

THE UNIVERSITY OF
SYDNEY

**From Stars to Sensors: Differentiable Methods for
Precision Astronomy**

Louis Desdoigts

This thesis is presented as part of the requirements for the conferral of the degree:

Doctor of Philosophy (PhD)

Supervisors:

Prof Peter Tuthill & A/Prof Benjamin Pope

The University of Sydney
School of Physics

2025

Declaration

I, *Louis Desdoigts*, declare that this thesis is submitted in partial fulfilment of 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. During the preparation of this thesis the author used ChatGPT for the purposes of text enhancement. The use of this generative AI tool includes improved paraphrasing, sentence structure, and spelling. The author confirms that where text was modified by generative AI, the content was reviewed for possible errors, inaccuracies, and bias. The author takes full responsibility for the submitted thesis and ensures the work is their own and has used generative AI within the parameters of use the University of Sydney generative AI guide for researchers. Each chapter is comprised of a first-author paper that is either accepted or under review that. They each include a title page that lays out the specific contributions made by myself and co-authors. This research was supported by the Research Training Program fee offset, the Faculty of Science Research Stipend Scholarship, and the Hunstead fund. This document has not been submitted for qualifications at any other academic institution.

Louis Desdoigts

January 20, 2026

Abstract

This thesis introduces novel methods in astronomical data analysis based on differentiable forward modelling, enabling physical processes and statistical inference to be treated within a unified computational framework. By embedding the full chain of optical propagation and detector physics into differentiable software, gradients can be propagated from photon arrival to detector readout, supporting principled calibration and inference. This approach is formalised in the “Pixels to Planets” philosophy: an end-to-end modelling framework that replaces traditional modular pipelines with continuous, interpretable models grounded in physics yet flexible enough for real data. It reframes data analysis as differentiable simulation and optimisation, enabling physical calibration and scientific inference with the same statistical model.

A key outcome is the development of ∂ Lux, an open-source library for differentiable optical modelling built in JAX. ∂ Lux supports efficient gradient computation through physically accurate models of diffractive optics and detectors, enabling uncertainty propagation and optimisation across the instrument lifecycle — from design to calibration and through analysis of data products. These tools are first applied to the TOLIMAN mission, a low-cost astrometric cubesat telescope targeting Earth-mass planets around the α Centauri star system by witnessing micro-arcsecond angular perturbations. By optimising diffractive pupil designs with respect to Fisher Information, we design an information-theoretically optimal experiment and allow instrument calibration directly from science exposures.

The same modelling philosophy is then extended to JWST’s NIRISS Aperture Masking Interferometry mode, which has not met pre-launch expectations for signal-to-noise, primarily due to limitations arising from uncorrected non-linear detector effects. The thesis presents a hybrid forward model, AMIGO, developed to jointly model optics, charge redistribution, and readout. It features a novel embedded neural network implemented within a wider physically-constrained ‘effective detector model’. Trained on raw JWST data, AMIGO recovers optical aberrations, emulates pixel-level systematics, and extracts stable interferometric observables from degraded exposures. These techniques enable the recovery of kernel amplitude and phase observables at precisions exceeding all prior methods, restoring the competitive scientific viability of AMI for high-contrast imaging at the diffraction limit.

More broadly, the thesis advocates for end-to-end differentiable models as a foundational framework for modern astronomical instrumentation. By unifying optical propagation, detector physics, and statistical inference within a single differentiable architecture, these models enable rigorous uncertainty quantification, principled calibration, and scalable optimisation. When confronted with systematics too complex or poorly understood to model explicitly — such as non-linear charge migration — this framework can be naturally extended with hybrid physical–machine learning components, embedding neural networks within interpretable physical structures. This approach preserves interpretability and causal structure meeting the challenges of real-world data. Differentiable modelling will be essential for next-generation observatories such as *Roman* and the *Habitable Worlds Observatory*, whose calibration demands exceed the capabilities of conventional pipelines.

Acknowledgements

Peter and Ben, I can hardly believe what we've achieved here. You have both gone above and beyond what anyone could expect of a supervisor. Your guidance, support, and mentorship extended far beyond academia. I truly could not have achieved this thesis with anyone else. Peter, you believed in me long before I ever did. Ben, your tireless support, no matter the hour or time zone, has been extraordinary. I think it's finally time, after all these years, that the beers are on me.

Mama and Papa, your support over the years has been immeasurable. You've been my silent cheer squad and my biggest fans. The Rossos, the lifts, the dinners, and your belief in me have meant more than you could know. You've helped me achieve things I never imagined, and this PhD is as much yours as it is mine.

Gabs and James, you've always been there keeping me sane, grounded, and with perspective. The gin nights, the games, the parties, the late night science rambles, thank you for enduring it all. Your support has meant the world.

Max, you've been more instrumental than you know. You've been such a kind presence over the years and made me want to actually come into uni after far too many years stuck at home on crutches. You've always been there to call on, especially at the 11th hour and have suffered through far too many of my manic rants about politics, maths, science, or my personal life. I'm blessed to have been able to work with you, and I hope we always do.

Becca, you came into my life at exactly the right time. I've never been able to lean on someone in the way I have on you. You truly helped me actually get over the line at the end of it all, and I will always be grateful for that.

Kez, Al, Giersz, Oli, Will, Kiesh, Loz, Bronte, Kat, Frances, Alex, Isaac, Emma, Camby, and the many Dans and Daniels. I could never list everyone who's been there. Your belief in me, the nights we've shared, and the memories we've made are among my most cherished. I can't wait to finally be able to spend the time you that you all deserve, you've all made this journey unforgettable.

To the PhD kiddies, Adam, Grace, Connor, Steph, Luci, Shrish, Coanire, Ryan, Hugh, Shashank, and Weeks. You all made this experience so much fun. You reminded me to enjoy the ride, let me drag you to the pub, and created moments I'll never forget. I can't wait to see all that you achieve.

I honour the Traditional Custodians of the lands on which this research was conducted and acknowledge their enduring resistance to dispossession and erasure. This research was completed during a time marked by devastating human suffering, reminding us that the pursuit of knowledge does not occur in a vacuum. May this work remain mindful of all peoples whose lives are shaped by violence, injustice, and the denial of freedom.

Contents

Abstract	iii
Acknowledgements	iv
List of Figures	vii
List of Tables	ix
1 Introduction	1
1.1 Exoplanets	3
1.2 Direct Imaging	5
1.2.1 High Contrast Imaging	6
1.2.2 Coronagraphs	7
1.2.3 Interferometry	8
1.2.3.1 Closure and Kernel Phases	9
1.2.3.2 Aperture Masking Interferometry	9
1.3 Modern Precision Astronomical Imaging Systems	10
1.3.1 The James Webb Space Telescope	11
1.3.1.1 NIRISS AMI Mode	12
1.3.1.2 H2RG Infrared Detectors	14
1.3.2 Toliman	15
1.4 Image Formation in Optical Systems	17
1.4.1 Physical Optics	17
1.4.2 Interferometry	19
1.4.2.1 Resolution Limits of Diffractive Systems	19
1.4.2.2 Development of Interferometric Methods	20
1.4.2.3 The van Cittert-Zernike Theorem	22
1.4.2.4 Interferometric Analysis & Kernel Spaces	24
1.5 Automatic Differentiation	25
1.5.1 Jacobians and The Chain Rule	27
1.5.1.1 The Adjoint Method: Reverse-mode	27
1.5.1.2 The Tangent Method: Forwards-mode	28
1.5.2 Differentiable Programming	28
1.5.3 Neural Networks	30
1.5.4 Differentiable Models	31
1.6 Modelling Paradigms for Scientific Inference	32
1.6.1 Inverse Modelling	33
1.6.2 Machine Learning	34
1.6.3 Forward Modelling	34

1.6.4	Hybrid Modelling	36
1.6.5	From Pixels to Planets	36
2	Deep Calibration of Flat Field and Phase Retrieval with Automatic Differentiation	38
2.1	Statement of Contribution	39
2.2	Introduction	40
2.2.1	Phase Retrieval	40
2.2.2	Detector Calibration	41
2.3	Differentiable Optical Models	42
2.4	Optical and Instrumental Modelling	43
2.5	Results	46
2.5.1	Optimisation	46
2.5.2	Phase Retrieval & Flat Field Estimation	47
2.5.3	Uncertainty Estimation	47
2.6	Noise & Performance Analysis	50
2.7	Conclusions & Future Work	53
2.8	Acknowledgements	54
3	Optical Design Maximising Fisher Information	55
3.1	Statement of Contribution	56
3.2	Introduction	57
3.3	Fisher Information and Bayesian Experimental Design	59
3.4	Comparison with Theory	60
3.5	Optical Design Case Study: the Toliman diffractive pupil	62
3.6	Discussion	71
3.7	Conclusion	71
3.8	Code, Data, and Materials	72
4	Amigo: a Data-Driven Calibration of the JWST Interferometer	73
4.1	Statement of Contribution	74
4.2	Introduction	75
4.2.1	JWST Aperture Masking Interferometer (AMI): A unique space-based interferometer	76
4.2.2	Differentiable Forward Models: From Pixels to Planets	78
4.3	The AMIGO Model & Pipeline	79
4.3.1	Calibration Data	80
4.3.2	Data Processing Pipeline	82
4.4	Optical & Visibility Model	84
4.4.1	Optical Model	84
4.4.1.1	Chromaticity and Spectral Model	84
4.4.1.2	Aperture Model	85
4.4.1.3	Wavefront Model	87
4.4.1.4	Propagation Model	87
4.4.2	Interferometric Visibility Model	89
4.4.2.1	Latent Visibility Model	90
4.4.2.2	Visibility Amplitudes & Phases Coupling	91
4.4.2.3	Kernel Amplitudes & Phases	91
4.5	The Effective Detector Model	95

4.5.1	Linear Detector Model	97
4.5.2	Non-Linear Ramp Model	97
4.5.2.1	Neural Network Implementation & Architecture	98
4.5.3	Electronics & Read model	99
4.5.3.1	Dark Current	99
4.5.3.2	Inter-Pixel Capacitance	101
4.5.3.3	Non-Linear Gain	101
4.5.3.4	1/ f Noise	101
4.6	Base Model Training and Calibration	101
4.7	Inference from Science Data	104
4.7.1	Observable Extraction Workflow	105
4.7.1.1	Data processing	106
4.7.2	Model Fitting	106
4.7.3	Visibility reduction	107
4.8	Results	108
4.8.1	Interferometric Model Fitting	109
4.8.2	COM 1093: Commissioning Data — AB Doradus AC	110
4.8.2.1	COM 1093: Analysis	110
4.8.3	GO 1843: HD 206893 at High Contrast	112
4.8.3.1	GO 1843: Analysis	115
4.9	Discussion	121
4.10	Conclusions and Future Work	122
4.11	Code and Data Availability	125
5	Conclusions & Future Work	127
5.1	Conclusions	127
5.2	Future Work	129
	Bibliography	135

List of Figures

1.1	Rendering of the Toliman and JWST spacecraft	2
1.2	Current known exoplanet population	4
1.3	JWST AMI’s non-redundant mask and associated PSF	13
1.4	The TOLIMAN pupil and PSF	16
1.5	Geometry of the diffraction integral	18
1.6	Visualisation of a monochromatic point source Airy disk PSF	20
1.7	Geometry of the double-slit experiment illustrating formation of fringes.	21
1.8	Interferometric relationships as a graph.	23
2.1	Optical Configuration Summary	45
2.2	Simulated data from which PRF is recovered.	47
2.3	Correlation plot of recovered astrometry and photometry.	48
2.4	Correlation plot and visualization of recovered OPD.	48
2.5	Recovered PRF, correlation, and residuals.	49
2.6	Recovered parameter error relative to photon noise expectation.	52
3.1	Analytic versus Autodiff Comparison	61
3.2	Geometric parameters of the Toliman astrometric problem	63
3.3	Toliman pupil design basis vectors	64
3.4	Diffractive Pupils and PSFs	66
3.5	Posterior Distributions for Unobstructed and Diffractive Pupils	67
3.6	Reduced Corner Plot Comparing Fisher Forecasts for Pupils Pre and Post Op- timisation	68
3.7	Relative Covariance Matrices for Pupils Before and After Optimisation	69
3.8	Optimised Diffractive Pupils and PSFs	70
4.1	7-hole non-redundant mask, associated PSF, and power-spectrum with labelled fringe components	76
4.2	High-level flow diagram of the AMIGO model and pipeline components	80
4.3	Residuals before and after sine-wave correction for ADC non-linearity	83
4.4	Residual mask distortion and PSF effects from optical aberrations	86
4.5	Flow chart of visibility signal injection into forward-modelled PSF	89
4.6	Latent visibility basis vectors and eigenvalues used for visibility fitting	92
4.7	Example DISCO basis vectors and their variances	93
4.8	Schematic of the EDM architecture with embedded neural network	96
4.9	Demonstration of the EDM BFE model and recovered pixel-level detector pa- rameters	100
4.10	Recovered OPD maps across filters showing recovered aberrations	102
4.11	Summary of AMIGO model fit diagnostics for F430M dithers	104

4.12 Schematic of the DISCO extraction workflow	105
4.13 Summary of AB Dor C fits: likelihood maps and posterior predictive checks . . .	111
4.14 Joint and per-filter MCMC posteriors for AB Dor C astrometry and photometry	113
4.15 AMI contrast sensitivity limits from 3σ detection curves for AD Dor	114
4.16 Summary of the fits to the HD 206893 companions.	117
4.17 Joint posterior distributions from an MCMC fit to the HD 206893 B companion.	118
4.18 Joint posterior distributions from an MCMC fit to the HD 206893 B companion.	119
4.19 Sensitivity limits derived using the Ruffio method for HD 206893.	120

List of Tables

4.1	Summary of JWST observations used for calibration	81
4.2	Allowed filters for AMI observations from JWST documentation	85
4.3	Summary of COM 1093 observations used to test the AMIGO model	108
4.4	Best-fit joint astrometry and per-filter photometry for AB Dor C	112
4.5	Summary of GO 1843 validation data	112
4.6	Best-fit relative joint astrometry and per-filter photometry for the HD 206893 B and companion.	114
4.7	Best-fit relative joint astrometry and per-filter photometry for the HD 206893 c companion.	115

Chapter 1

Introduction

Human curiosity about the cosmos predates civilisation itself (Norris 2016). While the question of whether life exists beyond Earth has ancient roots, only in recent decades have we gained the tools to explore it scientifically. This thesis contributes to that endeavour by developing new methods for the detection and characterisation of exoplanetary systems, those found outside the solar system, bridging fundamental questions with modern instrumentation.

The drive to find other worlds motivates this thesis, but I investigate the tools and technologies that make such discoveries possible. Astronomy is a science of observation: while theory helps guide the interpretation of our universe, it is our instruments and algorithms that set the limits on what we *can* understand. These technologies do more than augment our view of the universe — they shape the questions we are able to ask, defining the boundary of what can be understood. This thesis explores the interplay between hardware and software, where physical design and computational modelling combine to extend the frontiers of our understanding of, and place within the cosmos.

This thesis is anchored in two interwoven themes: advancing the capabilities of frontier instrumentation for exoplanet discovery, and the interrogation of implicit assumptions that guide our interpretation of the data they return. Although the outcomes are general in scope, the work was scaffolded upon two missions — the James Webb Space Telescope (JWST; Gardner et al. 2006) Aperture Masking Interferometer (AMI; Sivaramakrishnan et al. 2012; Soulain et al. 2020) and the TOLIMAN (Telescope for Orbital Locus Interferometric Monitoring of our Astronomical Neighbourhood; Tuthill et al. 2018) mission, show in Figure 1.1 — each designed to meet unprecedented demands on precision. Though they differ greatly in scale and in scope, both exemplify the evolving relationship between measurement, inference, and the pursuit of new frontiers in observational astronomy.

This thesis explores the use of differentiable simulations to bridge the gap between theory and data, empowering new ways to perform statistical inference. Holistic models are constructed to formalise the entire chain from measurement to parameter estimation as a differentiable graph, eroding the boundary between simulation and analysis. Conventional approaches often treat calibration and analysis as distinct stages stitched together by separate workflows. Differentiable modelling dissolves these divisions. This approach unifies the physical processes that generate our data with the statistical models used to interpret it, allowing inference to emerge




Figure 1.1: Rendering of the TOLIMAN and JWST spacecraft. Image credits: TOLIMAN¹ and NASA².

from a continuous dialogue between simulation and observation.

The TOLIMAN mission sets an extraordinary goal: to detect Earth-like planets through astrometry at a level of precision far beyond current practice. Meeting this challenge requires advances beyond incremental improvements — it catalysed a broader rethinking of how calibration, inference, and design could be unified into a single coherent problem through automatic differentiation. Chapter 2 introduces ∂ LUX, the differentiable optical modelling framework that underpins this thesis, and demonstrates how the traditionally separate processes of calibration and scientific inference can be formulated as a single optimisation problem within a unified computational graph. Chapter 3 extends this approach, applying information-theoretic tools to the joint design of instrumental metrology and scientific sensitivity, enabling a coherent framework for system-level optimisation. In many ways, the modelling principles developed throughout this thesis are a direct response to the demands — and possibilities — posed by TOLIMAN.

While TOLIMAN’s demands shaped the development of this modelling philosophy, its most consequential application to date lies in addressing the limitations of JWST’s AMI. The most powerful optical interferometer yet flown and the only one designed for science observations, AMI was intended to open a new window in precision interferometric measurement. However since launch, data fidelity has been limited by subtle, spatially structured detector systematics. These effects — non-linear, variable, and poorly characterised — interact with multiple aspects of the observation and degrade the precision of interferometric measurements. Despite its promise, AMI has fallen short of its anticipated pre-flight performance with significant implications for its scientific return during JWST’s (finite) operational lifetime.

Chapter 4 applies the differentiable modelling framework developed in the TOLIMAN context to understanding and interpreting AMI data. The full observation process — from starlight

¹toliman.space

²science.nasa.gov

to sensor — is modelled as a single computational graph, where the detector is no longer a black box with separable and invertible effect, but an integral part of the end-to-end signal chain. To capture the complex, spatially structured non-linearities in the detector response — dynamics that remain poorly understood and difficult to model from first principles — this work incorporates a machine-learned component into the detector’s forwards model, enabling the system to reproduce performance imperfections as part of the physical measurement chain. This approach does more than recover AMI’s lost performance — it reveals the broader potential of end-to-end differentiable modelling to resolve deeply coupled systematics that span optical, electronic, and observational domains. AMI thus becomes not only a scientific target, but a proving ground for a new mode of modelling in precision instrumentation.

The work presented in this thesis is therefore not an endpoint, but a foundation. The software tools and modelling philosophy developed here — centred on end-to-end differentiable inference — have reach that extends beyond the specific instrument use-cases for which they were developed. The demands of upcoming missions, for example in coronagraph design and calibration for NASA’s forthcoming Habitable Worlds Observatory, will require innovations that surpass the limits of hardware alone. The prospect of calibrator-free interferometry becomes feasible, as does the direct integration of machine-learned priors for joint inference of source structure and instrumental parameters during deconvolution. Further improvements to the models developed for AMI, enabled by new calibration datasets awaiting observation, have the potential to improve performance to meet and exceed pre-flight expectation and to push its utility beyond current observational bounds. These methods are well-positioned to shape coming generations of high-precision astronomical experiments and instruments.

1.1 Exoplanets

The discovery of planets orbiting stars beyond our Sun is one of the defining achievements of modern science and astronomy. The first confirmed detection of an exoplanet orbiting a main-sequence star, 51 Pegasi b, was made by Mayor and Queloz (1995) using Radial Velocity (RV) techniques (for a history of this discovery, see Mayor 2024). This breakthrough confirmed that planetary systems are not unique to our own Solar System, heralding in a new discipline in observational astronomy: exoplanetary science whose rapid growth has spawned corresponding new areas in astrophysical instrumentation. Over the last 30 years the pace of discovery has since accelerated dramatically: as of 2025, nearly 6,000 confirmed exoplanets have been catalogued across a diverse range of stellar hosts and orbital architectures, shown in Figure 1.2 using data retrieved from the NASA exoplanet Archive³. The overwhelming majority of these have been discovered through indirect methods, primarily transits (e.g. via the *Kepler* and *TESS* missions; Charbonneau et al. 1999; Borucki et al. 2010; Ricker et al. 2015) and RV surveys (Hara and Ford 2023).

While these indirect detection methods have yielded an extraordinarily detailed statistical picture of planetary populations, providing insights into planetary radii, masses, orbital peri-

³<https://exoplanetarchive.ipac.caltech.edu/>

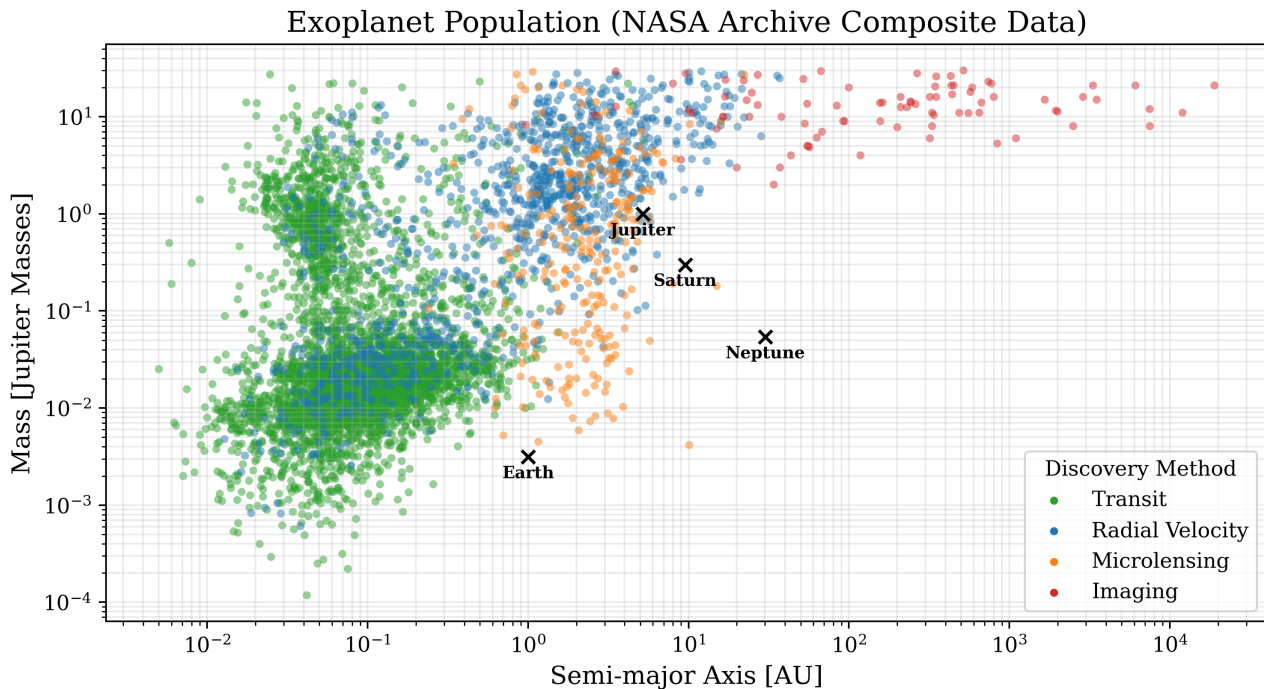


Figure 1.2: The known exoplanet population. Distribution of confirmed exoplanets as of 2025, plotted by Jupiter mass against semi-major axis in Astronomical Units (AU). The colour-coding illustrates the primary detection methods: transits (green), radial velocity (blue), microlensing (orange), and direct imaging (red). Solar System planets are marked with 'x' for scale.

ods, and even basic atmospheric properties (Fulton and Petigura 2018), the observations are fundamentally limited in spatial resolution (Perryman 2018). Transits reveal the silhouette of a planet against the disk of its host star, while radial velocities trace Doppler-induced shifts in stellar spectra. Neither yields a direct image of the planet, a picture of the system as a whole, or any spatially resolved imagery of any disk or structure in the circumstellar environment from which the planet may have formed.

This limitation is particularly acute for understanding the physics of planetary formation and evolution. Direct imaging opens a complementary observational window yielding information that is inaccessible to indirect methods. It enables the detection and characterisation of wide-separation gas giants and brown dwarfs, the measurement of orbital motion via astrometry, and studies of the physical structure of protoplanetary disks, circumstellar debris, and other extended features (Marois et al. 2008; Bowler 2016). With sufficient angular resolution and contrast, direct imaging permits detection and physical characterisation through photometry, spectroscopy, and dynamical analysis.

While the rewards for separating the faint planetary light from its bright host are great, the observational challenge to accomplish this task is profound. A Jupiter-mass exoplanet at 50-100 milliarcseconds separation from its host star may be 10^6 - 10^9 times fainter, depending on age and wavelength (Traub and Oppenheimer 2010). Extracting such signals requires instrumentation at the frontier of optical design, wavefront control, and system calibration. Coronagraphs must suppress starlight while preserving faint planetary signals (Kenworthy and Haffert 2025);

interferometric techniques must resolve structure within and below the diffraction limit (Monnier 2003); detectors must maintain exquisite linearity and stability (Rauscher et al. 2007). Even minor systematic errors in wavefront stability, detector non-linearities, or instrumental miscalibration can swamp the signal of a faint companion or lead to false detections (Mawet et al. 2012).

These challenges are not limited to exoplanetary science. The same tools can be used to investigate faint stellar companions, measure binary orbits, resolve circumstellar disks, and image the compact dusty environments of evolved stars (Tuthill et al. 2000a; Monnier et al. 2007; Kraus et al. 2016). On larger spatial scales, high angular resolution instruments contribute to the study of active galactic nuclei (Burtscher et al. 2013; Sturm et al. 2018), probing the host environments of supermassive black holes (GRAVITY Collaboration et al. 2018; GRAVITY Collaboration et al. 2020).

In this broader context, exoplanet imaging represents not just a demanding subfield of observational astronomy, but a crucible in which instrumentation, calibration, and inference methods are pushed to their limits. The need to detect and characterise faint, close-in companions drives innovation in optics, detectors, and data processing pipelines. The work presented in this thesis is situated within this technological and methodological frontier: while no explicit observational campaign to detect planets is presented, the aim is rather to develop statistically principled data analytic tools that make such detections possible despite the forbidding regime of simultaneous contrast and resolution demanded.

1.2 Direct Imaging

The most prolific techniques for exoplanet detection touched on previously — RV and transit photometry — have enabled valuable statistical insights into planetary architectures. However, their geometric and temporal selection effects result in meaningful blind spots in parameter space. For example neither method can probe face-on orbits, and both lose sensitivity for planets at wide separation from their host star and for lower mass (rocky) planets. Unfortunately many of these systematic biases run directly counter to what are arguably the most desirable of all parts of the exoplanetary canvas: rocky planets (like Earth) in wide (therefore temperate) orbits around main sequence host stars.

Other techniques fill some of these gaps. Microlensing can detect widely separated (\sim several AU) low-mass planets, but yields only single-epoch detections with limited information about host stars or orbital context. Followup of particularly promising discoveries is generally not possible. Astrometry (e.g. *Gaia*; Gaia Collaboration et al. 2016) probes intermediate-separation planets, but lacks contrast sensitivity. As a result, a large portion of planetary systems — particularly those with distant gas giants, embedded planets in disks, or complex multi-body dynamics — remain inaccessible to indirect methods.

Although to date sample sizes remain small, this is the regime in which direct imaging excels. It is uniquely sensitive to young, self-luminous gas giant planets at separations beyond ~ 10 AU, which are so far undetectable by other means. Moreover, it provides the only means to resolve the spatial context of planetary systems — enabling the study of circumstellar disks,

dynamical architectures, and planet-disk interactions. These capabilities enable the study of key open questions in planetary formation theory, including:

- The relative importance of core accretion and gravitational instability (Pollack et al. 1996; Mordasini et al. 2012);
- The frequency and architecture of Solar System analogues (Nielsen et al. 2019; Vigan et al. 2021);
- The origin of substructures observed in disks by ALMA and in scattered-light images (Andrews et al. 2018; Isella et al. 2016);
- Joint studies of the outer regions of solar systems combining direct imaging with astrometry (e.g. Currie et al. 2023);
- The formation, migration, and atmospheric properties of young giant planets (Bowler 2016).

Landmark detections such as HR 8799 bcde (Marois et al. 2008) and the protoplanets PDS 70 b and c (Keppler et al. 2018; Haffert et al. 2019) illustrate the unique scientific return of direct imaging. These systems are invisible to RV and transit surveys, yet are essential for testing formation scenarios and probing the early evolution of planetary atmospheres.

Ensuring sufficient signal fidelity to achieve this science requires extreme precision. Planet-to-star contrasts of 10^{-6} – 10^{-9} and angular separations of 50–200 mas set benchmarks that comprise a generational challenge for optical instrumentation and data analysis. Performance depends not only on coronagraphy and wavefront control, but also on precise calibration, detector modelling, and robust inference methods. The methodological challenges in this field form the core technical motivation of this thesis.

1.2.1 High Contrast Imaging

Faint planetary companions can be millions to billions of times fainter than their host stars and lie at angular separations of just tens to hundreds of milliarcseconds. Achieving detections under such conditions requires instrumentation capable of extreme contrast performance, exquisite optical stability, and with data handling augmented by sophisticated calibration techniques (Guyon 2005).

Specialised instruments have been developed to meet these challenges. Broadly speaking, high-contrast imaging systems combine starlight suppression techniques (such as coronagraphy) or self-calibration methods (like interferometry) with wavefront control, adaptive optics, and careful instrumental calibration. Each class of instrument targets a different regime in the contrast-separation space and is optimised for a specific scientific niche (Mawet et al. 2010; Guyon et al. 2006).

Coronagraphs operate by blocking or diffractively removing starlight in the focal or pupil planes, enabling the detection of faint off-axis sources such as exoplanets and circumstellar disks. Interferometric techniques, particularly aperture masking interferometry, on the other hand,

encode spatial information through interference fringes across a sparse pupil. Interferometry is especially powerful at resolving close-in companions near or within the diffraction limit, while also serving a role as optical calibration tools (Tuthill et al. 2000a; Sivaramakrishnan et al. 2012). Nulling interferometry (Bracewell and MacPhie 1979; Hinz et al. 2016) synthesises these ideas by rejecting starlight in interferometric measurements, and is a promising technique for exoplanet imaging missions in the longer term (e.g. the Large Interferometer For Exoplanets, *LIFE*: Quanz et al. 2022).

Ground-based high-contrast systems rely on extreme adaptive optics to correct for atmospheric turbulence in real time, enabling near-diffraction-limited imaging from the ground. Instruments such as Spectro-Polarimetric High-contrast Exoplanet REsearch (SPHERE; Beuzit et al. 2019), Gemini Planet Imager (GPI; Macintosh et al. 2014), and Subaru Coronagraphic Extreme Adaptive Optics (SCEXAO; Jovanovic et al. 2015) have demonstrated contrasts of 10^{-5} to 10^{-6} . In contrast, space-based systems like JWST and the upcoming Roman Space Telescope (Spergel et al. 2015) benefit from a stable thermal and mechanical environments, allowing them to push even deeper in contrast at small angular separations.

The most sought-after goal of the field remains the detection and characterisation of true *Earth-analogue* systems – a rocky Earth-mass object in a temperate orbit around a Sun-like star. Imaging such targets in reflected light requires contrasts better than 10^{-9} at typical angular separations inwards of 100 mas: well beyond the capabilities of current-generation instruments. However, this ambitious goal may finally be brought within reach by next-generation facilities such as the European Extremely Large Telescope (ELT; Gilmozzi and Spyromilio 2007), the Giant Magellan Telescope (GMT; Johns 2008), and the Habitable Worlds Observatory (Hab-Worlds; Gaudi et al. 2020). Far from being a niche field driven by isolated innovators, exoplanetary science is now entrenched as a central mission objective for coming generation flagship facility investment.

1.2.2 Coronagraphs

Coronagraphs are optical systems designed to suppress on-axis starlight by reshaping the electric field, often imposing wavefront modifications across multiple optical planes. A typical configuration includes a pupil-plane apodizer to sculpt the incoming wavefront, a focal plane mask to block light from the stellar core, then a so-called “Lyot stop” to remove residual diffracted light further downstream. Together, these components reshape the on-axis Point Spread Function (PSF), creating a deep central null and attenuating diffraction structures, thereby enhancing the detectability of faint off-axis sources (Kenworthy and Haffert 2025).

Originally developed for solar observations (Lyot and Marshall 1933), coronagraphy has since become a cornerstone of high-contrast imaging in exoplanetary science. Its capacity to isolate and characterise circumstellar material, wide-separation gas giants, and protoplanetary structures has made it a vital tool for testing theories of planet formation and system architecture (Follette 2023b).

Coronagraphs can achieve raw contrasts of approximately 10^{-6} to 10^{-9} , with inner working angles typically limited to $\gtrsim 2\lambda/D$, where λ is the observational wavelength and D is the

telescope diameter, depending on the optical design. However, reaching these performance levels requires exquisite stability. A range of physical and instrumental effects degrade coronagraphic contrast: residual wavefront errors scatter light into the dark zone; slowly evolving aberrations and misalignments introduce quasi-static speckles; and chromatic dispersion arising for broadband light reduces the effectiveness of starlight suppression. Sensitivity to low-order aberrations and imperfections in optical fabrication further compound these challenges. As a result, coronagraphic systems demand precise control over both optical alignment and temporal wavefront stability to perform (Kenworthy and Haffert 2025).

To mitigate the impact of these imperfections, coronagraphic images are typically post-processed using algorithms designed to suppress residual starlight. Techniques such as KLIP (Karhunen-Loève Image Projection; Soummer et al. 2012), Locally Optimised Combination of Images (LOCI; Lafrenière et al. 2007), and various forms of angular and spectral differential imaging (Sparks and Ford 2002; Marois et al. 2006; Follette 2023b) have become standard. These methods decompose image sequences into dominant modes of variability and attempt to isolate and subtract the stellar PSF. While often effective, these approaches can suffer from self-subtraction of the companion signal, reduced throughput near the star, and residual speckle noise in low signal-to-noise regimes.

Future gains in coronagraphic performance will rely on improvements to both hardware and software. Advances in optical design may reduce chromatic sensitivity and broaden usable bandwidths, while improved wavefront sensing and real-time control could offer better suppression of time-varying aberrations. More robust calibration and subtraction techniques — whether algorithmic or model-based — will be essential for pushing to deeper contrasts and smaller angular separations, ultimately enabling the detection of lower-mass planets and finer system features.

1.2.3 Interferometry

Interferometry is a foundational technique in high angular resolution astronomy, exploiting the wave nature of light to reconstruct spatial information at or beyond the diffraction limit. By combining light from spatially separated regions either across a single telescope pupil or between multiple telescopes, interferometers synthesise an effective aperture much larger than any individual subaperture. The resulting interference fringes encode the Fourier frequencies of the observed source brightness distribution, allowing reconstruction of spatial structure at angular scales $\lesssim \lambda/D$ (Monnier 2003; Lawson 2000).

These capabilities are particularly valuable for imaging close stellar binaries, probing the inner regions of circumstellar disks, mapping the surfaces of evolved stars, and resolving compact dusty environments surrounding young or massive stars (Tuthill et al. 2000a; Monnier et al. 2007; Kraus et al. 2016). Interferometric techniques are not restricted to exoplanetary science, but form a general-purpose tool for resolving faint and compact astrophysical structures.

Unlike coronagraphs, which act to suppress on-axis starlight in the image plane, imaging interferometers isolate the stellar signal in the Fourier domain. This makes them complementary: coronagraphs offer superior contrast at moderate-to-wide separations, while interferome-

try provides finer angular resolution (Guyon et al. 2006; Monnier 2003), often probing within the coronagraph’s Inner Working Angle (IWA).

The historical foundations of optical interferometry can be traced back to the work of Fizeau (1868) in the mid 19th century and later Michelson’s pioneering measurements of Jupiter’s moons and stellar diameters in the late 19th and early 20th centuries (Michelson 1891; Michelson and Pease 1921). Despite their successes, these early implementations were hampered by mechanical stability, atmospheric turbulence and lack of suitable detectors so that the field lay largely abandoned until the advent of intensity interferometry from the 1950’s onwards (Hanbury Brown and Twiss 1954; Hanbury Brown 1956; Hanbury Brown et al. 1974). Renewed interest went on to inspire the advent of speckle interferometry (Labeyrie 1970) and phase closure techniques (Jennison 1958). Many innovations and interplay of ideas arose from crossover with the burgeoning field of radio astronomy which confronted similar challenges in recovery of interferometric observables that are robust against phase noise and instrumental instabilities.

1.2.3.1 Closure and Kernel Phases

To overcome the problem of unknown or unstable phase in optical interferometry, robust observables such as closure phase were introduced. Closure phases are computed as the sum of the measured phase differences around a closed triangle of baselines. Their value lies in the fact that common-mode errors cancel, leaving a quantity that encodes only the intrinsic source structure (Jennison 1958; Baldwin et al. 1986).

More recently, kernel phase analysis was developed extending the idea of closure phase to filled apertures, applying to the special case where phase errors are sufficiently small to be in an approximately linear regime. This enables application of the self-calibrating properties of closure quantities to partially or fully redundant pupil geometries where the wavefront is nearly flat (to within ~ 1 rad Martinache 2010). This method linearises the mapping between pupil-plane aberrations and Fourier phases, projecting the phase space onto the kernel (null) space of the system’s response matrix. The result is a set of phase-like observables encoding source structure that are robust to small aberrations.

Kernel phase has enabled diffraction-limited science with instruments that lack traditional interferometric beam combination arrangements and has been particularly useful for detection of companions at small separation, astrometry, and studies of source morphology (Pope et al. 2013; Ireland 2013). This idea was later extended to kernel amplitudes, expanding the null space available to interferometric observables (Pope 2016); and to kernel nulling, applying the idea to nulling interferometry (Martinache and Ireland 2018). These developments form the mathematical and conceptual backbone of aperture masking interferometry. The mathematical formulation of closure and Kernel phase is explored more thoroughly in Section 4.4.2.3.

1.2.3.2 Aperture Masking Interferometry

AMI converts a conventional telescope into a sparse, non-redundant interferometric array by placing a carefully designed apodizing mask in the pupil plane. Baselines defined by each pair of holes in the mask deliver a unique Fourier coefficient obtained from analysis of fringe

spatial frequencies in the image. The non-redundancy condition imposed on the mask geometry guarantees a unique and unambiguous mapping between pupil geometry and Fourier-domain measurements (Baldwin et al. 1986; Frater et al. 1987; Tuthill et al. 2000a).

The result is a compact, self-calibrating interferometric instrument capable of achieving diffraction-limited resolution on a single telescope. Since it operates in the pupil plane, AMI is compatible with adaptive optics and can be implemented on both ground-based and space-based telescopes.

The method offers several advantages over conventional imaging:

1. It provides high angular resolution (the Michelson criterion for baseline length B is $\lambda/2B$) without requiring multiple telescopes or large baselines.
2. It is relatively insensitive to low-order aberrations and slowly evolving instrumental instabilities, due to the use of closure and kernel observables.
3. It enables reliable calibration through comparisons between science and calibrator targets, even in the presence of drifts in optical alignment (Ireland 2013).

Historically, AMI has achieved some of the highest-resolution optical images from the ground, revealing complex structures such as colliding-wind binaries (Tuthill et al. 1999), dynamic circumstellar dust shells (Monnier et al. 1999), and early disk structures (Tuthill et al. 2002). Its success at facilities such as the Keck, Palomar and VLT observatories firmly established AMI as a mature tool for high-resolution imaging.

The space environment offers advantages for interferometry, including aperture masking. In the absence of atmospheric turbulence a new regime is encountered where complex visibility information (calibrated visibilities and phases) can be measured with exceptional precision. However, small unstable instrumental aberrations in the wavefront remaining are amenable to correction with AMI in the drive to achieve the highest levels of signal calibration possible to detect faint companions. JWST hosts a dedicated AMI mode, the first science-dedicated optical interferometer in space. Designed with both instrumental calibration and science in mind, AMI delivers dual capability: direct imaging at a high angular resolution, while also offering unambiguous measurements of segment-to-segment phase differences for co-phasing and wavefront verification (Sivaramakrishnan et al. 2012; Soulain et al. 2020). While the theoretical performance of JWST AMI surpasses all ground-based analogues, realising this potential depends critically on understanding and correcting a host of residual instrumental systematics and nonlinearities — a challenge that forms the central theme to Chapter 4 of this thesis.

1.3 Modern Precision Astronomical Imaging Systems

The last two decades have witnessed a generational shift in the capabilities of astronomical imaging systems. Modern high angular resolution instruments now routinely achieve contrast ratios and spatial resolutions that were once considered aspirational, enabling the direct detection and characterisation of faint companions, circumstellar disks, and complex astrophysical

environments. These instruments are the product of sustained innovation across optics, control systems, thermal and electrical engineering, and calibration methodology. From ground-based facilities like SPHERE, GPI, and SCExAO to the space-based JWST, this new generation of observatories has transformed high-contrast imaging from a niche capability into a basic workhorse component of exoplanetary and stellar astrophysics.

Ground-based instruments typically go through multiple stages of hardware upgrade, with an equally important component being advancements in algorithms and codes to optimally recover information from the data. There is therefore a rapidly evolving ecosystem in which instrument design, calibration pipelines, and analysis methodologies are simultaneously developed and honed over time. As sensitivity and resolution have increased, so too have the challenges. Residual wavefront errors, segment misalignment, and detector systematics — especially prevalent in complex Infrared (IR) detectors — have emerged as critical performance bottlenecks. Data products have also grown in complexity, often comprising multi-dimensional datacubes tagged with time-resolved readouts and metadata. The interpretation of such data cannot be meaningfully decoupled from a physical understanding of the instrument that produced it. This has driven a shift in philosophy: calibration is no longer a fixed, pre-launch task but an ongoing, community-driven process of refinement.

1.3.1 The James Webb Space Telescope

The JWST represents the most ambitious space-based observatory ever launched, designed as the scientific successor to the Hubble Space Telescope (HST). Launched on 25 December 2021 aboard an Ariane 5 rocket from the Guiana Space Centre in Kourou, French Guiana, JWST was developed by the National Aeronautics and Space Administration (USA) in partnership with the European Space Agency (EU) and the Canadian Space Agency (CSA). With a total development cost of approximately ~ 10 billion USD, two decades of development in international collaboration, it stands as a landmark achievement in modern astrophysics.

After launch, JWST journeyed to the second Sun-Earth L_2 Lagrange point, situated roughly 1.5 million km from Earth, where it maintains a halo orbit to ensure thermal stability and uninterrupted solar shielding. This orbit keeps the spacecraft's sunshield continuously facing the Sun, Earth, and Moon, enabling passive cooling to approximately 40 K — an essential condition for high-sensitivity observations in the infrared.

At the heart of JWST is its 6.5 m segmented primary mirror, composed of 18 gold-coated beryllium hexagonal segments, giving it more than six times the collecting area of Hubble's mirror. The mirror was deployed via an intricate unfolding mechanism, operating through an array of precise actuators and wavefront sensing to achieve diffraction-limited image quality in the near-infrared. This large aperture, combined with the telescope's cryogenic design, enables unparalleled sensitivity in the 0.6-28.5 μm wavelength range (Rigby et al. 2023).

JWST's mission science goals are organised around four principal themes: (1) tracing the formation of the first stars and galaxies after the Big Bang; (2) understanding galaxy assembly and evolution over cosmic time; (3) studying the physics of star and planet formation; and (4) characterising the atmospheres of exoplanets, including potential biomarkers. These objectives

require both deep sensitivity and exquisite spectral and angular resolution, made possible by JWST’s precise optical system and suite of instruments.

The observatory houses four primary scientific instruments, all mounted within the Integrated Science Instrument Module together with the Fine Guidance Sensor:

- Near-Infrared Camera (NIRCam; Rieke et al. 2023) operates from 0.6-5 μm , offering wide-field imaging, coronagraphy, and wavefront sensing support. It is geared towards deep-field surveys, exoplanet transits, and characterising high-redshift galaxies.
- Near-Infrared Spectrograph (NIRSpec; Böker et al. 2023) delivers multi-object, integral-field, and fixed-slit spectroscopy over 0.6-5 μm , facilitating simultaneous observations of dozens of astronomical targets using a microshutter array.
- Mid-Infrared Instrument (MIRI; Wright et al. 2023) provides imaging, coronagraphy, and spectroscopy between 5 and 28 μm . It is equipped with a cryo-cooler to maintain its detectors below 7 K, enabling studies of cold dust, protoplanetary disks, and mature planetary systems.
- Near-Infrared Imager and Slitless Spectrograph (NIRISS; Doyon et al. 2023b) includes broadband imaging, slitless spectroscopy, and (most relevant to this thesis) a unique AMI mode, providing diffraction-limited imaging at thermal IR wavelengths and enabling high-angular-resolution observations from space.

