249–330. ISSN: 0922-6435.
- Readhead, A. et al. (Apr. 1988). “Diffraction Limited Imaging with Ground Based Optical Telescopes”. In: *The Astronomical Journal* 95, pp. 1278–1296. URL: [file:///Users/llil9854/Downloads/1988AJ\\_95\\_1278R.pdf](file:///Users/llil9854/Downloads/1988AJ_95_1278R.pdf).
- Richards, A. M. S. et al. (Sept. 1996). “Proper motions of water vapour masers and bipolar outflow from NML Cygni”. In: *Monthly Notices of the Royal Astronomical Society* 282.2, pp. 665–676. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/282.2.665>.
- Rosenblatt, F. (1958). “The perceptron: A probabilistic model for information storage and organization in the brain.” In: *Psychological Review* 65.6, pp. 386–408. ISSN: 1939-1471.
- Rucks, Melinda J. et al. (Apr. 2022). “Visible to Mid-Infrared Optical Constants of Orthopyroxenes”. In: *Earth and Space Science* 9.4. ISSN: 2333-5084.
- Ruggles, C (2015). *Handbook of Archaeoastronomy and Ethnoastronomy*. Ed. by Clive L.N. Ruggles. New York, NY: Springer New York. ISBN: 978-1-4614-6140-1.
- Rumelhart, David E. et al. (Oct. 1986). “Learning representations by back-propagating errors”. In: *Nature* 323.6088, pp. 533–536. ISSN: 0028-0836.
- Russell, R. W. et al. (May 1975). “Spectrophotometric observations of MU Cephei and the moon from 4 to 8 microns”. In: *The Astrophysical Journal* 198, p. L41. ISSN: 0004-637X. URL: <http://adsabs.harvard.edu/doi/10.1086/181806>.
- Sacuto, S. et al. (Mar. 2013). “The wind of the M-type AGB star RT Virginis probed by VLTI/MIDI”. In: *Astronomy & Astrophysics* 551, A72. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201220524>.
- Safonov, Boris et al. (Apr. 2019). “Differential speckle polarimetry at Cassegrain and Nasmyth foci”. In: *Monthly Notices of the Royal Astronomical Society* 484.4, pp. 5129–5141. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article/484/4/5129/5320684>.
- Sciicluna, P. et al. (Dec. 2015). “Large dust grains in the wind of VY Canis Majoris”. In: *Astronomy & Astrophysics* 584, p. L10. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201527563>.
- Selin, H (2000). *Astronomy Across Cultures*. Ed. by Helaine Selin et al. Vol. 1. Dordrecht: Springer Netherlands. ISBN: 978-94-010-5820-9.
- Setterholm, Benjamin R. et al. (Dec. 2020). “MIRC-X polarinterferometry at CHARA”. In: *Optical and Infrared Interferometry and Imaging VII*. Ed. by Antoine Mérand et al. SPIE, p. 24. ISBN: 9781510636798. URL: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11446/2562407/MIRC-X-polarinterferometry-at-CHARA/10.1117/12.2562407.full>.
- Sharpless, S et al. (1966). “An Analysis of the Light Curve of  $\mu$  Cephei”. In.
- Shenoy, Dinesh et al. (Feb. 2016). “SEARCHING FOR COOL DUST IN THE MID-TO-FAR INFRARED: THE MASS-LOSS HISTORIES OF THE HYPERGIANTS  $\mu$  Cep, VY CMa, IRC+10420, AND  $\rho$  Cas”. In: *The Astronomical Journal* 151.3, p. 51. ISSN: 1538-3881. URL: <https://iopscience.iop.org/article/10.3847/0004-6256/151/3/51>.
- Shenoy, Dinesh P. et al. (Sept. 2013). “ADAPTIVE OPTICS IMAGING OF VY CANIS MAJORIS AT 2-5  $\mu$ m WITH LBT/LMIRCam”. In: *The Astronomical Journal* 146.4, p. 90. ISSN: 0004-6256.
- Shenoy, Dinesh P. et al. (June 2015). “PROBING HYPERGIANT MASS LOSS WITH ADAPTIVE OPTICS IMAGING AND POLARIMETRY IN THE INFRARED: MMT-Pol AND LMIRCam OBSERVATIONS OF IRC +10420 AND VY CANIS MAJORIS”. In: *The Astronomical Journal* 150.1, p. 15. ISSN: 1538-3881.
- Shenoy, Dinesh Prabhakar (2016). “A Study of Hypergiant Mass Loss in the Near-To-Mid Infrared: VY CMa, IRC +10420, Mu Cep and Rho Cas”. PhD thesis. University of Minnesota.
- Shrestha, Manisha et al. (Dec. 2020). “Polarization simulations of stellar wind bow shock nebulae – II. The case of dust scattering”. In: *Monthly Notices of the Royal Astronomical Society* 500.4, pp. 4319–4337. ISSN: 0035-8711.
- Silvia D, S (2008). *Data Analysis, A Bayesian Tutorial*.

- Sivaramakrishnan, Anand et al. (Aug. 2014). “Non-redundant masking ideas on JWST”. In: ed. by Jacobus M. Oschmann et al., 91433S. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2056639>.
- Sloan, G. C. et al. (Aug. 2003). “A Uniform Database of 2.4–45.4 Micron Spectra from the *J* Infrared Space Observatory *J* Short Wavelength Spectrometer”. In: *The Astrophysical Journal Supplement Series* 147.2, pp. 379–401. ISSN: 0067-0049.
- Smartt, S. J. (Apr. 2015). “Observational Constraints on the Progenitors of Core-Collapse Supernovae: The Case for Missing High-Mass Stars”. In: *Publications of the Astronomical Society of Australia* 32, e016. ISSN: 1323-3580. URL: [https://www.cambridge.org/core/product/identifier/S132335801500017X/type/journal\\_article](https://www.cambridge.org/core/product/identifier/S132335801500017X/type/journal_article).
- Smartt, S. J. et al. (May 2009). “The death of massive stars - I. Observational constraints on the progenitors of Type II-P supernovae”. In: *Monthly Notices of the Royal Astronomical Society* 395.3, pp. 1409–1437. ISSN: 00358711. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2009.14506.x>.
- Smith, Nathan (Aug. 2014). “Mass Loss: Its Effect on the Evolution and Fate of High-Mass Stars”. In: *Annual Review of Astronomy and Astrophysics* 52.1, pp. 487–528. ISSN: 0066-4146.
- Snik, Frans et al. (2013). “Astronomical Polarimetry: Polarized Views of Stars and Planets”. In: *Planets, Stars and Stellar Systems*. Dordrecht: Springer Netherlands, pp. 175–221. URL: [http://link.springer.com/10.1007/978-94-007-5618-2\\_4](http://link.springer.com/10.1007/978-94-007-5618-2_4).
- Snik, Frans et al. (July 2018). “Review of high-contrast imaging systems for current and future ground-based and space-based telescopes III: technology opportunities and pathways”. In: *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation III*. Ed. by Roland Geyl et al. SPIE, p. 91. ISBN: 9781510619654.
- Soker, N. et al. (2002). “Stellar structure and mass loss during the early post asymptotic giant branch”. PhD thesis. Israel.
- Soulain, Anthony et al. (2020). “The James Webb Space Telescope aperture masking interferometer”. In: 1144611. December 2020, p. 62. ISSN: 1996756X.
- Srama, Ralf et al. (Mar. 2009). “Sample return of interstellar matter (SARIM)”. In: *Experimental Astronomy* 23.1, pp. 303–328. ISSN: 0922-6435.
- Stancliffe, Richard J. et al. (Feb. 2007). “Mass loss and yield uncertainty in low-mass asymptotic giant branch stars”. In: *Monthly Notices of the Royal Astronomical Society* 375.4, pp. 1280–1290. ISSN: 00358711. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2006.11363.x>.