Collectively, these instruments offer a unique combination of imaging, spectroscopic, and interferometric capabilities unmatched by any existing space observatory. JWST’s instrumentation enables not only deep surveys of the early universe but also detailed atmospheric characterisation of exoplanets and high-contrast observations of protoplanetary environments. In the context of this thesis, particularly Chapter 4, the focus is on the NIRISS instrument, especially its AMI mode, which offers a unique platform for pushing the boundaries of high-contrast imaging techniques in space-based exoplanet science.

1.3.1.1 NIRISS AMI Mode

The NIRISS instrument on JWST incorporates a dedicated AMI mode — the first science-focused implementation of non-redundant masking in space (Soulain et al. 2020). As detailed in Section 1.2.3.2, AMI transforms the telescope pupil into a sparse interferometric array, enabling stable, high-resolution measurements via self-calibrating kernel and closure phase observables. With JWST’s exquisite optical stability, AMI was founded on the promise to deliver the highest-precision interferometric observables yet obtained while also serving a dual role in wavefront metrology as well as science.

While masking was originally developed to mitigate atmospheric phase noise, AMI’s robustness also makes it ideal for space-based segmented telescopes like JWST. In particular, Non-Redundant Masking (NRM) configurations offer unambiguous recovery of inter-segment

phase differences — essential for verifying co-phasing and enabling diffraction-limited imaging in a segmented system (Sivaramakrishnan et al. 2012). Figure 1.3 shows the NRM mask configuration and resulting PSF for JWST AMI.

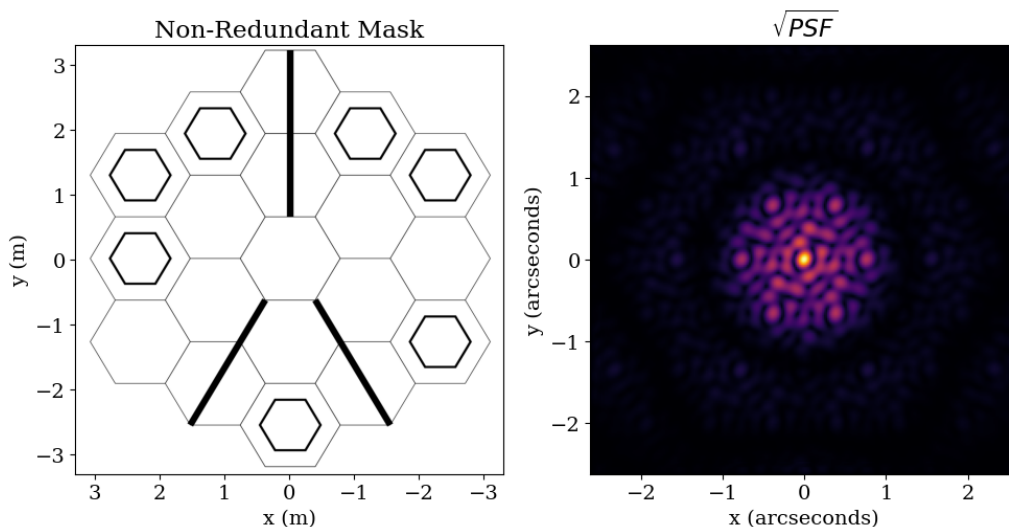


Figure 1.3: Left panel: Schematic diagram of the 7-hole NRM projected over the primary mirror. Right panel: The resulting PSF (i.e. interferogram) from the non-redundant mask, visualised on a square-root scale to highlight low-power features.

As with all major observatories, JWST’s calibration remains an ongoing communal effort, driving advances in both observing strategies and analytical methods. This process has remained particularly relevant for AMI, with its performance falling below the pre-flight specifications (Sivaramakrishnan et al. 2022; Sallum et al. 2024a; Ray et al. 2025). The standard JWST calibration pipeline (Bushouse et al. 2025) assumes a modular approach in which the optics, detector, and readout electronics are treated independently. These assumptions include ideal wavefront behaviour and detector linearity — neither of which were found hold to a sufficient degree for real data. Observations from commissioning and early science programs have shown that wavefront drift, pupil misalignment, and most importantly non-linear electronic effects within the sensor all introduce significant biases into the measured interferometric observables. These systematic errors degrade the calibration fidelity between science and calibrator targets, limiting contrast performance and weakening confidence in astrometric and photometric measurements derived from AMI data.

Many tools exist for JWST interferometric analysis — including AMICAL (Soulain et al. 2020)⁴, SAMpy (Sallum et al. 2022)⁵, ImPlaneIA (Greenbaum et al. 2015)⁶, and FOURIEVER⁷ — but all rely on a critical assumption: that the measured count rate in each pixel of the PSF is linearly related to the incident photon flux. This assumption breaks down in a real detector due to the presence of complex electronic effects, the foremost of which is non-linear charge migration (see Section 1.3.1.2). The resultant dynamics violate the separability of detector

⁴github.com/SAIL-Labs/AMICAL

⁵github.com/JWST-ERS1386-AMI/SAMpy

⁶github.com/anand0xff/ImPlaneIA

⁷github.com/kammerje/fouriever

and optical calibration, leading to biases in observables that, for the context of interferometric detection of faint companions, degrade contrast performance and introduce systematic error.

While standard calibration pipelines can address some kinds of non-linearity post hoc, the time-evolving and non-local nature of charge migration defies simple inversion. These effects introduce structure-dependent perturbations to recovered PSFs that are inseparable from the desired science outputs — for example the optical signature of faint companions — a critical limitation for interferometric precision.

As discussed further in Section 1.3.1.2, the root cause lies in the physics of the detector itself. To address these limitations, a new approach, detailed in Chapter 4, was required — one that would forgo the simplifying assumptions of standard calibration and instead forwards model the full imaging process. The result is Aperture Masking Interferometry Generative Observations (AMIGO), an end-to-end differentiable forwards model and analysis pipeline specifically developed for AMI. AMIGO fuses a physics-based model for the optics, detector, and electronics, with a learned neural component that captures non-linear charge-migration in the detector that is central to the problems plaguing AMI. This hybrid framework is combined with novel approaches to the extraction of interferometric kernel observables and operates directly from raw uncalibrated ramp data, bypassing the need for intermediate calibration products. This entire process is underpinned by the algorithmic differentiability of the code as implemented in the computer, a concept explored in further detail in Section 1.5.

This constitutes a reformulation of the process of obtaining scientific inference from data in which the physics informed forwards model takes primacy at the expense of more traditional ideas of “data reduction” which aim to deliver ad-hoc intermediate data products selected for perceived resilience to noise or error. Armed with these tools, the aim is to enable unprecedented contrast at small IWAs, exceeding limitations of existing pipelines for interferometric data reduction. The immediate science payoff targeted is the recovery of faint companions previously undetectable by standard methods along with improved astrometric precision and systematic error calibration. However the findings presented in this thesis and encapsulated within the AMIGO code can also be taken more generally: as an exemplar and illustration for next-generation modelling techniques, demonstrating that they can unlock the full potential of both ground- and space-based astronomy.

1.3.1.2 H2RG Infrared Detectors

All science instruments aboard JWST are equipped with IR detectors. Three of these, namely NIRISS, NIRCams, and NIRSpec, employ the Teledyne HAWAII-2RG (H2RG) variant (Rauscher et al. 2007) developed by Teledyne Imaging Sensors. These detectors are based on Mercury Cadmium Telluride photodiodes connected to a readout integrated circuit. A distinctive feature of H2RG detectors is their support for non-destructive readout, meaning that the accumulated charge in each pixel can be measured multiple times during a single exposure without being reset. This enables a sampling scheme known as up-the-ramp readout, where each pixel’s charge is sampled at uniform intervals over the course of the integration. By fitting a slope to the voltage ramp at each pixel, both the total flux and the effective uncertainty can be estimated, allowing for improved cosmic ray rejection, noise characterisation, and dynamic range. This

readout strategy plays a central role in JWST operations and underpins the standard calibration pipeline. However, several physical effects intrinsic to H2RGs — including electrostatic coupling and charge diffusion — complicate this simple picture and disrupt the assumptions underpinning PSF calibration in high-precision imaging.

A key motivation for rethinking JWST’s AMI calibration and analysis pipeline stems from the complex and poorly modelled behaviour of its detectors. Like most instruments aboard JWST, NIRISS employs H2RG detectors, which exhibit a host of non-linear effects that severely limit the assumptions underlying standard calibration models. Chief among these is the Brighter-Fatter Effect (BFE), wherein bright pixels accumulate charge that induces electrostatic repulsion, in turn causing redistribution of charge into neighbouring pixels (Antilogus et al. 2014; Guyonnet et al. 2015; Rowlands et al. 2018). This effect distorts the morphology of the PSF, is non-local, time-evolving, and couples to a host of other properties such as flux, sub-pixel positioning, PSF curvature over individual pixels, and variations in pixel-to-pixel sensitivities. This effect is problematic in many observing modes (Rowlands et al. 2018; Argyriou et al. 2023; Goudfrooij et al. 2024), but entirely deleterious for interferometric analysis (Sivaramakrishnan et al. 2022; Sallum et al. 2024a; Ray et al. 2025). The primary motivation for the work described in Chapter 4 of this thesis was to address this problem.

1.3.2 Toliman

Although the techniques developed over the course of the research program detailed in this thesis have wide applicability (with particularly critical application to JWST AMI) they owe their intellectual heritage to design and analysis challenges posed by the TOLIMAN space telescope mission. This is a cubesat-class satellite carrying a 125 mm aperture telescope currently being fabricated with launch targeted for 2026. The TOLIMAN has a straightforward statement of purpose: “To conduct a search revealing any terrestrial-mass planet in a temperate orbit circling either component of our binary nearest neighbour α Centauri AB star system”. Despite its brightness and proximity as one of the most prominent naked-eye stars in the night sky, astronomers possess no existing instruments or technologies capable of revealing such objects. TOLIMAN aims to recover such signals through *relative* binary astrometry. If the relative separation of the two stars comprising the binary can be measured precisely over time, the gravitational reflex motion due to any (unseen) orbiting planets will impose cyclic variations on the expected orbital locus of the stellar pair. Therefore binary separation data — if recovered with sufficient (extreme) precision — can betray the existence of orbiting exoplanets (Tuthill et al. 2018), as well as yielding information on the period and mass.

Historically, astrometric science (cataloguing exact positional information for stars) has been concerned with *absolute astrometry*, i.e. the position and motion of stars with respect to a “stationary” frame of reference taken to be ultimately provided by objects in the distant universe, and more conveniently in the local context by galactic background stars. Recovery of sufficiently precise absolute astrometric data poses major instrumental challenges, primarily that in order to register motion against background field stars — the ruler establishing measurement stability — a large (\sim meter-class) aperture is required to collect sufficient photons

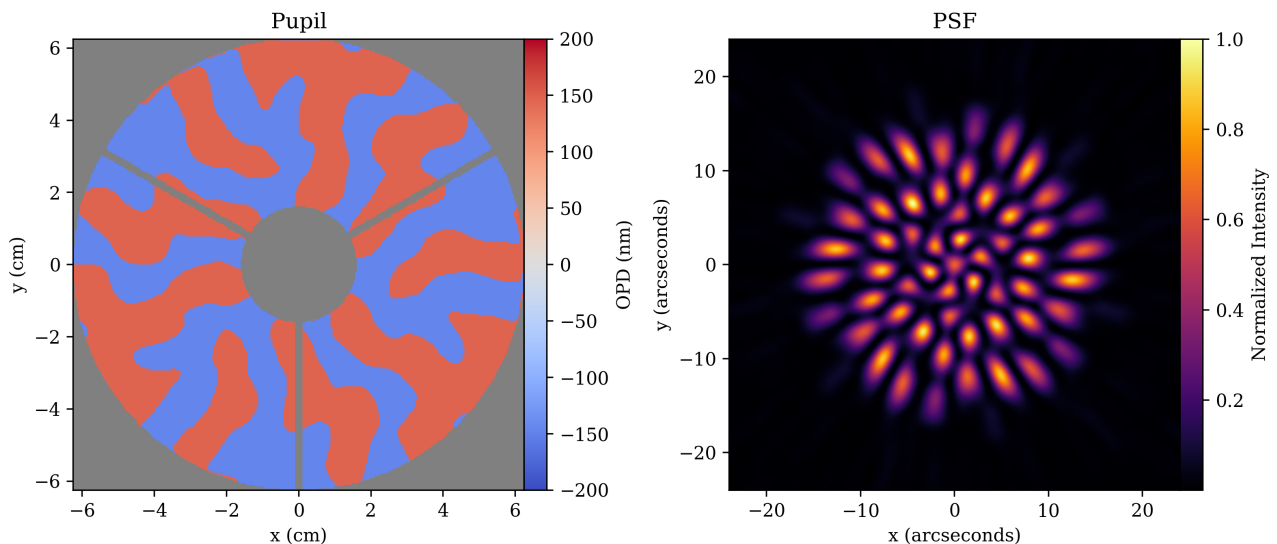


Figure 1.4: The TOLIMAN telescope’s optical configuration. (Left) The diffractive pupil mask featuring a binary phase pattern etched onto the entrance pupil. The binary valued design uses two unique delay values in anti-phase to modulate the incoming light. (Right) The resulting PSF, which appears as a complex set of overlapping interference fringes. This coded PSF acts as a high-precision ruler to measure the relative separation between the stars in the α Cen binary system, providing stability against optical distortions and enabling the detection of gravitational wobbles from orbiting exoplanets.

because conveniently located background reference stars are typically of order 10 000 times fainter. Furnishing an aperture of this size aboard a space platform is well beyond the financial scope of the program. TOLIMAN instead leverages the fact that α Cen is a binary system: if the relative separation of the two stars (both very bright) can be recovered with sufficient precision, then the gravitational perturbation of the planet can be witnessed without the need for imaging of any faint background sources. To make this approach viable, TOLIMAN employs a novel optical system featuring a diffractive pupil. By etching an engineered binary phase pattern onto the entrance pupil of the telescope, the PSF is sculpted into a carefully designed, sharply structured pattern. Consisting of a complex set of overlapping interference fringes, this coded PSF acts as a ruler that can be engineered to yield measurements invariant to unstable optical distortions and aberrations (Tuthill et al. 2018). Furthermore, by enforcing a binary valued pupil (the phase pattern is made up of only two values of delay that are in anti-phase) then it turns out that with further bespoke design the PSF can also act as a spectrometer and continuous metrology device encoding the state of the optical system and sensor along with the desired science signal. Figure 1.4 presents the TOLIMAN pupil and PSF.

Despite this innovative approach, the challenges associated with achieving the primary mission goal cannot be overstated. The expected astrometric signal amplitude for an earth analog is of order $\sim 1 \mu\text{as}$, which translates in the optical system to an image displacement of about one millionth of a single detector pixel. In calibrating systematics at 10^{-6} levels of precision, a host of high order terms and small effects normally negligible must all be carefully examined. The scale of the technical challenge is thrown into sharp relief as TOLIMAN is a low-cost project deployed on a cubesat platform in low Earth orbit. A mission attempting scientifically

audacious goals yet fitting within an aggressive cost cap cannot simply engineer its problems away by invoking flagship-class infrastructure. This core tension within the TOLIMAN program created an environment in which radical innovation was a basic requirement to push beyond the existing limits of signal fidelity, calibration and data analysis. It is this context that drove examination of the problem of inference at the most basic level leading to a reformulation of the process of extracting meaning from data.

Therefore as a direct response challenges posed by the TOLIMAN mission, a synthesis of ideas from machine learning, statistical methods and computational optimisation were forged into a new library of codes for modelling of optical systems: ∂ LUX. From its basic philosophy through to the construction, utility and notable use-cases such as the JWST data calibration, these codes form the enabling theme of this thesis. Chapter 2 explores the use of the differentiable models permitted by construction in ∂ LUX for instrumental calibration of a TOLIMAN-like instrument directly from science data, and Chapter 3 details the methods of information maximising design used for the diffractive pupil employed by TOLIMAN.

1.4 Image Formation in Optical Systems

Optical imaging is one of the oldest scientific disciplines with mathematical foundations tracing back to antiquity. Euclid’s early geometrical description of light treated rays as straight lines obeying simple laws of reflection — an idea that laid the groundwork for what is now formalised as geometric optics. Far from obsolete, geometric optics remains a foundational and widely used framework in modern optical design. By neglecting the wave nature of light, it enables tractable analysis of complex systems through ray tracing, and is particularly powerful when modelling the reflective and refractive behaviour of optical elements. This approximation underpins the design of lenses, mirrors, and apertures in most instruments, where it is used to model the trajectory of photons to a focus, maximising the spatial information recovered about the observed scene (Ewart 2019).

1.4.1 Physical Optics

While powerful, geometric methods ignore the diffractive properties of light arising from its wave nature — a characteristic essential to the work of this thesis. This wave theory of light was first proposed by Huygens, describing the interferometric, polarimetric and diffractive properties exhibited by light waves by envisioning a wavefront comprised of a superposition of small spherical wavelets (Huygens 1690). Huygens principle was formalised by Kirchhoff (1883) and shown to be valid under a paraxial assumption (on or close to the optical axis of a lens or mirror), describing the resulting field at any point as the superposition of small spherical waves from the input plane as

$$E(x, y, z) = \frac{1}{i\lambda} \iint E(x', y') \frac{e^{ikr}}{r} \cos(\theta) dx' dy'. \quad (1.1)$$

The resulting field at at point $E(x, y, z)$ due to the original planar field $E(x', y')$ is a function

of the path length r , angle $\cos(\theta) = z/r$, k is the wavenumber $2\pi/\lambda$, and wavelength λ . The geometry of this system is shown in Figure 1.5.




Figure 1.5: Geometry of the diffraction integral showing the resulting field $E(x, y, z)$ as a superposition of spherical waves from an input planar field $E(x', y')$.

While this equation is general, particularly relevant to optical modelling is the Fraunhofer, or far-field, regime ($z \gg D^2/\lambda$) where D is the diameter of the aperture or wavefront. This yields the Fraunhofer diffraction integral (Born et al. 1999)

$$E(x, y, z) = \frac{e^{ikz}}{i\lambda z} \exp\left[\frac{ik}{2z}(x^2 + y^2)\right] \iint_{-\infty}^{\infty} E(x', y') \exp\left[-\frac{ik}{z}(xx' + yy')\right] dx' dy', \quad (1.2)$$

which is valid at the focus of a lens or mirror. Remarkably, this integral can be directly expressed as the 2D Fourier Transform of the initial field $E(x', y')$ via

$$E(x, y, z) = \frac{e^{ikz}}{i\lambda z} \exp\left[i\frac{\pi}{\lambda z}(x^2 + y^2)\right] \cdot \mathcal{F}\{E_0(x', y')\}\left(\frac{x}{\lambda z}, \frac{y}{\lambda z}\right), \quad (1.3)$$

which enables the simple and direct computation of the diffractive principles that govern light at the focus of an optical system. This relationship in which wavefront propagation can be understood by way of a Fourier transform is foundational to the field of physical optics, as explored throughout this thesis. The behaviour of light in the near-field can also be expressed as a similar integral or Fourier transform, however the more general treatment arises from the angular spectrum method (Morse and Feshbach 1953; Born et al. 1999; Goodman 2005) as

$$E(x, y, z) = \iint E(x', y') \frac{1}{(2\pi)^2} \iint \exp[i(k_x \hat{x} + k_y \hat{y} + k_z z)] dk_x dk_y dx' dy' \quad (1.4)$$

where \hat{x} , \hat{y} are $x - x'$ and $y - y'$ respectively, $k_z = \sqrt{k^2 - k_x^2 - k_y^2}$, and k_x, k_y correspond to the angular spectrum (spatial frequencies) of the wave. This equation enables the description of propagating plane waves at all distances. This equation can also be expressed through Fourier transforms as

$$E(x, y, z) = \mathcal{F}^{-1}[\mathcal{F}[E(x', y')](k_x, k_y) \cdot e^{ik_z z}], \quad (1.5)$$

placing the angular spectrum method as a general framework for modelling scalar wave propagation in optical systems. Together, these equations enable the rigorous evaluation of diffractive

behaviour — fundamental to the formation of the PSF, which describes the response of an optical system to a point source. Much of this thesis centres on how optical systems can employ phase and amplitude masks to engineer PSFs with tailored properties for the detection and characterisation of astronomical signals.

1.4.2 Interferometry

Among the methods used to achieve high angular resolution in astronomy, interferometry stands out for its ability to resolve structure below the diffraction limit. It does so not through standard imaging as it is classically understood, but by forming fringe patterns that provides access to information that conventional imaging struggles to recover.

1.4.2.1 Resolution Limits of Diffractive Systems

The diffractive principles that govern the wave-like behaviour of light set fundamental limitations on the achievable spatial resolution of an optical system. For a clear, circular aperture imaged in the Fraunhofer, i.e. far-field, regime, a monochromatic point source produces the Airy disk PSF, first shown by Airy (1835) and visualised in Figure 1.6. This familiar patterns consists of an intense central peak, followed by a radial null and a series of concentric rings with diminishing brightness with radius. The first dark ring is found at an angle of

$$\theta = 1.22 \frac{\lambda}{D} \quad (1.6)$$

where λ is the wavelength of light and D is the diameter of the circular aperture. When identified as the angle between two point objects, Equation 1.6 defines the Rayleigh criterion (Rayleigh 1880) and is commonly used as a quantitative proxy for the resolution limit of an imaging system. Although commonly misunderstood as a hard limit on optical performance, this condition is ultimately nothing more than a handy rule-of-thumb. With a high quality calibration and sufficient signal-to-noise, resolutions beyond this limit can be practically achieved.

In practice, the question of whether two sources are ‘resolved’ or not is not one that is answered by a simple equation, but rather by how much information about them is recoverable from some observation. In the case of a binary source we commonly parametrise them by their separation and contrast. When inferring these parameters from data, we do not suddenly lose all knowledge as to the existence of the companion crossing some limit threshold, but rather gradually lose statistical certainty about the the exact position and relative brightness between the two. Taking this idea to its limits, we imagine two point sources with zero separation — for any brightness we chose for the primary source we can find a contrast for the companion such that the total brightness remains unchanged, rendering the brightness ratio entirely degenerate.

In an interferometric system (which I argue later applies to all pupil-focal optical imaging systems), the minimum and maximum resolvable angle is generally expressed as a function of the baseline length B , defined as the separation between the various apertures that contribute to the interferometric array (Labeyrie 1975), through the equations

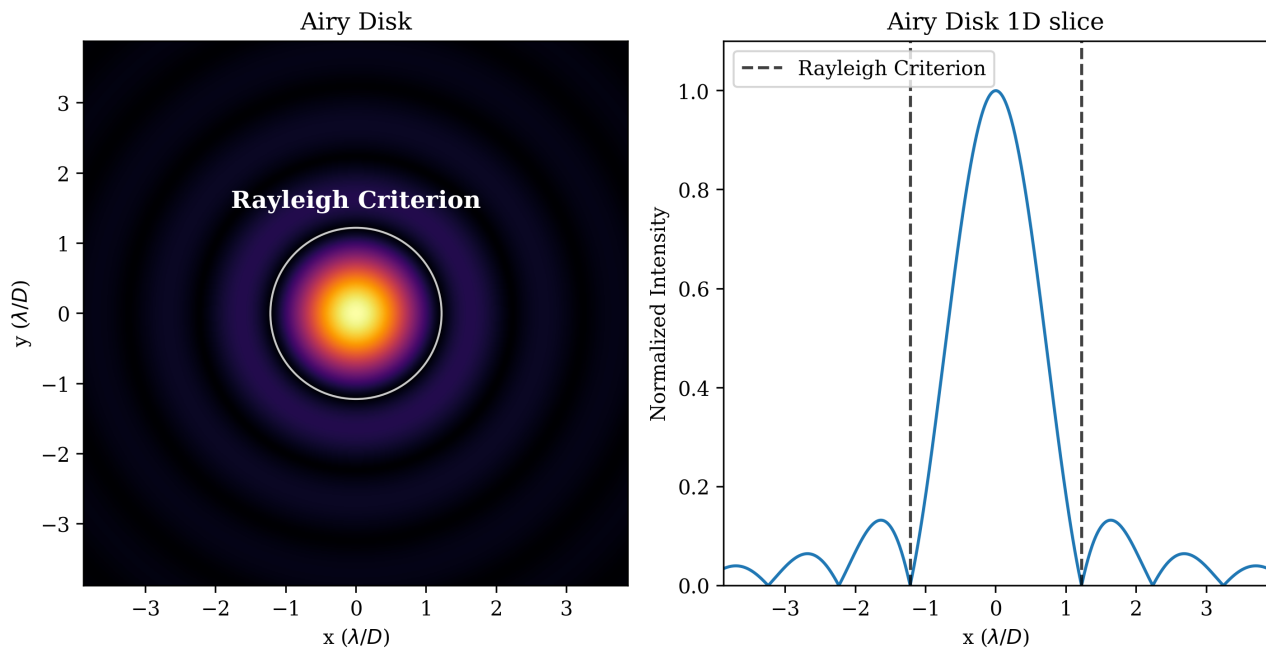


Figure 1.6: Visualisation of a monochromatic point source Airy disk PSF. The right panel displays the 2D diffraction pattern while the left panel shows a 1D cross-section (slice) of the square root intensity pattern. The first radial null, i.e. the Rayleigh criterion, is indicated by the white circle and dashed line at $1.22\lambda/D$.

$$\theta_{\min} = \frac{\lambda}{2B_{\max}}, \quad \theta_{\max} = \frac{\lambda}{2B_{\min}}. \quad (1.7)$$

The same principles apply to this criteria: it is not a fundamental limit, but a simply boundary at which the information we wish to recover becomes more challenging for statistical inference. As I show in later chapters of this thesis, that these the true graduated form of these “resolution limits” can be understood through principled treatment of the underlying physics with advanced modelling techniques.

1.4.2.2 Development of Interferometric Methods

Although the wave nature of light had been introduced since Huygens’ first founded the field of physical optics (Huygens 1690; Kirchoff 1883), the sub-field of interferometry began with the double-slit experiment (Young 1804). By passing light through two narrow slits in an opaque screen, Young engineered the first set of interference fringes, providing the first experimental demonstration of interference and an emphatic illustration of the wave nature of light. Calculation of the intensity pattern from two slits in the far field is straightforward. For the purposes of the derivation, Figure 1.7 presents a conceptual diagram of Young’s celebrated experiment.

Two slits, separated by a distance d , are illuminated by an incident monochromatic wavefront of wavelength λ . At some point P in the far field, offset from the centre of the two slits by an angle θ , we can find the relative path difference between the two rays $\Delta r = r_2 - r_1$ from each aperture in order to find the resulting intensity at P . Since P is in the far-field, r_1 and r_2 are approximately parallel and the path difference is therefore $\Delta = d \sin(\theta)$. Our light is coherent




Figure 1.7: Geometry of the double-slit experiment: an incident monochromatic plane wave from the left arrives at an angle θ at the plane of a screen creating a path difference $\Delta r = d \sin \theta$ between the rays from slits A_1 and A_2 at the far-field point P . The resulting interference pattern produces sinusoidal intensity fringes whose spacing depends on both the wavelength λ and the slit separation d . This basic scheme underpins the resolution formula $\theta = \lambda/2d$ and forms the conceptual foundation for modern optical interferometry, where fringe visibility encodes angular information about extended sources.

we can find the resulting phasor, ie the complex electric field value, at this point P as the superposition of the individual phasors from the two slits as a function of the path difference Δr , ignoring the global phase. Assuming the amplitude of the input wave is A at both slits, we can describe this field as

$$A_P = A \exp[ik\Delta r/2] + A \exp[-ik\Delta r/2] \quad (1.8)$$

where k is the wavenumber $2\pi/\lambda$. Using Euler's formula $e^{i\phi} = \cos(\phi) + i \sin(\phi)$ the resulting field is

$$A_P = 2A \cos(k\Delta r/2). \quad (1.9)$$

Substituting in $k = 2\pi/\lambda$ and $\Delta r = d \sin(\theta) \approx d\theta$ under the small angle approximation we find its resulting intensity at any point is then

$$P(\theta) = 4A^2 \cos^2(d\pi\theta/\lambda), \quad (1.10)$$

taking sinusoidal form. Since the first null of a squared cosine is found when its argument is equal to $\pi/2$, we find the interferometric equivalent of the Rayleigh criterion defined in Equation 1.6, known as the *Michelson* resolution limit given by $\theta = \lambda/2d$, equivalent to Equation 1.7 and smaller than that of the equivalent clear aperture. The foundations of interferometry as a high-angular resolution method originate from this very simple derivation showing that the ability of light to interfere with itself can be leveraged to produce nulls at smaller angles than is possible with a filled aperture of the same diameter as the baseline.

Following the derivation of the far-field interference fringes from a pair of slits, we might ask how this can be used to extract useful astrophysical information. To understand this we we must consider that, while incredibly small, all astrophysical sources have nonzero angular extent, meaning that different parts of the source subtend slightly different angles at the slits. Each point on the surface of these objects produces can be considered to produce a unique fringe pattern, offset by some angle relative to other points.

The final observed intensity pattern is then produced by an *incoherent* sum (i.e. adding in through intensity) of multiple slightly shifted sets of fringes, which is mathematically identical to a convolution. We can therefore represent any source as a convolution between the point source fringe patterns and intensity distribution of the observed source. This convolution has a clear observational effect on fringes in the image plane — it reduces the contrast between the peaks and nulls of the fringe pattern. The more extended the source, the more these fringes overlap and wash out, lowering the contrast of the interference pattern. This idea led to the development of the concept of *fringe visibility*, defined as the contrast between the maximum and minimum intensity of the observed fringes

$$V = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}. \quad (1.11)$$

For an unresolved point source, as we have just seen, this results in a visibility of unity. However, once the angular extent of the observed object begins to approach the interferometric resolution, the visibility contrast, V , decreases. The first measurements of the angular diameter of stars was performed by Michelson and Pease (1921) by measuring the decrease in fringe visibility from an opaque mask with two holes placed at the aperture of a telescope.

The visibility contrast at a given slit separation corresponds to the amplitude of the spatial Fourier component at that baseline — a key insight that van Cittert-Zernike theorem (following section) formalises. In this sense, the fringe visibility becomes a directly observable property that is able to constrain the one dimensional angular structure of some imaged source.

1.4.2.3 The van Cittert-Zernike Theorem

Following the insight that fringe contrast encodes spatial information of our source, we introduce the framework that emerged to generalise the mathematics that underpins interferometric imaging. This is the van Cittert-Zernike Theorem: the visibility pattern observed by an interferometric array is directly proportional to the Fourier transform of the source intensity distribution (Cittert 1934; Zernike 1938; Goodman 2005).

Figure 1.8 presents a diagram of this process, illustrating how an incident wavefront is transformed as it passes through an optical system. The top half of the figure describes how each coherent monochromatic wavefront is transformed by the optical components of an imaging system, and eventually measured. A wavefront in the pupil plane, denoted as E_{pupil} , diffracts through an optical system to the focal plane — which can be mathematically modelled directly though a Fourier transform as shown in preceding sections — producing the complex field E_{focal} at the focal plane. However, with the exception of some special cases (for example imaging from radio telescope data), we do not have data recovering this resulting complex field, only the resulting intensity, i.e. the PSF \mathcal{I} , found via the squared modulus of E_{focal} .

This transformation is not invertible, as we lose the phase information of the wavefront in the process. However, we can gain further insight by considering the autocorrelation of the wavefront in the pupil plane, E_{pupil} . This operation yields the Optical Transfer Function (OTF), \mathcal{O} , acting as a filter in the frequency domain, limiting the spatial frequencies an optical system is responsive to. This limitation of spatial frequencies is what drives the resolution limits of an

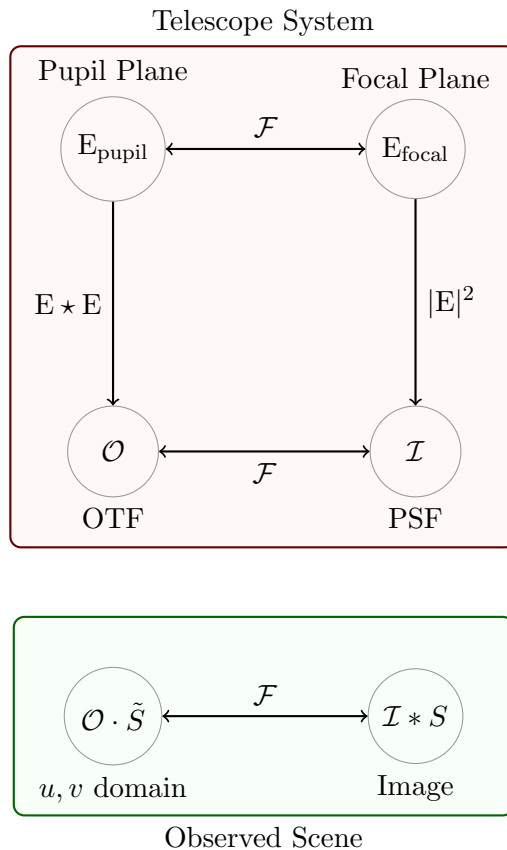


Figure 1.8: Diagram showing the relationships between wavefront propagation, intensity formation, and interferometric observables in an optical system. The telescope (red box) maps the pupil-plane wavefront E_{pupil} to the focal-plane wavefront E_{focal} via a Fourier transform. The measurable quantity is the PSF \mathcal{I} , formed by taking the square modulus of the complex field. The OTF \mathcal{O} is the autocorrelation of the pupil wavefront $E_{\text{pupil}} \star E_{\text{pupil}}$ (where \star is a correlation), and is related to \mathcal{I} via a Fourier transform. When the system observes a source with angular brightness distribution S , the final image is a convolution $\mathcal{I} \star S$ (where \star is a convolution), whose Fourier transform is $\mathcal{O} \cdot \tilde{S}$. The van Cittert—Zernike theorem corresponds to the operations in the green box. This figure has been adapted from Pope (2017).

imaging system — a delta function (ie, a perfect point source) is composed of infinite spatial frequencies, whereas an imaging system can only respond to those that its aperture allows. Crucially, the Fourier transform of the PSF \mathcal{I} from a point source, i.e. the PSF, is equal to the OTF \mathcal{O} , providing a map between measured intensity as the frequency response of an optical system and plays a central role in understanding how phase errors affect imaging resolution (Roddier 1981).

The bottom half of Figure 1.8 introduces the idea of an astrophysical scene. When a telescope observes an extended source, its intensity distribution S is convolved with the instrumental PSF \mathcal{I} . This produces the final observed image $\mathcal{I} \star S$. By the convolution theorem, the Fourier transform of the image is simply the element-wise product of the OTF with the Fourier transform of the source \tilde{S} (Goodman 2005). This duality lies at the heart of the van Cittert-Zernike theorem — the interferometric signal is not an image, but a set of filtered spatial frequency measurements that can be represented entirely in the Fourier domain.

This perspective clarifies the operation of an interferometer: while it does not produce a direct image of a source, rather it samples structure at the specific spatial frequencies allowed by its aperture configuration (Labeyrie 1975). In this formulation, interferometry is distinct from an image formation problem via its analytical representation in the frequency domain. Spatial frequencies that would be degenerate through clear aperture systems become unambiguously resolved through the visibilities.

1.4.2.4 Interferometric Analysis & Kernel Spaces

As discussed in the previous section, optical imaging systems only recover information on the intensity of a wavefront, losing phase information in the process. In systems with a redundant aperture configuration, each Fourier frequency is sampled multiple times across the pupil, causing the measured visibility phase to be the sum of multiple different values. Because any given complex visibility recovered in the Fourier plane is the sum of a number of indistinguishable pairs of patches in the incident wavefront, this makes recovery of the wavefront information degenerate; mathematically it is an ill-posed problem. By contrast, non-redundant aperture configuration enforces a *unique* mapping between the incident wavefront and the Fourier phase and amplitude, entirely avoiding the degeneracy (Baldwin et al. 1986; Haniff 2007).

The unique phase mapping provided by non-redundant configurations can be leveraged to further improve the resilience of interferometric observables to wavefront error by extracting an observable known as the *closure-phase* (Jennison 1958). For a pupil consisting of at least three subapertures, we can form a number of closing triangles between hole triplets. For three baselines, i , j , and k , the closure phase is defined

$$\Phi_{ijk} = \phi_{ij} + \phi_{jk} + \phi_{ki}, \quad (1.12)$$

where ϕ_{ij} is the measured baseline phase for the two aperture holes i and j . The closure phase observable Φ_{ijk} is invariant to any piston phase errors across these holes, since each baseline records the relative phase between its two contributing holes. Any non-zero closure phase value can therefore only represent the angular structure intrinsic to the observed source, at least for an ideal optical system in which aberrations occur in the pupil plane. The ability to extract these phase-error invariant observables has made ground based interferometers invaluable to the recovery of high-angular resolution structures when observing from ground based sites that must contend with astrophysics phase corruption.

This self-calibrating observable produced by interferometers has proven remarkably powerful in recovery of real astrophysical structure, however the basic idea is not limited purely to non-redundant apertures. Closure phases are simply a self-calibrating subspace of an optical system: a property that exists and can also be calculated for either fully or partially redundant apertures. The observable that extends capability to redundant pupils is known as *kernel phase* (Martinache 2010), and (unlike closure phase) it relies on the fact that in the small aberration regime the mapping between pupil and visibility phases is approximately linear. By calculating the phase-transfer matrix, the linear mapping between phases in the pupil and Fourier plane (which is simply a Jacobian), a phase-invariant (kernel) nullspace can

be extracted. This kernel matrix enables the extraction of phase-invariant observables from arbitrary pupil geometries, no matter how redundant. By considering all optical systems as an interferometer, we can find self-calibrating observables resilient to the generally ill-posed problem of phase-reconstruction mandated by the measurement of the intensity of the incident wavefront.

Crucially, this formulation is fully compatible with the van Cittert-Zernike theorem, which relates interferometric measurements to the spatial Fourier transform of the source. Instrumental effects manifest as multiplicative filters in this domain, and the ability to project onto calibrated subspaces allows us to subtract or suppress such effects at the level of the data themselves. This perspective enables interferometers to infer high-resolution spatial information without needing to reconstruct the full pupil phase — a particularly important feature for systems where direct wavefront calibration is difficult or impossible.

While kernel phase was originally derived under the assumption of phase-only errors, real optical systems inevitably exhibit coupling between phase and amplitude, especially in the presence of defocus or atmospheric turbulence (Pope 2016). These effects violate the idealised separation between visibility amplitude and phase assumed in conventional models. However, because the kernel formulation is a linear algebraic projection, it is naturally extensible: any system where visibility perturbations can be written as a differentiable function of instrumental parameters can, in principle, be linearised and projected onto a new self-calibrating subspace.

This generalisation plays a pivotal role in Chapter 4 of this thesis. By embedding forwards-modelled interferometric signals into a differentiable computational framework, we will demonstrate that the kernel approach can be extended beyond phase-only models, enabling calibration and inference even in the presence of non-trivial optical effects.

1.5 Automatic Differentiation

Automatic Differentiation (autodiff; Margossian 2019) is a foundational algorithm to modern computer science and Machine Learning (ML). In analogy to the way the fast Fourier transform turned the mathematical elegance of the Fourier transform into a practical computational tool, automatic differentiation brings the chain rule to life as an efficient and exact algorithm. Together, they exemplify how foundational mathematical concepts can be transformed into algorithmic pillars of modern computation. Autodiff enables the calculation of (mostly) arbitrary floating point function derivatives with respect to floating point arguments. While there are many ways to find the gradients of some function, symbolic differentiation and finite differences being the two most familiar to physicists, autodiff offers machine-precision exact derivatives and works on arbitrarily complex algorithms, even when faced with branching, looping, or intermediate numerical solvers. Despite being foundational to the field of ML, the power of autodiff as applied to the field of physics beyond the context of Neural Networks (NNs) is only beginning to be realised. This thesis explores a unification of physical modelling with autodiff, and hopefully sheds light on the path towards next generation tools to enable the ambitious goals of physics and astronomy over the coming decades.

Historically, autodiff was formalised during the 1960s and 1970s under the names “analytic

differentiation” or “computational differentiation” (Wengert 1964; Griewank 1988). Early implementations were used in scientific and aerospace computing applications, where accurate and efficient derivative calculations were critical (Unger and Hall 1994; Barthelemy and Hall 1995). The technique gained renewed prominence in the 1980s with the rise of backpropagation in neural networks, culminating in modern deep learning frameworks. Geoffrey Hinton, upon award of the 2024 Nobel Prize in physics for his foundational work on neural networks (Royal Swedish Academy of Sciences 2024), acknowledged autodiff as an essential component of the field’s success. As Graphics Processing Unit (GPU) hardware became capable of training complex NNs, the profit potential of ML was realised by major corporations like Google and Facebook (for better or worse). Consequently, major investments were put into the development of autodiff frameworks with strong focus on efficiency and usability. This set the stage for the modern landscape of the fast, user-friendly, open-source tools `PYTORCH` (Paszke et al. 2019) and `JAX` (Bradbury et al. 2018), developed by Facebook and Google respectively.

To appreciate the innovation that autodiff represents, it is instructive to contrast it with the two other primary methods for computing derivatives:

- **Finite difference methods** approximate derivatives numerically using small perturbations via

$$\frac{df}{d\theta} \approx \frac{f(\theta + h) - f(\theta)}{h}. \quad (1.13)$$

These methods are simple to implement but suffer from a number of issues: numerical instability, sensitivity to the choice of step size h resulting in truncation error, and most importantly, computational complexity that scales with the number of parameters differentiated, making it essentially useless for everything except small models with a few parameters.

- **Symbolic differentiation** manipulates expressions algebraically using the rules of formal calculus. While exact in principle, it can lead to exponential growth in expression size (expression swell) and does not apply easily to programs involving control flow, recursion, or external numerical solvers.

Automatic differentiation circumvents all of these issues: it is as accurate as symbolic differentiation, can be implemented entirely algorithmically, and has computational complexity agnostic to the number of differentiated parameters. It works by decomposing a computation into primitive operations and systematically applying the chain rule. In what follows, we build up reverse-mode autodiff from first principles. Beginning with the multivariate chain rule and the Jacobian matrix, and then develop the concepts of the core mathematical engine behind autodiff, the Vector-Jacobian Product (VJP) and Jacobian-Vector Product (JVP) — defined as the dot product between a vector and Jacobian or Jacobian and vector respectively.

1.5.1 Jacobians and The Chain Rule

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable scalar-valued function. Its gradient is defined as the vector of partial derivatives, i.e. its Jacobian

$$\nabla f(\boldsymbol{\theta}) = \mathbf{J}_f(\boldsymbol{\theta}) = \left[\frac{\partial f}{\partial \theta_1}, \dots, \frac{\partial f}{\partial \theta_n} \right] \in \mathbb{R}^{1 \times n}. \quad (1.14)$$

This gradient describes how much the function changes with respect to each parameter $\boldsymbol{\theta}$. If f is composed from a sequence of intermediate functions u , e.g.

$$f(\boldsymbol{\theta}) = u_k(u_{k-1}(\dots u_1(\boldsymbol{\theta}))) \quad (1.15)$$

or as a compute graph

$$\boldsymbol{\theta} \mapsto u_1(\boldsymbol{\theta}) \mapsto u_2(u_1(\boldsymbol{\theta})) \mapsto \dots \mapsto u_k(u_{k-1}(\dots u_1(\boldsymbol{\theta}))). \quad (1.16)$$

Given our universe is well described by a Markovian process — one where the next state depends only the one that precedes it — it is almost always possible to construct a differentiable mathematical representation of some physical processes in this way. The symmetry between the chain rule and physics simulation make autodiff the natural computational tool for our universe (Lin et al. 2017).

In practice, these functions are often multivariate during their calculation and cannot be expressed as simply as in Equation 1.15. Nonetheless, the same mathematical structure applies: complex programs can be decomposed into sequences of primitive operations, each contributing a partial derivative via the chain rule. These local sensitivities can then be composed to compute the total derivative of the output with respect to the inputs through

$$\frac{df}{d\boldsymbol{\theta}} = \sum_i \frac{\partial f}{\partial u_i} \cdot \frac{\partial u_i}{\partial \boldsymbol{\theta}}. \quad (1.17)$$

This formulation captures the essential insight of automatic differentiation: the derivative of a complex function can be built up from the composition of simpler derivatives. Depending on how this composition is carried out — either from outputs back to inputs, or from inputs forward to outputs — we obtain the two complementary modes of autodiff: reverse-mode and forward-mode. Each mode follows directly from this formulation but traverses the computational graph in opposite directions, with different tradeoffs in efficiency and memory usage. We now examine each of these methods in turn.