- Stark, Christopher C. et al. (Sept. 2024). “Paths to robust exoplanet science yield margin for the Habitable Worlds Observatory”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 10.03. ISSN: 2329-4124. URL: <https://www.spiedigitallibrary.org/journals/Journal-of-Astronomical-Telescopes-Instruments-and-Systems/volume-10/issue-03/034006/Paths-to-robust-exoplanet-science-yield-margin-for-the-Habitable/10.1117/1.JATIS.10.3.034006.full>.
- Stokes, G. (1851). “On the Composition and Resolution of Streams of Polarized Light from different Sources”. In: *Transactions of the Cambridge Philosophical Society*.
- Tang, Ya-Wen et al. (May 2017). “Planet Formation in AB Aurigae: Imaging of the Inner Gaseous Spirals Observed inside the Dust Cavity”. In: *The Astrophysical Journal* 840.1, p. 32. ISSN: 0004-637X.
- Tazzari, Marco et al. (June 2018). “GALARIO: a GPU accelerated library for analysing radio interferometer observations”. In: *Monthly Notices of the Royal Astronomical Society* 476.4, pp. 4527–4542. ISSN: 0035-8711.
- Ten Brummelaar, T (1999). “A Pictorial Essay of the CHARA Array”. In: *Proceedings from ASP Conference Vol. 194*.
- The CASA Team et al. (Nov. 2022). “CASA, the Common Astronomy Software Applications for Radio Astronomy”. In: *Publications of the Astronomical Society of the Pacific* 134.1041, p. 114501. ISSN: 0004-6280.
- Thompson, Madeleine B. (Oct. 2010). “A Comparison of Methods for Computing Autocorrelation Time”. In: URL: <http://arxiv.org/abs/1011.0175>.
- Trippe, Sascha (Feb. 2014). “POLARIZATION AND POLARIMETRY: A REVIEW”. In: *Journal of The Korean Astronomical Society* 47.1, pp. 15–39. ISSN: 1225-4614.

- Tsuji, T. (Aug. 2000). “Water in Emission in the ISO Spectrum of the Early M Supergiant Star  $\mu$  Cephei”. In: URL: <http://arxiv.org/abs/astro-ph/0008058><http://dx.doi.org/10.1086/312879>.
- Turner, David G (2014). “Why Visual Observations of  $\mu$  Cephei are Important Why Visual Observations of  $\mu$  Cephei are Important”. In: May 2012.
- Tuthill, P (1994). “Peter Tuthill’s PhD Thesis”. In.
- Tuthill, P. et al. (Apr. 2000a). “Michelson Interferometry with the Keck I Telescope”. In: *Publications of the Astronomical Society of the Pacific* 112.770, pp. 555–565.
- Tuthill, P. G. et al. (Apr. 2000b). “Michelson Interferometry with the Keck I Telescope”. In: *Publications of the Astronomical Society of the Pacific* 112.770, pp. 555–565. ISSN: 0004-6280.
- Tuthill, Peter et al. (June 2006). “Sparse-aperture adaptive optics”. In: 62723A. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.672342>.
- Tuthill, Peter et al. (July 2010a). “Sparse aperture masking (SAM) at NAOS/CONICA on the VLT”. In: *Ground-based and Airborne Instrumentation for Astronomy III*. Ed. by Ian S McLean et al. SPIE.
- Tuthill, Peter et al. (June 2010b). “Sparse Aperture Masking (SAM) at NAOS/CONICA on the VLT”. In: URL: <http://arxiv.org/abs/1006.2586><http://dx.doi.org/10.1117/12.856806>.
- Ulyanov, Dmitry et al. (Nov. 2020). “Deep Image Prior”. In: *International Journal of Computer Vision* 128, pp. 1867–1888. URL: <http://arxiv.org/abs/1711.10925><http://dx.doi.org/10.1007/s11263-020-01303-4>.