1.5.1.1 The Adjoint Method: Reverse-mode

Importantly, this equation can be understood as a VJP. The left term $\partial f / \partial u_i$, commonly expressed as \bar{u}_i , is known as the adjoint and describes the gradient of the original function f with respect the intermediate value u_i . The right term $\partial u_i / \partial \boldsymbol{\theta}$ is simply the Jacobian of the intermediate value u_i with respect to the original input variable $\boldsymbol{\theta}$. Extending this to the more common multivariate case we find

$$\nabla f(\boldsymbol{\theta}) = \sum_i \frac{\partial f}{\partial u_i} \cdot \nabla_{\boldsymbol{\theta}} u_i = \nabla_{\mathbf{u}} f \cdot \mathbf{J}_{\mathbf{u}}(\boldsymbol{\theta}) \quad (1.18)$$

where $\nabla_{\mathbf{u}} f \in \mathbb{R}^{1 \times m}$ is the gradient of the scalar output with respect to the vector of intermediates $\mathbf{u} \in \mathbb{R}^m$, and $\mathbf{J}_{\mathbf{u}}(\boldsymbol{\theta}) \in \mathbb{R}^{m \times n}$ is the Jacobian of the intermediate with respect to the input values $\boldsymbol{\theta}$. This equation is the general form of reverse-mode autodiff, providing a compute graph able to evaluate the vector of parameter gradients with respect to arbitrary functions.

1.5.1.2 The Tangent Method: Forwards-mode

While reverse-mode autodiff traces influence backward from outputs to inputs, forward-mode autodiff moves in the opposite direction. It tracks how perturbations to the input affect the computation as they flow forward through the graph. Where reverse-mode accumulates *adjoints* — quantities representing how changes in intermediate variables affect the output — forwards-mode accumulates *tangents*, which represent how changes in the inputs propagate to intermediate computations.

In reverse mode, Equation 1.17 is interpreted as propagating the gradient of f with respect to the intermediates (the adjoints $\partial f / \partial u_i$) back through the Jacobian $\partial u_i / \partial \theta$. In forward-mode, we flip our viewpoint: we push a perturbation in θ forward, computing how each intermediate variable u_i changes under that perturbation, and then how those changes affect the output. To formalise this, we introduce a small scalar parameter ϵ and define a directional perturbation of the input

$$\boldsymbol{\theta}(\epsilon) = \boldsymbol{\theta}_0 + \epsilon \cdot \mathbf{v}. \quad (1.19)$$

The directional derivative of f along \mathbf{v} , $df/d\epsilon$, describes how f changes when $\boldsymbol{\theta}$ is nudged in the direction \mathbf{v} . Using the chain rule, we can express this in a form directly analogous to the reverse-mode formulation of Equation 1.18. Letting $\mathbf{u} \in \mathbb{R}^m$ denote the intermediate variables, we obtain

$$\frac{df}{d\epsilon} = \left(\frac{d\mathbf{u}}{d\epsilon} \right)^\top \cdot \nabla_{\mathbf{u}} f = (\mathbf{J}_{\mathbf{u}}(\boldsymbol{\theta}) \cdot \mathbf{v})^\top \cdot \nabla_{\mathbf{u}} f. \quad (1.20)$$

This is a JVP: the Jacobian of the intermediate variables with respect to the input, applied to the input perturbation \mathbf{v} , then contracted with the gradient of the output with respect to the intermediates. This is the general form of forwards-mode autodiff, and is efficient when computing the gradient of a small number of inputs with respect to a large number of outputs — making it complementary to reverse-mode autodiff.

1.5.2 Differentiable Programming

To truly understand what makes a framework ‘differentiable’ we need to understand how to make Equations 1.18 & 1.20 efficient for computation and memory use. The key lies in how low-level mathematical operations are defined in the language. For every function that a dif-

ferentiable language implements, it holds the rules for how to evaluate it, but also the rules for computing its derivative, VJP, JVP, and some other higher-order terms are not relevant here. By storing these function transformation rules, operations like the VJP can be evaluated *directly* without ever materialising the (memory inefficient) full Jacobian.

Let us examine the mechanics of reverse-mode directly to understand this. To maintain computational efficiency, any differentiated input function performs both a ‘forwards’ pass where all the required intermediate values are cached, followed by a ‘reverse’ pass where the function transformation like the VJP can be applied. To see this in action we will construct a small example. Let

$$f(\theta_1, \theta_2) = (\theta_1 \theta_2)^2 + 3 \sin(\theta_1), \quad (1.21)$$

and define intermediate variables

$$u_1 = \theta_1 \theta_2, \quad u_2 = \sin(\theta_1), \quad (1.22)$$

which allows us to re-express f in terms of u_1, u_2

$$f(u_1, u_2) = u_1^2 + 3u_2. \quad (1.23)$$

From this expression it is trivial to then populate the terms of Equation 1.18

$$\nabla_{\mathbf{u}} f = [2u_1, 3], \quad \mathbf{J}_{\mathbf{u}}(\boldsymbol{\theta}) = \begin{bmatrix} \frac{\partial u_1}{\partial \theta_1} & \frac{\partial u_1}{\partial \theta_2} \\ \frac{\partial u_2}{\partial \theta_1} & \frac{\partial u_2}{\partial \theta_2} \end{bmatrix} = \begin{bmatrix} \theta_2 & \theta_1 \\ \cos(\theta_1) & 0 \end{bmatrix}. \quad (1.24)$$

The evaluation of the parameter Jacobian now becomes straight forward

$$\nabla f(\boldsymbol{\theta}) = [2u_1, 3] \cdot \begin{bmatrix} \theta_2 & \theta_1 \\ \cos(\theta_1) & 0 \end{bmatrix} \quad (1.25)$$

$$= [2u_1 \theta_2 + 3 \cos(\theta_1), 2u_1 \theta_1] \quad (1.26)$$

$$= [2\theta_1 \theta_2^2 + 3 \cos(\theta_1), 2\theta_1^2 \theta_2]. \quad (1.27)$$

On the forwards pass of the function, the values of $\theta_1, \theta_2, u_1, u_2$ can be cached with minimal overhead as they are scalars. Then in the reverse pass, provided the VJP rules of Equation 1.18, the final expression for the parameter gradients can be evaluated directly without ever materialising any intermediate Jacobians. More complex functions can have this processes applied recursively to each operation, returning exact parameter gradients. Remarkably, this computation only requires the equivalent of ~ 2 evaluations of the original function (forwards and reverse).

While the discussion thus far has focused on reverse-mode autodiff, a comprehensive autodiff system also implements forward-mode differentiation. Where reverse mode is optimal for functions with many inputs and a small number of outputs — such as scalar loss functions

— forward mode is ideal when the situation is reversed: few inputs, many outputs. Forward-mode autodiff is based on the Jacobian-vector product, which efficiently propagates directional derivatives through a computational graph. For functions that map $\mathbb{R}^k \rightarrow \mathbb{R}^n$ with small k , forward mode evaluates derivatives at a cost proportional to the number of inputs, making it a more efficient choice for Jacobians with respect to a small set of parameters or when computing sensitivities in simulation.

Crucially, modern autodiff frameworks like JAX and PYTORCH expose both modes in a composable way, allowing users to nest and combine them to compute both stable and exact higher-order derivatives. By applying forward- and reverse-mode transformations recursively, one can compute Jacobians, Hessians, Hessian-vector products, and even higher-order curvature information with remarkable efficiency. This is enabled by viewing differentiation as a sequence of linear transformations applied to a computational graph, and exploiting the associativity and distributivity of these operations using standard linear algebra tools. Rather than symbolically differentiating expressions, these libraries generate low-level code that propagate derivatives directly through the program’s control flow and memory layout.

Together, these capabilities provide a unified framework for efficiently differentiating through any model: analytical, numerical, neural, or hybrid. For physicists and astronomers, the implications are profound. Differentiation is the engine behind most forms of inference, optimisation, and sensitivity analysis. Historically, its application to physics-based models has been stymied by the complexity of those models: simulations involving partial differential equations, iterative solvers, or time-stepping methods could not be differentiated using classical calculus tools or were rendered intractable by the curse of dimensionality inherent to finite-differencing approaches. With autodiff, these limitations disappear. Provided the model can be expressed as code within a supported language, its derivatives become accessible, efficient and exact.

This synergy between physical modelling and autodiff is a central theme of this thesis. As a primarily observational science, astrophysical research is often driven through simulation: from the propagation of light through an optical system, to the redistribution of charge within a detector, to the inference of planetary orbits from projected astrometry. Each of these processes is inherently differentiable, or can be reformulated to be so. By embedding these models in autodiff-aware frameworks, we gain not just gradients, but the ability to calibrate, design, and infer from complex systems in a statistically coherent and computationally efficient way. It is this philosophy that underlies the models developed in subsequent chapters, where the boundary between simulation and inference dissolves, and autodiff becomes the glue that connects physical insight to data.

1.5.3 Neural Networks

Automatic differentiation is often perceived primarily as an enabler of machine learning — a utility developed to make training neural networks tractable. Indeed, the rise of deep learning owes much of its success to the fact that autodiff allows efficient and exact computation of gradients, even in models with millions or billions of parameters. Without autodiff, the optimisation of such high-dimensional, non-linear systems via gradient descent would be entirely

intractable.

However, this framing misses a broader truth: autodiff is not inherently tied to machine learning, but is a general-purpose computational tool. It operates at the level of programs, not disciplines. Neural networks just happen to be a particularly clear-cut example of non-linear functions that are differentiable by design and whose usefulness hinges on the ability to compute derivatives efficiently.

At a high level, neural networks are differentiable function approximators constructed from compositions of affine transformations and elementwise non-linearities. Their structure is both modular and expressive, and they are capable of modelling a wide range of functional relationships. The Universal Approximation Theorem guarantees that, under mild assumptions, even shallow networks can approximate any continuous function to arbitrary accuracy (Cybenko 1989; Hornik 1991). In practice, deeper networks offer more compact and structured representations of complex mappings.

Neural networks are fundamentally compositional. Each layer transforms an input space into a new, typically higher-dimensional representation, and the final output emerges from a series of these transformations. Because each operation is differentiable, the entire network becomes differentiable end-to-end — enabling gradients to be propagated efficiently with autodiff. This synergy between structure and differentiability is what makes neural networks so scalable and so effective in modern computational pipelines.

Numerous architectural variations exist — fully connected (dense) networks, convolutional networks, recurrent structures — but all share the core property of being differentiable, learnable functions. It is this property, not their connection to learning per se, that makes them compelling tools for abstract and complex problems.

1.5.4 Differentiable Models

While autodiff can be understood from a high level as ‘computational models with efficient derivative calculations’, its true power is more expansive. Beyond the ability to accelerate a wide range of classical and modern inference and optimisation techniques, access to leading edge statistical tools combined with efficient Jacobians and Hessians calculations makes differentiable models substantially more capable than their classical counterparts. It facilitates the acceleration of gradient based optimisation strategies like gradient descent, BFGS (Broyden 1970; Fletcher 1970; Goldfarb 1970; Shanno 1970), Levenberg-Marquardt (Levenberg 1944; Marquardt 1963), and Newton’s method (Ypma 1995) by providing curvature information without manual derivation. Importantly, these second-order and quasi-Newton methods, which were previously limited by derivative complexity, now become efficient and broadly applicable at essentially no additional computational cost when using autodiff.

In Bayesian inference, autodiff enables gradient-based sampling algorithms like Hamiltonian Monte Carlo (HMC; Betancourt 2017; Duane et al. 1987; Brooks et al. 2011). Unlike traditional single-site or random-walk Markov Chain Monte Carlo (MCMC) methods (Hastings 1970; Metropolis et al. 1953; Brooks et al. 2011), which break down in moderate- to high-dimensional parameter spaces (typically above 10-12 dimensions) (Huijser et al. 2022), gradient-informed

samplers can explore the posterior effectively in high dimensions. This capability is crucial for forward models in astronomy and physics, where parameter spaces can easily have tens to hundreds of dimensions. Furthermore, having direct access to Hessians allows for Laplace approximations (Kass et al. 1991) and computation of Fisher information matrices (Fisher 1925) — which are only made stable through autodiff (Bhandari et al. 2021) — essential for Bayesian experimental design, parameter uncertainty estimation, and metric-based search strategies.

Autodiff also supports techniques that respect the intrinsic geometry of parameter spaces. By computing Fisher metrics or identifying null- and orthogonal-subspaces, autodiff enables natural-gradient optimisation (Amari 1998; Martens 2020) and manifold-aware parameter transformations (Benner et al. 2015), improving convergence and maintaining essential dynamics in physical systems via projections to lower-dimensional, statistically independent spaces.

A major benefit of autodiff is its ability to handle differentiation through iterative solvers, such as differential equation integrators or root-finders (Kidger 2021). This capability makes dynamic and simulation-based forward models fully differentiable “black-box” components. It opens doors to gradient-based sensitivity analysis, model calibration, control, and solver-in-the-loop optimisation, enhancing both the interpretability and efficiency of scientific models.

Together, these capabilities can be woven into a consistent narrative: by offering exact, efficient derivatives through arbitrary numerical code, autodiff transforms forward models into richer, more tractable objects. This enables optimisation, uncertainty quantification, design, and simulation workflows that were previously cumbersome or computationally infeasible. This interplay between autodiff and modelling sits at the heart of this thesis, forming the foundation for the hybrid physical-learned models and inference methods developed in later chapters.

1.6 Modelling Paradigms for Scientific Inference

Scientific models are not neutral tools — they encode assumptions about causality, uncertainty, and the nature of the data-generating process. In astronomy, these assumptions are often implicit: inverse reductions, correction chains, and surrogate models are widely used without recognition as distinct methodological choices. While machine learning is frequently highlighted, the paradigm within which it is used — as an emulator, an inverse map, or a physical surrogate — is rarely stated. This lack of explicit framing can obfuscate the limits of what models can meaningfully infer.

Because this thesis adopts a modelling framework that departs from standard practice — integrating physical realism, differentiability, and inference into a unified structure — it is important to first outline the dominant paradigms of scientific modelling. The aim is not philosophical, but practical: the form of a model fundamentally shapes what can be asked of data, and what answers can be trusted.

“*All models are wrong, but some are useful.*” (Box 1976)

This phrase captures a central truth of scientific inference: models are not faithful replicas of reality, but structured approximations, crafted for utility within specific regimes. They reflect both our understanding of physical systems and the constraints under which we operate, be they computational, observational, or conceptual.

At the highest level, the task of inference can be expressed as

$$F(\boldsymbol{\theta}_{\text{true}}) = \mathbf{X} \quad (1.28)$$

where F denotes the true physical process mapping parameters $\boldsymbol{\theta}$ to data \mathbf{X} . In practice, we only have access to an approximate model f , and must use it to reason about $\boldsymbol{\theta}$. Whether f is inverted, sampled, learned, or simulated categorises the modelling paradigm.

In observational astronomy, the dominant paradigm has been inverse modelling. Pipelines such as those used by JWST apply a sequence of modular corrections — for example linearity, flat-fielding, dark subtraction, ramp-fitting — to transform raw detector voltages into calibrated science products. Each step aims to undo some known effect of the measurement chain, implicitly assuming that those effects are separable and invertible.

This modularity has practical appeal: it is tractable, interpretable, and often sufficient. But it also encodes a view of inference in which systematics can be sequentially corrected in isolation, and where uncertainties are only understood in an isolated context. This thesis, in response to the exacting precisions demanded by exoplanetary science, challenges that view. It explores models that simulate the system as a whole, preserve physical structure, and propagate uncertainty coherently. To contextualise this shift, we begin by outlining the four major modelling paradigms used in contemporary scientific inference: inverse, machine learned, forward, and hybrid.

1.6.1 Inverse Modelling

Inverse models form the dominant framework for inference in observational astronomy. Rather than simulating how parameters $\boldsymbol{\theta}$ produce data \mathbf{X} , they begin with the data and attempt to recover $\boldsymbol{\theta}$ by inverting an approximate model

$$\hat{\boldsymbol{\theta}} = f^{-1}(\mathbf{X}). \quad (1.29)$$

This approach underlies both formal data reduction pipelines and many common modelling tasks, even when not explicitly labelled as such. In both cases, the system is treated as a sequence of independent transformations to be reversed, with each step targeting a specific effect assumed to be separable and invertible. A canonical example is the JWST calibration pipeline, which applies dark subtraction, non-linearity correction, flat-fielding, and ramp fitting. These models are favoured for their efficiency and interpretability, but they rest on strong assumptions: that instrumental effects are stable, decoupled, and correctable in isolation. Uncertainties are typically propagated locally, if at all, and errors introduced early in the chain accumulate without feedback or reassessment.

These limitations become critical in regimes that demand precision and robust error control.

Inverse models break down when systematics are non-linear, coupled, or information-destroying. In such cases, it is the assumptions built into the model — not the information content of the data — that set the limits on inference.

A representative failure appeared during early JWST transit observations, when a mirror segment tilt shifted the PSF and introduced a discontinuity in the aperture photometry (Rigby et al. 2023, Figure 10). Because the model had no representation of the optical state, it could not detect or account for the change directly. The result was a spurious instrumental feature masquerading as an astrophysical signal.

1.6.2 Machine Learning

Unlike inverse models, ML approaches rely on flexible architectures such as NNs to learn a mapping from observed data \mathbf{X} to parameters $\boldsymbol{\theta}$, without explicitly modelling the underlying process f . Once trained, these models can be highly efficient, scale well with data, and capture complex relationships that evade traditional methods.

However, their strength lies in prediction rather than explanation. As Hogg and Villar (2024) point out, most ML models are built around an ontology where only data relationships matter, and validation is based on performance on held-out sets — not on capturing underlying mechanisms. This can conflict with the goals of scientific inference, where understanding causal structure is often as important as making accurate predictions.

Consider ML-based PSF emulators for wavefront sensing in adaptive optics. A typical pipeline simulates PSFs using an optical model with injected wavefront errors, then trains a NN to invert these outputs (Andersen et al. 2019; Weinberger et al. 2024). While effective within the training regime, these models implicitly learn both the diffractive physics and the instrumental configuration, despite both being known. Because they do not encode physical parameters or constraints directly, they lack interpretability and often extrapolate poorly when conditions shift, for instance, due to changes in alignment, turbulence, or thermal state (Liu et al. 2020).

As Hogg and Villar (2024) caution, ML models that replace physical components can amplify confirmation bias. A network trained on nominal system conditions may perform well in typical settings but fail silently when unexpected systematics arise. Without encoded structure or uncertainty modelling, such failures can propagate into downstream inference, leading to biased or invalid results.

While ML offers exceptional predictive power, this is often traded against interpretability, uncertainty quantification, and robustness to out-of-distribution data. In common with inverse models, ML depends on implicit assumptions — here about training data representativeness and model inductive biases — and lacks internal mechanisms to identify or correct failures when those assumptions break.

1.6.3 Forward Modelling

Forward models simulate the generative process by which parameters $\boldsymbol{\theta}$ produce observations \mathbf{X} . Unlike inverse or ML approaches, which aim to undo or approximate this process, forward

models seek to reproduce it directly

$$f(\boldsymbol{\theta}) = \mathbf{X}_{\text{pred}}. \quad (1.30)$$

Inference proceeds by adjusting $\boldsymbol{\theta}$ to minimise the mismatch between the predicted and observed data, typically via

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{X}_{\text{obs}}, f(\boldsymbol{\theta})), \quad (1.31)$$

where \mathcal{L} is a loss function, such as a likelihood or residual norm. This reframes inference as optimisation within a constrained generative model.

Historically, forward modelling has been underutilised, mainly due to computational limitations. Evaluating complex models across high-dimensional parameter spaces is expensive, and many physical simulations were non-differentiable or locked in legacy code, making gradient-based inference impractical. Inverse approaches, with modular corrections and local approximations, were preferred for tractability.

In recent years this landscape has shifted. Modern autodiff frameworks enable gradients through complex simulations, allowing physical forward models to be optimised efficiently and embedded in probabilistic inference workflows. This enables differentiable forward models: physically grounded simulations that support scalable optimisation, uncertainty propagation, and rigorous inference. These models offer several key advantages:

- They simulate the full data-generating process, preserving causal relationships across system components.
- They enable coherent uncertainty propagation, tracing parameter uncertainty through to observables.
- They retain interpretability: model parameters correspond to physical quantities.
- They can incorporate learned submodels where physics is unknown, without compromising overall structure.

Unlike inverse or ML models, forward simulations generate predictions from explicit physical assumptions. This makes residuals between model and data informative: they reflect missing physics, miscalibration, or unmodelled effects. In contrast, errors in inverse or ML pipelines can be harder to interpret or diagnose.

Recent works in astronomical optics have explored differentiable forward models. Liaudat et al. (2023a) presented a diffraction-constrained differentiable optics model, but situates it within an inverse framework reliant on pre-corrected inputs. As noted in Liaudat et al. (2023b), this limits its utility and accuracy. Stone et al. (2023) used autodiff to model flexible PSFs, enabling excellent photometry, but their models are not physically constrained and thus lack generalisability beyond photometric tasks.

1.6.4 Hybrid Modelling

The emergence of differentiable programming frameworks has enabled a new class of models that combine physical simulations with data-driven components. These hybrid models integrate mechanistic structure where available with machine-learned components where necessary. This allows models to remain grounded in domain knowledge while extending their expressivity to capture unknown or computationally intractable processes.

A canonical example is the physics-informed neural network (Raissi et al. 2019), which augments a NN’s loss function with penalties for violating physical laws, typically expressed as differential equations. This bias toward physically consistent solutions allows them to generalise better, learn from limited data, and retain interpretability. Similar principles underpin neural ordinary differential equations (Chen et al. 2019), where the dynamics are parameterised by a learnable function. This approach has been extended to stochastic and controlled differential equations (Kidger et al. 2020; Kidger 2022a), enabling hybrid models to represent complex time-series processes with embedded structure.

Hybrid modelling has gained traction across the physical sciences. In fluid mechanics, NNs are used to learn closure terms while preserving the Navier-Stokes formulation (Ling et al. 2016). In climate modelling, hybrid systems approximate unresolved sub-grid processes within physical solvers (Stengel et al. 2020). In molecular dynamics, learned potentials accelerate simulation while conserving energy and momentum (Schütt et al. 2017). In computational imaging, the Deep Optics framework (Sitzmann et al. 2018a) jointly optimises physical phase mask parameters and a neural reconstruction module in a fully differentiable optical pipeline. These examples show how hybrid models can embed learning into structured systems, aligning with the goals of inference rather than prediction alone.

1.6.5 From Pixels to Planets

This thesis adopts a unified modelling framework, hereby referred to as the ‘Pixels to Planets’ philosophy, developed in response to the limitations of conventional methods in precision astronomical imaging. It models the entire measurement and inference process — from photon arrival to detector readout — as a single, differentiable computational graph grounded in physical principles. Simulating the full forward chain enables coherent uncertainty propagation and preserves causal structure throughout the system.

The framework centres on differentiable forward models that balance physical realism with flexibility. Rather than enforcing idealised or invertible models, the structure is shaped by what the data can meaningfully constrain. When parts of the system are too complex or uncertain to simulate directly, hybrid components such as neural submodules can be integrated without sacrificing interpretability or coherence.

A central advantage is lifecycle continuity. The same model supports instrument design, calibration, commissioning, and scientific analysis. This avoids the need for separate models or stage-specific corrections. Instrument design is explored in Chapter 3; calibration and analysis are addressed in both Chapters 2 & 4.

This approach is particularly effective for modelling non-linear or coupled systematics. By

simulating how these effects distort observations, rather than attempting to correct them after the fact, the model captures complex behaviour coherently. Uncertainties remain traceable from input to output, enabling principled error propagation. Residuals between prediction and data become scientifically meaningful. They reflect gaps in the model — such as unmodelled physics, miscalibrated parameters, or structural assumptions — rather than arbitrary errors.

The Pixels to Planets framework emerged directly from the modelling challenges addressed in this thesis. It enabled solutions where inverse methods failed, without relying on black-box ML or rigid pipelines. The differentiable optics library ∂ LUX and the AMIGO calibration framework embody this approach, offering a flexible and physically grounded path for next-generation inference in astronomical instrumentation.

Chapter 2

Deep Calibration of Flat Field and Phase Retrieval with Automatic Differentiation

The precision limits of modern space telescopes are increasingly dictated not by fundamental noise processes, but by residual systematics in optical propagation and detector response. Aberrations in the Point Spread Function (PSF) and spatial variations in pixel-level sensitivities both contribute systematic errors that degrade the scientific return of high-precision observations. Conventional strategies rely on static pre-flight calibrations or specialised hardware, which remain vulnerable to non-common path effects and time-evolving systematics. This chapter introduces ∂ LUX, the differentiable optical modelling framework built in JAX, which enables high-dimensional inference and offers native hardware-acceleration. We demonstrate how differentiable modelling with ∂ LUX unifies instrumental calibration and scientific inference into a single, jointly solvable optimisation problem. Under traditional approaches, the dimensionality of this problem would mandate a separation of calibration and analysis; here, we show that by treating the full image formation chain as a differentiable model, we can recover wavefront aberrations, pixel sensitivity variations, and astrophysical parameters simultaneously at scale. This chapter lays the foundation for the end-to-end modelling philosophy developed throughout this thesis, extended in the next chapter to hardware design, and culminating in the restoration of James Webb Space Telescope (JWST; Gardner et al. 2006) Aperture Masking Interferometer (AMI; Sivaramakrishnan et al. 2012; Soulain et al. 2020)’s performance to pre-flight specifications — an outcome that has remained elusive since launch.

2.1 Statement of Contribution

This chapter faithfully reproduces material from Desdoigts et al. (2023), which was accepted for publication in the *Journal of Astronomical Telescopes, Instruments, and Systems*, Vol. 9, Issue 2, 028007 (June 2023), with minor adaptation to fit into the thesis style. I am the first author, with Benjamin Pope, Jordan Dennis, and Peter Tuthill as co-authors. The research problem was conceived by both Peter and Ben, originated by the challenges presented by the calibration of the TOLIMAN (Telescope for Orbital Locus Interferometric Monitoring of our Astronomical Neighbourhood; Tuthill et al. 2018) instrument. Jordan aided in the early development of ∂ LUX with continual advice provided by Ben and Peter. All of the code, models, simulations, and inference for the paper was done by myself. I wrote the paper, with co-authors contributing feedback on the scientific results, reviewing the manuscript, and assisting with minor improvements to the code and interpretation. All major analytical and computational work was performed by me.

2.2 Introduction

At the vanguard of space-based astronomical imaging, photometry, and spectroscopy, imaging precision is limited by systematics introduced by aberrations in the optics and noise processes in the detector. While problems are ubiquitous in astronomy, we are motivated by several core examples. In exoplanet direct imaging, we want to achieve high resolution and high contrast simultaneously with instruments like JWST Coronagraphy (Girard et al. 2022), Aperture Masking (Sivaramakrishnan et al. 2022), or Kernel Phase (Martinache 2010; Kammerer et al. 2022) modes. The astrometric mission Toliman (Tuthill et al. 2018) aims to measure precise relative positions of the binary stars α Centauri AB to reveal the gravitational influence of unseen planets; this will require micro-arcsecond, micro-pixel astrometric precision. We may also want to perform high precision photometry, whether with dedicated missions like *Kepler* (Borucki et al. 2010) and TESS (Ricker et al. 2015), or as ancillary science with Toliman.

In each case, we face serious limitations from an imperfect knowledge of the Pixel Response Function (PRF) — the map of the intra- and inter-pixel variations in detector sensitivity; and also of the PSF — the diffraction-limited pattern by which light from a point source like a star spreads across a detector.

In this series of papers, we present a new software package, ∂ LUX, for fitting high-dimensional parametrised physical optics models to astronomical data. By using the Python library JAX (Bradbury et al. 2018), we obtain Graphics Processing Unit (GPU) hardware acceleration, as well as automatic differentiation or ‘autodiff’ (Margossian 2018) features that enable high-dimensional optimisation and inference with gradient descent or Hamiltonian Monte Carlo (HMC). In the present paper, we focus on the particular problem of estimating the PRF and PSF simultaneously — i.e. joint flat field calibration and phase retrieval. In Paper II, we will show how autodiff enables improvements to hardware *design*, directly calculating and optimising the Fisher information and therefore fundamental figures of merit of an optical system end-to-end. And in subsequent papers, we will use this for end-to-end deconvolution of high angular resolution imaging with an unknown PSF. We make ∂ LUX available as open-source software on GitHub¹, and encourage interested readers to use and contribute to this package as a community resource.

2.2.1 Phase Retrieval

The PSF depends on the successive planes through which light passes from the entrance pupil of the telescope through to the detector. This can be calculated from physical optics: for a simple camera bringing light from the pupil to focus at a detector plane (the regime of Fraunhofer diffraction), the PSF is the Fourier transform of the input wavefront. In the more general Fresnel regime where the detector is slightly out of focus, there are additional quadratic phase factors before and after this Fourier transform. In more complicated instruments like coronagraphs (Bowler 2016), the light might pass through multiple re-imaged pupil and focal planes, being operated on in each. In either case, the trouble is that the PSF is distorted by unknown

¹github.com/LouisDesdoigts/dLux

aberrations — distortions in the wavefront, which can be represented as a spatially-varying phase map across the optical plane.

Phase retrieval is the problem of inferring these aberrations from data (Shechtman et al. 2014), which is in general ill-posed (Barnett et al. 2020): because of the Hermitian symmetry of the Fourier transform, and because we measure *intensity* and not electric field in optical astronomy, there is a large space of aberrations that would generate the same intensity PSF. Fortunately, this space can be restricted to physically-realistic solutions and readily solved by algorithms such as the Gerchberg-Saxton algorithm (Gerchberg and Saxton 1972), using ideas from compressed sensing (Candes et al. 2011), or by machine learning (Metzler et al. 2018; Işil et al. 2019; Nishizaki et al. 2020). Phase retrieval was memorably performed to infer and correct the serious aberration on the *Hubble Space Telescope* mirror at launch (Fienup et al. 1993). Earlier work in the vein of this paper has shown that the phase retrieval problem can be efficiently solved even in the case of detector nonlinearity by taking a forwards model of physical optics, obtaining partial derivatives with autodiff, and optimising its parameters by gradient descent (Jurling and Fienup 2014; Wong et al. 2021), including a conference presentation of an early version of the present work (Desdoigts et al. 2022).

2.2.2 Detector Calibration

High-precision, high-cadence time series of space photometry from *Kepler* (Borucki et al. 2010), K2 (Howell et al. 2014), TESS (Ricker et al. 2015) and CHEOPS (Broeg et al. 2013) space telescopes have been revolutionary for exoplanetary science (Zhu and Dong 2021) and stellar astrophysics (Aerts 2021; Jackiewicz 2021).

In many practical cases, photometric precision is limited by systematic errors due to variations in sensitivity within and between pixels. Changes in telescope pointing or variations in the PSF with focus couple to these inter- and intra-pixel variations in the flat field to produce changes in overall measured flux. This was particularly severe for the K2 mission where periodic thruster firings introduced a saw-tooth systematic on 6 hour timescales, comparable to the durations of exoplanet transits and asteroseismic signals of interest.

Much work has been put towards developing data-driven self-calibration of the flat field in K2, which are of general interest and applicability for other space photometry missions. In Self Flat-Fielding (SFF; Vanderburg and Johnson 2014), the systematics are modelled with a spline regression of raw flux versus detector position, together with a linear combination of principal components of all light curves on the detector, and *k2sc* (Aigrain et al. 2016) uses a gaussian process similarly. The Causal Pixel Model (CPM; Wang et al. 2016) constrains instrumental effects by measuring correlations in the light curves of spatially distant, causally-disconnected pixels. In Pixel Level Decorrelation (PLD; Luger et al. 2016) the light curve is de-trended by a linear combination of regressors formed by the ensemble of normalised pixel time series contributing to the light curve, together with their higher-order products. Halo photometry (White et al. 2017; Pope et al. 2019) extracts light curves from a weighted sum of pixels where the weights are learned by minimising the total variation of the resulting light curve, a convex optimisation which is tractable because of autodiff.

Accurate measurement of background noise is also a serious issue for imaging science, where for example bad pixels, background noise, and uncertainties in the detector systematics limit sensitivity of aperture masking interferometry and kernel phase (Kammerer et al. 2019), and for low-surface-brightness science such as with *Euclid* (Collaboration et al. 2022). In each case it is normal for detectors to be rigorously calibrated on the ground before flight, and with dedicated observations in space; but at the levels of contrast and resolution required for aperture masking on JWST (ideally 10^{-4} at a few λ/D) even small residual mis-calibrations can seriously affect performance (Sivaramakrishnan et al. 2022) and we require some method of self-calibration from science data themselves.

2.3 Differentiable Optical Models

Rather than working with reduced data products like visibilities and light curves, an alternative is to directly model the pixel-level images with a parametrised model of the optics and detector. For example, the popular package POPPY (Perrin et al. 2012) can accurately simulate Fresnel and Fraunhofer propagation, with a WebbPSF extension for preset JWST instrument models (Perrin et al. 2014). In principle, it is possible to fit such a model to data with Markov Chain Monte Carlo (MCMC), fitting the positions of stars and a modal basis representation of aberrations; but common samplers can be very slow, or fail to converge in high dimensions (Huijser et al. 2022). This restraint is lifted by using HMC, which requires computation of the gradient of the log-likelihood with respect to parameters.

The core technology that enables this for physics also underpins the last decade’s revolution in deep learning (LeCun et al. 2015): autodiff. By application of the chain rule, it is possible to calculate the partial derivatives of almost arbitrary numerical functions with respect to their floating-point arguments, so long as they are executed in an appropriate software framework. Considerable private-sector and open-source investment has gone into building autodiff frameworks to enable machine learning, such as Theano (Team 2016), TensorFlow (Martín Abadi et al. 2015), PyTorch (Paszke et al. 2019), native support in the Julia language (Bezanson et al. 2012), and the framework we apply in this paper, JAX (Bradbury et al. 2018). JAX has an Application Programming Interface (API) mirroring the common Python array library NUMPY (Harris et al. 2020), while supporting GPU acceleration, Just-In-Time (jit) compilation, and autodiff.

In this paper we introduce a new optical simulation package, ∂ LUX, available under an open-source BSD 3 license. It significantly extends the capabilities of *morphine* (Pope et al. 2021; Wong et al. 2021), a previous JAX optics package from our team based on *poppy*, being rewritten from the ground up taking full advantage of more-recent JAX features and libraries. It is built in an object-oriented JAX framework called ZODIAX which extends EQUINOX (Kidger and Garcia 2021a) for scientific programming. An optical system consisting of a ZODIAX module wrapping a stack of ‘layers’ which are themselves ZODIAX modules applying phase or amplitude screens, or performing Fourier or Fresnel transforms on a wavefront. This allows us to re-frame the idea of a computational optical model by analogy to a ‘parametrised neural network’, with each operation performed on the wavefront analogous to a neural network layer encoding the

physics of the transformation. The object-oriented nature of ZODIAX allows for larger and more complex models to be built and seamlessly integrated than would be possible within the purely functional framework default to JAX. ZODIAX is also built to be directly integrated into the JAX framework, allowing for all JAX functions to seamlessly interface with any ZODIAX models and vice-versa. We use the optimisation library OPTAX (Hessel et al. 2020) for gradient-descent optimisation, and the probabilistic programming library NUMPYRO (Phan et al. 2019a) for HMC.

There are several other frameworks for automatically-differentiable optics simulation: the most similar is WaveBlocks (Page and Favaro 2020), a PyTorch package for object-oriented modelling of fluorescence microscopy. The DeepOptics package (Sitzmann et al. 2018a) has been developed in TensorFlow for camera design and computational imaging, and WaveDiff in TensorFlow (Martín Abadi et al. 2015) has been used for detector calibration (Liaudat et al. 2021) and PSF modelling (Liaudat et al. 2022). We believe that ∂ LUX fills a unique niche in this ecosystem as the most general modelling framework for physical optics. This has driven a design focus on flexibility and ease of open-source contribution and development. The goal is to provide a general object-oriented framework that enables a new user to harness the advantages of autodiff and hardware acceleration provided by JAX with its user-friendly NUMPY-like API.

2.4 Optical and Instrumental Modelling

Within the ∂ LUX framework we can construct an optical model with only a pupil and focal plane like so:

```
import jax.numpy as np
import jax.random as jr
import dLux as dl
import dLux.utils as dlu

# Define wavelengths
wavels = 1e-9 * np.linspace(545, 645, 3)

# Basic Optical Parameters
diameter = 0.5 # meters
wf_npix = 256

# Construct a simple aperture
oversample = 3
coords = dlu.pixel_coords(oversample * wf_npix, diameter)
outer_circle = dlu.circle(coords, diameter / 2)
inner_circle = dlu.circle(coords, diameter / 10, invert=True)
transmission = dlu.combine([outer_circle, inner_circle], oversample)

# Construct Zernike aberrations
noll_inds = np.arange(4, 11)
coeffs = 2e-8 * jr.normal(jr.PRNGKey(1), (len(noll_inds),))
coords = dlu.pixel_coords(wf_npix, diameter)
basis = dlu.zernike_basis(noll_inds, coords, diameter)

# Detector Parameters
det_npix = 1024
```

```

det_pixsize = 12e-3 # arcseconds

# Build the layers
layers = [
    ("aperture", dl.layers.BasisOptic(basis, transmission, coeffs, normalise=
        True)),
    ("mask", dl.Optic(phase=np.load("mask.npy"))),
]

# Combine into Optics object
optics = dl.AngularOpticalSystem(wf_npix, diameter, layers, det_npix,
    det_pixsize)

```

The aperture object defines the pupil support as well as low-order phase errors parameterised by coefficients on Zernike polynomials. We include the first 4 radial terms, ignoring global piston, tip and tilt for a total of 7 terms. Coefficients were randomly drawn from a normal distribution of mean 0 and standard deviation 20 nm. These aberrations applied to the pupil mask and resulting PSF are shown in the middle panel of Figure 2.1, where we see this induces a large amount of distortion in the PSF.

The mask object models the binary phase plate diffractive pupil from Toliman (Tuthill et al. 2018), with a half-wave total Optical Path Difference (OPD) between the binary regions. We use this as an information-preserving way to spread the PSF across a large number of pixels which encodes the PRF information for a large number of pixels simultaneously. Alternative methods such as defocusing or using a diffuser (Stefansson et al. 2017) come at the cost of eliminating Fourier information which constrain the wavefront and positions. An important consideration which is alleviated by the Toliman pupil is that focal-plane wavefront sensing suffers a sign ambiguity for all phase modes that are inversion-symmetric about the origin, and for unambiguous phase retrieval we therefore need an asymmetric or odd-mode symmetric pupil (Martinache 2013). This pupil and corresponding aberrated PSF are shown in the left panel of Figure 2.1.

This is then propagated to the 1024×1024 pixel focal plane by a two-sided matrix Fourier transform (Soummer et al. 2007; Martinache et al. 2020). The detector PRF is modelled by multiplying the resulting PSF by a pixel sensitivity map which is drawn from a normal distribution with mean 0 and standard deviation 0.1. The histogram of pixel sensitivities and total pixel sensitivity map is shown in the right panel of Figure 2.1. Here we show how to create this within the ∂ Lux framework and create an instrument object to model this:

```

# Pixel response
pix_response = 1 + 0.1 * jr.normal(jr.PRNGKey(2), [det_npix, det_npix])

# Create Detector object
detector = dl.LayeredDetector([dl.ApplyPixelResponse(pix_response)])

```

This operation is then performed over 20 stars with uniform 100 nm bandpass spectra sampled at three wavelengths 545 nm, 595 nm and 645 nm. This total model requires the calculation of 60 individual PSFs over a 1024×1024 detector from a 512×512 pixel wavefront. This calculation is vastly dominated by the propagation of the wavefront using a 2-sided MFT, which requires two matrix multiplications for each individual PSF of sizes $[N_o \times N_i] \cdot [N_i \times N_i]$

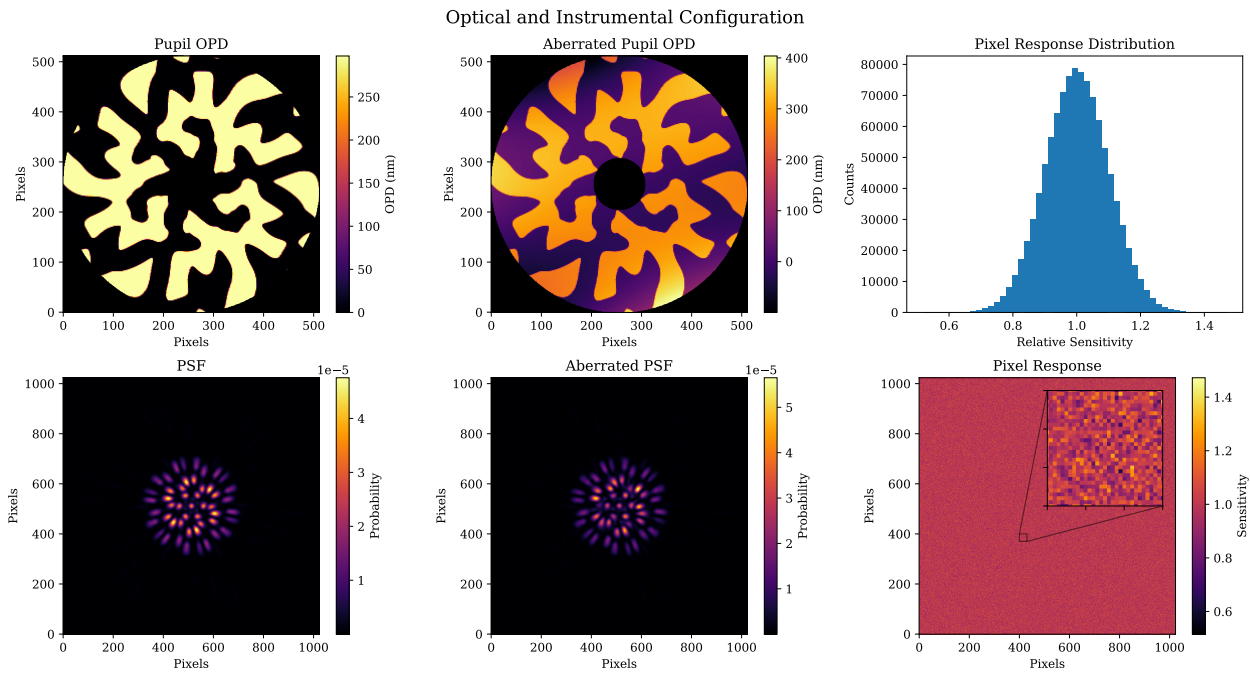


Figure 2.1: Summary of the optical configuration. Left panel: The top plot show the OPD of the pupil at the aperture of the telescope. The binary values create a half-wave step at the mean observation wavelength. The bottom plot shows the resulting large single-star PSF without aberrations applied. Middle panel: The top plot shows the total OPD of the pupil with the optical aberrations applied using low-order Zernike polynomials that is used to generate the data. The bottom plot shows the resulting single-star PSF with the aberrations applied. Clearly these aberrations have a large effect on the PSF and would make recovering information very difficult without appropriate calibration. Right panel: The top plot shows the histogram of the PRF that is applied in the focal plane. These values have a large spread and would greatly affect any results without calibration. The bottom plot shows the full PRF across the whole detector, with a small zoomed region used to show the fine detail that can not be seen when examining the full detector.

and $[N_o \times N_i] \cdot [N_i \times N_o]$ where $N_i = 512$ is the input size of the wavefront and $N_o = 1024$ is the output size of the wavefront. This takes a total of ~ 1.25 seconds to evaluate on an Apple M1 CPU, with all operations performed using FP32 precision default to JAX. FP32 was chosen as this problem does not require precision over a large dynamic range, however FP64 calculations can be performed as required. We also apply 4 integer-pixel dithers of ± 20 pixels in the (x,y) directions, equivalent to a $1 \lambda/D$ dither for a total of 5 images with the total flux spread evenly across each image. This helps the over-parametrisation consequential from having a free parameter for each pixel. Using pixel-integer dithers similarly allows for a single PSF calculation on a slightly larger detector plane (1064x1064 in this case) and then sampling the smaller 1024x1024 regions to produce the full set of 5 images.

The data set is generated by taking this set of images and applying Poisson photon noise to the PSFs. Using this image we recover stellar positions and fluxes, optical aberrations, and individual pixel sensitivities. This gives us 40 stellar positional parameters (20 RA and DEC each), 20 stellar flux parameters, 7 optical aberration parameters, and 1,048,576 detector parameters: the dimensionality of the problem is vastly dominated by the PRF parameters, which is why this problem is made tractable only by using autodiff.

2.5 Results

2.5.1 Optimisation

To recover these parameters we initialise a naive unaberrated optical model with a uniform PRF. Stellar positions were perturbed by adding a random value drawn from a normal distribution mean 0 deviation 1 pixel and fluxes were perturbed by multiplying their value by a random value drawn from a normal distribution mean 1 deviation 0.1. We will use this model to recover the true values. Figure 2.2 shows the data and the initial and final residuals.

In order to recover these input aberrations, we perform gradient descent using Adam (Kingma and Ba 2014) as implemented in OPTAX (Hessel et al. 2020), minimizing the posterior of a per-pixel Poisson log-likelihood with a χ^2 log prior on the PRF values, matching the true distribution of mean 1 deviation 0.1. We take the gradient of this loss function with respect to the position and flux of all stars, the response in all pixels, and the Zernike coefficients. We also use a staged optimisation strategy, initially optimising the positions, fluxes and optical aberrations in order to recover a good PSF model before learning the PRF. This is necessary to circumvent the high degree of covariance that each individual pixel experiences with the parameters of the PSF. To achieve this the learning rate of the PRF is set to zero for the first 100 epochs of optimisation, after which the PRF learning rate is initialised while the position, flux and aberrations parameters are frozen for the remaining 50 epochs. This helps to circumvent the global covariance between the mean PRF and flux which can result in an overestimation of the flux parameters as the model tries to fit to the noise.

The optimisation is performed for a total of 150 epochs, with each value-and-gradient calculation taking ~ 2.75 seconds to evaluate on an Apple M1 Max CPU, for a total optimisation time of ~ 6.5 minutes. Faster convergence can be achieved by careful tuning of the optimisation




Figure 2.2: Left: The sum of the images from which the PRF is recovered. By eye it is clearly difficult to disentangle the astrophysical information. The large number of overlapping PSFs is chosen in order to spread light across the majority of the detector so that we encode the PRF information for as much of the detector as possible. Middle: The sum of the residual of the data and the initial uncalibrated model. Clearly these residuals are large, showing that there is a large amount of calibration required. Right: This same residual after the model has been optimised. The residual values are much smaller with no discernable structure remaining. A small zoomed region is shown so that the individual pixel-level residuals can be seen.

hyper-parameters such as learning rate and momentum, however we take a slower approach in this work in order to get good performance across a range of fluxes and PRFs which is analysed in Section 2.6.

2.5.2 Phase Retrieval & Flat Field Estimation

This optimisation strategy performs well, with Figure 2.2 showing the final residual image with no visually discernible structure. A further analysis of parameter recovery over a range of fluxes and PRFs is shown in Section 2.6. Correlations and residual plots of the positions, fluxes and aberrations after the optimisation are shown in Figure 2.3 and Figure 2.4. Figure 2.4 shows the recovery of the optical aberrations, with a final residual aberration of 0.0282 nm RMS.

Finally, Figure 2.5 shows a correlation plot of the true and recovered sensitivity values and the corresponding residual histogram. The points of the correlation plot are colour-coded with the total counts for each pixel, showing as expected that pixels with higher photon counts having a better recovery of the PRF since these pixels have a greater signal-to-noise ratio (SNR). Overall the parameters are very well recovered down to the noise level.

2.5.3 Uncertainty Estimation

We can also leverage the differentiability of these models in order to estimate the errors in these parameters under the Laplace approximation (Bard 1974), by taking the inverse Hessian of the Poisson likelihood. Using this method marginalised errors can be calculated incorporating parameter covariance. There are some limitations using this method as it requires the calculation of the covariance matrix over all parameters — which in this model is over a million, giving a covariance matrix of over a computationally-intractable trillion values. Instead

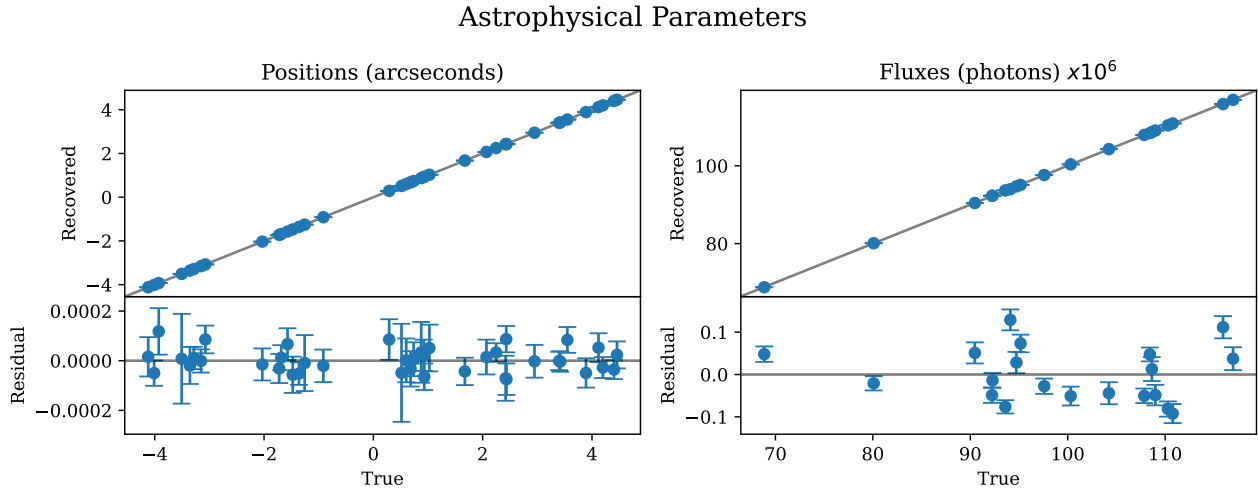


Figure 2.3: Left: recovery of the individual RA-Dec positions of each star, with the top section showing the correlation between the true and recovered values in units of arcseconds. The bottom section shows the individual residuals for each. Clearly these parameters have been well recovered through the optimisation. Right: the recovery of the flux parameters in units of photons. The top section shows the correlation between the true and recovered values, and the bottom showing the resulting residuals. The error bars are the $1\text{-}\sigma$ deviations calculated from the covariance Hessian matrix marginalised over positions, fluxes and aberrations.




Figure 2.4: Recovery of the optical aberrations after optimisation. Left: the amplitudes of the Zernike polynomial coefficients, with the top section showing the correlation plot of the true and recovered values, and the bottom showing the individual residuals. Middle and right: the true total Optical Path Difference (OPD) and the residuals generated by these values respectively. The error bars are the $1\text{-}\sigma$ deviations calculated from the covariance Hessian matrix marginalised over positions, fluxes and aberrations.

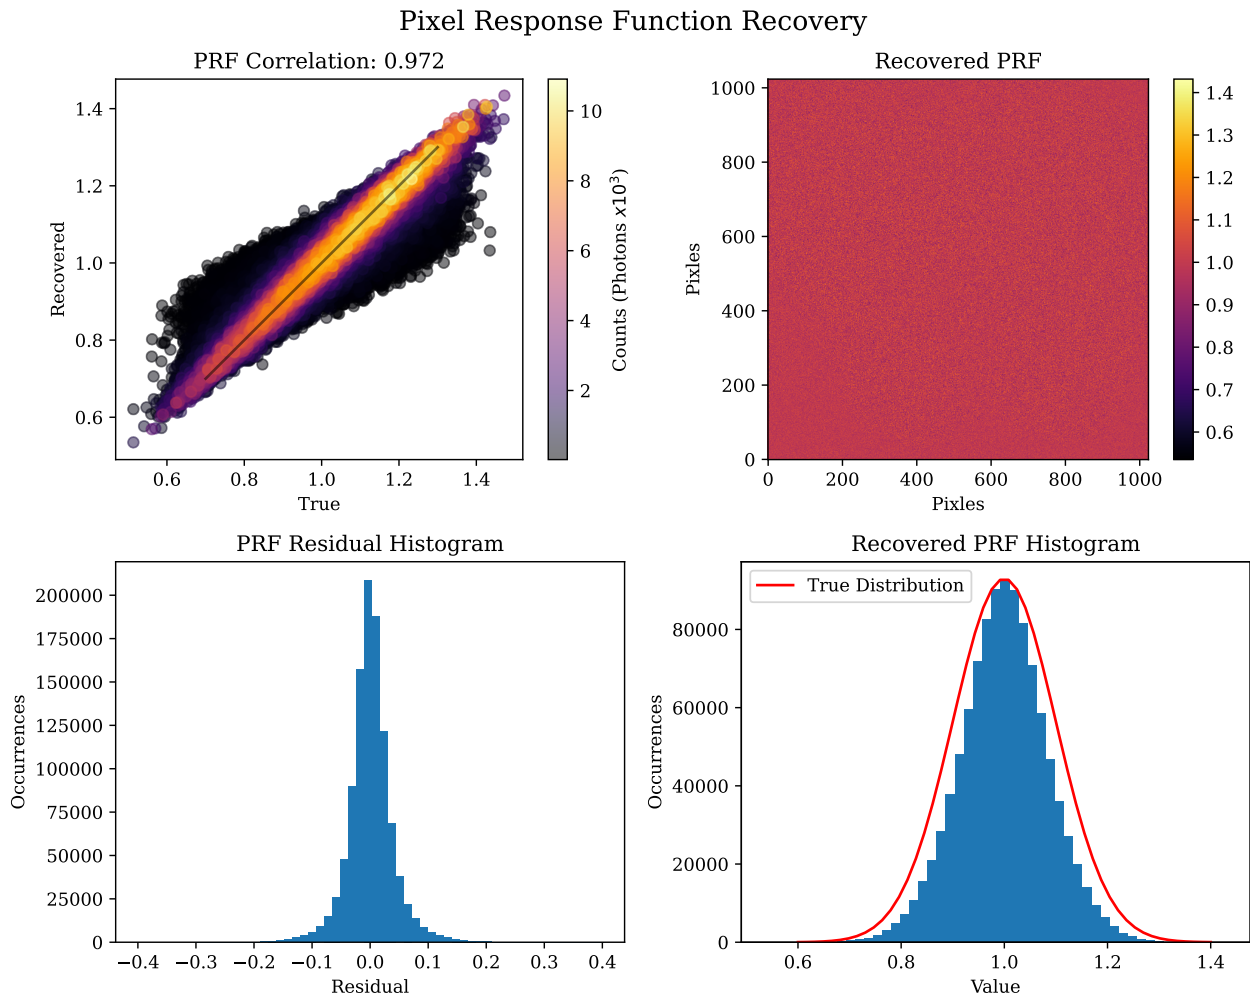


Figure 2.5: PRF values after optimisation. Left: Correlation between the pixel-level true and recovered values. Each point is color-coded with the total flux incident to each pixel, ie its SNR. The pixels with a greater signal are recovered better as expected. Right: Histogram of these residuals, with the majority of values being very well recovered and small symmetric about a residual of zero.

we calculate the errors marginalising *only* over stellar positions, fluxes and optical aberrations, with pixel sensitivities held constant. This provides us with an accurate calculation of the true errors when photon noise is dominant, but it will remain an underestimate when the PRF is the dominant noise source. We chose to estimate our errors in this manner as opposed to the conventional photon-noise scaling relations in order to better capture covariance between parameters. In Section 2.6 we provide a comprehensive analysis of our results with respect to these photon-noise scaling relations.

The $1\text{-}\sigma$ error bars are generated by autodiff, calculating the covariance matrix under the Laplace approximation as the negative matrix inverse of the Hessian of the likelihood function evaluated at the maximum likelihood estimate. These error bars are clearly an underestimate: this is understandable as they are marginalised over positions, fluxes and aberration, but not the PRF residuals which are the dominant noise source. The million-dimensional Hessian required to estimate these uncertainties is at present intractable even in JAX. These parameters are recovered to the expected distribution when performed without modelling the PRF.

The optical propagation algorithm employs two-sided matrix Fourier transforms (Soummer et al. 2007), which can be greatly improved if deployed onto GPU hardware. In our a test case, we initialise from a very naive state far from the truth, requiring a larger number of optimisation epochs than would be required in a real-world example. Furthermore as a calibration algorithm across a whole detector the complexity of the model is greatly increased, applying this to real data would likely be focused on small detector regions resulting in less computationally intensive models.

The use of a forwards model clearly requires a good match between the true and model PSF. The benefit of using differentiable optical models however is that it provides the perfect platform upon which to directly learn this model from the data, with the ability to scale the complexity as required for each problem. Similarly since we are working with direct image data as opposed to reduced data-products, mis-calibrations of the system are simple to diagnose through direct examination of the residuals of the data and model. Figure 2.2 demonstrates that in this example we can be sure of a good fit because the final residual contains no visually identifiable structure.

2.6 Noise & Performance Analysis

This section examines how well this method performs across a range of both PRF and flux values, which are our two dominant noise sources. Recovering parameters in a Poisson-noise-dominated regime for simple optical systems should have the following relationships:

$$\sigma_{\text{flux}} = 1/\sqrt{N_{\text{phot}}} \quad (\text{photons}), \quad (2.1)$$

$$\sigma_{\text{position}} = \frac{1}{\pi} \sqrt{2/N_{\text{phot}}} \frac{\lambda}{D} \quad (\text{radians}), \quad (2.2)$$

$$\sigma_{\text{zernike}} = 1/\sqrt{N_{\text{phot}}} \quad (\text{RMS radians}). \quad (2.3)$$

Comparing the errors in the recovered parameters across a range of fluxes provides a good way to examine when we are in the photon noise dominated regime. After optimisation the errors are calculated for each parameter and compared to these expected values

$$\sigma_{\text{relative}} = \sum_{i=1}^N \frac{1}{N} \frac{\sigma_i \text{ recovered}}{\sigma_i \text{ expected}}. \quad (2.4)$$

The top-left panel of Figure 2.6 shows the weighted Pearson correlation coefficient (Costa 2011) as a function of both flux and PRF. The transition where the PRF is recoverable denotes two regimes: one where photon noise dominates and another where PRF noise dominates. We have distinct results in these different regimes. Astrophysical and optical parameters are recovered to the photon noise expectation, while there is poor recovery of the PRF when photon noise is the dominant source and vice versa when the PRF is the dominant noise source. The top-right and bottom panels of Figure 2.6 show how σ_{relative} changes as a function of both flux and PRF. These σ_{expected} values are expected to be underestimates of the true values because they are not marginalised over all of the fitted parameters in the forwards model. By examining Figure 2.6, we can see that there is a transition between two regimes: photon noise dominated and PRF noise dominated. In the photon noise dominated regime, we can see that the parameters are recovered to the level expected by these relationships and diverge from these once the PRF becomes the primary noise source.

It is important to note in this section that there is a degree of scatter in the recovered results. Due to the gradient descent optimisation strategy there is a degree of sensitivity to the optimisation hyper-parameters, such as learning rates which must be tuned based on the flux, momentum, and distance to the noise floor. These parameters are simple enough to tune for individual data sets with similar scales, but it is non-trivial to automate this process for data with a vast range of scales like those being analysed here. The hyper-parameters chosen were tuned on the $1e8$ flux data set and roughly scaled to work with the other data sets, which is why we observe the best recovery for that set of parameters. These figures are designed to demonstrate the overall trends in the different noise regimes, rather than to be a precise analysis of the fundamental limits.




Figure 2.6: Recovered parameter error relative to photon noise expectation. Left: Mean relative per-mode Zernike error. Middle: Mean relative flux error. Right: Mean relative positional error.

2.7 Conclusions & Future Work

In this first paper of a planned series, we have demonstrated photometry, phase retrieval, and detector calibration using autodiff to make tractable gradient descent in very high dimensions. Our software ∂ LUX integrates with a rich and user-friendly ecosystem of JAX packages for optimisation and sampling, which point to a number of natural extensions.

The most obvious is to produce calibrated photometric time series, making use of correlations in time rather than the ‘bag of images’ approach taken in this work. Instead of independently optimising each frame, they could be treated hierarchically, with regularisation imposed on each star’s light curve. Gaussian process models are widely used in astronomy, and JAX packages such as `tinygp` (Aigrain and Foreman-Mackey 2022) and `GPJax` (Pinder and Dodd 2022) permit differentiable optimisation and sampling that integrate efficiently with ∂ LUX.

Photometry of crowded fields is often of poor quality even from space telescopes, and there has been considerable investment in algorithms for de-blending light curves in the K2 galactic plane campaigns (Zhu et al. 2017) and in TESS with its large pixels (Nardiello et al. 2019; Hedges et al. 2021; Higgins and Bell 2022). High dimensional optimisation with autodiff in principle renders de-blending tractable even close to the diffraction limit with unknown PSFs. Prior knowledge of a structured PSF, such as with the diffractive pupil architecture of Toliman (Guyon et al. 2012; Guyon et al. 2013), may further improve this de-blending and retrieval.

High-precision photometric measurements often need to also consider variations in the intra-pixel sensitivity. This work only examined the inter-pixel sensitivity for simplicity, however the PSFs examined here were sampled at 10x the Nyquist limit, well beyond the typical 1-1.5x Nyquist used in practice. The ability to recover sensitivity variations over these sub-fringe scales show that this method can be directly applied to recovery of both inter and intra-pixel sensitivity fluctuations.

High contrast imaging with the recently-launched JWST (Gardner et al. 2006) and the upcoming Roman Space Telescope coronagraph (Zellem et al. 2022) will in each case demand a new generation of data-driven PSF and detector calibrations beyond those used fruitfully on ground-based data (Cantalloube et al. 2021). While we have not attempted this in the present work, ∂ LUX can flexibly model aberrations in the multiple intermediate planes of these coronagraphs, and in future work we intend to assess the practical and fundamental limits of multi-plane phase retrieval and calibration to improve coronagraphic modelling. Uncertainties in the flat field are also the leading contributor to noise in the exoplanet transmission spectroscopy with JWST (Rustamkulov et al. 2022), and a promising future direction with ∂ LUX will be to build models not just of imaging data, but of spectrophotometry.

In Paper II in this series, we will tackle the isomorphic problem of hardware design. While we have tackled phase design in previous work (Wong et al. 2021), autodiff permits calculation and optimisation of the Fisher information directly (Coe 2009), which means we can optimise towards the fundamental limits of parameter sensitivity. In subsequent papers, we intend to explore parametric and nonparametric deconvolution of images at high angular resolution and contrast.

2.8 Acknowledgements

We thank Laurent Pueyo and Tim White for their helpful discussions.

We acknowledge and pay respect to the Gadigal people of the Eora Nation, upon whose unceded, sovereign, ancestral lands the University of Sydney is built; and the traditional owners of the land on which the University of Queensland is situated. We pay respects to their Ancestors and descendants, who continue cultural and spiritual connections to Country.

This research made use of NUMPY (Harris et al. [2020](#)); Matplotlib (Hunter [2007](#)); JAX (Bradbury et al. [2018](#)); equinox (Kidger and Garcia [2021a](#)); and OPTAX (Hessel et al. [2020](#)).

Chapter 3

Optical Design Maximising Fisher Information

The design of high-precision astronomical hardware operating at the diffraction limit requires the optimisation of physically accurate optical simulations of the instrument with respect to desired figures of merit, such as photometric or astrometric precision. System design entails many parameters, some of which introduce strong nonlinear dependencies between instrument properties and the fidelity of science observables. This chapter explores how differentiable optical models built in ∂ LUX can leverage numerically stable second- and higher-order derivatives to design hardware that maximises the information content of scientific parameters, marginalised over nuisance instrumental parameters. By directly calculating the Fisher information matrix of an observation — including optical and detector metrology — we show how statistically principled, Bayesian experimental design targeting abstract figures of merit, such as precision on science observables, becomes tractable through complex forward models. We validate the method by comparing parameter forecasts from Fisher matrix calculations against analytic theory and inference algorithms. Finally, we extend this to the design of a real-world instrument for exoplanet detection, optimising a binary phase plate to maximise astrometric precision while marginalising over instrumental nuisance parameters.

3.1 Statement of Contribution

This chapter faithfully reproduces material from Desdoigts et al. (2025), which was submitted to the Open Journal of Astrophysics on the 24th of June 2025, with minor adaptation to fit into the thesis style. I am the first author, with Benjamin Pope, Michael Gully-Santiago, and Peter Tuthill as co-authors. The original research problem has been a long-term project with Peter for the design of the TOLIMAN (Telescope for Orbital Locus Interferometric Monitoring of our Astronomical Neighbourhood; Tuthill et al. 2018) space telescope. Ben and Peter both provided guidance over the methodology developed over the project. All of the code, models, simulations, and inference for the paper was done by myself. I wrote the paper, with co-authors contributing feedback on the scientific results, reviewing the manuscript, and assisting with minor improvements to the code and interpretation. All major analytical and computational work was performed by me.

3.2 Introduction

Advances in contemporary observational astronomy are driven by progress in instrumentation and hardware, enabling measurements with improved sensitivity and lower noise. This is particularly true for exoplanetary science where signals can be many orders of magnitude smaller than the noise. When designing instrumentation for this goal, it is essential to accurately model the pattern, or Point Spread Function (PSF), governing the spread of starlight on the detector to enable the engineering of system hardware to optimise science deliverables. For example, apodizing phase plate coronagraphs (Guyon et al. 2006) suppress light from on-axis sources by modifying the wavefront in multiple conjugate planes, creating spatially-varying PSFs. This transformation allows faint nearby companions to be detected against the glare of their bright host star. This can be achieved with an apodizing phase plate (Codona et al. 2006), a phase mask placed in a telescope pupil plane that reshapes the PSF, the design of which poses non-linear optimisation challenges (Por 2017). Astrometric exoplanetary detection (Sozzetti 2005) similarly requires careful consideration of instrumental design such that PSFs are stable and can be used to characterise the time-varying instrumental imperfections, achieving measurement precisions required for tiny (micro-arcsecond) signals. This can be achieved with widefield space-based surveys such as *Gaia*, or with smaller telescopes that employ a diffractive pupil to engineer PSFs with favourable properties (Guyon et al. 2012; Tuthill et al. 2018). In either case and throughout observational astronomy we find escalating demands for stability and precision in data-driven calibration of diffraction effects in the PSF.

Advances in software and algorithms are therefore essential for both data analysis and to facilitate hardware design. A range of open-source physical optics simulation codes (Perrin et al. 2012; Dube 2019; Por et al. 2018) have been developed to model imaging systems end to end, providing a framework to optimise design and/or data analysis schemes by way of grid-based or Markov Chain Monte Carlo (Metropolis et al. 1953) methods. For models with many parameters (the norm for realistic situations with, for example, many modes of phase aberration, multiple sources in the field of view and/or pixel-level imperfections in the detector) then reliable sampling or optimisation can become computationally intractable (Huijser et al. 2022) unless it is also possible to evaluate not just the value but also the *gradient* of the objective function.

Algorithms to implement automatic differentiation or ‘autodiff’ (Margossian 2018) comprise the foundational technology used in artificial intelligence and machine learning (LeCun et al. 2015), and are now rapidly becoming a critical enabling technology in software for the physical sciences. Autodiff makes it possible to evaluate the partial derivatives of a computer program’s floating-point outputs with respect to floating-point arguments. Autodiff computes exact derivatives using successive applications of the chain rule, with computational cost scaling similarly to that of the original model evaluation, rather than increasing with the number of parameters. Considerable industry and academic effort has gone into developing performant and user-friendly numerical software libraries with autodiff capability: PYTORCH (Paszke et al. 2019), TENSORFLOW (Martín Abadi et al. 2015), Julia (Bezanson et al. 2012), or our preferred library used in this work, JAX (Bradbury et al. 2018). Importantly, autodiff removes the need

to manually derive or symbolically express gradients or Hessians — they emerge directly from the differentiable structure of the implemented forward model. This is a central departure from traditional inference workflows, and it underpins our ability to evaluate and optimise realistic optical systems without requiring closed-form expressions for likelihood derivatives.

Differentiable forward models in physics permit optimisation and inference with very many parameters, and the inclusion of flexible nonparametric models (Lavin et al. 2022) jointly with deterministic physics. Autodiff has been used recently in optical and imaging science for phase retrieval and PSF modelling (Jurling and Fienup 2014; Wong et al. 2021; Liaudat et al. 2021; Liaudat et al. 2023a). In previous work (Desdoigts et al. 2023) we introduced ∂LUX , a differentiable physical optics framework designed to address a variety of high dimensional optical calibration tasks. Using the accelerated numerical computation methods and gradient calculations enabled by its JAX back-end, efficient solutions to complex optical problems such as end-to-end phase retrieval and pixel-level calibration from simulated imaging data were demonstrated.

In this manuscript we build on this previous work and focus on scenarios where the posterior distribution over model parameters is well described by a multivariate normal near its peak, known as the Laplace approximation (Kass et al. 1991; MacKay 2002). We show how numerically stable higher-order gradients permit the calculation of covariance and Fisher matrices (Bhandari et al. 2021; Coulton and Wandelt 2023) — a key object in statistical design that quantifies how precisely model parameters can be inferred from observations — under the Laplace approximation. By constructing differentiable optical models we can differentiate a likelihood function with respect to any astrophysical or instrumental model to compute parameter covariance matrices. This means that neither the likelihood function nor its derivatives need to be written in closed form. Instead, they are evaluated programmatically via autodiff directly through the forward model, a key methodological shift enabled by this approach. This enables *Fisher forecasting*: a statistical technique used to estimate the precision of parameter estimates from a future experiment, and enables computation of the covariance matrix. The covariance matrix obtained through the Fisher matrix gives us the Cramér-Rao Lower Bound (CRLB; Radhakrishna Rao 1945; Cramé, Haraldr 1947) — the information-theoretical limit on parameter constraints achievable with an experiment. This statistically principled approach is common for the planning of new instruments and surveys; for example in cosmology the forecasting of how well observations of the cosmic microwave background (Liu and Bunn 2016) or spectroscopic surveys (D et al. 2023) will constrain bulk cosmological parameters, under varying choice of settings in a preliminary instrument design, and marginalised over nuisance parameters.

The ability to calculate these Fisher matrices using autodiff then enables *Bayesian experimental design* (Fedorov 1972; Chaloner and Verdinelli 1995; Ryan et al. 2016) via gradient descent. This might involve optimising the uncertainty of a particular parameter, or a norm such as the determinant or trace of the CRLB. In Section 3.3 of this Paper we briefly describe the underlying theory, and in Section 3.4 we validate calculations made with ∂LUX against analytic theory. By way of illustration with a concrete example, Section 3.5 applies this framework to optimise the design of a phase mask for the astrometric mission TOLIMAN (Tuthill et al. 2018),

a small telescope aiming to detect planets around α Cen AB by measuring micro-arcsecond perturbations in the relative separation of the binary.

3.3 Fisher Information and Bayesian Experimental Design

The key feature of autodiff exploited in this paper is the ability to efficiently calculate a multivariate normal approximation to a probability distribution. Consider an imaging system and sources, parametrised by $\theta \equiv (\theta^1, \dots, \theta^N)$, that generate data d . In solving an inverse problem, we want to infer θ given d and any prior information. We can use Bayes' rule to update our prior knowledge of θ conditioned on our prior knowledge I , given d

$$\underbrace{p(\theta|d, I)}_{\text{posterior}} = \frac{\underbrace{p(d|\theta, I)}_{\text{likelihood}} \cdot \underbrace{p(\theta|I)}_{\text{prior}}}{\underbrace{p(d|I)}_{\text{evidence}}}. \quad (3.1)$$

It is convenient computationally to express this in logarithmic units

$$\log p(\theta|d, I) = \underbrace{\mathcal{L}(d|\theta, I)}_{\text{log-likelihood}} + \underbrace{\Pi(\theta|I)}_{\text{log-prior}} - \underbrace{\log Z}_{\text{log-evidence}}. \quad (3.2)$$

At the maximum likelihood, the *score* — the gradient of the log-likelihood with respect to the parameters — vanishes, and $\partial_i \mathcal{L}(\theta_0) = 0 \forall_i$ where $\partial_i \equiv \partial/\partial\theta^i$ is the maximum likelihood estimate of these parameters. We then motivate consideration of a Taylor expansion of \mathcal{L} about this point, dropping the prior Π^1 , to estimate the distribution near the maximum. Going up to second-order is equivalent to approximating \mathcal{L} as a multivariate normal — ie Laplace's Method

$$\mathcal{L}(\theta^i) \approx \underbrace{\mathcal{L}(\theta_0^i)}_{=0} + (\theta - \theta_0)^i \underbrace{\frac{\partial \mathcal{L}}{\partial \theta^i}}_{\theta=\theta_0} + (\theta - \theta_0)^i (\theta - \theta_0)^j \underbrace{\frac{\partial^2 \mathcal{L}}{\partial \theta^i \partial \theta^j}}_{\equiv F_{ij}} \bigg|_{\theta=\theta_0} + \underbrace{\dots}_{\text{higher-order terms}}$$

and we identify the negative Hessian of \mathcal{L} as the Fisher Information Matrix (FIM; Fisher 1925)

$$\mathbf{F}(\boldsymbol{\theta}) \equiv F_{ij}(\boldsymbol{\theta}) = -\frac{\partial^2 \mathcal{L}}{\partial \theta^i \partial \theta^j}. \quad (3.3)$$

This is formally defined as the variance of the score, and for appropriately regular functions is given by this Hessian. This second-order approximation treats the log-likelihood as a quadratic function near the maximum, and is often used in practice when full posterior sampling is computationally prohibitive. The parameter covariance matrix \mathbf{C} is calculated as the inverse of the FIM

¹We can do so without loss of generality; to include the prior, we can replace \mathcal{L} throughout this calculation with $\mathcal{L} + \Pi$ instead.

$$\mathbf{C} = \mathbf{F}^{-1} \quad (3.4)$$

that fully describes the best-fit multivariate normal; an estimator of the behaviour of a likelihood distribution around its peak. In classical applications, this would require symbolic or manual calculation of second derivatives of the log-likelihood. However, in our framework, these quantities are obtained directly via automatic differentiation of the implemented model, without the need for algebraic derivation.

The FIM has some convenient properties (Coe 2009), firstly the Fisher matrices for two independent experiments add ($\mathbf{F}_{1,2} = \mathbf{F}_1 + \mathbf{F}_2$) allowing for the FIM to be calculated independently for any observations that add linearly, such as dithered images. Secondly, since parameter marginalisation happens through the matrix inversion of \mathbf{F} , row and column i can be deleted from \mathbf{F} in order to remove its contribution to the resulting covariance matrix. Importantly, the covariance matrix of a model with parameters $\boldsymbol{\theta}$ calculated using the observation produced by that model given parameters $\boldsymbol{\theta}$ is the CRLB: the lower bound on the variance of an unbiased frequentist estimator of $\boldsymbol{\theta}$. It is not possible to recover parameters better than the CRLB, and so it is useful for forecasting the sensitivity of an experiment to the parameters of a model under consideration, irrespective of how data analysis will be carried out.

In the remainder of this paper, we will assume that the multivariate normal approximation to our likelihood holds, and that the Fisher and covariance matrices calculated from this Hessian are accurate to within a reasonable tolerance. We also assume that our measurements are photon-noise dominated and therefore a Poissonian likelihood is used exclusively throughout the remainder of the work. Furthermore, all calculations of, and references to, the covariance matrix are of the parameter covariance matrix, calculated against simulated data without noise realisations, and therefore completely describe the CRLB under the Laplace approximation (which is assumed valid throughout).

3.4 Comparison with Theory

In especially simple cases and under a number of simplifying assumptions, the Fisher forecast of parameters of interest can be obtained by analytic theory. In this Section we compare the estimated CRLB uncertainties via three different methods: an analytic derivation, through the Fisher matrix calculated via a differentiable model (the methodology explored in this paper), and direct sampling of the posterior via an Hamiltonian Monte Carlo (HMC) algorithm. The posterior samples generated via the HMC serve as a numerical benchmark we can compare the Fisher and analytic methods against, given that the HMC directly samples the full posterior.

We start with an illustrative toy problem: obtaining a position measurement from the image plane of a simple telescope. Given a circular pupil support and a monochromatic point source, the PSF is the well known Airy disk. The analytic CRLB (i.e. best achievable precision, independent of the analysis method used) on the localisation of a point source through this system as a function of photon count is given by the expression




Figure 3.1: Comparison of the posterior positional uncertainty, σ_r , via three estimators: HMC posterior samples (used as a numerical benchmark), a Fisher-information estimate obtained via autodiff, and a closed form analytic derivation. Top: absolute estimated σ_r versus photon count for the three estimators, all showing strong agreement. Bottom: ratio of σ_r to the HMC benchmark (unity = perfect agreement), revealing a systematic difference between the analytic method and the HMC benchmark of $\sim 6\%$, indicative of an consistent overestimate of system performance. The Fisher estimate is in strong agreement with the HMC benchmark at all values.

$$\sigma_r = \frac{1}{\pi} \sqrt{\frac{2}{N_{\text{phot}}} \frac{\lambda}{D}}, \quad (3.5)$$

where σ_r is the radial uncertainty in the positional measurement in radians, N_{phot} is the number of photons, λ is the wavelength of light and D is the diameter of the aperture (Guyon et al. 2012). Figure 3.1 now compares σ_r as calculated from this analytic equation and the σ_r values calculated using the Fisher methods described in Section 3.3, both registered against results of the posterior samples generated via the HMC algorithm, used as the benchmark. We assumed that the signal was photon noise limited, and therefore chose a per-pixel Poisson distribution as the likelihood function for both the Fisher and HMC methods, which is the same as that used to derive Equation 3.5.

The top panel of Figure 3.1 shows consistency between the three methods across the full range of the fluxes, however the bottom panel that compares the Fisher and analytic method to the HMC numerical benchmark shows a consistent overestimate of performance with the analytic approach. This highlights the limitations of the idealised assumptions typically required to find closed-form solutions to these equations. In this case, the finite pixel and image plane size form two such complexities not addressed by analytic derivation. Furthermore, any real-

world imaging system is likely to entail significantly more complexity than this toy model. Instruments operate over a non-zero passband and feature chromatic performance, surfaces are imperfect, and the clear aperture can be cluttered with blockages from the secondary mirror and spiders. All these, and many more non-ideal elements will directly influence the PSFs and in general can only be modelled numerically. In these cases, differentiable forward models provide an effective way to accurately calculate the CRLB of a given observation without resort to the simplifying assumptions required for analytic expressions.

3.5 Optical Design Case Study: the Toliman diffractive pupil

These methods allow the evaluation of the covariance matrix at fixed parameter values and the direct optimisation of its components using higher-order derivatives. This enables gradient-based experimental design, targeting properties of the covariance matrix such as entropy or marginalised variances, using a differentiable forward model. This opens up the possibility of *Bayesian experimental design*, allowing for gradient descent directly on such properties and functions of the covariance matrix, or as we show in this example, on individual components of the covariance matrix marginalised over the rest of the model. In keeping with the central theme of this paper, accomplishing this requires only a forward model of the system in a differentiable framework.

The illustrative optical problem chosen is the design of the diffractive pupil required for the TOLIMAN mission (Tuthill et al. 2018), a ~ 100 mm class aperture space telescope which aims to measure the micro-arcsecond scale angular perturbation induced on the host star by the orbital gravitational reflex motion from an (unseen) low-mass planet. Such a formidable challenge requires the engineering of a PSF that maximally encodes both instrument metrology and scientific signals of interest simultaneously on the sensor. The solution to these challenges incorporates a binary-valued diffractive pupil and a spectrometer, all integrated into the PSF of the telescope by engraving phase patterns onto the entrance pupil. Specifically crossed sinusoidal gratings are used to produce a spectrometer (not discussed further in the present manuscript), which are overwritten onto the primary diffractive pupil pattern which enables the optical and detector calibration while also allowing astrometric science measurement. Here we focus solely on this latter pattern which consists of regions of zero/ π phase which are separated by sharp transitional boundaries. The requirement for the binary pattern arises as a primary design constraint of the TOLIMAN diffractive pupil (Tuthill et al. 2018). The optical signal encoding and associated recovery methods for system state metrology for the TOLIMAN mission present complex requirements, providing an ideal robust test case for evaluating the practical capabilities and limitations of the Fisher-information-based design methods discussed in this paper. While full details of the mission-specific constraints are not essential here, this context highlights the degree of design complexity and abstracted figures of merit that can be effectively managed through differentiable optical modelling.

Here, we take the design of the TOLIMAN diffractive pupil as a complex yet achievable




Figure 3.2: Geometric parameters and Toliman astrometric problem. Left: telescope aperture geometry with random phase errors generated on the aperture. Optical aberrations Z_i are represented with Zernike polynomials in units of nm. Right: parametrisation of the binary star geometry; mean position (x, y) (arcseconds), relative separation r (arcseconds), position angle θ (degrees).

demonstration of analysis by our autodiff frameworks, and further posit that arriving at any comparably performing outcome by way of analytic or heuristic approaches would present a formidable task. In framing our optimisation problem, the objective can be stated find a binary-valued diffractive pupil that maximally constrains the relative angular separation of a binary star r after marginalising over the remaining astrophysical parameters, namely the mean position (x, y) in arcseconds on the sensor, the position angle θ in degrees, the total flux F in photons, a dimensionless contrast ϕ as the ratio of the binary component fluxes, the mean wavelength λ in nm, as well as the remaining optical parameters — the pixel scale γ in arcsec/pixel and optical aberrations Z_i modelled as a sum of normalised Zernike polynomials with coefficients in nm. Figure 3.2 shows the basic geometry of both the telescope aperture and the astrophysical system, without any diffractive pupil in the system.

TOLIMAN’s requirement for binary phases $\in \{0, \pi\}$ presents a challenge for gradient-based optimisation methods as the distribution is inherently discontinuous. To map differentially between a continuous set of basis vectors and binary valued mask we employ the CLIMB algorithm, first introduced for the Toliman mask design problem in Wong et al. (2021). By using a set of $3 \times$ oversampled basis vectors, a binary mask can be constructed with a single pixel boundary between the two regions with continuous values that enable gradients to propagate smoothly. In brief, any 3×3 oversampled pixel with all values above or below zero can be assigned to one or zero respectively. Any remaining pixels form the boundary region, where a plane can be fit via least-squares to approximate the fraction of the pixel above zero, yielding a single pixel boundary of continuous values between zero and one. More details and full psuedo-code can be found in Wong et al. (2021).

The basis vectors used can, in principle, be chosen freely. However, it is advantageous to tailor their selection to the physical symmetries of the problem. In this work, we employ basis




Figure 3.3: Top two panels: selection of the the ortho-normalised continues basis vectors built from log-harmonics and radial sine and cosine functions. Colour scales are arbitrary. Bottom panel: Binary masks generated by random coefficients. The dark and light regions correspond to zero and π respectively.

vectors with three-fold rotational symmetry, which naturally enhances sensitivity to even-mode Zernike aberrations that suffer sign degeneracies in conventional optical systems. By improving the sensitivity to these modes we enable the diffractive pupil to both constrain the astrophysical geometry and act as a wavefront sensor. Ultimately the choice of basis vectors is arbitrary, provided it has the appropriate expressiveness for the problem as the differentiable engine propagates gradients from PSFs back to the basis of choice. In this work we build the basis vectors from a set of log harmonic radial and radial sine and cosine functions orthogonalised using the Gram-Schmidt algorithm (Björck 1994). We then keep the top 100 basis vectors and discard the rest. A selection of these basis vectors are shown in Figure 3.3, along with realisations of the binary masks generated by randomly generate coefficients.

We construct a polychromatic model of the TOLIMAN optical system featuring:

- A clear 125 mm aperture diameter.
- A polychromatic bandpass modelled at three wavelengths: 530, 585, and 640 nm.
- PSF pixel scale of 0.375 arcseconds/pixel, $\sim 1.5\times$ Nyquist.
- The 4th - 10th Zernike modes (ignoring the lowest 3: piston, tip, and tilt).

We then model the expected signals from the binary star α Centauri AB assuming a representative projected separation of 10 arcseconds.

To transform the Fisher information into a concrete optimisation objective, we construct a loss function that minimises the uncertainty of a specific scientific parameter of interest under a full forward model — the binary separation r in this case

$$\text{loss}(\boldsymbol{\eta}) = [\mathbf{C}(\boldsymbol{\psi}_0)]_{rr} = \left[-\nabla_{\boldsymbol{\psi}}^2 \mathcal{L}(f(\boldsymbol{\eta}, \boldsymbol{\psi})) \Big|_{\boldsymbol{\psi}=\boldsymbol{\psi}_0} \right]_{rr}^{-1} = \text{Var}(r \mid \boldsymbol{\psi}_0, \boldsymbol{\eta}), \quad (3.6)$$

where

- $\boldsymbol{\eta}$ are the design parameters, the diffractive pupil basis vector coefficients.
- $\boldsymbol{\psi}$ are the astrophysical and instrumental parameters over which we marginalise (e.g. binary separation r , flux, aberrations).
- $f(\boldsymbol{\eta}, \boldsymbol{\psi})$ is the differentiable forwards model simulating the system.
- \mathcal{L} is the log-likelihood function (Poissonian, consistent with the assumption of photon noise-limited observations).
- $\nabla_{\boldsymbol{\psi}}^2$ is the Hessian operator applied to the $\boldsymbol{\psi}$ parameters of the model.
- $[\cdot]_{rr}$ denotes the component of $\boldsymbol{\psi}$ corresponding to the binary separation.

We note that the likelihood is only a function of the model, as we generate the data at the input at $(\boldsymbol{\eta}, \boldsymbol{\psi})$ in order to ensure our Hessian computation produces the Fisher matrix and its inverse is therefore the CRLB. This loss function computes the covariance matrix of the parameters through the forward model, extracting the variance component corresponding to r , which serves as the optimisation target. Other reductions of the covariance matrix such as its trace or entropy can also be optimised if desired. Gradients of the loss with respect to the design parameters $\boldsymbol{\eta}$, ie the coefficients of the diffractive pupil basis vectors, are computed via automatic differentiation. Optimisation is performed using 50 epochs of the Adam algorithm (Kingma and Ba 2017).

The initial and final pupil-PSF pairs are shown in Figure 3.4, where we can see that the final PSF is concentrated into a small number of brighter peaks near the centre, plus a series of dimmer peaks surrounding it. Figure 3.5 shows the parameter posteriors for both the initial and final pupils, as well as an optical system with a clear pupil producing an Airy-disk like PSF as a benchmark. We see the greatest improvement for the optimisation metric, the binary separation r , as well as the pixel scale γ and mean wavelength λ . Using this Fisher optimisation approach, we improve the CRLB of the binary star separation from ~ 32 mas to ~ 22 mas, approximately a 30% improvement from this single seed. The full corner plot is shown in Figure 3.6. Note that the optical aberrations have been omitted from these plots for visualisation purposes, although they were present for the optimisation and have been marginalised over.