- Vandeportal, Julien et al. (Mar. 2019). “Polarization of stars with debris discs: comparing observations with models”. In: *Monthly Notices of the Royal Astronomical Society* 483.3, pp. 3510–3528. ISSN: 0035-8711.
- Vermeulen, O. et al. (Aug. 2025). “Impact of a binary companion in AGB outflows on CO spectral lines”. In: *Astronomy & Astrophysics* 700, A85. ISSN: 0004-6361.
- Vink, Jorick S. et al. (Oct. 2023). “Exploring the Red Supergiant wind kink”. In: *Astronomy & Astrophysics* 678, p. L3. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/202347801>.
- Vlemmings, Wouter et al. (Sept. 2024). “One month convection timescale on the surface of a giant evolved star”. In: *Nature* 633.8029, pp. 323–326. ISSN: 0028-0836. URL: <https://www.nature.com/articles/s41586-024-07836-9>.
- Wahhaj, Z. et al. (July 2024). “PDS 70 unveiled by star-hopping: Total intensity, polarimetry, and millimeter imaging modeled in concert”. In: *Astronomy & Astrophysics* 687, A257. ISSN: 0004-6361.
- Walmswell, Joseph J. et al. (Jan. 2012). “Circumstellar dust as a solution to the red supergiant supernova progenitor problem”. In: *Monthly Notices of the Royal Astronomical Society* 419.3, pp. 2054–2062. ISSN: 00358711. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2011.19860.x>.
- Wang, Zihao et al. (May 2023). “Theoretical Analysis of Inductive Biases in Deep Convolutional Networks”. In: *arXiv*. URL: <http://arxiv.org/abs/2305.08404>.
- Weingartner, Joseph C. et al. (Feb. 2001). “Dust Grain-Size Distributions and Extinction in the Milky Way, Large Magellanic Cloud, and Small Magellanic Cloud”. In: *The Astrophysical Journal* 548.1, pp. 296–309. ISSN: 0004-637X.
- Wilson, R. W. et al. (Nov. 1997). “The changing face of Betelgeuse”. In: *Monthly Notices of the Royal Astronomical Society* 291.4, pp. 819–826. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/291.4.819>.
- Winters, J.M et al. (1997). “Circumstellar dust shells around long-period variables”. In: *Astronomy and Astrophysics* 326, pp. 305–317.
- Wit, S. de et al. (Sept. 2024). “Investigating episodic mass loss in evolved massive stars”. In: *Astronomy & Astrophysics* 689, A46. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/202449607>.
- Wit, W J de et al. (Sept. 2008a). “A Red Supergiant Nebula at 25  $\mu$ m: Arcsecond-Scale Mass-Loss Asymmetries of  $\mu$  Cephei”. In: *The Astrophysical Journal* 685.1, pp. L75–L78.
- (Sept. 2008b). “A Red Supergiant Nebula at 25  $\mu$ m: Arcsecond-Scale Mass-Loss Asymmetries of  $\mu$  Cephei”. In: *The Astrophysical Journal* 685.1, pp. L75–L78. ISSN: 0004-637X. URL: <https://iopscience.iop.org/article/10.1086/592384>.
- Wittkowski, M. et al. (May 2011). “The extended atmospheres of Mira variables probed by VLTI, VLBA, and APEX”. In: URL: <http://arxiv.org/abs/1105.3386>.

- Wittkowski, M. et al. (Apr. 2012). “Fundamental properties and atmospheric structure of the red supergiant VY Canis Majoris based on VLTI/AMBER spectro-interferometry”. In: *Astronomy & Astrophysics* 540, p. L12. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201219126>.
- Wittkowski, M. et al. (June 2014). “On the atmospheric structure and fundamental parameters of red supergiants”. In: *Proceedings of the International Astronomical Union* 9.S307, pp. 280–285. ISSN: 1743-9213. URL: [https://www.cambridge.org/core/product/identifier/S1743921314006930/type/journal\\_article](https://www.cambridge.org/core/product/identifier/S1743921314006930/type/journal_article).