Interestingly, the pre-optimisation model already outperforms the Airy system, implying a good choice of basis vectors for the diffractive pupil. Despite the improved performance in the binary separation metric, the diffractive pupil model is inferior in recovery of both the (x, y) position and position angle θ . Figure 3.7 compares the full performance of these models to each

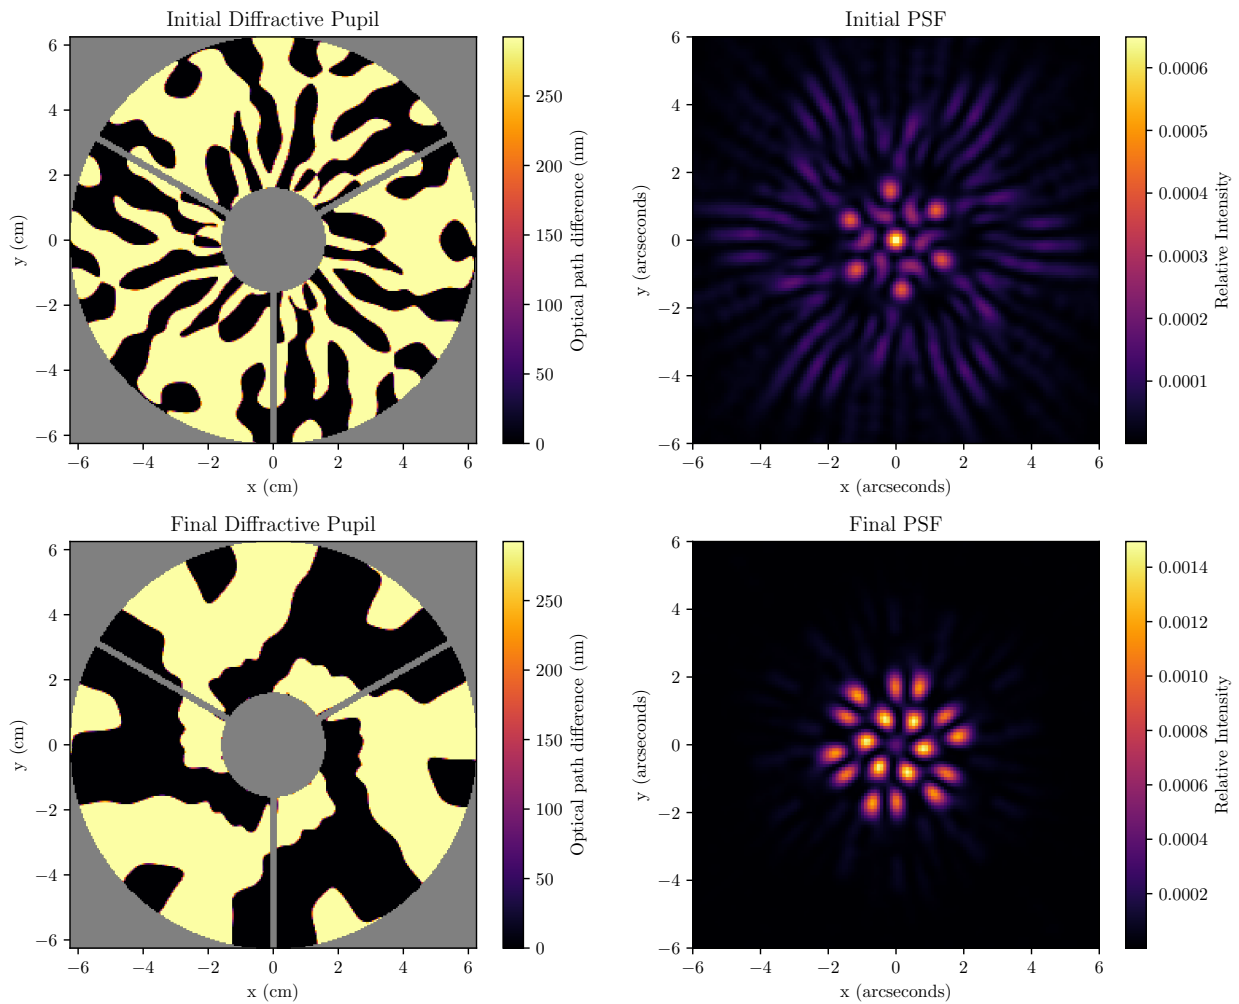


Figure 3.4: Diffractive Pupil - PSF pairs before and after Fisher optimisation. Top panels: Randomly initialised pupil pattern (left) consisting of binary-valued $0/\pi$ phases within the circular support of primary/secondary mirrors and 3 opaque spiders, together with its corresponding PSF (right). Bottom panels: Final pupil after optimisation (left) and its corresponding PSF (right).

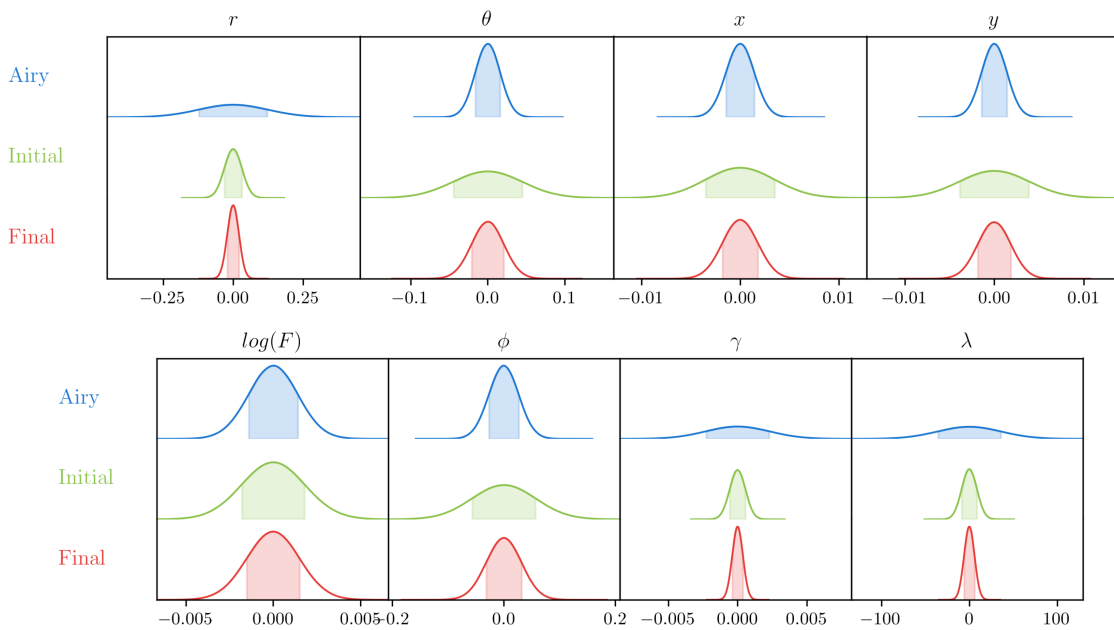


Figure 3.5: CRLB posteriors for the parameters of each model. Note the optical aberration posteriors have been omitted for plotting purposes, but have been marginalised over in the calculation. The ‘Airy’ model is an optical system without a diffractive pupil, ‘Initial’ is the pre-optimisation diffractive pupil and ‘Final’ is the post-optimisation diffractive pupil model. The optimisation process has improved the performance of all displayed parameters, with most gain coming from the separation r , pixel scale γ and the mean wavelength λ . Interestingly, while the optimised model greatly outperforms the ‘Airy’ system in the minimised parameter separation r , pixel scale γ , and wavelength λ (as desired), its performance is worse for all the remaining displayed parameters.

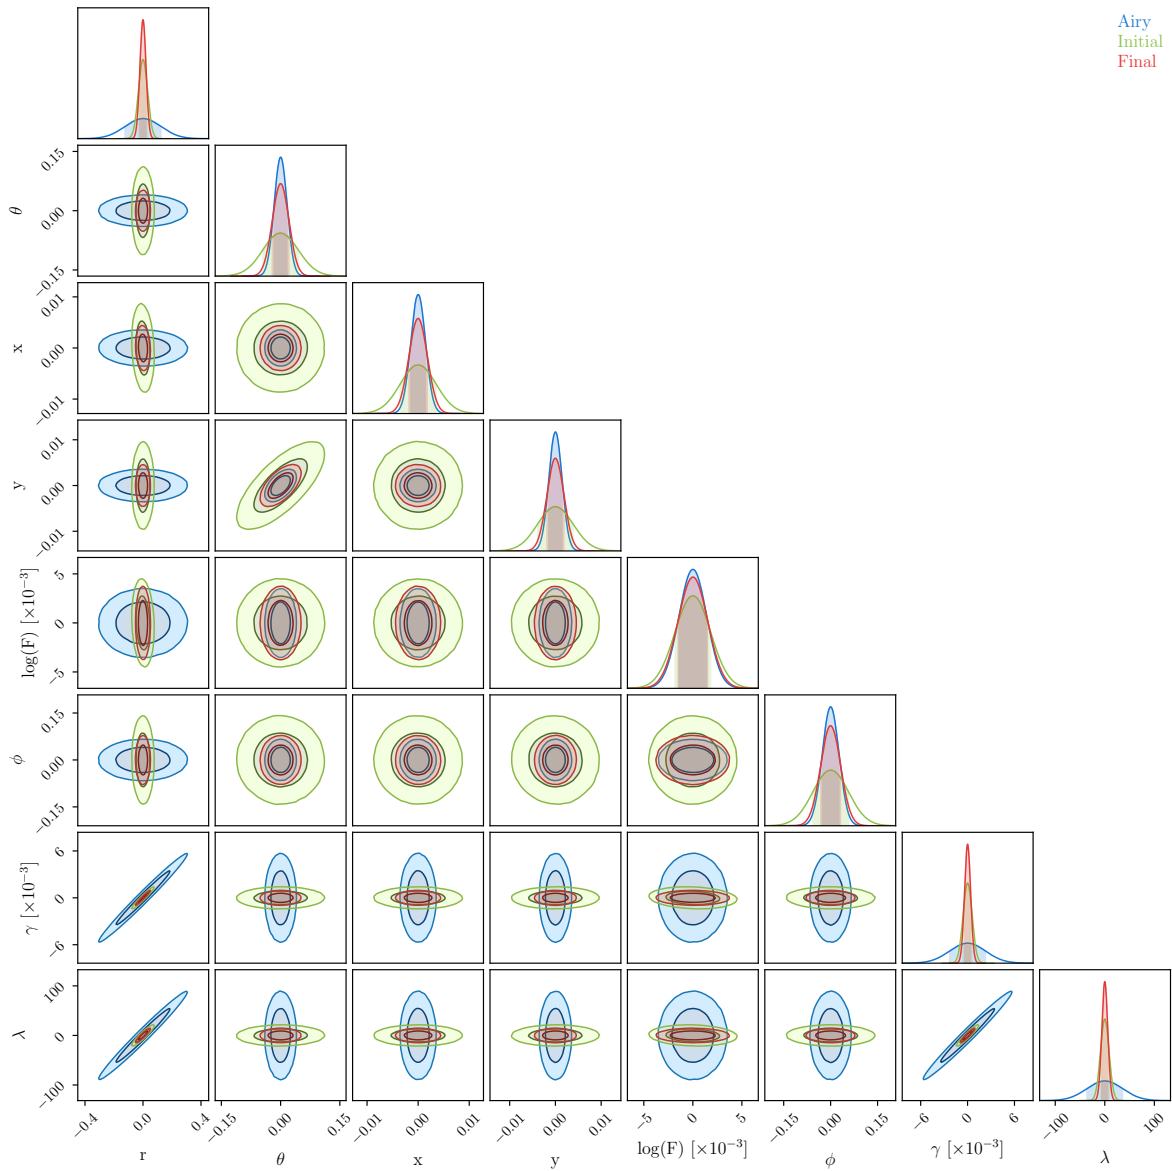


Figure 3.6: Reduced corner plot comparing the model performance of an Airy-like model (blue), the pre-optimisation diffractive pupil (green), the post optimisation diffractive pupil (red). While still marginalised over all parameters in the model, this plot does not show the optical aberration components as the resulting figure becomes too large, and all of the models show the same behaviour for those model parameters.

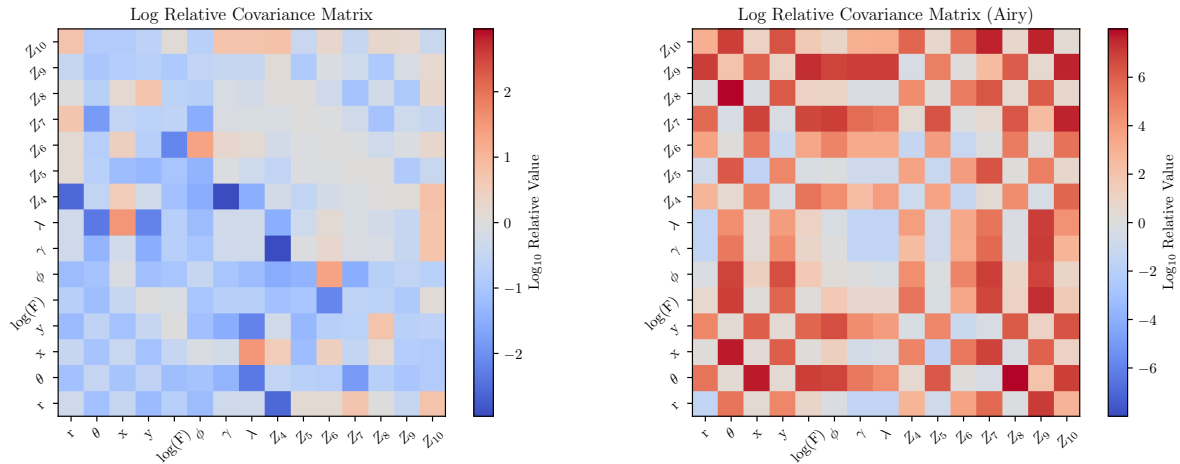


Figure 3.7: Logarithm of absolute element-wise relative covariance matrices, i.e. $\log_{10}(|C_{ij,1}/C_{ij,2}|)$. Negative values are blue and indicate a better constraint of that parameter while red flags the opposite. Left panel: The post-optimisation covariance matrix relative to the pre-optimisation covariance matrix. Right panel: The post-optimisation covariance matrix relative to the Airy-disk covariance matrix.

other by visualising their relative covariance matrices ($\log_{10}(|C_{ij,1}/C_{ij,2}|)$). This Figure reveals some interesting properties of this optimisation process, we can see that while most elements of the covariance matrix show improved performance, some exhibit degraded performance. This is explained through the marginalisation process inherent when inverting Fisher matrices to get covariance matrices — parameters that have little to no covariance with the parameter of interest (the binary separation r) can have their precision decreased in order to gain improved precision over those that are. This demonstrates an optimisation process that is fully coherent of the complex relationships between different parameters within a system. The pupil mask design approach here offers a significant improvement in the astrometric performance of TOLIMAN, and these methods will form the basis of its design and analysis, presently in progress.

Given the high-dimensional nature of this problem and the complexity of the loss space introduced by the mask binarisation process, this is a non-convex problem where only local minima can be found. To ensure that the solutions found are sufficient, we repeat the optimisation process from five randomly initialised pupil patterns. Figure 3.8 shows the resulting pupil designs along with their associated PSFs and relative covariance matrices, expressed as $\log_{10}(|C_{ij,1}/C_{ij,2}|)$. Each of these five random seed improves the CRLB of the binary star separation over the optimisation by approximately 33%, 48%, 35%, 36% and 21%. These results are presented as illustrative examples of the method’s capabilities, without detailed interpretation of the relative performance across different seeds.




Figure 3.8: Summary of diffractive pupil optimisations for five different random seeds. Each row is a different seed, left panel: Final diffractive pupil, middle panel: Resulting PSF, right panel: \log_{10} relative covariance matrix, normalised by the initial pre-optimisation covariance matrix.

3.6 Discussion

The methods to calculate Fisher Information presented in this work provide a simple and straightforward way to explore, analyse and optimise models, however it is important to be cognisant of limitations. Firstly the Laplace approximation assumes that all posteriors are well described by multivariate normals about their peaks. This assumption will not hold in a general sense for systems with complex parameter degeneracies. While an understanding of the local topology of the posterior about the maximum likelihood estimate can be found, this can not function as a *general* replacement for sampling of the posteriors via Monte Carlo methods which serve to understand the global topology of the posterior. In many cases, the posterior distribution is either well-approximated by a multivariate normal, or a global understanding of its structure is unnecessary. In such scenarios, the second-order Laplace approximation provides sufficient insight and this is often the case for well-constrained optical instruments where local sensitivity is the dominant driver of performance. Furthermore, this work only explored the optimisation of the Fisher information about a single realisation of model parameters. In practice for an instrument like Toliman the data will span a varying range of both instrumental and astrophysical states. Accounting for this variation can be done during the optimisation process by sampling the parameters of the covariance matrix over their expected range of values, ensuring solutions robust to all states are found.

3.7 Conclusion

Progress in astronomical science depends not only on advances in hardware and observational platforms, but also on innovations in software that enhance instrument design and data analysis. Exoplanetary science, in particular, places some of the most stringent demands on both instrumentation and inference due to the intrinsically faint, noise-dominated signals that require careful calibration and characterisation. Autodiff has already demonstrated profound utility in improving astrophysical measurements and calibration techniques (Liaudat et al. 2023a). Modern frameworks such as JAX enable a new class of scientific software that builds on autodiff to offer powerful extensions to existing tools. This manuscript aims to showcase the modelling strategies made possible by adopting differentiable software paradigms such as ∂ LUX that treat autodiff not as an afterthought, but as a foundational design principle. By replacing hand-derived analytic equations with programmatic differentiation, these tools allow complex models that incorporate realistic optics and detector effects to be optimised and explored using the same physical principles, but with greater flexibility and computational tractability.

In this manuscript we have explored new avenues for analysis and design of optical systems harnessing the stable calculation of Fisher matrices empowered by autodiff. This method yields results in agreement with analytically derived expressions when identical assumptions are enforced, and furthermore it can be used to explore extensions to a more relaxed set of assumptions. The framework provides new ways to probe the effects of different instrumental architectures on the recovery of science outcomes, and straightforward algorithms to optimise experimental design targeting constraint of specific astrophysical parameters without detailed

pen-and-paper analysis.

Differentiable methods also enable inference to be performed on complex posteriors. In cases where distributions are non-Gaussian due to complex parameter degeneracies, approximate inference methods become essential. Stochastic variational inference (Hoffman et al. 2013) enables parameter degeneracies resulting in non-Gaussian posteriors to be approximated by minimising the divergence between the inferred and true posterior. Extending this, automatic differentiation variational inference (Kucukelbir et al. 2016) utilises the derivatives of a model to explore the posterior parameter space efficiently, and can be applied directly to differentiable optical models in order to address unavoidable parameter degeneracies present in complex data sets.

These methods offer a new framework not only for analysing observational data, but also for informing the design of future instruments employing statistically grounded optimisation criteria described by the Fisher information. As one particularly promising use case, future coronagraphic instrument design based on autodiff may be made more robust to realistic noise processes such as low-order wavefront error (Currie et al. 2018) or the low wind effect (Milli et al. 2018). With these tools, novel architectures can be rapidly explored and optimised with respect to evidence ratios, robustness to optical aberrations, or spectral parameter estimation. The same principles can be applied to other instruments such as spectrographs for radial velocity planet detection, enabling joint end-to-end optimisation of optical systems and observational regimes.

The methods presented were implemented using ∂ LUX, however they are not inherently tied to any specific software package. The central contribution of this work lies in the application of automatic differentiation for optical system design and uncertainty forecasting. We anticipate that similar approaches can be adapted within other autodiff-capable frameworks, and future work may benefit from comparative benchmarking across such platforms.

3.8 Code, Data, and Materials

As part of our commitment to open science, we have released ∂ LUX as an open source package under a BSD three-clause at github.com/LouisDesdoigts/dLux. Furthermore an accompanying Jupyter notebook that produces all results and figures in this paper is publicly hosted at github.com/LouisDesdoigts/FIM_tutorial. The accompanying code repository allows for replication of results and adaptation of the methods to a variety of optical modelling tasks.

Acknowledgements

This research made use of NUMPY (Harris et al. 2020); Matplotlib (Hunter 2007); JAX (Bradbury et al. 2018); NUMPYRO (Phan et al. 2019a); equinox (Kidger and Garcia 2021a); OPTAX (Hessel et al. 2020); and ChainConsumer (Hinton 2016).

Chapter 4

Amigo: a Data-Driven Calibration of the JWST Interferometer

The James Webb Space Telescope (JWST; Gardner et al. 2006) hosts a non-redundant Aperture Masking Interferometer (AMI; Sivaramakrishnan et al. 2012; Soullain et al. 2020) in its Near-Infrared Imager and Slitless Spectrograph (NIRISS; Doyon et al. 2023b) instrument, providing the only dedicated interferometric facility aboard — orders of magnitude more capable than any interferometric experiment previously flown. However, the performance of AMI (and other high resolution approaches such as kernel phase) in recovery of structure at high contrasts has not met design expectations. A major contributing factor has been the presence of uncorrected detector systematics, notably charge migration effects in the H2RG sensor, and insufficiently accurate mask metrology, previously hidden beneath precision limits imposed by the atmosphere. Here we present AMIGO, a data-driven calibration framework and analysis pipeline that forward-models the full JWST AMI system — including its optics, detector physics, and readout electronics — using an end-to-end differentiable architecture implemented in the JAX framework and in particular exploiting the ∂ LUX optical modelling package. AMIGO directly models the generation of up-the-ramp detector reads, using an embedded neural sub-module to capture non-linear charge redistribution effects, enabling the optimal extraction of robust observables, for example kernel amplitudes and phases, while mitigating systematics such as the brighter-fatter effect. We demonstrate AMIGO’s capabilities by recovering the AB Dor AC binary from commissioning data with high-precision astrometry, and detecting both HD 206893 B and the inner substellar companion HD 206893 c: a benchmark requiring contrasts approaching 10 magnitudes at separations of only 100 mas. These results exceed outcomes from all published pipelines, and re-establish AMI as a viable competitor for imaging at high contrast at the diffraction limit.

4.1 Statement of Contribution

This chapter faithfully reproduces material from Desdoigts 2025b (in review), which was submitted to the Publications of the Astronomical Society of Australia on the 29th of June 2025, with minor adaptation to fit into the thesis style. I am the first author, with Benjamin Pope, Max Charles, Peter Tuthill, Dori Blakely, Doug Johnstone, Shrishmoy Ray, Anand Sivaramakrishnan, Kevin Volk, Jens Kammerer, Deepashri Thatte, and Rachel Cooper as co-authors. The original research problem was proposed by Peter. All of the co-authors contributed guidance throughout the project. All of the code, models, simulations, and inference for the paper was done by myself. I wrote the paper, with co-authors contributing feedback on the scientific results, reviewing the manuscript, and assisting with minor improvements to the code and interpretation. Max produced all the flow diagrams used throughout the manuscript, and all major analytical and computational work was performed by me.

4.2 Introduction

Direct imaging of exoplanets and their environments against the glare of the host star makes extreme demands on precision in calibration of optics, electronics, and computational data analysis. Recovering signals requires simultaneous high contrast and angular resolution with performance levels dictated by the planet’s relative faintness and proximity to their host-stars (Follette 2023a), being limited mainly by systematics from optical aberrations and detector electronics. Coronagraphic methods can achieve contrast ratios of 10^{-4} to 10^{-6} , and are particularly favoured for deployment on modern space observatories like JWST (Gardner et al. 2006; Gardner et al. 2023) that avoid speckle noise arising from atmospheric turbulence. Despite proven performance at high contrasts, coronagraphs remain limited by their Inner Working Angle (IWA), with best performance usually found beyond $> 2\lambda/D$ (Guyon et al. 2006). This inability to study the inner structures of extra-solar systems, regions crucial to the understanding of exoplanetary formation (Wagner et al. 2019), leaves a glaring observational gap at high angular resolutions.

Aperture masking interferometers provide imaging capabilities at and beyond the classical diffraction limit $\sim \lambda/D$ (Monnier 2003), filling the observational gap left at high angular resolutions. Non-Redundant Maskings (NRMs) propagate interferometric phase information from each mask hole to a unique location in the Fourier plane, enabling precise calibration and subtraction of instrumental effects. Furthermore, closing triangles and squares in the aperture yield closure-phases (Jennison 1958) and amplitudes (Twiss et al. 1960; Readhead et al. 1980) respectively, observables that are robust to low-order wavefront and amplitude errors. When combined with short exposures, they enable high angular resolution images through the phase corruption of atmospheric turbulence (Baldwin et al. 1986; Robertson et al. 1991). Further advances have produced images in the Near Infrared (NIR) at the diffraction limit (Monnier et al. 1999; Tuthill et al. 2000b), cementing these methods as the only way to simultaneously image at both high contrasts and angular resolutions through wavefront phase errors.

The JWST NIRISS (Doyon et al. 2012; Doyon et al. 2023a) hosts a NRM, providing an AMI observational mode — the first of its kind aboard a space observatory (Sivaramakrishnan et al. 2023; Sivaramakrishnan et al. 2012; Soulain et al. 2020). While the Hubble Space Telescope (HST) hosts an interferometer in each arm of its fine guidance system (Jefferys et al. 1985), AMI on JWST is the first of its kind capable of complex imaging. NRM systems have hitherto been used to mitigate atmospheric turbulence, prompting the question as to their value on a space observatory without a turbulent atmosphere corrupting its wavefronts. However, despite its optical stability, JWST’s segmented aperture demanded precise phasing of the mirrors after launch, a problem not well addressed by in-focus, clear-pupil imagers due to the redundancy of wavefront phase information, but well suited to non-redundant apertures (Cheetham et al. 2012). The NRM configuration allows AMI to act a proven alternate wavefront calibration device that also provides unique scientific capabilities — and was added to NIRISS’ suite of observing modes after JWST’s preliminary design review. AMI provides a complementary role to the various coronagraphic modes and explores a search space within the IWAs of JWST’s coronagraphs, reaching the snow line of nearby exoplanet systems (Ray et al. 2023).

4.2.1 JWST AMI: A unique space-based interferometer

JWSTs AMI mode employs a 7-hole NRM to form a stable interferometric Point Spread Function (PSF) suitable for classic interferometric analysis methods. Its configuration, shown in Figure 4.1 provides approximately even sampling in the Fourier plane, and primarily observes in three medium band filters from $\sim 3\text{-}5\ \mu\text{m}$.

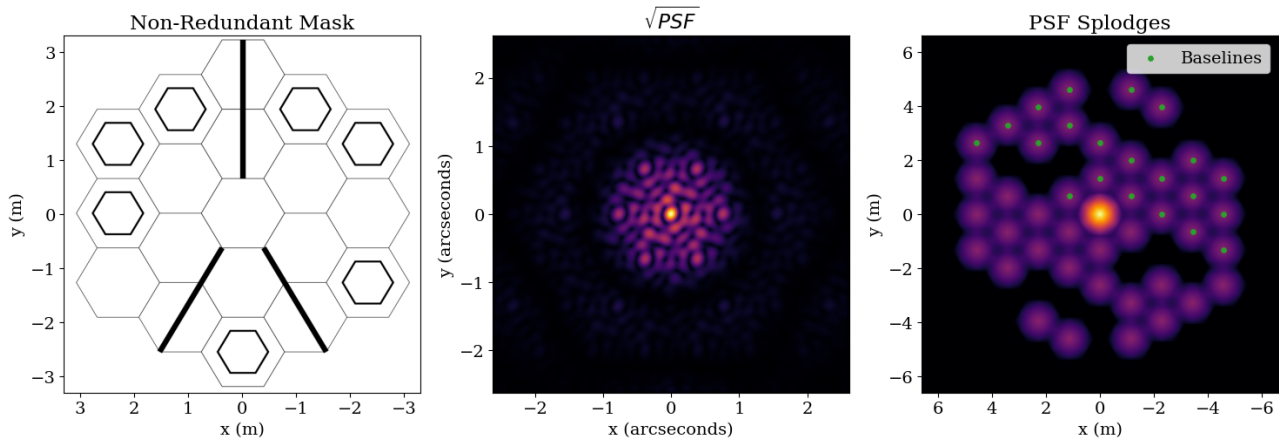


Figure 4.1: Left panel: Schematic diagram of the 7-hole NRM projected over the primary mirror. Middle panel: The resulting PSF (i.e. interferogram) from the non-redundant mask, visualised on a square-root scale to highlight low-power features. Right panel: The power-spectrum of the PSF featuring baseline-specific regions of fringe power known in the literature as *splodges* that can be conveniently found by Fourier transform of the PSF. The 21 discrete non-redundant baselines are indicated by the overlaid green dots.

Interferometric data is analysed in the Fourier plane — hereby referred to as the uv -plane — by examining its complex Fourier coefficients described as its *complex visibilities*. Adjacent observations of calibrator stars enable the subtraction of both the host star and instrumental signals via a division of these visibilities in complex form, ideally revealing the minute signals of the near-stellar environment.

Various pipelines exist to extract the interferometric observables from calibrated JWST images. **AMICAL** (Soulain et al. 2020) and **SAMpy** (Sallum et al. 2022) both perform analysis on the Fourier transform of the calibrated JWST images and have been used in the context of ground-based interferometry to great success (Sallum et al. 2023; Vides et al. 2023; Blakely et al. 2022; Blakely et al. 2025; Lucas et al. 2024). **FOURIEVER** (Kammerer et al. 2023) harnesses similar ideas and extends them for kernel phase (Martinache 2010). Other pipelines such as **ImPlaneIA** (Greenbaum et al. 2015) perform analysis in the image plane, using an analytic forward model of the PSF as a way to better address piston phase errors and pixel-level miscalibrations (Lau et al. 2023; Greenbaum et al. 2019). This work draws much of its intellectual heritage from these earlier pipelines, in particular ImPlaneIA through its use of an analytical forward model for the interferogram, but greatly extended to include the instrument as a whole including the detector and a high-order model of the optical aberrations.

Notwithstanding the rigour and success of these analytical tools, alongside much work from researchers in the field — AMI mode has failed to provide its promised fringe stability and precision, with both found to be up to an order of magnitude worse than expectation in the

worst cases (Sivaramakrishnan et al. 2023; Ray et al. 2025; Sallum et al. 2024b). Despite this degraded performance, careful calibration has been able to recover images of both the circumstellar disk and companion planets around the PDS 70 system (Blakely et al. 2024), albeit at the cost of discarding very large portions of the data. Similar approaches have been used to resolve the dusty environments around WR 137 (Lau et al. 2023), though visible PSF miscalibration remains persistent. Efforts have been made to fix these problems but follow the approach of simply discarding data without addressing the underlying issues (Goudfrooij et al. 2024).

The most significant factor contributing to the degraded performance of AMI arises from the Brighter-Fatter Effect (BFE), or charge migration between pixels (Antilogus et al. 2014; Guyonnet et al. 2015; Rowlands et al. 2018; Argyriou et al. 2023; Goudfrooij et al. 2024). Electrostatic interactions within the substrate push excited photoelectrons into neighbouring pixels. This results in an effective distortion of the measured PSF, typically seen as a broadening of bright sources that fill neighbouring pixels with more charge than dim sources — hence ‘brighter-fatter’. While this effect is troublesome for various imaging modes, it presents a uniquely challenging problem for interferometric analysis dependent on precise inference of the PSF to calibrate observations. To identify why this effect plagues AMI in particular we must examine the JWST data processing pipeline, not the interferometric analysis methods themselves.

These other pipelines adopt *inverse modelling*: a series of transformations of the data to extract summary statistics (a calibrated image, or parameters of that image) that we can fit astrophysical models to. The [JWST pipeline](#) applies fixed pixel-wise linearity corrections to the data, and the AMI pipelines assume images are formed linearly downstream of that. This assumption breaks down significantly in the presence of the BFE, which is a local nonlinear convolution — at least when considered in the context of the PSF precision requirements found within interferometry. There is not a known accurate expression or simulator for the BFE, and certainly not an inverse operator to restore the un-blurred ideal pixel response.

This induces a nonlinear change in PSF shape between target and calibrator stars, so that different pipelines return different complex visibilities that do not calibrate in the Fourier domain by simple division. While the BFE was known prior to launch (Rowlands et al. 2018), its seriously harmful effect on interferometric analysis was only realised post launch.

The BFE, together with imperfect gnosys of the AMI metrology, have resulted in AMI underperforming. Until now, fringe stability and precision has not been much better than ground-based observations and AMI proposals have fallen short in the competitive environment of available JWST General Observer (GO) observing time.

In this paper, we present a new approach: Aperture Masking Interferometry Generative Observations (AMIGO), a pipeline in which we jointly train a physical model of the optical system with a hybrid forwards and machine-learned Effective Detector Model (EDM), and extract rich visibility information with a generalisation of kernel phase. We fit the optical and electronic models simultaneously to on-sky calibration data that allow us to separate these effects for the first time, and apply this base instrument model to extract interferometric observables from several science targets, which have hitherto resisted AMI pipelines, achieving

near-photon-noise-limited performance in detecting faint companions.

4.2.2 Differentiable Forward Models: From Pixels to Planets

AMIGO uses an end-to-end differentiable forward model of the entire end-to-end chain of physics based on Automatic Differentiation (autodiff; Margossian 2019) — the foundational algorithm of machine learning (LeCun et al. 2015). By decomposing functions into a sequence of function primitives and applying the chain rule programmatically, autodiff enables an algorithmic computation of machine-precise derivatives. Importantly, autodiff does not harness finite differences nor symbolic differentiation, instead computing the exact derivative of its input function directly. Its success can be attributed to its computational complexity, scaling with the model itself — even for high-dimensional or nested models — *rather* than the number of parameters being differentiated. Two primary algorithms underpin autodiff: ‘forwards’ mode (or the tangent method; Siskind and Pearlmutter 2008) and ‘reverse’ mode autodiff (or backpropagation; Griewank and Walther 2008; Rumelhart et al. 1986). These can be composed to efficiently implement operators such as Jacobians and Hessians through arbitrary computational programs. In particular, such partial derivatives are necessary for optimisation and sampling in high dimensions, for example by stochastic gradient descent (SGD; Ruder 2016) or Hamiltonian Monte Carlo (HMC). This native computational efficiency and accuracy has enabled the training of Machine Learning (ML) models with *billions* of parameters, giving rise to much of the modern world.

In this paper we build on the growing body of work on differentiable modelling in optics (Page and Favaro 2020; Wong et al. 2021; Desdoigts et al. 2023; Desdoigts et al. 2025; Sitzmann et al. 2018b; Liaudat et al. 2023a) and astronomy more generally (e.g. and non-exhaustively, Gully-Santiago and Morley 2022; Campagne et al. 2023; Hattori et al. 2024; Foreman-Mackey 2023; Dholakia et al. 2024; Horta et al. 2025; McDougall et al. 2025). AMIGO takes the leap to a true end-to-end approach, something yet to be comprehensively explored until now.

Beyond enabling optimisation and sampling with gradients, current autodiff frameworks like JAX (Bradbury et al. 2018) and PYTORCH (Paszke et al. 2019) offer substantial benefits over standard numerical processing libraries. Built and designed for ML research, almost all autodiff libraries offer many highly optimised tools that ease development and efficiency. Native deployment to hardware accelerators such as GPUs & TPUs, efficient compilers, function vectorisation and parallelisation, and higher order derivatives all give access to a toolbox that can accelerate research and development. Furthermore, many of the bleeding-edge optimisation and inference algorithms rely on the efficient derivatives offered by autodiff. Any new software tools built within autodiff frameworks offer strict benefits over the standard numerical libraries found throughout the astronomical software landscape.

The AMIGO model is built using ∂ LUX (Desdoigts et al. 2023) & ZODIAX are used as the base framework for differentiable optics and a user-friendly interface for scientific forward models respectively, which are built on JAX & EQUINOX (Kidger and Garcia 2021b) to ensure it can act as a single end-to-end differentiable system. JAX provides the core autodiff engine with EQUINOX providing the framework for object-oriented programming as well as the tools required

to implement ML models.

The most significant innovation within AMIGO is found in the detector model. Presented with the challenge of forward modelling the charge migration that manifests as the BFE — ultimately governed by differential equations — a novel solution was found by directly integrating a Neural Network (NN) *inside* the detector. This formulation breaks from the three commonly understood modelling paradigms: forwards, inverse and machine-learned, but is part of an emerging body of work known as ‘hybrid models’ (Akhare et al. 2023; Karniadakis et al. 2021; Willard et al. 2022).

Differentiable forward models show a path towards the next generation of precision calibration and data analysis. For variable nuisance processes that must be inferred directly from science data, which might be affected by heteroskedastic noise and unknown nonlinearities, it is better to fit a forward model to data with Bayesian methods (e.g. discussion in Hogg and Villar 2024). Such models can have very many parameters, necessitating autodiff; and may have to represent noise processes for which a physical simulation may be inadequate or unavailable, but which a neural network can adequately predict (e.g. discussion in Hogg and Villar 2024). In this case, we may not care about interpretability of the model for nuisance processes — in our case, an EDM for the BFE — but only that they perform well and do not damage recovery of the physics we *do* care about (the astrophysical scene).

The hybrid modelling approach presented in this paper connects forward models with machine-learned ones, gaining the best attributes of both approaches. Because the NN is embedded within the overall forward model, it cannot be trained in isolation on a well-curated, diverse dataset, and must instead be evaluated solely through its influence on the end-to-end model predictions. This approach is likely to cause discomfort for many, in particular in its generalisation to datasets very different to those on which it was trained; however its effectiveness in the regime of high contrast imaging will be thoroughly demonstrated in this work.

4.3 The Amigo Model & Pipeline

The AMIGO model and pipeline consists of a digital twin of AMI, trained only on high-quality in-flight point source data and flat-field calibration data; and then the application of this pre-trained ‘base model’ to Bayesian inference from science data of the instrument state (principally wavefront and Fresnel defocus) and astrophysical observables (visibilities, flux, spectrum, though with the option to fit more complex models), holding most of the model’s other parameters fixed.

While most data analysis pipelines work on processed and calibrated data, AMIGO is designed to generate predictions of the *uncalibrated* JWST data, fully independent of any other software, including the JWST official pipeline. While the embedded optical model resembles those found in WEBBPSF (Perrin et al. 2014; Perrin et al. 2012), the fiducial STScI-supported physical-optics simulator for JWST, its direct integration of a visibility forward model, and a detector model designed to produce the 3-dimensional time-evolving pixels found in uncalibrated data are significant departures from existing models.

AMIGO consists of five distinct and modular sub-models of: the optical system, the complex

visibilities, a linear detector, a non-linear ramp, and the read electronics. A flow diagram of the full model is shown in Figure 4.2, with each component detailed in later sections. While these components are modular and can operate independently, they are trained and behave as a single cohesive system.

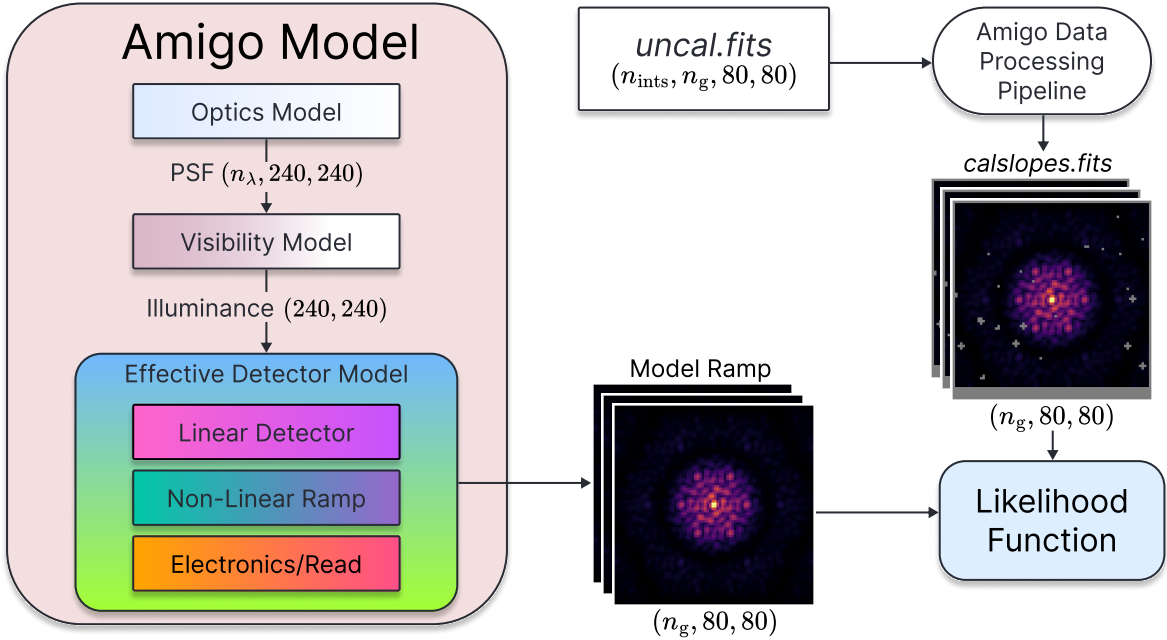


Figure 4.2: High level flow diagram of the AMIGO model and pipeline, showing the input and output product and shapes passed between each modular component. n_λ is the number of wavelengths modelled by the optics, n_g is the number of groups in the data, and n_{ints} is the number of integrations. Each of these model and pipeline components are discussed in detail in their own section.

4.3.1 Calibration Data

High quality calibration data is of the utmost importance for the ambitious calibration goals of the AMIGO model. Due to a late switch out of the detector hardware on JWST, there is not much publicly available, high-quality, pre-flight calibration data, necessitating an approach exclusively using on-sky data. While a number of calibrator stars have been observed by AMI mode, a requirement for interferometric calibration, most of these programs select bright targets in order to reduce required observation times. Consequently, almost all of the existing data is very coarsely sampled up-the-ramp, making it insufficient for accurate inference of the dynamics of the BFE.

To aid on-going calibration efforts of AMI and to seek a deeper understanding of the BFE, program [CAL 4481](#) (PI: Sivaramakrishnan) was proposed, accepted, and finally observed on 5 May 2024. Designed purely for calibration, it employs a 5-point sub-pixel dither, an uncommon approach for AMI data, with 2×10^9 photons at each dither position. Seeking to better

understand the detector systematics, HD 41094 was chosen (with the aid of [SearchCal](#); Bonneau et al. 2012) as our calibrator for building the digital twin. It is an appropriately-bright ($W2 = 0.69$ Jy) target for a good number of groups up-the-ramp, with 11, 20, and 30 groups per integration in the F380M, F430M and F480M filters respectively. As a K0 giant beyond 300 pc, with no known companions, we can be confident that it is a point source in the AMI band. In order to capture the full dynamics of the BFE without dealing with the compounding complexities that arise from deeper exposures, a peak pixel depth of $\sim 50 ke^-$ was chosen — approximately half way to saturation. A summary of the observing program is presented in Table 4.1.

A future calibration data-set designed to explore the dynamics of the BFE further up-the-ramp and with more complex illuminance patterns, GO 8330 has been already been accepted and awaits observation. Ideally, this program should allow for the AMIGO model to remain performant for brighter targets and deeper observations, expanding the observational capabilities of AMI into the future.

The CAL 4481 proved indispensable to the calibration of AMIGO. Diversity of both wavelengths and dithers, deep exposures with high temporal resolution, with balanced signal across filters all combine to provide an excellent training and testing ground for all AMI pipelines both at present and into the future. All proceeding discussion and presentation of calibration products in this work originate exclusively from this dataset, combined with in-flight flat-fielding data which disentangle the pixel sensitivities and nonlinearity from the spatially-varying effects of the BFE.

In order to ensure model generalisation and avoid over-fitting, a validation dataset was built from existing public programs [GO 1843](#), [GTO 1242](#), [ERS 1386](#), and [COM 1093](#). This set of calibrator star exposures were chosen with the goal of having as much validation data as possible, while keeping a balance of total signal in each filter, detailed in Table 4.1.

Table 4.1: Summary of JWST CAL 4481 observations for HD 41094, used for model calibration. All targets are point sources. Pixel well depth values have dimensionless integer units (‘digital number’ or DN), of which each count corresponds to $\sim 1.6e^-$.

Type	Program	Star	Filter	Groups	Integrations	Dithers	Well Depth	Photons
Calibrator	CAL 4481	HD 41094	F480M	30	3800	5	30,000	9.98×10^9
Calibrator	CAL 4481	HD 41094	F430M	20	4525	5	29,000	9.63×10^9
Calibrator	CAL 4481	HD 41094	F380M	11	5300	5	30,000	9.39×10^9
Validator	COM 1093	HD 36805	F480M	9	65	1	24,000	0.10×10^9
Validator	COM 1093	HD 36805	F430M	6	82	1	24,000	0.10×10^9
Validator	COM 1093	HD 36805	F380M	3	122	1	22,000	0.10×10^9
Validator	GO 1843	HD 205827	F480M	10	641	1	20,000	0.86×10^9
Validator	GO 1843	HD 205827	F430M	7	1885	1	20,000	2.15×10^9
Validator	GO 1843	HD 205827	F380M	3	7800	1	18,000	4.75×10^9
Validator	GTO 1242	HD 18638	F480M	8	4869	1	16,000	3.93×10^9
Validator	GTO 1242	HD 18638	F430M	5	6256	1	15,000	3.62×10^9
Validator	ERS 1386	HD 116084	F380M	3	10000	1	17,000	5.86×10^9
Validator	ERS 1386	HD 116084	F380M	3	6000	1	17,000	3.52×10^9

4.3.2 Data Processing Pipeline

JWST employs Teledyne HAWAII-2RG (H2RG) near-infrared detectors across various instruments, including NIRISS which hosts AMI mode. These detectors use a non-destructive read-out pattern known as ‘up-the-ramp’ sampling, where the voltage is measured multiple times as it accumulates charge producing a time evolving measurement in each pixel. This method provides better read noise characteristics and makes identifying and rejecting cosmic-rays easier (Rauscher et al. 2007). Ideally these pixels can be treated independently with linear fits to the resulting ramp solving for the incident flux.

Standard data calibration approaches based on inverse models seek to subtract these effects sequentially from data, ideally returning a clean signal representative of the input photon distribution. However, non-linear effects like the BFE are self-interacting, non-local, and couple to various properties not addressed at a pixel-level such as PSF shape, curvature across individual pixels, and sub pixel positioning. As a result, these effects do not calibrate straightforwardly by division in the uv plane.

The AMIGO data processing pipeline takes a very different approach, seeking to preserve as much of the physics as possible in the output product. The ‘up-the-ramp’ readout of the H2RG detectors provides 4D uncalibrated data: Two spatial dimensions, one time dimension for the read at each group, and a second time dimension for each integration, or ‘image’, taken.

The AMIGO pipeline performs a single calibration to the data, correcting for the Analogue to Digital Converter (ADC) integral non-linearity. There is a strong periodic residual in uncalibrated data that, until we identified this, made interpreting trends in the group-level data challenging as it induces unique per-pixel periodic signals greater than the intrinsic noise. This behaviour is thought to arise from a lack of power being supplied to the amplifiers. It is periodic in *raw counts*, and we estimate it simply by performing a least-squares fit to the average cleaned ramp values. Plotting the residual to this fit against the data value (i.e. including the pixel bias) reveals a strong sinusoidal signal. While we are not able to infer a functional form accurate at all count levels, we find a sufficiently accurate fix for our purposes by subtracting off a sinusoid with period 1024 and amplitude 2. This correction leaves some periodic residuals in the ramp-level data, however the existing correction is currently sufficient. We also found this signal in other observing modes, with a more comprehensive treatment and correction to come in subsequent work (Dholakia, in prep). Figure 4.3 shows this residual and the post-correction residuals.

Next, simple outlier rejection is performed on the ADC corrected ramps, rejecting the 3σ outliers from each group. This same outlier rejection is performed at 3σ on the data slopes, found by taking the differences between each group read. This outlier detection process is aimed at removing cosmic rays and other spurious jumps in the data.

From this outlier cleaned data the mean and covariance matrix of each pixel is calculated along the integration axis for both the ramp and slope data. This produces the final `calslope` output product. At present the AMIGO model is designed to fit the slope data, although the ramp data is still preserved with the intention of extending the model to predict ramp level, rather than slope level data in order to produce higher fidelity over the pixel gain.

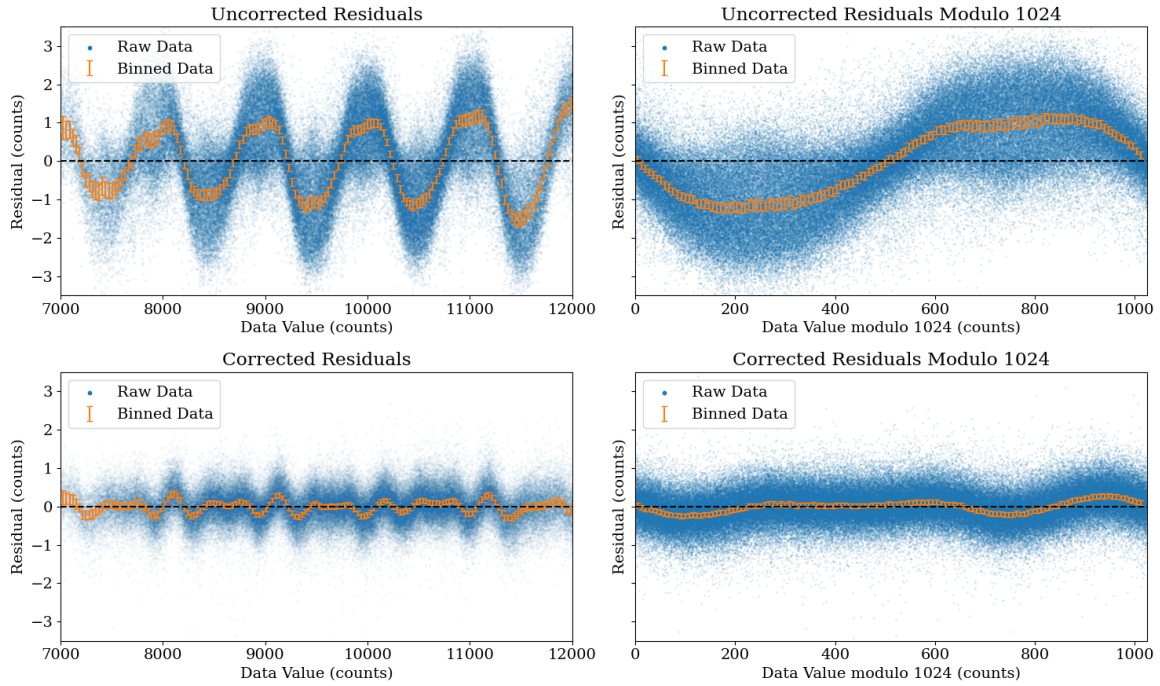


Figure 4.3: Residuals from second-order polynomial fits to ramp data, shown before (top row) and after (bottom row) applying a sine-wave-based correction for ADC integral non-linearity. The left panels plot residuals as a function of the ramp value, while the right panels show the same residuals folded over a modulo 1024 pattern, revealing periodic structure. Prior to correction, the residuals exhibit a strong sinusoidal modulation. The applied correction consists of a 1024-period sine wave with fixed amplitude, significantly reducing both the overall residual structure and the folded periodicity (bottom panels). Orange points and error bars represent binned data mean and standard error in each bin, highlighting the improved uniformity of residuals post-correction.

An important consequence of fitting to the slope data is the loss of the first group read. While this had the benefit of avoiding predicting the pixel-to-pixel bias, it also results in a loss of $1/n_{\text{groups}}$ worth of data. As a result, the AMIGO model at present can not predict observations made with a single group read, and loses a higher fraction of the total photons for observations made with a small number of groups.

As most AMI observations deliberately do not go above half pixel depth, in order to avoid BFE, it is not possible with current resources to train the AMIGO model to cope with saturated sources, and we obtain worse fits with increasing well depth. This may be mitigated in future as deeper datasets become available and the model is re-trained to include these.

4.4 Optical & Visibility Model

AMIGO combines a diffractive physical optics model with a novel interferometric visibility forwards model. This enables the injection of observed source brightness distributions into PSFs in an optically coherent manner and encoded via its complex visibility.

4.4.1 Optical Model

The first stage of AMIGO uses a differentiable physical optics model based on ∂LUX (Desdoigts et al. 2023; Desdoigts et al. 2023), which has similar features to other Python physical optics packages like `WEBBPSF` (Perrin et al. 2012; Perrin et al. 2014) and `PRYSM` (Dube 2019) but provides automatic differentiation and hardware acceleration and parallelisation through `JAX`.

To circumvent these issues the AMIGO optical model has a strong focus towards flexibility, an approach facilitated by `autodiff`. With direct calibration from on-sky data, it emerges as the most precise physical optics model of AMI, with dynamically generated, non-linearly distorted aperture geometries, persistent wavefront sensing, broadband PSFs, and Fresnel diffraction.

The optical model can be broken into four components:

1. The spectral model
2. The aperture model
3. The wavefront model
4. The propagation model

Each component enables the coherent flow of gradients from residuals of the predicted PSF back through to the underlying parameters, enabling precise calibration and resulting in a highly accurate and physically principled PSF describing the instrumental response to incoming wavefronts.

4.4.1.1 Chromaticity and Spectral Model

AMI observations can be taken with four different filters, detailed in Table 4.2, although most observations avoid F277W due to these wavelengths being significantly under-sampled by the

NIRISS pixels. Other potential problems arise with the F277W filter, since the BFE is also believed to have wavelength dependence, generalising the implementation of the AMIGO model would require greater complexity and more calibration data. For these reasons, the CAL4481 program (used for training AMIGO) did not observe using the F277W filter, performance in this band is not well characterised and not discussed further.

Table 4.2: Allowed Filters for AMI observations. Values taken from the [JWST documentation](#). Full tabulated curves used in propagation.

Filter	λ central	$\Delta\lambda$
F480M	4.815 μm	0.289 μm
F430M	4.285 μm	0.203 μm
F380M	3.825 μm	0.205 μm
F277W	2.771 μm	0.717 μm

All three of the primary AMI filters are relatively narrow in comparison to the size of the generally expected spectral features, allowing a simple linear spectral energy distribution model to be used as the default via

$$F(\lambda) = 1 + m(\lambda - \lambda_c) \quad (4.1)$$

where λ_c is the filter’s central wavelength and m is the spectral slope parameter. This default model is unlikely to capture the complexities present in some science cases. However, the forward modelling framework enables the inclusion of user-defined differentiable spectral models, tailored to the specifics of any observation.

By default AMIGO propagates 9 monochromatic wavelengths through the optical systems to capture the appropriate spectral diversity. Each wavelength is weighted by the integrated filter bandpass and the spectral weights and summed to a broadband illuminance pattern.

4.4.1.2 Aperture Model

The AMIGO aperture model is similar to, but significantly more flexible than existing aperture models used throughout the field. Each aperture mask hole is modelled as a soft-edged hexagon rendered dynamically on the aperture coordinates. The soft-edges of the aperture enable stable gradient propagation through to the coordinate grid, despite the output array being dominated by the binary ones and zeros. The mask is parameterised by two sets of 2D polynomial distortion coefficients: one that controls relative positions of each aperture mask hole, and one for each of the 7 aperture mask holes, applied to the hexagon coordinate array.

These aperture polynomial distortions provide the flexibility required to model the astrometric distortion seen in real optical systems. Typically, these distortions are modelled by applying a polynomial distortion to the PSF in the image plane, as opposed to the aperture in a pupil plane. The choice to apply these distortions in the pupil plane enables the prediction of a PSF that remains governed by diffractive physics. Additionally, this approach accounts for any relative shears/rotations between the wavefront and the aperture mask when encountered in

the filter wheel.

The use of 2D polynomial distortions is important as it enables the removal of two problematic degrees of freedom: the global x - and y -positional shift. Since the optical model is diffractive, the wavefronts are propagated from pupil to focal plane via a Fourier transform, and the PSF is the squared modulus of this focused wavefront — hence the phase information encoding the global position of the aperture is lost. Consequently, the model is invariant to x, y aperture shifts; however, noise in the gradients produced by autodiff can result in undesired drifts. By removing these degrees of freedom from the model, we can pin down both the global aperture mask position and the relative positions of each hole, avoiding this problem entirely. Figure 4.4 presents the residual between the recovered distorted aperture mask and its undistorted counterpart along with its effect on the PSF.

We identified an unexpected discrepancy in the reported size of the AMI mask holes across different online resources and software packages. The [JDox](#) documentation states that the face-to-face diameter of each hole, as projected onto the primary mirror, is 0.80 m. However, the [mask definition file](#) used by both `WEBBPSF` and `IMPLANEIA` adopts a slightly larger value of 0.82 m. Several other analysis pipelines in the AMI ecosystem were also found to prefer one of these two values. This discrepancy in the manufactured size of the mask holes was resolved in two ways: first, by allowing them to be inferred directly during the AMIGO calibration process, and second, by cross-checking the result against the original manufacturing specification file used. The AMIGO model confidently recovered a hole diameter of 0.80 m, consistent with both the `JDox` documentation and the manufacturing specification file, thereby validating the smaller of the two reported values.

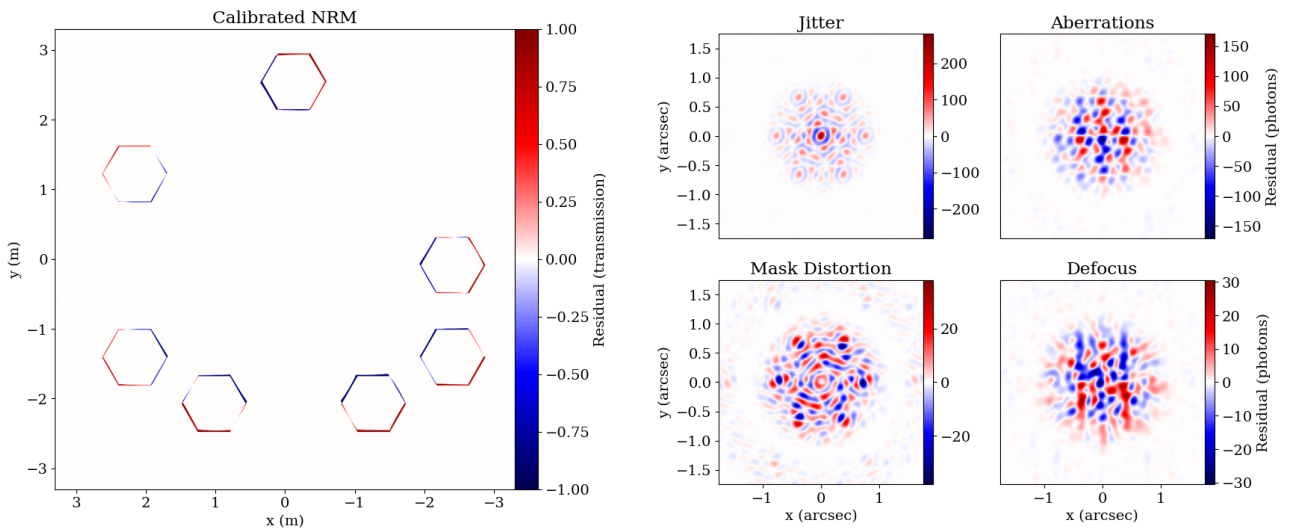


Figure 4.4: Left panel: Residual between the calibrated aperture mask and its idealised undistorted counterpart. Right panel: PSF residuals of the four primary optical effects on the PSF. Top left: Instrumental jitter, applied through a convolution with a Gaussian kernel. Top right: Primary mirror aberrations, modelled using Zernike polynomials on the primary. Bottom left: Aperture mask distortions, modelled by applying a distortion to the coordinates over which the aperture mask is calculated. Bottom right: Fresnel defocus modelled using a Fresnel propagation algorithm. All effects are shown for the F430M filter, using a PSF with 10^6 total photons.

4.4.1.3 Wavefront Model

Accurate PSF modelling of any real system must also account for phase errors in the wavefront, accumulated by both the primary mirror and through the optical train. Despite its unprecedented optical stability, JWST is no exception. Observations made during its commission phase revealed two important time-dependent optical degradations: mirror tilt-events and micro-meteoroids (Rigby et al. 2023). These effects necessitate the recovery of wavefront phases between, and sometimes within, any given observing program, a task made far simpler by both the non-redundant aperture mask and the gradients provided by the differentiable model. While most optical systems suffer from a sign degeneracy within certain phase modes, NRMs offer unambiguous recovery of all phase modes, making them an ideal calibration tool (Cheetham et al. 2012; Pope et al. 2014).

To recover the Optical Path Difference (OPD) state of the cumulative optical surfaces we employ Zernike polynomials, a set of orthogonal functions defined over the unit disk, commonly used to represent wavefront aberrations in optical systems. Their orthogonality and correspondence with classical aberration types (e.g., defocus, astigmatism, coma) make them particularly useful for decomposing and quantifying optical distortions. Due to these properties, Zernike expansions provide a compact and physically interpretable basis for modelling and correcting wavefront errors across a wide range of optical applications (Lakshminarayanan and Fleck 2011).

Given that JWST has a segmented primary mirror, we model a unique set of moderate order Zernike polynomials over each hole within the aperture which are then fit to any given observation. Choosing a total of 4 radial orders, we get 10 Zernike modes per hole for a total of 70 phase modes across the aperture. The inclusion of higher-order effects is trivial but was found to be unnecessary, as these terms remained statistically insignificant in the calibration data, up to $\sim 10^{10}$ photons. By default, AMIGO will recover wavefront phases for all observed data, providing a potential way to do long-term wavefront sensing independent of dedicated observations. The effect these aberrations have on the resulting PSF as found in the calibration data is shown in Figure 4.4.

Somewhat surprisingly, slightly different wavefront phases are recovered across the three filters, with the most significant deviation found in the F430M filter. These differences, visualised in Figure 4.10, we believe to arise from imperfections in the optical surfaces of the filters. The differences found in F430M can also be found through a significant deviation in the recovered Fresnel defocus, discussed in the next section.

4.4.1.4 Propagation Model

Under the Fraunhofer approximation, optical systems are typically modelled using two conjugate planes: the pupil and the focal plane. This approach assumes that light propagates as planar wavefronts, which is valid when the observation point is in the far-field or at the focal point. However, many practical optical systems exhibit complexities such as a misalignment along the optical axis which necessitate modelling wavefront propagation to intermediate planes. In these scenarios, Fresnel diffraction theory (Morse and Feshbach 1953; Born et al.

1999; Hecht 2002; Goodman 2005) becomes essential, as it accounts for the coupling between wavefront phase and amplitude variations, providing a more accurate representation of PSF behaviour in these regimes.

The Fresnel diffraction integral, best expressed through Fourier Transforms, which describes the complex field E at point (x, y, z) , is given by

$$E(x, y, z) = \mathcal{F}^{-1} [\mathcal{F}[E(x', y')](k_x, k_y) \cdot e^{ik_z z}] \quad (4.2)$$

where $E(x', y')$ is the field in the pupil plane, $k_z = \sqrt{k^2 - k_x^2 - k_y^2}$, and k_x, k_y correspond to the angular spectrum (spatial frequencies) of the wave. A notable application of Fresnel diffraction modelling is in the analysis of the HST optical performance. Thermal fluctuations in the HST structure cause "breathing" modes, leading to temporal variations in the telescope's focus. By employing Fresnel-based models, researchers have been able to accurately characterise and correct these focus variations, thereby enhancing the quality of the scientific data obtained from HST (Krist et al. 2011).

Fresnel diffraction algorithms were found to be a necessary component to recover accurate PSF morphology with the AMI observing mode, with significant differences found across filters. The recovered defocus values were found to be $0.017 \mu\text{m}$, $0.050 \mu\text{m}$, and $0.010 \mu\text{m}$ in the F480M, F430M, and F380M filters respectively. F430M was found to have the largest defocus, with more than double that found in the other filters. The unique effect induced to the PSF in the F430M filter is shown in Figure 4.4. It is worth noting that due to the need to recover both primary mirror aberrations and Fresnel defocus values, which have a high degree of covariance in the small defocus regime, that these values can not be recovered fully independently. Considerable testing was done and accurate PSFs could only be recovered when considering both of these effects in tandem.

The coupling of amplitude and phase effects in wavefronts introduced by Fresnel effects is typically not important to downstream analysis, however careful consideration of its implication on interferometric observables is essential to recovering well-calibrated and high-precision visibilities. This is discussed more in Section 4.4.2.2.

We enhance the fidelity of the modelled PSF by applying a cubic spline interpolation to up-sample the image from a $3\times$ to a $9\times$ resolution. The result is then down-sampled back to the original $3\times$ grid. While this may appear redundant, the process improves realism by more accurately accounting for the finite area over which detector pixels integrate light — an effect that is not captured by simple point sampled PSF models. This method was benchmarked against a reference PSF generated at $30\times$ oversampling and was found to match the fidelity of a $\sim 12\times$ oversample, while requiring only a $3\times$ optical propagation. The choice of $9\times$ for the up-sampling factor reflects a practical compromise between accuracy and computational efficiency. The resulting increase in spatial resolution reveals finer PSF structure, which is especially important given the BFE's strong coupling to PSF curvature.

4.4.2 Interferometric Visibility Model

Optical interferometry is typically conceived of as a purely inverse problem: each image has a set of complex visibilities that can be reduced down to the non-redundant baselines and analysed. This paradigm is insufficient for forward modelling approaches because in reality the visibility signal is defined across the entire support of the Optical Transfer Function (OTF), something that must be captured by our modelling approach. A visibility model without the flexibility to reproduce the full behaviour of the PSF as observed through the optical system will introduce biases and non-physical signals in the predicted PSF that further couples through the downstream non-linear detector model. This problem has mandated a re-think of the concept of a visibility as it applies to forward modelling, as well as the analysis methods used on the recovered observables.

In order to achieve the required behaviour our model must produce a continuous array of interferometric amplitudes and phases, appropriately conjugated about the origin, that can be multiplied by the PSF in the uv -plane. This is done by defining a set of knots across the OTF, whose values (one amplitude and one phase each) can be interpolated to the appropriate uv -coordinates as defined by the Fast Fourier Transform (FFT) of each monochromatic PSF. The resulting complex visibility map is then multiplied by the complex PSF splodges and transformed back to the image plane. The resulting PSF can now express complex instrumental and astrophysical effects while remaining firmly grounded in the diffractive physics that governs imaging systems. Figure 4.5 presents a diagram of this process from the wavefront to the final interferogram.

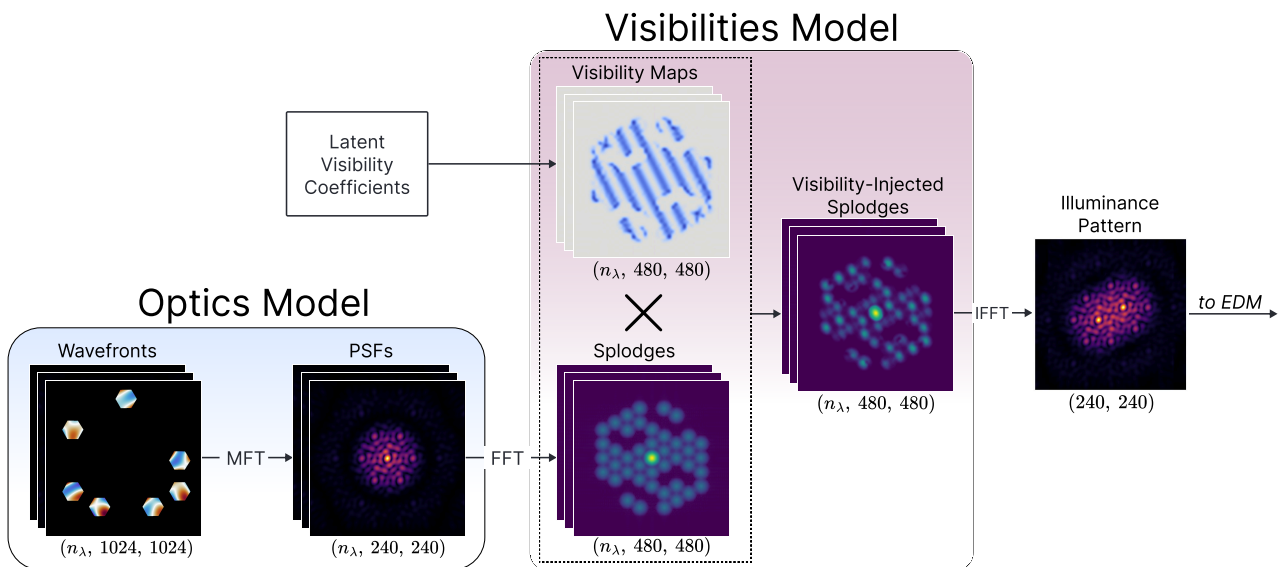


Figure 4.5: Flow chart of the injection of visibility signals to forwards-modelled PSFs. This demonstrates how high-resolution visibility signals can be directly injected into any PSF model provided the appropriate set of visibility basis vectors. An example binary-star signal is injected as a demonstrator.

4.4.2.1 Latent Visibility Model

This simple visibility modelling method satisfies the requirements for use in a forward model, but comes with some draw-backs. The pixel basis used for the visibilities results in high-dimensional and covariant parameters, with many pixels falling out outside the OTF and therefore being unconstrained by the PSF. Highly dimensional problems present no inherent issues for differentiable models, but strongly covariant or unconstrained parameters can cause problems for optimisation algorithms. This mandated a deeper and more principled consideration of forward modelled visibilities, underpinned by Reduced Order Modelling (ROM; Benner et al. 2015).

ROM is a set of methods that can reduce parameter dimensionality in a way that preserves the essential dynamics or structure of the original system. We apply this concept to our visibility model, using autodiff to design an orthogonal low-dimensional latent set of parameters that fully capture the dynamics of the pixelised visibility model. However, ROM methods rely on linear model assumption, a property exhibited by visibility phases, but not amplitudes. This mandates a treatment of the *logarithm* of the complex visibilities

$$\ln(Ae^{i\phi}) = \ln(A) + i\phi \quad (4.3)$$

reformulating the mapping between the parameters in the uv and focal plane to be linear though the real and imaginary components of the logarithmic complex visibility (Pope 2016). Using this construction, we now seek a matrix \mathbf{V} that can map between our full set of log amplitudes and phases

$$\ln(\mathbf{A}), \boldsymbol{\phi} \in \mathbb{R}^N \quad (4.4)$$

and a latent set

$$\mathbf{A}_l, \boldsymbol{\phi}_l \in \mathbb{R}^M \quad (4.5)$$

where $M < N$.

To construct the matrix \mathbf{V} , best envisioned as a set of basis vectors, we compute the Jacobian matrix $\mathbf{J}_{\text{psf}} \in \mathbb{R}^{N \times L}$ of the PSF, where L is the number of pixels in the PSF, with respect to the pixelised log visibilities $\ln(\mathbf{A}), \boldsymbol{\phi}$. This Jacobian captures the local sensitivity of each pixel in the PSF to each visibility log amplitude and phase parameter. Using the Gauss-Newton Hessian approximation

$$\mathbf{H} = \mathbf{J} \cdot \boldsymbol{\Sigma}^{-1} \cdot \mathbf{J}^{-1} \in \mathbb{R}^{N \times N} \quad (4.6)$$

we can estimate the Hessian of the pixelised log visibility parameters under a log-likelihood. An eigen-decomposition of this Hessian matrix

$$\mathbf{H} = \mathbf{Q}\boldsymbol{\Lambda}\mathbf{Q}^T \quad (4.7)$$

returns eigenvectors \mathbf{Q} that form an orthonormal basis for the uv -plane, ordered by their

influence on the interferogram. By selecting out the top M eigenvectors corresponding to the largest eigenvalues Λ (i.e. ones with high Fisher information, that are constrained by data), we tailor a low dimensional latent set of orthogonal basis vectors $\mathbf{V} \in \mathbb{R}^{N \times M}$ that exhibit favourable properties for modelling.

This method is very powerful in many modelling regimes and enables customisation of the properties of the latent basis vectors through the choice of Σ in Equation 4.6. In the context of interferometric visibilities we consider two choices for Σ . The first is the PSF itself, providing basis vectors that best constrain point sources — ideal for high-contrast companion recovery. This nevertheless could introduce difficulties when recovering extended, resolved, or medium-high contrast sources, where the noise distribution may be very different from the simulated point source used to build the basis. The second option is the identity matrix, providing an equal weighting to all pixels in the image plane. This provides higher expressiveness in the resulting basis vectors that remain capable of the high-contrast sources often targeted by interferometric observations.

In order to maximise generality, we chose an identity pixel covariance matrix to produce a set of latent visibility vectors, visualised in Figure 4.6. Interestingly but perhaps unsurprisingly, the low order basis vectors appear to select out the classic interferometric baselines. Examining the eigenvalues in Figure 4.6, we can see there is a knee around index 600 in both the amplitudes and phases. These reflect the visibility values *outside* of the OTF, where the PSF is unresponsive to any signal - as it is outside of the range of frequencies that the optical system is responsive to. We therefore model only those modes that are actually constrained by the optical configuration; an approach which implies a general method for describing the information content of diffraction-limited images, generalising the speckle statistics described by Mawet et al. (2014); a full exploration of this idea is beyond the scope of the present work.

4.4.2.2 Visibility Amplitudes & Phases Coupling

An interesting consequence of the observed Fresnel effects discussed in Section 4.4.1 is the coupling of amplitude and phase effects in the wavefront. This paradigm further translates to the complex visibilities — meaning that the visibility amplitudes and phases do not cleanly separate and act independently on the PSF, instead having non-insignificant covariances. This was directly observed in the estimated Hessian of the pixelised visibilities as calculated by Equation 4.6. This implies the true orthonormal visibility basis vectors live in *complex* space, and is therefore composed of both visibility amplitudes *and* phases. However keeping the classically understood visibility amplitudes and phases separated eases both their interpretation and comparisons to existing methods, and is a convention adopted through the rest of this work.

4.4.2.3 Kernel Amplitudes & Phases

Kernel phase analysis offers a powerful method to extract interferometric observables robust to residual wavefront error, even through clear aperture optical system with redundant uv baselines (Martinache 2010), and conversely to infer wavefront error from science data by the same approximation (Martinache 2013; Pope et al. 2014). It generalises the idea of closure

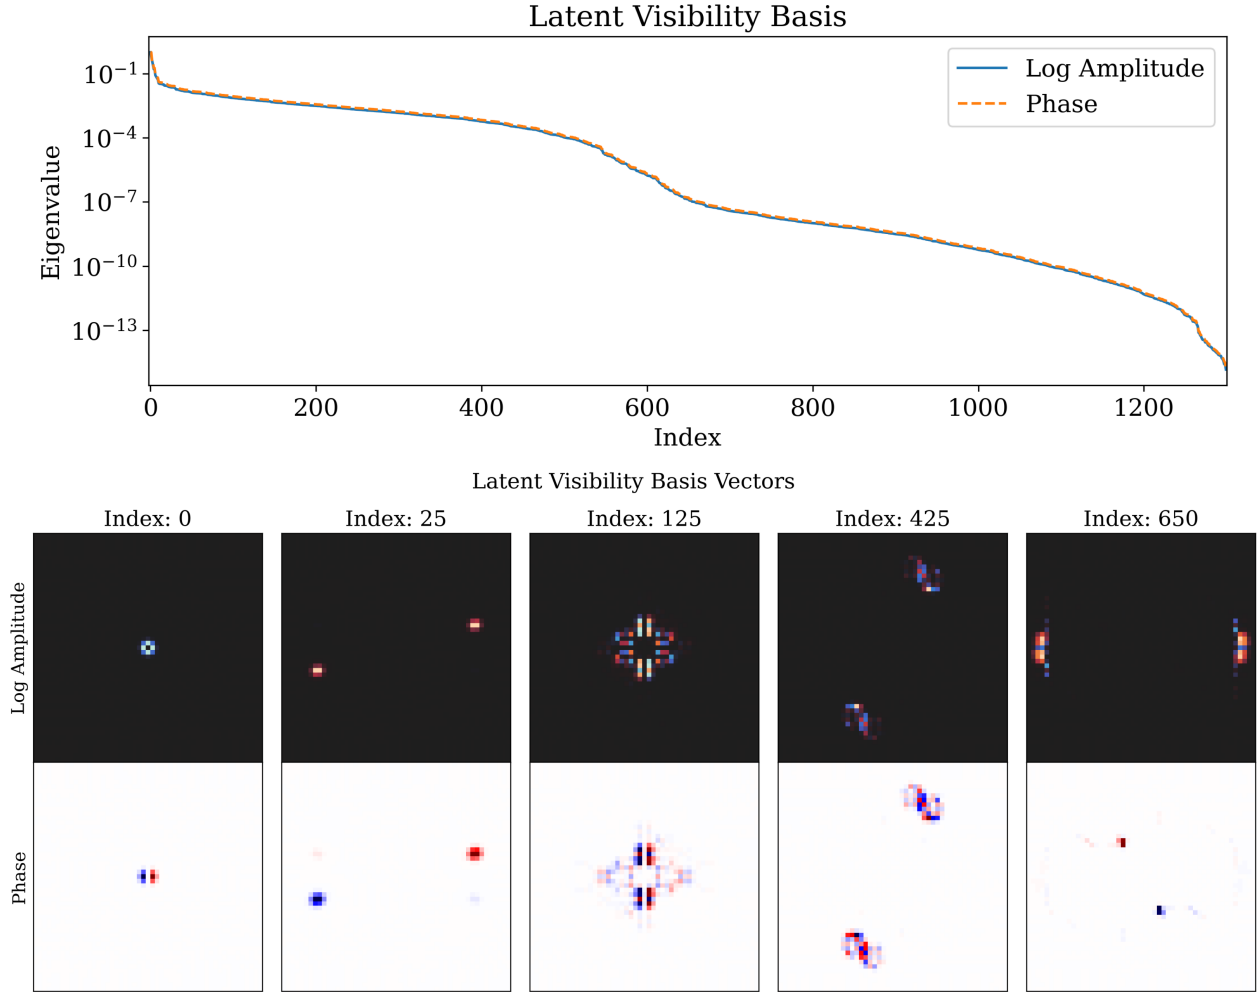


Figure 4.6: Demonstration of the produced latent visibility basis used for model-fitting. Top panel: The normalised eigenvalues for each visibility basis vector, ordered by their impact on the PSF in the image plane. Bottom: Representative log amplitude and phase basis vectors over a range of indexes. We can see that higher basis indices have increasing spatial resolution over the OTF, with low order ones picking out the classical interferometric baselines and their conjugates. Indices above ~ 600 start to put power outside the OTF, are un-sensed by the optical system, and are excluded from the model, but shown here to demonstrate how the basis can be restricted to inside the OTF.

phases (Jennison 1958), by approximating the effect of phase aberrations on pupil elements ϕ_j on the phases measured on baselines Φ_i , using a Jacobian matrix $\mathbf{J} \equiv \partial\Phi_i/\phi_j$

$$\Phi^{\text{meas}} = \Phi^{\text{sky}} + \mathbf{J} \cdot \phi. \quad (4.8)$$

The kernel phase idea is then to use singular value decomposition to find a kernel matrix \mathbf{K} such that $\mathbf{K} \cdot \mathbf{J} = 0$, so that

$$\mathbf{K} \cdot \Phi^{\text{meas}} = \mathbf{K} \cdot \Phi^{\text{sky}} + \underbrace{\mathbf{K} \cdot \mathbf{J} \cdot \phi}_{=0} \quad (4.9)$$

and therefore recovering ‘kernel phases’ $\mathbf{K} \cdot \Phi^{\text{meas}}$ which are self-calibrating with respect to

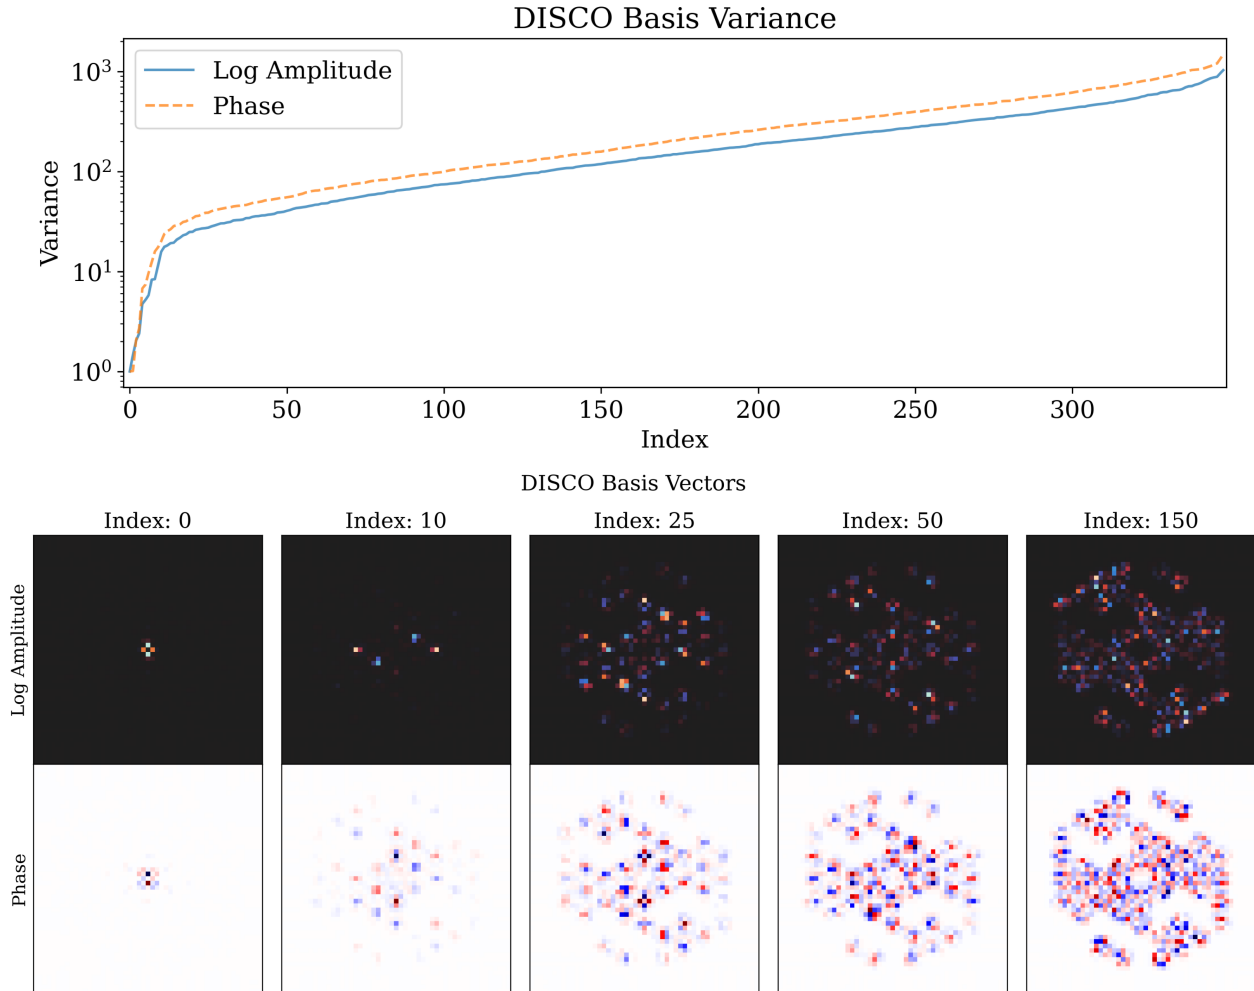


Figure 4.7: Example Delay-Invariant Subspace of Calibrated Observables (DISCO) basis vectors found from the GO 1843 observation, discussed in Section 4.8.3. Top panel: Normalised log amplitude and phases basis vector variances. Bottom panel: Selection of representative Delay-Invariant Subspace of Calibrated Observables (DISCO) basis vectors. Matching with the latent visibility basis vectors, low index vectors are better constrained and have lower spatial fidelity.

wavefront noise.

In its first implementations (Martinache 2010) this is calculated analytically by autocorrelation of a grid of points, and later by autodiff (Pope et al. 2021). This idea can be extended to kernel amplitudes (Pope 2016), by performing analysis on the log-visibility, a practice we adopt here, which has the advantage of making gain errors additive rather than multiplicative and therefore putting phase and log-amplitude on an equal footing as the complex logarithm of a wavefront or visibility. The kernel phase and amplitude relations are approximate and rely on linearisation for redundant apertures, but are exact and recover closure phases as a special case for non-redundant apertures.

Even though we infer Zernike coefficients for each mirror fairly precisely in the workflow above, the kernel phase idea is still required here: instrumental degrees of freedom (here, Zernike coefficients) are linearly indistinguishable from a subspace of astrophysical degrees of freedom

(i.e. linear combinations of complex visibilities that are sensitive to Zernikes, and cannot be inferred separately). Furthermore, some observations may require fitting a unique wavefront or visibilities to different detector positions, over observations made at different observational epochs, or may need to account for tilt-events during or between different exposures, as occurred between imaging AB Dor and its calibrator.

Further building on these ideas and leveraging autodiff we expand the null-space from pure wavefront phases on a Zernike basis to also null over small defocus, flux and spectral miscalibrations. Trivial to calculate using autodiff (Pope et al. 2021), this formulation enables fine-grained control over the order of wavefront error, and other nuisance optical effect over which to null the observables. In this work we only null over the same Zernikes modelled in the optical system (i.e. the first 10 modes, as discussed in Section 4.4.1), preserving more astrophysical information in the interferometric outputs than a pixel-basis would allow. In principle, this same approach could be used to generate kernels invariant with respect to any instrumental or nuisance degrees of freedom, such as mask rotation or jitter.

Kernel amplitudes must be used for a different set of reasons. As discussed in Section 4.4.1, accurate PSF metrology can only be recovered by modelling wavefront behaviour outside of the Fraunhofer regime, i.e. using a Fresnel propagation, which projects phase aberrations into effects both in phase and amplitude in the final uv -plane. Furthermore, the finite field of view causes spatial correlation of complex visibilities, convolving the real and imaginary parts with a window function and therefore mixing amplitude and phase. In developing AMIGO, we notice miscalibration in amplitudes unless we account for this effect.

This projection of our extracted visibilities into a kernel space provides one final statistical hurdle: the output kernel visibilities do not preserve the original statistical independence of the basis vectors. To circumvent this final issue, an eigen-decomposition is performed on the measured visibility covariance matrix to restore statistical independence (following Ireland 2013). If we were to keep the full covariance matrix and use this in the likelihood, this would not change the information content, but projecting once to a basis of statistically independent observables does allow for accelerated downstream analysis in evaluating the likelihood without matrix inversion each time.

Each operation in this process from pixelised visibilities through to statistically independent observables is linear, enabling each projection to be mapped into a single matrix, the Delay-Invariant Subspace of Calibrated Observables (DISCO) matrix. Figure 4.7 shows the DISCO basis vectors re-projected onto the pixel basis with their corresponding variances. Interestingly, many of the produced basis vectors look similar to combinations of the latent visibility basis vectors, implying that much of the information is preserved even through projection to the null space.

We believe that this representation of our data is close to optimal: using latent visibilities derived from the Fisher matrix eigenvectors ensures inclusion of *all and only* the information passed by the optical system, so that unconstrained modes are ignored but able to express any detectable visibility pattern across the uv plane. Then by projecting to a kernel space these are protected from miscalibrated instrumental degrees of freedom; and the full information from correlated noise is preserved in the final representation. By using an accurately trained pupil

model and an exact Fourier sampling to construct this basis, we alleviate the issue of model misspecification which has rendered the original HST/NICMOS kernel phase results suspect (Pope et al. 2013; Martinache et al. 2020), by instead directly solving for the instrument metrology first.

4.5 The Effective Detector Model

AMIGO treats the detector, with its complicated pattern of sensitivity and cross talk, with a non-parametric EDM. The JWST H2RG detectors in NIRISS are subject to the BFE, driven by electrostatic interactions within the detector substrate, which induces flux-dependent charge redistribution across neighbouring pixels. These effects evolve over time and are entangled with other systematics, including variations in the PSF and inter-pixel sensitivity. Importantly, they non-linearly act on a patch of pixels in the core of the PSF, depending sensitively on the illumination pattern over these pixels, which will differ from star to star. As a result, pipelines depending on calibration in the Fourier domain may be inadequate for recovering signals at the precision levels required for high-contrast exoplanet imaging.

The EDM circumvents these issues by modelling the detector as a coherent physical system. Its architecture mirrors key components of the standard JWST pipeline but employs the pixel-to-planets philosophy of end-to-end differentiable *forward* models. This construction enables further innovations ideally situated for complex or poorly understood physical processes through the direct integration of a NN.

The defining feature of the EDM is its integration of machine-learned components into a physics-based forward model, forming what is known as a *hybrid model* (Akhare et al. 2023; Karniadakis et al. 2021; Willard et al. 2022). In particular, a NN is embedded *inside* the detector model as a differentiable transformation, trained to capture the charge migration behaviour of the BFE. This breaks from common NN usage in astronomy, where models are trained independently and used as surrogates for entire processes. Here, the NN is treated as one operator in a long Markovian chain, calibrated through gradients propagated from the raw detector data all the way back to the astrophysical parameters.

This design yields several critical advantages:

- The NN can leverage accurate upstream physical predictions, reducing its complexity and improving interpretability.
- Training is conducted *entirely* on real, on-sky data, without the need for lab-based calibrations or curated training sets.
- Physical constraints can be imposed on the NN outputs, such as flux conservation and locality of charge migration, drawing inspiration from the field of physics-informed machine learning (Raissi et al. 2019; Kidger 2022b).

The hybrid approach of this model introduces new challenges. Since the NN is not trained in isolation, the quality of its predictions is dependent on both the accuracy of the upstream op-

tics and visibility models as well as the downstream electronics model. If other components are poorly calibrated, the NN may erroneously learn to compensate for them, degrading generalisation. This coupling mandates joint optimisation of all model components, enforcing physical realism at each stage.

The EDM is comprised of three main stages:

1. Linear Detector Model
2. Non-linear Ramp Model
3. Electronics Model

Together, these components form a structured pipeline that mirrors the physical data generation and readout process of H2RG detectors, as outlined in Figure 4.8. Because the BFE is spatially non-local, non-linear, and self-interacting, it poses the greatest challenge to conventional approaches out of the effects in this chain. Additionally, its interaction with other effects (such as flat field variations) renders modular correction ineffective. For example, an insensitive pixel will accumulate fewer electrons, which biases the electric field measurement and thus affects the migration behaviour of surrounding pixels.