- Wittkowski, M. et al. (Mar. 2016). “Near-infrared spectro-interferometry of Mira variables and comparisons to 1D dynamic model atmospheres and 3D convection simulations”. In: *Astronomy & Astrophysics* 587, A12. ISSN: 0004-6361. URL: <http://www.aanda.org/10.1051/0004-6361/201527614>.
- Wong, P. (2023). “Artificial Intelligence for Astronomical Imaging”. PhD thesis. URL: <https://ses.library.usyd.edu.au/handle/2123/30068>.
- Woodruff, H. C. et al. (Jan. 2008). “The Keck Aperture Masking Experiment: Multiwavelength Observations of Six Mira Variables”. In: *The Astrophysical Journal* 673.1, pp. 418–433. ISSN: 0004-637X.
- Wolf, N.J. (1969). “Circumstellar Infrared Emission from Cool Stars”. In: *Journal of Chemical Information and Modeling* 155.9, pp. 1689–1699. ISSN: 1098-6596.
- Wootten, A. et al. (Aug. 2009). “The Atacama Large Millimeter/Submillimeter Array”. In: *Proceedings of the IEEE* 97.8, pp. 1463–1471. ISSN: 0018-9219.
- Yang, Haifeng et al. (Nov. 2017). “Scattering-produced (sub)millimetre polarization in inclined discs: optical depth effects, near–far side asymmetry and dust settling”. In: *Monthly Notices of the Royal Astronomical Society* 472.1, pp. 373–388. ISSN: 0035-8711.
- Yang, Ming et al. (Aug. 2023). “Evolved massive stars at low-metallicity”. In: *Astronomy & Astrophysics* 676, A84. ISSN: 0004-6361. URL: <https://www.aanda.org/10.1051/0004-6361/202244770>.
- Yasuda, Yuki et al. (Feb. 2012). “FORMATION OF SiC GRAINS IN PULSATION-ENHANCED DUST-DRIVEN WIND AROUND CARBON-RICH ASYMPTOTIC GIANT BRANCH STARS”. In: *The Astrophysical Journal* 745.2, p. 159. ISSN: 0004-637X.
- Young, J. S. et al. (July 2000). “New views of Betelgeuse: multi-wavelength surface imaging and implications for models of hotspot generation”. In: *Monthly Notices of the Royal Astronomical Society* 315.3, pp. 635–645. ISSN: 0035-8711. URL: <https://academic.oup.com/mnras/article/315/3/635/972357>.
- Ysard, N. et al. (Sept. 2018). “The optical properties of dust: the effects of composition, size, and structure”. In: *Astronomy & Astrophysics* 617, A124. ISSN: 0004-6361.
- Zapartas, E. et al. (Oct. 2024). “The effect of mass loss in models of red supergiants in the Small Magellanic Cloud”. In: URL: <http://arxiv.org/abs/2410.07335>.
- Zeidler, S. et al. (May 2013). “Optical constants of refractory oxides at high temperatures”. In: *Astronomy & Astrophysics* 553, A81. ISSN: 0004-6361.
- Zhang, Manxuan et al. (Oct. 2023). “Characterizing the instrumental polarization of SCExAO VAMPIRES”. In: *Techniques and Instrumentation for Detection of Exoplanets XI*. Ed. by Garreth J. Ruane. SPIE, p. 29. ISBN: 9781510665743.
- Zhao-Geisler, R et al. (2012). “Dust and molecular shells in asymptotic giant branch stars - Mid-infrared interferometric observations of R Aql, R Aqr, R Hya, W Hya and V Hya”. In: *Arxiv*.
- Zurlo, A (2024). “Direct imaging of exoplanets”. In: *Arxiv*. URL: <https://arxiv.org/abs/2404.05797>.