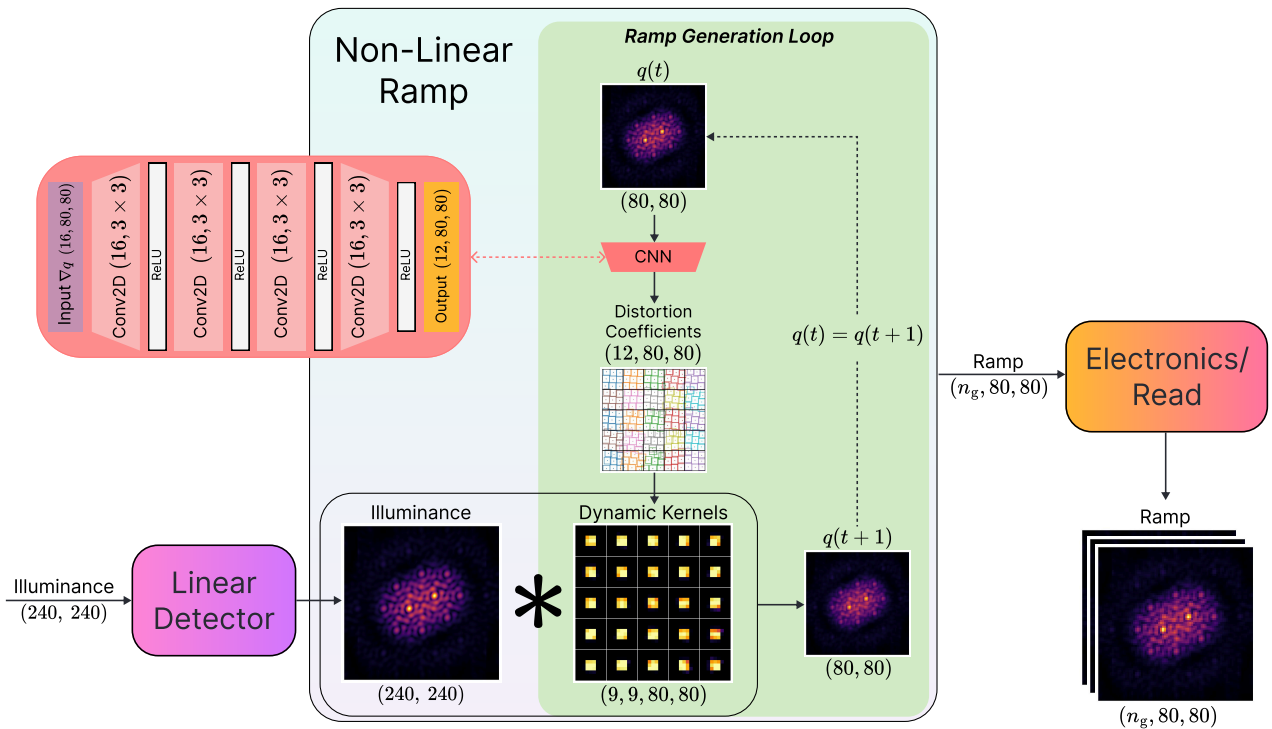


Figure 4.8: Schematic diagram of the EDM architecture, showing the three major components and the embedded neural network used to capture non-linear charge migration.

To our knowledge, the EDM presented here is a novel tool in astronomy, analogous to the recent use of ML for subgrid physics in otherwise physically-rigorous weather simulations (Kochkov et al. 2024). The result is a unified, physics-informed, and data-driven model capable of meeting the stringent requirements of exoplanet imaging.

4.5.1 Linear Detector Model

The first component of the EDM is the linear detector model which handles the transfer of the predicted PSF onto the pixels of the detector. This process includes two effects: instrumental jitter and PSF resampling.

Starting with the PSF predicted by the upstream optics and visibility model, instrumental jitter is applied at the calculated $3\times$ PSF oversample by a Gaussian convolution. The instrumental jitter discussed here is distinct from the more commonly understood *pointing* jitter. Instrumental jitter is a blur arising from imperfect stability of all components in an optical train, as opposed to telescope pointing instability. JWST has an exquisite pointing accuracy of ~ 1 mas, but the instrumental jitter recovered by AMIGO is a much larger at ~ 21 mas. This blur can arise from vibrational modes at frequencies much greater than a single integration in any of the optical components, such as the secondary mirror, or from electron diffusion in the detector substrate. It is not known which of these two potential sources is the origin of this blur. Instrumental jitter causes a substantial reduction in the sharpness of the PSF; this is visualised in Figure 4.4.

The next step requires resampling the predicted PSF from the optical coordinates to the detector pixels. The NIRISS detectors exhibit a [documented anisotropy](#), as well as a small $\sim 0.5^\circ$ rotation with respect to the optical axis. Both of these operations are simultaneously applied via a cubic interpolation of the oversampled PSF, while preserving the $3\times$ oversample. Note that AMIGO does not use existing distortion solutions for NIRISS, since the PSF distortion effects in AMIGO are modelled via the pupil rather than focal plane. The final output is the oversampled illuminance pattern incident to the detector pixel.

4.5.2 Non-Linear Ramp Model

The greatest challenge in the entire AMIGO pipeline is accurately capturing the dynamics of the BFE, handled by the non-linear ramp component of the EDM. The dynamics of the BFE should be expressible as a differential equation that can be evaluated and solved directly (as in MIRI, at longer wavelengths: Argyriou et al. 2023), but this is at present not tractable for the BFE in HgCdTe detectors. The approach used in that work, while accurate and crucial to understanding the true physics behind the BFE, falls short of the key requirements in this project: that it be differentiable and computationally efficient. Instead, we build a hybrid model, using an NN to capture the BFE dynamics without explicitly solving the complex differential equation.

The non-linear ramp model transforms the static 2D oversampled illuminance pattern into a 3D time-evolved charge accumulation ramp. Its construction is informed by the work on modelling the MIRI BFE (Argyriou et al. 2023), aiming to predict an evolving pixel area describing how charge transports from the excited photo-electrons in the photosensitive region down into the individual depletion layers of each pixel.

The time evolution of the charge accumulation is addressed by using a recurrent architecture, adding charge to each pixel in a fixed number of small time-steps. The measured charge at each group-read is found by interpolating the fixed number of time steps to the correct time

stamp of each group read.

The evolving pixel areas are modelled with a dynamic filter (Brabandere et al. 2016), with each individual pixel in the 80×80 sub-array used for AMI observations using its own predicted unique convolutional kernel at each time step. Further complexities arise from the unique sensitivity of each pixel in this process (the flat-field). These variations result in either more or less charge accumulating in any one pixel, which then influences the resulting collecting area of both it and its neighbouring pixels. This mandates that both the inter- and intra-pixel sensitivity be modelled separately.

4.5.2.1 Neural Network Implementation & Architecture

The non-linear ramp employs a recurrent architecture to mirror the time evolution of the BFE. Provided with the oversampled illuminance I and a charge distribution q_t , it seeks to predict the charge distribution at the time step q_{t+1} .

Using a Convolutional Neural Network (CNN) applied to the current charge distribution q_t , a set of polynomial distortion coefficients are predicted for each individual pixel. These distortion coefficients describe how each individual oversampled illuminance pixel is distorted as it travels through the electric fields of the detector substrate. Starting with a 3×3 rectilinear coordinate grid over each output pixel the distortions coefficients are applied, producing a local spatial transformation. These distorted coordinates are then used to calculate the overlap fraction from the new position with the neighbouring detector pixels. These overlap fractions then form the individual weights of the dynamic per-pixel kernels used to transport charge from the illuminance down to the detector pixels.

This construction is very important to ensuring physically constrained predictions by the CNN, since the overlap fractions from the illuminance pixels can only sum to one — directly enforcing flux conservation from the input illuminance. By wrapping the NN outputs in a kernel model that enforces these physical constraints the NN has greatly favourable properties to the system as a whole, easing the calibration process of both it and the wider forward model.

Inter- and intra-pixel sensitivity variations are also simple to include in this architecture. Since the model predicts where excited photo-electrons will be measured by the sensor, the quantum efficiency of this location can be baked directly into the kernels through a multiplication. A unique sensitivity value is used for each pixel in the detector, i.e. the flat-field. Intra-pixel sensitivity variations are parameterised by a simple quadratic function, using the same value for all pixels. Figure 4.9 shows the recovered intra- and inter-pixel sensitivity variations. These flat-field values are only calibrated using the publicly available in-flight flat-fielding data. This is a necessity in this process, as even though the AMI PSF covers many more pixels than clear aperture PSFs, many pixels in the AMI calibration data have little to no incident flux.

Using the dynamic per-pixel kernels that include pixel sensitivity variations, a dynamic convolution with the illuminance array I is performed to transport the charge from the detector surface down to the individual pixels. This output charge distribution can then be directly added to the present charge distribution q_t in order to predict the next charge state q_{t+1} and the processes repeated in order to predict the time evolving charge distribution in each pixel.

This recurrent processes is visualised in Figure 4.8, showing how the dynamical kernels are generated and applied. Figure 4.9 shows the various steps in the process. It presents the resulting pixel distortions and dynamic kernels as produced from an example input charge realisation. It also shows the final effect on the final measured PSF and how the charge bleeds between neighbours around the brightest pixels in the detector.

The embedded CNN in the non-linear ramp model is small with a simple architecture. Given that the charge bleeding between pixels is dominated by the electric field between each pixel, the input charge array is first convolved with 16 different spatial gradient kernels which are then fed to the CNN. These kernels are kept static and not trained. The CNN employs a total of 3 convolutional layers of width 16, followed by a final layer of width 12, and relu activations. All layers use a 3×3 kernel size and do not use bias terms, giving a total of 8150 parameters — very small in the context of modern deep-learning. This structure is shown in Figure 4.8. The choice to avoid biases ensures that the NN returns zero migration for an input array with zero charge, matching the inductive bias of our problem since zero charge results in zero electric field gradients and no charge bleeding. This further aids training as the CNN does not need to learn that no charge results in no bleeding. The 12 output channels of the final layer are the 2D polynomial distortion coefficients which are used to generate the normalised convolution kernels used to transport charge into each pixel. This hybrid model approach demonstrates the power that is gained by offloading the known physics to an encompassing forwards model that is able to provide both high-quality inputs and later transformations to outputs, while leaving the complex or unknown physics to a learned component.

4.5.3 Electronics & Read model

The non-linear ramp model is able to produce highly-accurate charge evolution prediction in each pixel, but there remains one last set of transformations that are applied, arising from the electronic devices that measure the voltage in each pixel. This model handles four primary effects:

1. Dark current
2. Inter-Pixel Capacitance (IPC)
3. Non-linear gain
4. Amplifier noise

4.5.3.1 Dark Current

Dark current effects are quite simple to implement. Individual contacts for each pixel heat up as they measure voltage and consequently can emit photons from the back of the detector. In practice the probability of this event varies from pixel to pixel, however AMIGO uses a constant value for all pixels both for simplicity and to avoid the including dark-current data in the calibration processes. This is desirable as dark-current calibration requires large files with

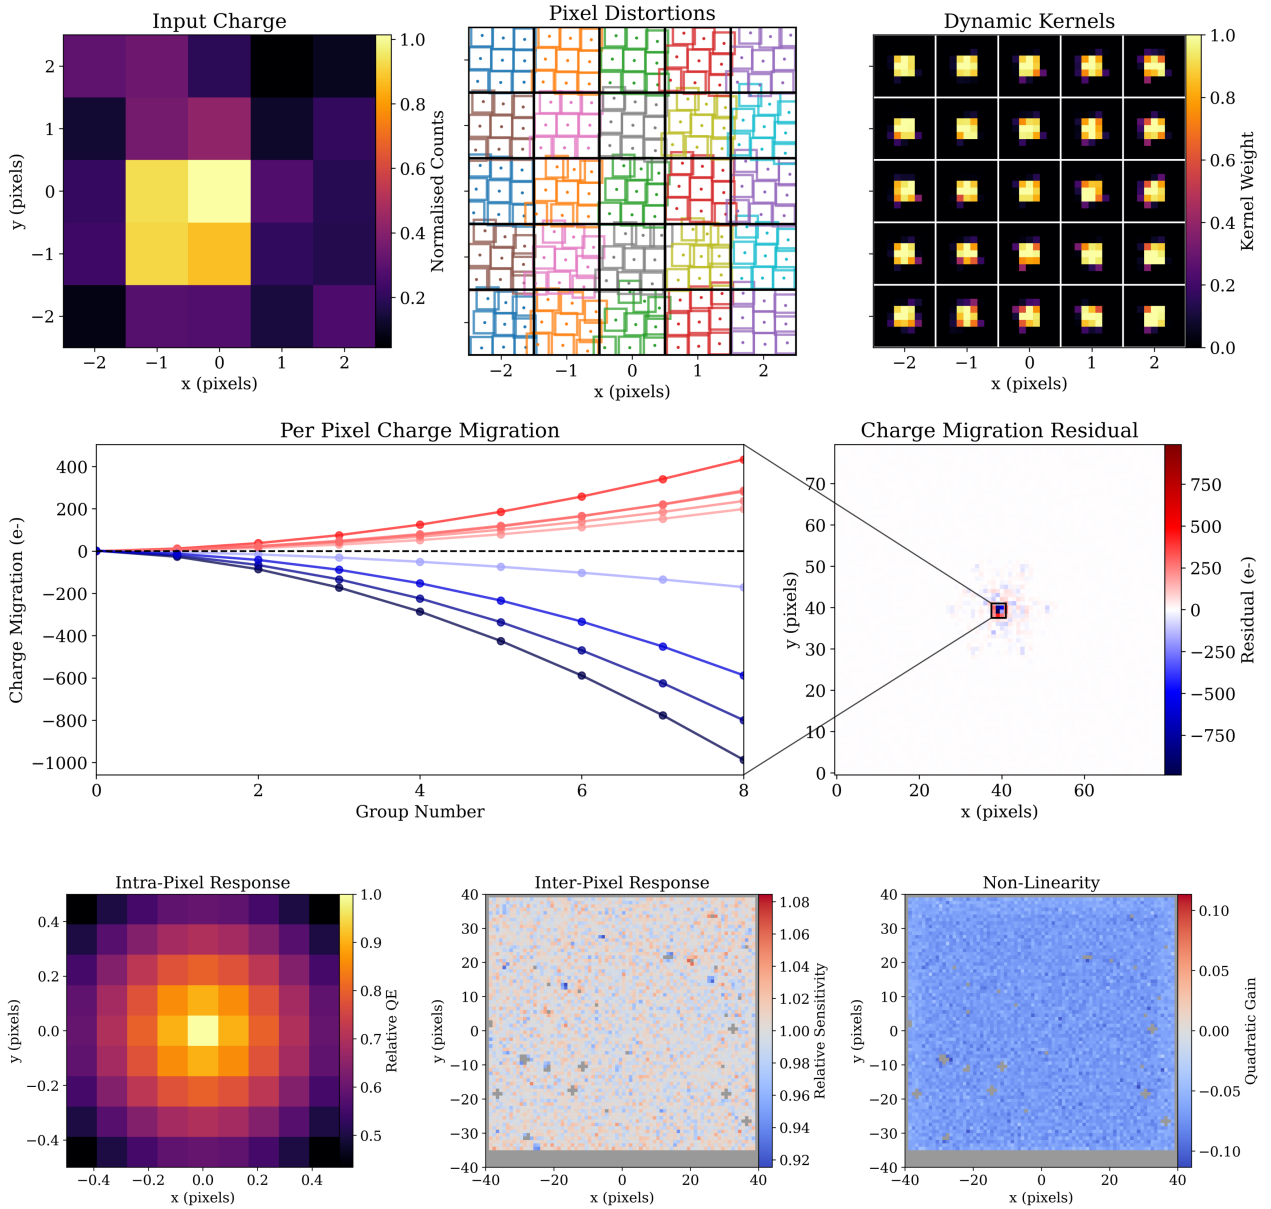


Figure 4.9: Demonstration of the EDM BFE model and recovered detector parameters. Starting with the normalised input charge distribution (top left), the CNN predicts a series of distortion coefficients that are applied to a $3\times$ oversampled set of coordinates for each pixel. The distorted output pixel positions are visualised in the top middle panel. These output positions have their overlap fractions with each neighbouring pixel calculated in order to produce a set of flux conservative convolutional kernel for each pixel, shown in the top right panel. This dynamic convolution is applied to the predicted input illuminance pattern for a fixed number of time steps. The middle left panel shows the charge migration between a set of pixels around the brightest region of the PSF as charge is accumulated up the ramp. The middle right panel shows the cumulative charge migration between pixels for each individual pixel. The bottom row shows the various recovered pixel-level effects. The bottom left panel shows the recovered intra-pixel sensitivity variations, applied through a simple quadratic function. The AMIGO model only uses a $3\times$ oversample, but is visualised here with $9\times$ oversample. The bottom middle panel shows the inter-pixel sensitivity variations, ie the flat-field. The resulting distribution of sensitivities shows similar properties to the existing JWST pipeline calibrations. The bottom right panel shows the recovered quadratic non-linear gain term applied through the electronics model. Every pixel shows the expected negative gain as a function of pixel depth, in line with existing calibrations.

many group reads in order to get sufficient signal on the singular photons emitted. AMIGO applies this dark current additively to each group read and a final recovered value of $\sim 0.45 e^-$ is found.

4.5.3.2 Inter-Pixel Capacitance

IPC is a generally well understood process in detectors — capacitive coupling between pixels results in the charge of a pixel influencing the measured value of its neighbours. IPC is conventionally modelled as a convolution; AMIGO employs this by convolving each group read with a static kernel. This is the only step where AMIGO uses ground-based calibration data products as the IPC is expected to be highly covariant with the charge migration effects of the BFE and therefore difficult to isolate. Using the ground-based calibration eases the training of the NN embedded in the EDM. The 5×5 IPC kernel, primarily composed of coupling to the directly adjacent neighbours, used is taken from the [calibration-reference data system](#) (CRDS).

4.5.3.3 Non-Linear Gain

NIRISS’ H2RG detectors also exhibit a non-linear gain. This is modelled with a unique per-pixel quadratic polynomial response to the input charge. Typically, a much higher order polynomial is used (4th or 5th order), however higher order effects are small until pixel charges exceeds half their full well depth, a regime not recommended for AMI observations and that we have no calibration data for. As each pixel has a set of unique coefficients, it is only calibrated from the in-flight flat-fielding data, the same processes used for the intra-pixel sensitivity variations described in Section 4.5.2. This helps to avoid any biases introduced by uncalibrated optical or charge-migration systematics. The recovered per-pixel quadratic term is presented in Figure 4.9.

4.5.3.4 $1/f$ Noise

$1/f$ noise, or red noise, is the thermal drift of the amplifiers as they read along the pixel columns (Rauscher et al. 2017). As the amplifier temperature drifts over time, the measured voltage drifts in kind, producing a visible vertical striping in recovered images. The magnitude of $1/f$ noise is larger for observations with fewer integrations and groups as the effect averages out to zero with an increasing number of reads. AMIGO has the ability to model $1/f$ noise with a low order polynomial added to each column for each group, however in practice this is unnecessary and usually skipped.

4.6 Base Model Training and Calibration

The AMIGO model is calibrated using gradient-based optimisation, in which all model parameters (including both physical parameters and the embedded NN weights) are jointly fit. The validation datasets are simultaneously fitted (but not used for model calibration) to ensure model generalisation; these datasets are detailed in Section 4.3.1. The model is implemented

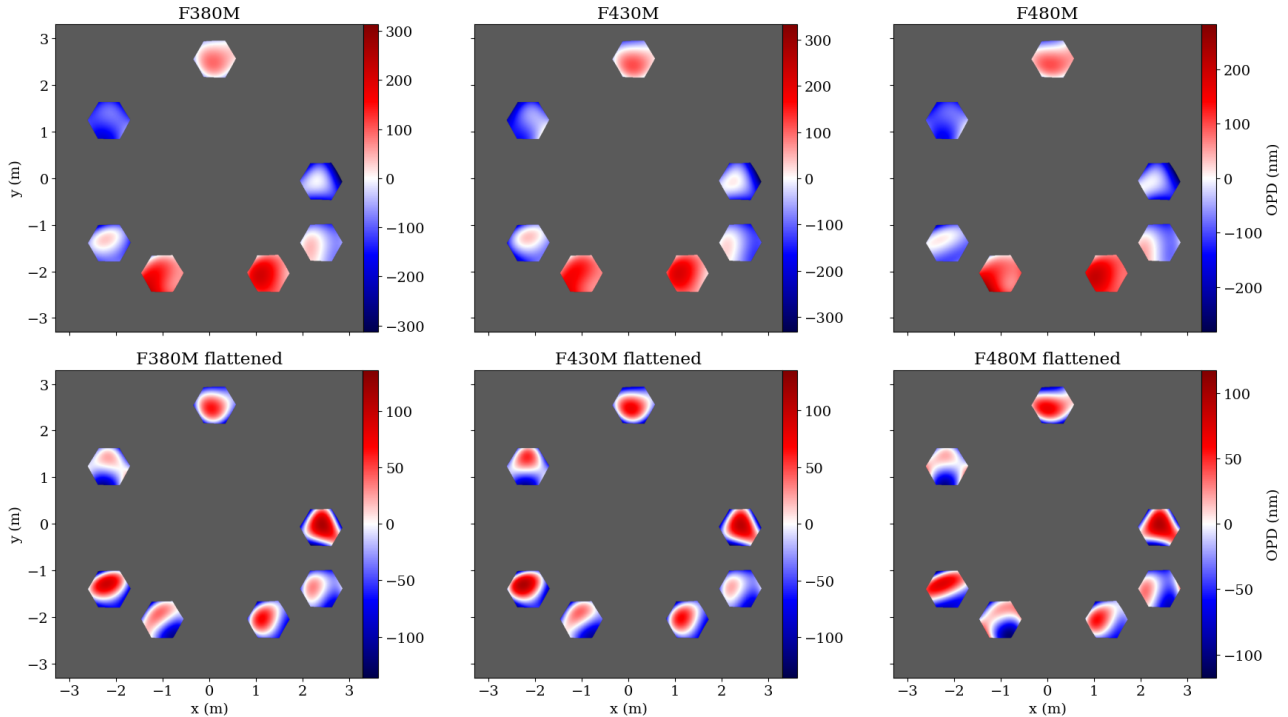


Figure 4.10: Recovered OPD maps from the calibration data for each filter. The top row shows the full OPD maps recovered for the F380M, F430M, and F480M filters, revealing large-scale piston, tip, and tilt terms across the segmented aperture. These low-order aberrations are expected given the off-axis subarray placement used in AMI mode. The bottom row displays the same OPD maps with piston, tip, and tilt removed from each sub-aperture, enhancing the visibility of higher-order surface errors across individual segments. Residual aberrations at the $\sim 50\text{--}100$ nm level are evident, and are consistent across wavelengths, suggesting a static contribution from optical surface figure errors or segment alignment offsets. Small differences are observed across filters, consistent with the observed Fresnel defocus discussed in Section 4.4.1.

in a fully differentiable framework, enabling end-to-end training of all components. This approach allows the full model — including optical system and EDM — to act as a single coherent system, representative of the true chain of physics leading to any observation.

Native gradient descent methods can be troublesome for data that has not been normalised, particularly when different datasets have different levels of total signal. Rather than normalising the data, AMIGO circumvents this issue by using a natural gradient descent approach (Martens 2020) which involves approximating the Fisher matrix of the parameters under the Laplace approximation (Kass et al. 1991; MacKay 2002). This enables the model *gradients* to be normalised, as opposed to the data. Fisher matrices are a natural tool for AMIGO as they can be efficiently calculated with autodiff for each dataset, added in order to reflect hierarchically constrained parameters, and approximated through Jacobians. The use of natural gradient descent speeds up convergence and transforms most parameter learning rates to order \sim unity, all while preserving the Bayesian weighting between disparate datasets.

Each trained non-NN parameter uses a momentum-based optimiser and is assigned a unique learning rate. Complex visibilities are not fit during training. This is a deliberate choice: visibilities are highly covariant with optical parameters and “soak up” residual systematics.

All model calibration is done without ever transforming to the uv -plane.

The NN component is optimised with the Adam optimiser (Kingma and Ba 2014). Training is performed in batches to enable stochastic updates and efficient GPU usage. The calibration data was split into five batches, and their order randomised each epoch. Validation data is not batched because it does not influence the calibration of the model and stochastic behaviour has no effect. The NN is strongly covariant with physical model components, particularly during early stages of training where it may compensate for optical/detector effects. To mitigate this, the model is retrained multiple times from different initial weights while preserving the best-fit instrumental parameters. This staged optimisation process improves convergence and helps to separate the roles the NN and physical parameters play in capturing systematics.

Detector-specific parameters that are *unique* per-pixel are fit in parallel with, but independently from, the rest of the model. These components are trained exclusively on flat-field calibration data. This avoids over-fitting arising from any miscalibration elsewhere in the model, as well as ensuring that each pixel has approximately the same total signal. The resulting flat-field maps as shown in Figure 4.9 exhibit smooth variations consistent with standard JWST pipeline calibrations, while the per-pixel non-linear gain terms match expected trends with signal depth.

All model training was performed on a NVIDIA RTX 4080 GPU, with convergence reached after $\sim 20,000$ steps across multiple re-trainings. Figure 4.10 presents the recovered wavefront state after the training process from the calibration data. There are only small differences between filters, which is expected given the different defocus values found in each filter (see Section 4.4.1.4). The resulting model shows a statistically good fit to the calibration data. Figure 4.11 shows a summary of the fits to the F430M filter for all five subpixel dither positions. Little to no PSF structure is present in the residuals. The residuals are close to being distributed as a unit-normal, indicating neither an under- nor over-fit.

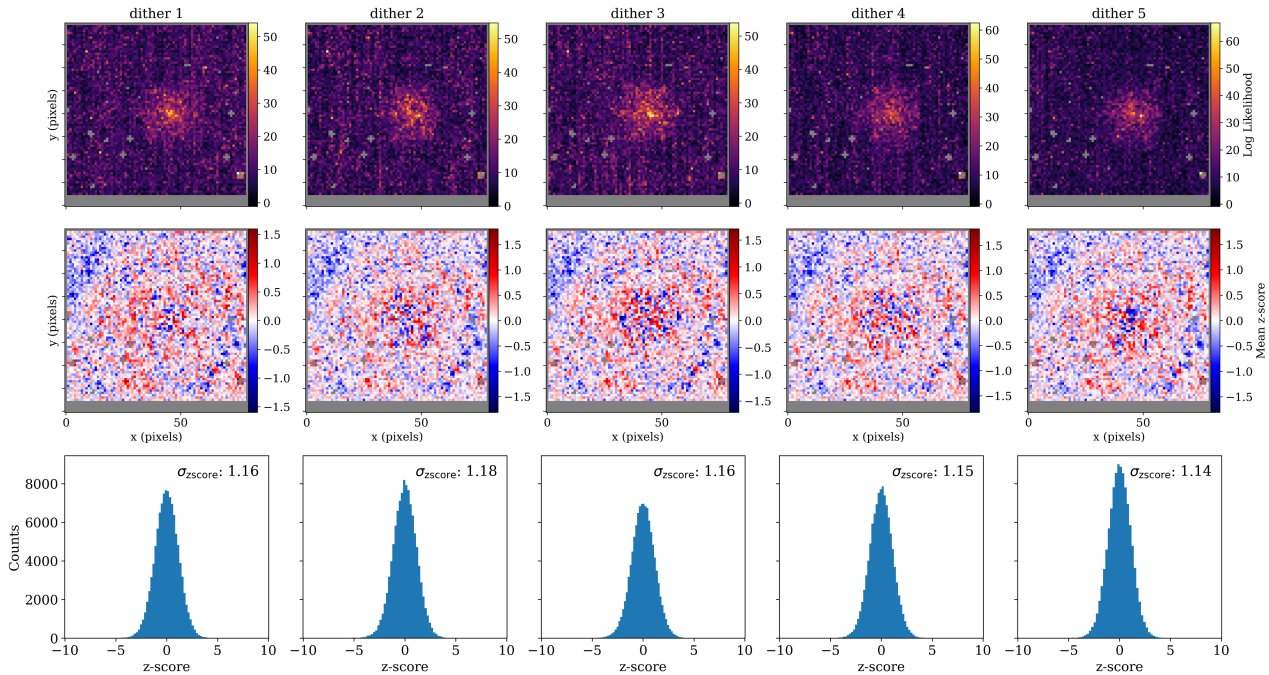


Figure 4.11: Summary of AMIGO model fit diagnostics across all five dithers for the F430M calibration data. Each column corresponds to a single dither position. The top row shows the per-pixel log-likelihoods from the final fit, highlighting the location of the target PSF. The middle row displays the average residual z-score per pixel, computed by averaging the uncertainty-normalised residuals across all groups in the ramp, revealing the spatial structure of any systematic model misfit. The bottom row shows histograms of all z-scores across pixels and groups for each dither, without any averaging over the groups. A perfect fit would have a standard deviation in the z-score be 1; we recover values between 1.1-1.2 in all three filters, indicating a good fit that has not learnt any noise characteristics. We note that the full likelihood is described with a covariance matrix that accounts for the anti-correlation between adjacent group-reads seen in slope data. Consequently, these summary statistics are only an helpful approximation and correct performance can only be described through the likelihood.

4.7 Inference from Science Data

The AMIGO framework is built around a fully differentiable forward model, which enables efficient gradient-based inference on both astrophysical and instrumental parameters as a digital twin of AMI. First we build this twin from high-quality calibration data; then, this same *base model* can be used in Bayesian inference by holding most parameters fixed and only fitting parameters of astrophysical interest like wavefronts, spectra, and source intensity distributions by gradient descent or HMC. A serialised version of this trained base model is included in the AMIGO repository and does not need to be retrained by users.

In order to make this efficient for the end user as a black box, AMIGO caches large Jacobians and estimates per-dataset Fisher matrices which are used to normalise gradients, implementing natural gradient descent that improves convergence speed and robustness across diverse observations. For situations where further stability or precision is required, AMIGO also provides access to second-order optimisation routines and the ability to control learning rates over

each recovered parameter. Users can choose to infer parameters either jointly or independently across multiple exposures — for instance, sharing optical aberrations across a sequence of observations while allowing astrophysical parameters to vary — and disentangling instrumental and scientific effects over multiple epochs.

By default, AMIGO only recovers a focused set of parameters: source positions, brightnesses, spectra, optical aberrations, and complex interferometric visibilities. Calibration parameters, such as pixel sensitivities, detector non-linearities, or the weights of the neural network within the EDM, fixed when fitting science data. However, thanks to its modular and generative structure, AMIGO allows users to selectively activate inference on any of these parameters when needed.

For example, as discussed in Section 4.5.3, AMIGO can model $1/f$ detector noise, but by default does not. Another example involves the [documented physical alignment of the NIRISS aperture mask](#). Small differences in the mask rotation between calibration and science observations, especially when taken across different epochs, can introduce characteristic residuals. These are directly visible in the pixel-level residuals of the model fit. When present, AMIGO can infer a unique mask rotation angle per exposure, allowing this systematic effect to be characterised and marginalised as part of the forward model.

In general, as a forwards model that infers parameter via Bayesian statistics, AMIGO is capable of placing priors over any set of parameters during its inference workflow. This can be achieved via a simple modifications of the likelihood functions used for optimisation through the inclusion of a prior over any model parameters. This also applies to model-fitting to the extracted visibilities — AMIGO is an end-to-end forwards model, enabling Bayesian methods to be applied widely to any component of the system or its outputs.

4.7.1 Observable Extraction Workflow

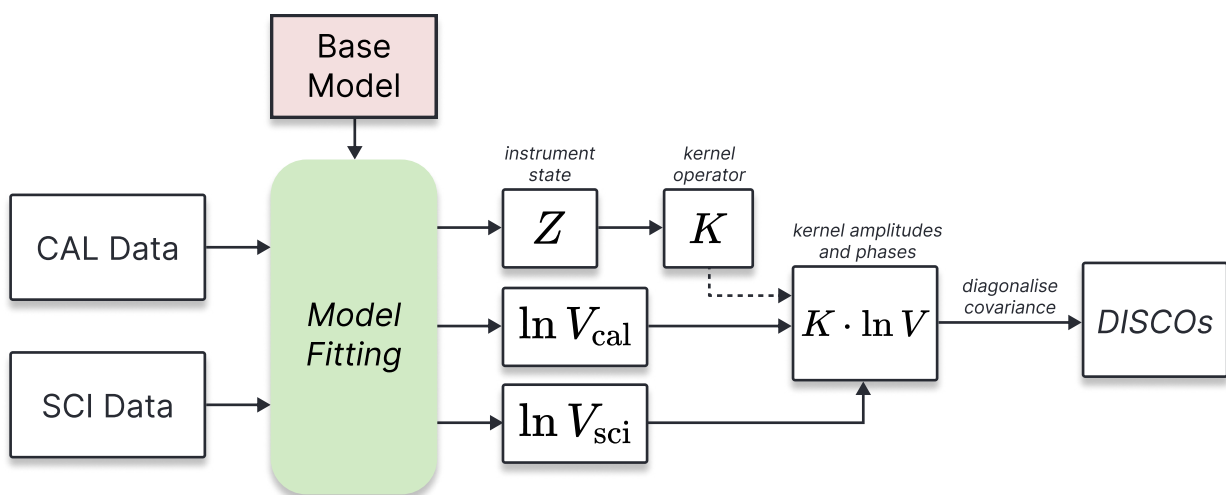


Figure 4.12: Schematic diagram of the basic workflow of extracting the DISCOs from data.

The AMIGO pipeline is designed to extract the DISCO observables discussed in Section 4.4.2.3,

starting from uncalibrated JWST data. Here we detail the basic workflow used for this process, comprised of three steps:

1. Data processing
2. Model fitting
3. Visibility reduction

Each of these steps are detailed in the following sections, with Figure 4.12 detailing the overall workflow from data to DISCOs. As a forward-modelling pipeline, AMIGO operates differently to most existing tools. While default operation with little user input should suffice for the majority of use cases, the model provides great flexibility over how observables are extracted and is able to directly solve for multiple non-standard effects when they influence the data.

Despite modelling the full system, AMIGO still requires a PSF reference star in order to provide high-quality calibrated observables. As such, standard interferometric observing programs are recommended following the [recommendation in JDox](#). Note that a further set of recommendations are provided in Section 4.9 designed to mitigate any potential model miscalibrations.

4.7.1.1 Data processing

The initial data processing step is to reduce the 4D `uncal` JWST data to the `calslope` data produced used by AMIGO as described in Section 4.3.2. This is an automated procedure, and unlike the standard JWST data processing pipelines does not require any external calibration files. Users are able to chose the $N\sigma$ threshold used for outlier rejection, but defaults to 3σ . Default operation will also apply the ADC correction as described in Section 4.3.2, which should always be performed. Both the PSF reference and target stars should use the same data processing configuration. A [notebook detailing this process](#) is provided.

4.7.2 Model Fitting

The model fitting step is the most important to the recovery of high-quality observables. As a forwards model of the instrument, the AMIGO pipeline can fit a number of different source (i.e. astrophysical) models to any provided `calslope` data, such as a point source or visibility model. Standard operation for interferometric observable recovery should use the `SplineVisFit` class on both the PSF reference and target stars, which take the `calslope` file as it input and produces an `Exposure` object. The `Exposure` class operates on the AMIGO model in order to produce a predicted observation, and we maximise the log-likelihood of this model with respect to data. Flags in this class control which parameters in the base model are fixed and which are recovered from the data, for example the recovery of the $1/f$ stripe signal.

As discussed in the previous section, under default settings we recover source position, flux, spectral slope, visibilities, and wavefront phase. The AMIGO pipeline will not recover the wavefront phases from non-PSF reference stars, as the wavefront phase and interferometric signal are strongly covariant. The wavefront phase is recovered only from the PSF reference stars,

and fixed in inferring visibilities from the science target and in generating the kernel observable operator. Users have full control over what is recovered from data, and for example can infer more complex spectral models; or different wavefront states can be used for observations taken in different epochs.

With these model fit objects, users now need to construct the AMIGO base model used by the fit to predict observations by loading a pre-calculated calibration file (produced as described in Section 4.6), which is released for each version of AMIGO.

With the calibrated base model and the model fit objects the observables can be recovered from the data. This is done via natural gradient descent, minimising the log-likelihood, using the Fisher matrix (which can be constructed approximately from pre-calculated Jacobians) to normalise each parameter of the model so that we require only one gradient descent hyperparameter. Model fits to both the target and PSF reference stars must be done at the same time to ensure accurate wavefront phases are used on the target stars. These can be done in batches to reduce both computational and memory overheads. The model should be optimised for ~ 200 epochs for good convergence, however this can vary depending on the specifics of the input data. There is no fixed termination criterion, and convergence must be determined by the user. In some cases tweaks to the parameter learning rates is required. While it is currently experimental, AMIGO also provides an interface to the BFGS optimiser (Broyden 1970; Fletcher 1970; Goldfarb 1970; Shanno 1970) as implemented in OPTIMISTIX (Rader et al. 2024) which provides improved convergence guarantees.

With the fitted model, the parameter uncertainty covariance matrices are estimated via the Hessian of the log-likelihood evaluated at the best-fit parameters. For computational efficiency this is only done over the recovered visibility parameters. These parameters and their uncertainties are then saved to an intermediate data product which is next used for interferometric observable reduction and calibration. A [notebook detailing this process](#) is provided.

4.7.3 Visibility reduction

The next step projects the recovered observables to the kernel space and calibrates them. The recovered wavefront state is loaded into the AMIGO model and we calculate the Jacobian of the visibilities with respect to not just pupil Zernike phase modes but also flux, spectral slope, and Fresnel defocus. A singular value decomposition is performed on the Jacobian and used to construct a kernel basis and used to project the complex visibilities to kernel observables. This projection is done on both the PSF reference and target stars, where the target star visibilities are calibrated by a subtraction of the reference star visibilities. This operation is subtraction, not a division, as AMIGO recovers the log complex visibilities (see Section 4.4.2.1). The same projection is performed on the parameter covariance matrices to propagate uncertainties correctly.

The final step in this procedure is the projection to the statistically-independent DISCO space, as detailed in Section 4.4.2.3. This is found via an eigendecomposition of the calibrated kernel visibility uncertainty covariance matrix, where the eigenvectors are used to project both the visibilities and covariance matrices into a statistically-independent space (Ireland 2013),

yielding the DISCO observables. Finally, a `calvis` data product is produced that reduces all of the remaining relevant information needed for downstream interferometric analysis. It provides the uv coordinates of the observation, mean wavelength, a single DISCO matrix that maps from the pixelised uv coordinates back to the DISCO space, as well as all the intermediate data products. A [notebook detailing this process](#) is provided. This is considered the end of the basic AMIGO pipeline, however AMIGO also provides a number of classes and interfaces useful for fitting interferometric models to the DISCOs, although they are experimental and not yet fully supported for users.

4.8 Results

Table 4.3: Summary of COM1093 program observations used to test the AMIGO model. The number of photons is an estimation from the raw data, and details the number of usable photons after accounting for the 1/groups fractional loss from discarding the first group of an integrations. The percentage loss of photons is also shown.

Observation	Type	Star	Filter	Groups	Integrations	Photons	Loss (%)
12	SCI	AB Dor	F480M	5	69	92×10^6	20.0%
12	SCI	AB Dor	F430M	4	82	94×10^6	25.0%
12	SCI	AB Dor	F380M	2	160	97×10^6	50.0%
15	CAL	HD 36805	F480M	12	61	85×10^6	8.3%
15	CAL	HD 36805	F430M	9	78	104×10^6	11.1%
15	CAL	HD 36805	F380M	4	118	100×10^6	25.0%

This section offers a quantitative assessment of the AMIGO model’s performance on both medium and high contrast interferometric imaging data from JWST NIRISS AMI mode. Rather than a reduction to calibrated visibilities and closure phases, AMIGO infers DISCOs (described in Section 4.4.2). These latent observables are robust to both low- and moderate-order wave-front miscalibration and capture the complex visibilities over the full OTF and sources across the FOV. As such, direct comparisons using conventional metrics like closure phase scatter are not applicable. Instead, performance is evaluated empirically via the recovery of known companions in representative and publicly available archival data. Two observing programs are used for this comparison: COM1093 (PI: Thatte), a preliminary exploration of performance on the AB Dor system during the commissioning phase, and GO 1843 (PI: Kammerer), a deep observation pushing the limits of the instrument seeking to characterise the red substellar companions around HD 206893.

The theoretical performance of an AMI system is commonly benchmarked using the expected contrast floor for detecting a point source near a host star. This is determined by the closure phase uncertainty σ_{CP} (Ireland 2013)

$$\sigma_{\text{CP}} = \frac{N_h}{N_p} \sqrt{1.5N_p} \quad (4.10)$$

where N_p is the number of photons collected and N_h is the number of holes in the NRM. In

the high-contrast regime σ_{CP} is used as an approximation to the $1\text{-}\sigma$ per-observable detection limit in *contrast* of a companion (Sivaramakrishnan et al. 2023). For a 7-hole mask like AMI this provides an approximate contrast limit cited in JDox as a recommendation for exposure time calculations

$$\text{contrast limit} \approx \sqrt{100/N_p}. \quad (4.11)$$

This value does not directly translate to a principled ‘ $N\sigma$ detection’ limit as it is not cognisant of mask or signal geometry, nor analysis methodology; it also does not appear in this form in the Ireland (2013) paper. Nevertheless as shown below, this rule of thumb for the contrast limit proves to be accurate.

4.8.1 Interferometric Model Fitting

Model fits to the calibrated observables are performed with the software DRPANGLOSS, an under-development JAX accelerated interferometric analysis package (Blakely et al. 2024). While the final output product of the official AMIGO pipeline are the DISCOs, AMIGO also provides basic interfaces for integration with DRPANGLOSS given its natural synergy as JAX + ZODIAX based software and to facilitate seamless analysis of its novel output products.

To fit interferometric models to the data, first the calibrated DISCOs are loaded into the `AmigoOIData` class which holds the recovered observables, the uncertainties, the uv coordinates, and the DISCO matrix that maps between the uv plane and the DISCO space. There are two main procedures used to fit high-contrast models to the DISCOs, which are used in both of the observing programs analysed in this manuscript. To identify the existence of any companions, a grid-search is performed over a 1 arcsecond radius of the primary star. This grid search is done by solving for the best-fitting contrast at each position in the grid by minimising the log-likelihood of the predicted binary via a BFGS algorithm. This fit is performed via a vectorised operation in JAX, providing orders of magnitude speed up over comparable software like CANDID (Gallenne et al. 2015), taking ~ 20 seconds in each filter on an NVIDIA RTX 4080 Graphics Processing Unit (GPU). This produces a log-likelihood detection map that can be used to identify companions through the maximum likelihood estimate. Both the best fit companion position and contrast at the maximum likelihood estimate can then be used to infer its properties via an Markov Chain Monte Carlo (MCMC) sampler: we use HMC as implemented in the probabilistic programming language NUMPYRO. HMC achieves much better performance on high-dimensional problems than other MCMC samplers by taking advantage of derivative information.

For companion astrometry, we adopt a uniform prior over position in RA and Dec with width 150 mas, and a uniform prior over log-contrast with width 2 orders of magnitude, in both cases where the means of the priors are determined by the best-fit values found from the grid search. The HMC chain also fits a multiplicative error term σ_{scale} to both the DISCO amplitude and phase uncertainties. This has the effect of pushing the χ^2_ν to ~ 1 during the MCMC sampling. This operates under the assumption that the model being fit (a point source in this case) is accurate, allows for the sampler to operate more efficiently, and provides an understanding over

the quality of the resulting fit, the uncertainties of the recovered parameters, and the quality of the estimated errors on the data. As our uncertainties are derived from propagating photon noise uncertainty through the pipeline, σ_{scale} is approximately the ratio of our data scatter to the ideal photon noise limit. This process is performed both jointly across all filters, and uniquely per filter.

In the case of multiple companions, as with the GO 1843 program, the best-fit primary companion found via MCMC can be subtracted from the data. Then the grid-search can be repeated in order to find the position and contrast of any other companions. Next the MCMC fit is performed simultaneously using a joint model of both companions in order to recover precise parameters for both companions.

4.8.2 COM 1093: Commissioning Data — AB Doradus AC

The AB Dor system was the first observed by AMI as part of the COM 1093 program during JWST's commissioning phase. This program aimed to demonstrate the detection of point source companions at moderate contrasts and test target placement precision through primary and sub-pixel dithering operations. Preliminary analysis of this data produced results inconsistent with pre-flight estimates (Sivaramakrishnan et al. 2023). The recovered contrast limits fell short of expectations by ~ 0.5 -1 magnitudes, making this program ideal to benchmark AMIGO.

To test AMIGO's performance on point-source companion recovery, we use a subset of the data in line with AMI's recommended observing strategy of single exposures with no dithers. Table 4.3 details the subset of exposures used for analysis.

Default operation of AMIGO takes likelihood statistics over *slopes* rather than ramps. This aids problems that arise from predicting the pixel-to-pixel bias level that can result in model over-fitting as discussed in Section 4.3.2. While this is not a fundamental problem and future releases of AMIGO plan to directly address this, presently AMIGO discards the first group read from an exposure resulting in a loss of data. Exposures with few groups are more affected as the first group forms a larger fraction of the overall data. As such, AMIGO as currently working operates on *less* overall signal than conventional pipelines. Improving use of low group count data is an important step in AMIGO development. Table 4.3 also presents the approximate photon counts in the COM 1093 dataset used for analysis, along with the fraction that AMIGO uses. Examination reveals that each filter has a total of $\lesssim 10^8$ photons each, which yields an estimated contrast floor of $\lesssim 7.5$ magnitudes from Equation 4.11.

4.8.2.1 COM 1093: Analysis

The raw `uncal` exposures were downloaded from Mikulski Archive for Space Telescopes (MAST) and processed into the `calslope` format using the default operation of the AMIGO data processing pipeline, described in Section 4.3.2. The calibrated AMIGO model, described in Section 4.3, was then fit to the processed data to recover the latent visibilities described in Section 4.4.2.1. The recovered wavefront phases were used to calculate the DISCO basis used for the output observables of the AMIGO model.

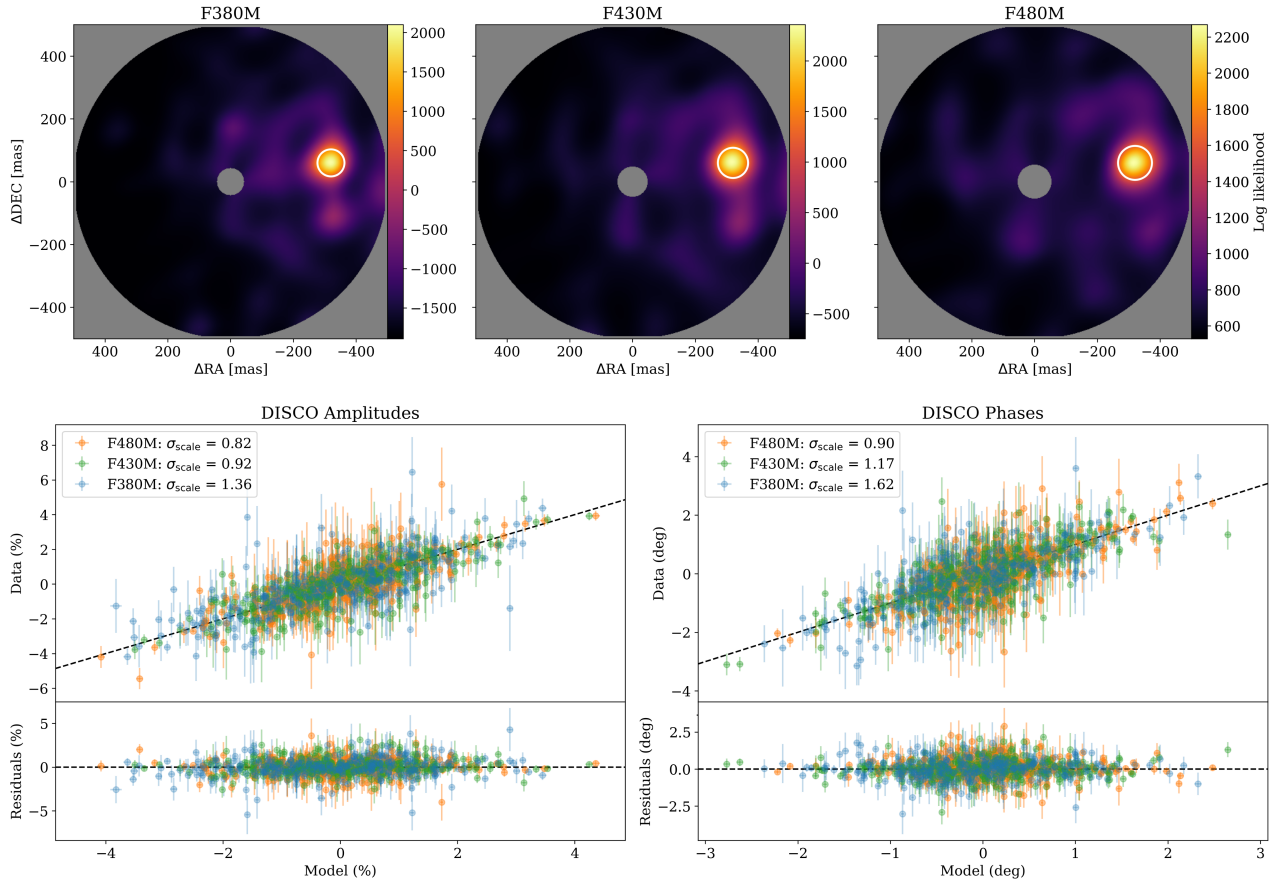


Figure 4.13: Summary of the fits to the AB Dor C companion. The top panels show the log-likelihood detection maps over companion position for each filter (F380M, F430M, F480M) from a grid search over the recovered DISCOs. Each panel shows the marginalised log-likelihood surface as a function of companion offset in ΔRA and ΔDEC , revealing a strong and consistent peak in all filters. The greyed central region denotes the IWA masked by the interferometric null. The recovered peak likelihood location is consistent across filters, indicating the reliability of the model prediction. The bottom panels show the predictive posterior check comparing AMIGO recovered DISCO amplitudes (left) and phases (right) against the predicted values from the MCMC samples, for the three filters: F380M, F430M, and F480M. Each panel shows individual measurements with 1σ error bars, with the top panels plotting model predictions against data and the lower panels showing residuals. The black dashed line represents the 1:1 line for a perfect model prediction. Scale factors σ_{scale} are applied to the observational uncertainties with the effect of ensuring $\chi^2_{\nu} \approx 1$ during MCMC fitting and are quoted in the legend for each filter. The agreement across all filters and observables confirms both the validity of the forward model and appropriate uncertainty quantification in the recovered posterior.

Figure 4.13 shows the recovered log-likelihood detection maps in each filter, calculated with a binary fit to contrast at each RA, Dec with the models provided by DRPANGLOSS. These show a confident and independent detection of the known companion in all three filters.

To recover posterior samples of the detected companion we run an HMC sampler using NUMPYRO (Phan et al. 2019b; Bingham et al. 2019) and DRPANGLOSS. While the companion is detected independently in all filters, we perform a broadband joint-fit to the astrometry of the companion to ensure the best constraints on both the astrometry and photometry. To assess

Table 4.4: Best-fit relative joint astrometry and per-filter photometry for the AB Dor C companion, showing the $\pm 1\sigma$ bounds. Reported quantities include the on-sky separation (Sep), position angle (PA), and Δ -magnitude in each of the F380M, F430M, and F480M filters. Results from two different fit types are compared: A joint fit in all three filters and a fit to each filter uniquely.

Fit Type	Sep (mas)	PA ($^\circ$)	Δ mag (F380M)	Δ mag (F430M)	Δ mag (F480M)
Joint	$328.85^{+0.48}_{-0.85}$	$169.482^{+0.020}_{-0.022}$	$4.313^{+0.013}_{-0.023}$	$4.354^{+0.016}_{-0.012}$	$4.381^{+0.017}_{-0.016}$
F380M	$329.9^{+1.3}_{-1.2}$	$169.521^{+0.045}_{-0.033}$	$4.301^{+0.021}_{-0.016}$	—	—
F430M	$328.58^{+0.98}_{-1.18}$	$169.468^{+0.033}_{-0.036}$	—	$4.357^{+0.014}_{-0.015}$	—
F480M	$327.5^{+1.3}_{-1.1}$	$169.447^{+0.041}_{-0.036}$	—	—	$4.377^{+0.018}_{-0.015}$

Table 4.5: Summary of GO 1843 program observations used to test the AMIGO model. The number of photons is an estimation from the raw data, and details the number of usable photons after accounting for the 1/groups fractional loss from discarding the first group of an integrations. The percentage loss of photons is also shown.

Obs	Type	Star	Filter	Groups	Integrations	Photons	Loss (%)
1	SCI	HD 206893	F480M	11	720	1.05×10^9	9.09%
1	SCI	HD 206893	F430M	8	2177	2.47×10^9	12.5%
1	SCI	HD 206893	F380M	4	7046	5.31×10^9	25.0%
2	CAL	HD 205827	F480M	10	641	0.86×10^9	10.0%
2	CAL	HD 205827	F430M	7	1885	2.15×10^9	14.3%
2	CAL	HD 205827	F380M	3	7800	4.74×10^9	33.3%

the quality of the fit and the recovered parameters we perform a predictive posterior check — visualising the correlation and residual between the recovered and predicted DISCOs found with the MCMC samples and shown in Figure 4.13. Multiplicative error inflation factor σ_{scale} is fit to account for and quantify uncertainty in the AMIGO estimated errors. These terms remain close to one, indicating both a good fit and approximately accurate uncertainty quantification. The full posterior samples are shown in Figure 4.14, and the resulting fit in Table 4.4.

To quantify the overall detection limits of both the model and the data the best-fit companion is subtracted from the recovered DISCOs, while applying the error jitter term σ_{scale} . Using these cleaned DISCO values we calculate the azimuthal Bayesian upper limits following the methods introduced by Ruffio et al. (2018), shown in Figure 4.15. This demonstrates that the AMIGO model performs at the expected contrast limits of the instrument, indicating it is well calibrated and free of the detector systematics that cause issues for AMI analysis pipelines, up to the total signal level provided by this data.

4.8.3 GO 1843: HD 206893 at High Contrast

To demonstrate the performance of AMIGO on a deep AMI dataset, we analyse archival images of the HD 206893 exoplanetary system at three wavelengths (F380M, F430M, F480M). These images are part of GO 1843 program with the ultimate goal of detecting the brown dwarf

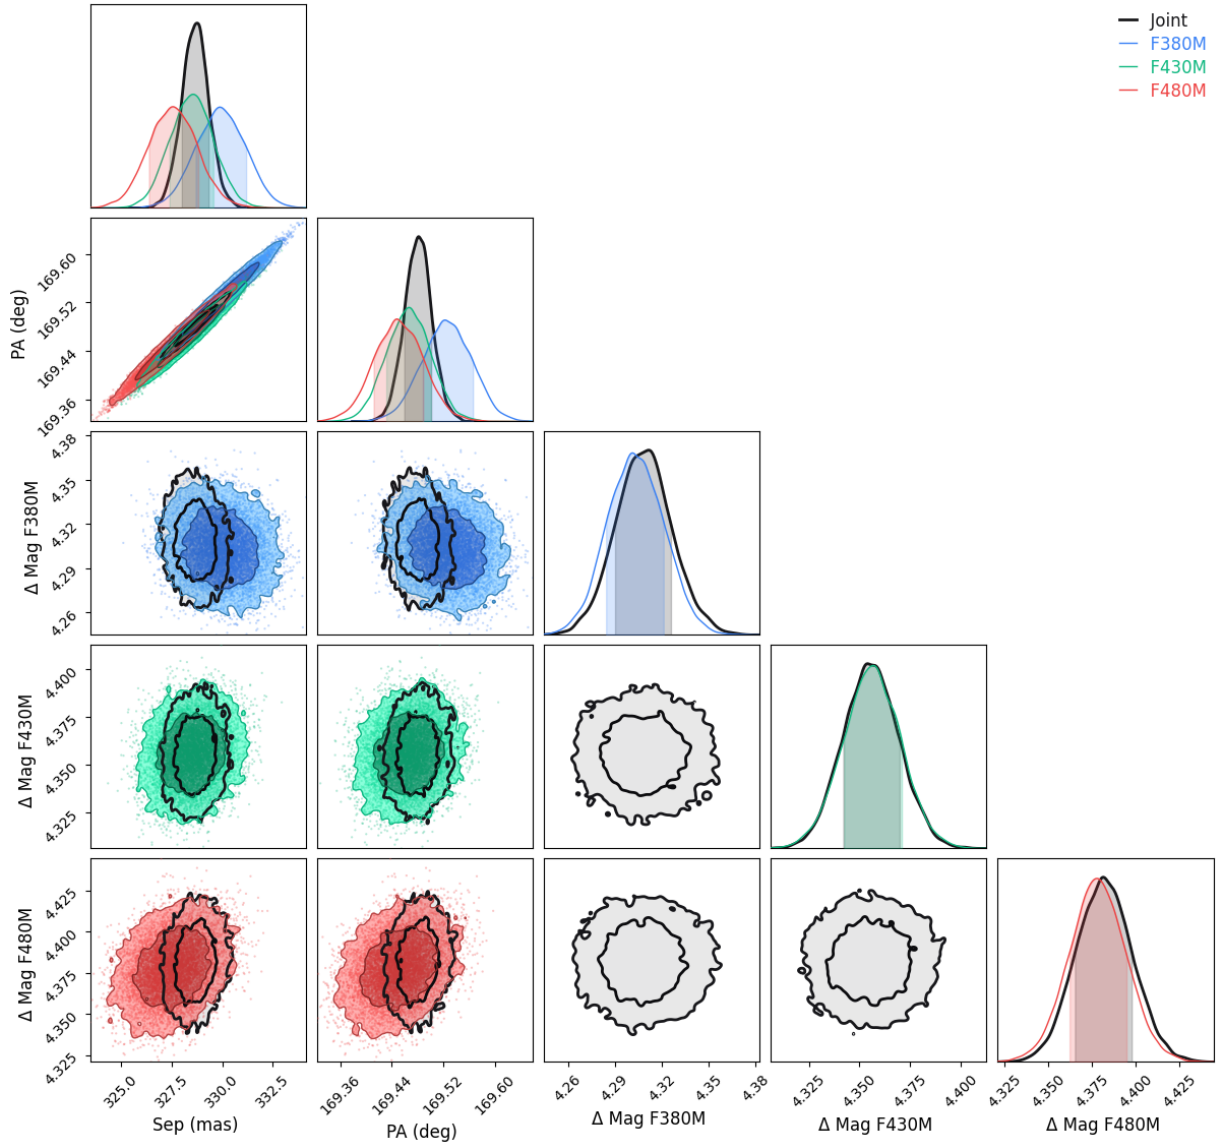


Figure 4.14: Joint posterior distributions from an MCMC fit to the AB Dor C companion. Two fit types are shown: Joint — simultaneously modelling astrometry and photometry across all three filters — and per-filter fits. The parameters shown include the separation (mas), position angle (degrees), and contrasts (Δmag) in the **F380M**, **F430M**, and **F480M** bands. The joint-fit samples are shown in black. One- and two- σ credible regions are shown in dark and light shades, respectively. Strong correlation is observed between separation and position angle due to the projection geometry, while photometric parameters are weakly correlated and independently constrained in each filter. The tight constraints in both astrometry and photometry reflect the high signal-to-noise of the companion detection.

HD 206893 B and planet HD 206893 c at angular separations of ~ 200 and ~ 100 mas, respectively. The three NIRISS AMI filters chosen for this program probe the presence of CO and CO₂ at 3–5 μm wavelengths. Together with the data at shorter wavelength measured from the ground, this constrains the carbon and oxygen chemistry in the atmospheres of the two substellar companions and provides insights into their formation history. Additionally, photometric constraints at the AMI wavelengths can constrain the origin of the extremely red near-infrared colour of HD 206893 B, which is also the reddest known substellar companion.




Figure 4.15: Sensitivity limits derived using the Ruffio et al. (2018) method for calculating 3σ contrast upper limits as a function of angular separation, applied to the F380M, F430M, and F480M filters. The shaded regions denote the $\pm 1\sigma$ variation across azimuthal annuli. The dashed black line indicates the approximate contrast limit as calculated from Equation 4.11, confirming that model performance is consistent with the expected limits. The grey dotted line shows the size of a single pixel for reference. The recovered companion 3σ contour plots are shown in their respective colours but are so well constrained they only appear as single dots on the figure. These curves quantify the AMI contrast performance and establish detection limits for additional companions in the field.

Table 4.6: Best-fit relative joint astrometry and per-filter photometry for the HD 206893 B companion, showing the $\pm 1\sigma$ bounds. Reported quantities include the on-sky separation (Sep), position angle (PA), and Δ -magnitude in each of the F380M, F430M, and F480M filters. Results from two different fit types are compared: A joint fit in all three filters and a fit to each filter uniquely.

Fit Type	Sep (mas)	PA ($^\circ$)	Δ mag (F380M)	Δ mag (F430M)	Δ mag (F480M)
Joint	$197.1^{+4.3}_{-4.1}$	$102.3^{+1.2}_{-1.6}$	$8.33^{+0.18}_{-0.16}$	7.98 ± 0.15	$7.49^{+0.11}_{-0.12}$
F380M	199.3 ± 6.2	$101.9^{+2.7}_{-2.6}$	$8.33^{+0.21}_{-0.14}$	—	—
F430M	$192.1^{+9.4}_{-8.4}$	$99.4^{+2.4}_{-3.5}$	—	$7.96^{+0.16}_{-0.14}$	—
F480M	$199.4^{+9.0}_{-8.4}$	$104.0^{+2.0}_{-2.1}$	—	—	$7.47^{+0.142}_{-0.096}$

This system is ideal for testing performance because prior observations with VLTI/GRAVITY provide highly-precise constraints on the companion orbits (Hinkley et al. 2023). This paper presents the detection of the HD 206893 B/c companions with AMIGO as a demonstration of its performance close to the theoretical contrast detection limits based on photon noise. A future publication will present the scientific analysis of the JWST photometry in the context of the atmospheric carbon chemistry and the dust cloud properties (Kammerer et al., in prep.).

This dataset represents a far greater challenge to AMI, seeking the recovery of companions close to the theoretical limit of recoverable contrasts and pushing beyond the diffraction limit for the inner companion. A classical analysis of the HD 206893 images with AMICAL (Soulain et al. 2020) provides a tentative detection of the brighter B companion in all three bands

(without *c*), but at a position that is not compatible with the orbit of the object — known to ~ 1 mas from GRAVITY observations — by several tens of mas. This is likely due to uncorrected charge migration and distortion effects. When fixing the position of the companion in the fit, reasonable photometric constraints in all three bands can be obtained, but at much lower signal to noise than what would be expected without detector systematics. The photometry is also biased by these systematic effects. The fainter *c* companion remains undetected below the systematic noise when using the existing analysis pipelines.

Table 4.7: Best-fit relative joint astrometry and per-filter photometry for the HD 206893 *c* companion, showing the $\pm 1\sigma$ bounds. Reported quantities include the on-sky separation (Sep), position angle (PA), and Δ -magnitude in each of the F380M, F430M, and F480M filters. Results from two different fit types are compared: A joint fit in all three filters and a fit to each filter uniquely.

Fit Type	Sep (mas)	PA ($^\circ$)	Δ mag (F380M)	Δ mag (F430M)	Δ mag (F480M)
Joint	103_{-14}^{+16}	$-80.7_{-5.1}^{+4.1}$	$9.14_{-0.30}^{+0.43}$	$8.32_{-0.29}^{+0.23}$	$8.07_{-0.28}^{+0.34}$
F380M	125_{-20}^{+30}	$-88.0_{-8.9}^{+10.4}$	$9.17_{-0.30}^{+0.48}$	—	—
F430M	102_{-19}^{+24}	$-79.6_{-7.3}^{+7.4}$	—	$8.36_{-0.34}^{+0.28}$	—
F480M	91_{-36}^{+26}	$-83.8_{-8.4}^{+9.6}$	—	—	$8.03_{-0.70}^{+0.51}$

Analysis with AMIGO yields significantly better results than classic approaches. We find a confident detection of both companions in agreement with GRAVITY data, with well constrained photometry even when analysed per filter.

The targeted observational approach of the GO 1843 program provides a standard template for high-contrast companion recovery: one target and one calibrator in the three primary filters, F380M, F430M, F480M, without any dithers. Here we use all of the exposures, summarised in Table 4.5. The GO 1843 data targets are dimmer than those in the COM 1093 program (with more photons at shorter wavelengths), providing more groups per integration and preserving more photons. Table 4.5 also provides an overview of the approximate photon counts in each exposure as well as the fraction that is lost through the present AMIGO pipeline. Shorter wavelengths lose more photons due to having fewer groups but most of the signal is preserved. From these values and Equation 4.11 this gives an expected contrast limit of ~ 9.66 , 9.24 , and 8.78 Δ mags in the F380M, F430M and F480M filters respectively.

4.8.3.1 GO 1843: Analysis

The same process used in Section 4.8.2.1 was used for the GO program, taking `uncal` exposures from MAST and processing into `calslopes` with the default AMIGO data processing pipeline, then fitting and recovering the instrument state, latent visibilities, and DISCOs.

As a high-contrast companion observation we use the same methods for analysis of the DISCOs, expanded to account for the second companion. Figure 4.16 shows the log-likelihood detection maps for the full data, along with the data after the best-fit B companion has been subtracted off. Subtracting the B companion helps to reveal the dimmer *c* companion that can be obscured by the signal of the B companion. This shows a strong detection of both

companions with statistically consistent astrometry (as compared against VLTI/GRAVITY constraints provided in Hinkley et al. (2023)), independently in all filters.

Full inference is performed with a joint MCMC fit to both the B and c companions with astrometry constrained across all filters to ensure the most accurate recovery of photometry. The quality of this fit is assessed in Figure 4.16 via a correlation plot between the measured and predicted DISCOs. The recovered error scaling terms σ_{scale} are all close to unity, and reasonably consistent with the values found in from the fit to the AB Dor companion, indicating consistent performance in the high contrast regime and a good fit to the data. While both the B and c companions are simultaneously fit in the HMC, we present the samples for their parameters independently to prevent clutter. Examination of the full samples reveal insignificant correlation between the parameters of the two companions. Figure 4.17 shows the parameter posteriors for the B companion and Figure 4.18 shows the same for the c companion. A summary of the recovered parameters is presented in Table 4.6 and Table 4.7.

Quantification of the detection limits in this dataset is found by subtracting out the best-fit values from the MCMC while also applying the error scaling term to the data uncertainty. The azimuthal upper-limits are calculated on the cleaned DISCOs via the Ruffio et al. (2018) method and shown in Figure 4.19. This reveals that the performance of the AMIGO model remains in strong agreement with the expected contrast limits even down to ~ 9.5 magnitudes, inside the average diffraction limit of of all three filters, ~ 120 mas.

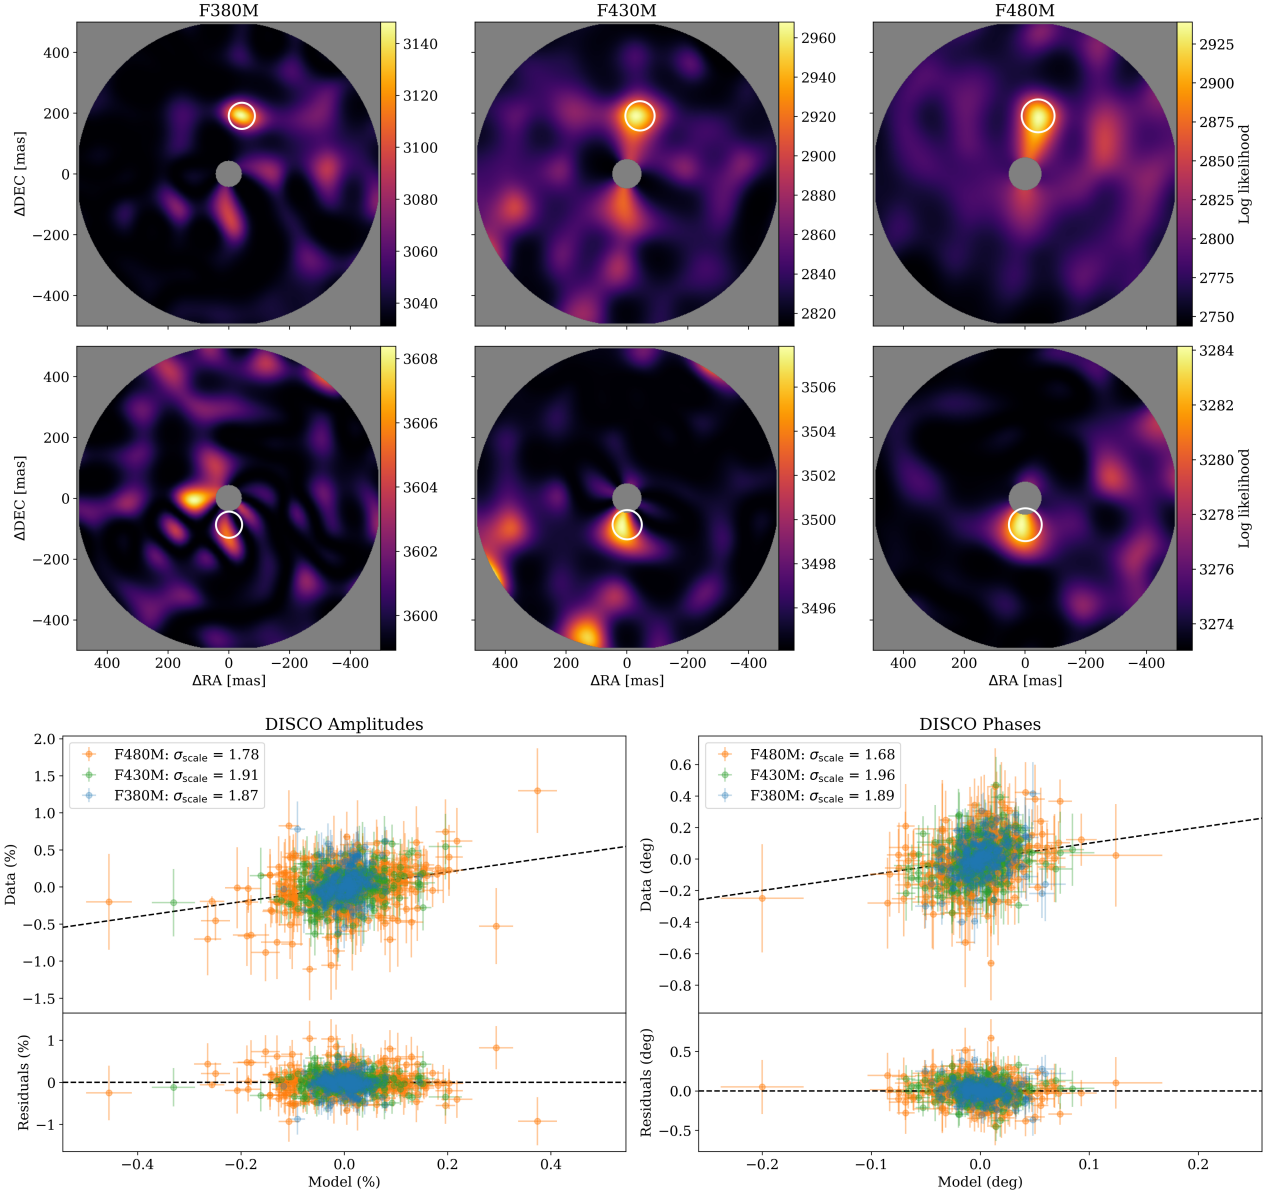


Figure 4.16: Summary of the fits to the HD 206893 companions. The top two panels shows the marginalised log-likelihood surface as a function of companion offset in ΔRA and ΔDEC , revealing a strong and consistent peak in all filters. The top row shows the detection maps for the full data, with the GRAVITY prediction for the B companion shown as a white circle. The middle row shows the detection map after the best-fit B companions has been subtracted from the data, revealing the inner c companion being consistently detected in all three filters, with the GRAVITY prediction overlaid with a white circle. The peaks in each filter for both companions can be seen matching the expected positions. The greyed central region denotes the IWA masked by the interferometric null. The bottom panels show the predictive posterior check comparing AMIGO recovered DISCO amplitudes (left) and phases (right) against the predicted values from the MCMC samples, for the three filters: **F380M**, **F430M**, and **F480M**. Each panel shows individual measurements with 1σ error bars, with the top panels plotting model predictions against data and the lower panels showing residuals. The black dashed line represents the 1:1 line for a perfect model prediction. Scale factors σ_{scale} are applied to the observational uncertainties in MCMC with the effect of ensuring $\chi^2_{\nu} \approx 1$ and are quoted in the legend for each filter. The agreement across filters and observables confirms the validity of the forward model and appropriate uncertainty quantification in the recovered posterior.

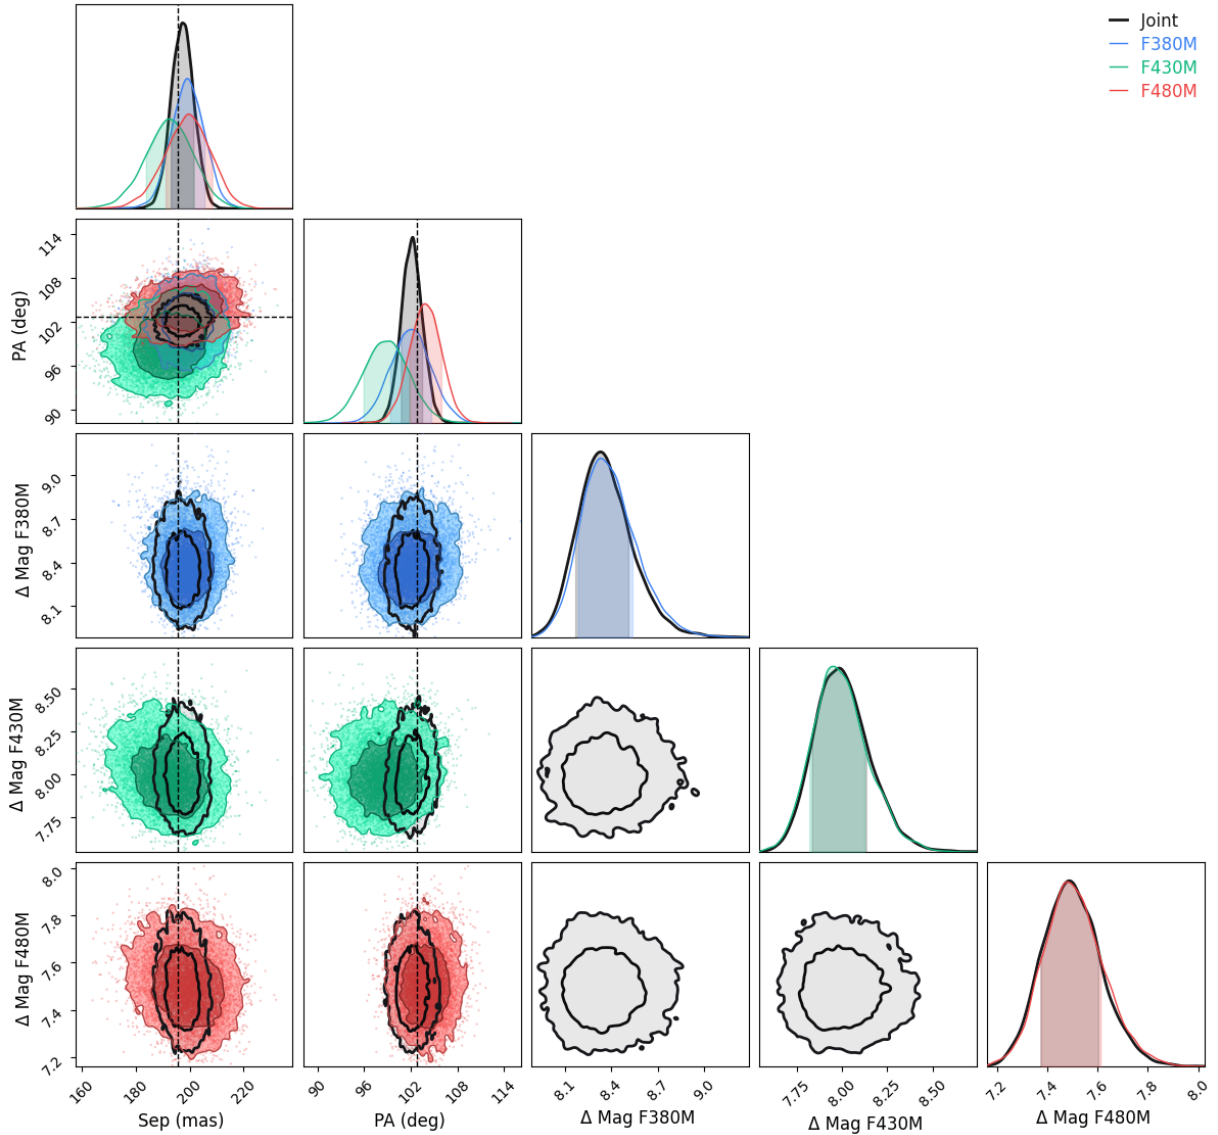


Figure 4.17: Joint posterior distributions from an MCMC fit to the HD 206893 B companion. Two fit types are shown: Joint — simultaneously modelling astrometry and photometry across all three filters — and per-filter fits. While both companions are fit simultaneously, only the B companions samples are shown here for clarity. The parameters shown include the separation (mas), position angle (degrees), and contrasts (Δmag) in the **F380M**, **F430M**, and **F480M** bands. The joint-fit samples are shown in black. One- and two- σ credible regions are shown in dark and light shades, respectively. The tight constraints in both astrometry and photometry reflect the high signal-to-noise of the companion detection. The expected position predicted by GRAVITY orbit fits (with a precision of ~ 1 mas) are overlaid in a black dashed line, showing strong agreement with AMI, both in each filter and in the joint-fit.

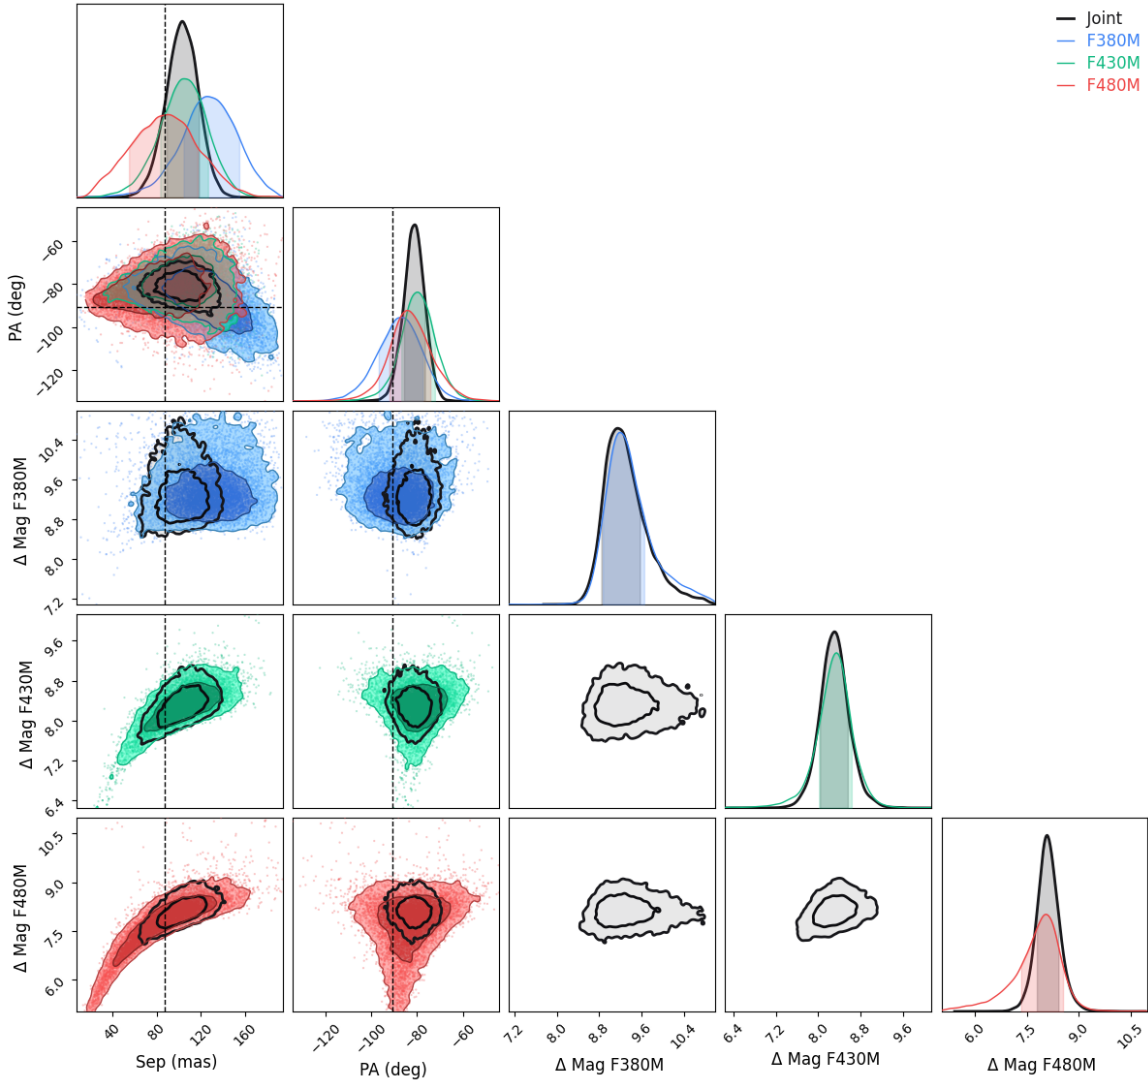


Figure 4.18: Joint posterior distributions from an MCMC fit to the HD 206893 c companion. Two fit types are shown: Joint — simultaneously modelling astrometry and photometry across all three filters — and per-filter fits. While both companions are fit simultaneously, only the B companions samples are shown here for clarity. The parameters shown include the separation (mas), position angle (degrees), and contrasts (Δmag) in the F380M, F430M, and F480M bands. The joint-fit samples are shown in black. One- and two- σ credible regions are shown in dark and light shades, respectively. The degeneracy between separation and contrast inside the diffraction limit is clearly shown in the F480M filter, demonstrating how photometry can be improved with a multi-band fit. The fit to the F380M filter shows that it falls right on the border of detection, with chains becoming poorly constrained above ~ 10 mags. The expected position predicted by GRAVITY orbit fits are overlaid in a black dashed line, showing decent agreement with AMI. Deviations between the expected and recovered positions could arise from either statistical noise or coupling to non-linear distortion arising from an imperfect calibration of the BFE, however given the relatively large astrometric uncertainty, we expect this to arise largely from low signal from the dim companion.

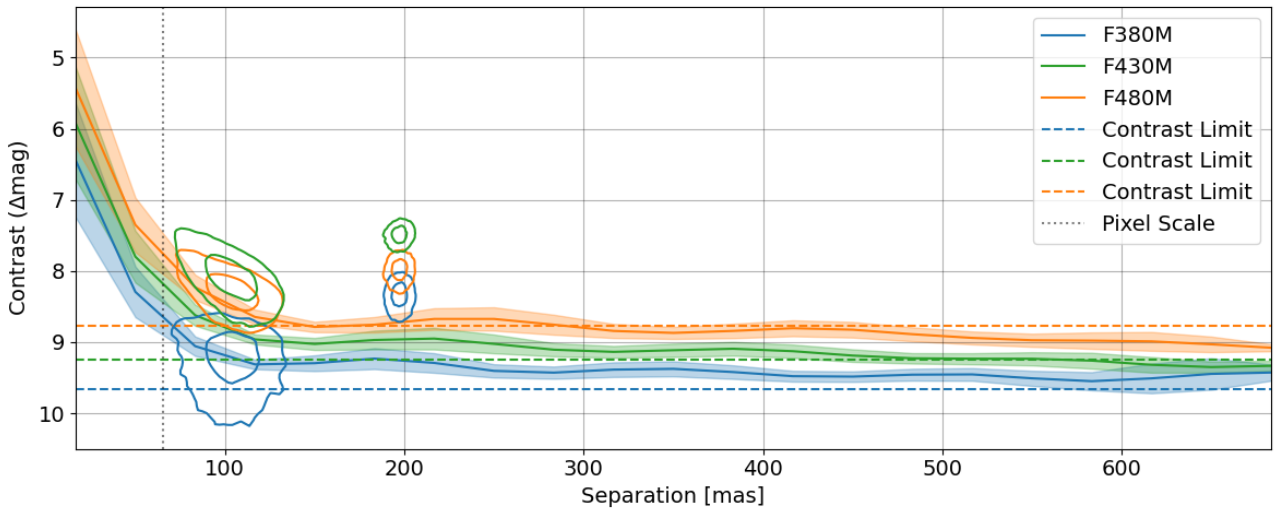


Figure 4.19: Sensitivity limits for HD 206893 derived using the Ruffio et al. (2018) method for calculating 3σ contrast upper limits as a function of angular separation, applied to the F380M, F430M, and F480M filters. The shaded regions denote the $\pm 1\sigma$ variation across azimuthal annuli. The dashed black line indicates the approximate contrast limit as calculated from Equation 4.11, confirming that model performance is consistent with the expected limits. The grey dotted line shows the size of a single pixel for reference. The recovered companion 1 and 2σ contour plots are shown in their respective colours. Recovery of the dim inner companion near the IWA and inside of the diffraction limit quantifies the AMI contrast limits as well calibrated and close to the theoretical limits of performance for AMI.

4.9 Discussion

The results presented in this work demonstrate that a robust analysis of JWST AMI data through differentiable forward model provides a new way to calibrate and extract observables. These methods provide insight into a new set of observing strategies better suited to these approaches. Prior investigations by Sallum et al. (2024b) have led to a set of three recommendations for future observations to mitigate the effect of charge migration:

1. Maximise the number of groups per integration.
2. Observing calibrator and science targets to a similar well depth.
3. Selecting calibrators with a similar brightness to science targets.

We agree with these recommendations. We also provide two further recommendations for using AMIGO, while mitigating its weaknesses:

1. Remaining comfortably below half pixel well depth ($\sim 50k e^-$, $\sim 30k$ counts): this is already a recommended strategy for AMI observations given the worsening charge migration above these values. Adding to this the AMIGO model is only trained on exposures below these values, making data with deeper pixel wells fall into a regime where model behaviour is unexplored and unlikely to perform well. Future calibration data with deeper exposures should provide a path towards precise observables with forward modelling methods, but until calibrations are performed in this regime, observations should remain below these levels.
2. A 5-point sub-pixel dither pattern: provided precise telescope pointing, conventional interferometric calibration should remove the effect of pixel-level miscalibrations making dithering an unnecessary complication to analysis. In practice non-linear and pixel to pixel detector effects can couple strongly to measured PSF shapes through the BFE. While AMIGO provides a direct model of these effects, nonlinearity makes precise calibration challenging without a more thorough treatment and better calibration data. A 5-point sub-pixel dither increases computational costs but provides a direct way to decouple pixel-level miscalibration from science observables and is expected to improve precision through the AMIGO model and pipeline.

These items provide a generally achievable set of guidelines to aide overall precision in output observables. Faster read-out modes for AMI are also being developed that will make it easier to select acceptable targets that fit within these guidelines.

The case study of HD 206893 highlights the strengths of AMIGO and also provides insight into potential limitations. The B companion is recovered with high confidence in both the astrometric and photometric precision. There is very strong agreement between the AMI astrometry and the predictions made by GRAVITY (Hinkley et al. 2023) indicating that AMI is now capable of confident results in the high-contrast regime up to $\Delta \text{mag} \sim 9$. There is also a confident detection of the c companion, however it shows a $\sim 2\sigma$ deviation in the astrometry.

This companion presents the greatest challenge to the AMIGO model, with its harsh contrast and separation well beyond the capabilities of any other JWST observing modes, and placement inside the diffraction limit of all three filters. These deviations could be consistent with an astrometric bias induced from this minute separation, given its position well inside of the pixels most affected by charge migration from the BFE. This could also explain the significantly larger uncertainty in its recovered position. Given these considerations the possibility of biases in the recovered photometry will remain until further study can be performed on other well constrained systems. A promising sign is that both of these companions remain detectable independently in each filter indicating that these results can be still be considered credible.

AMIGO has shown it can recover high-contrast companions at the diffraction limit. We recover a 3σ contrast limit through the Ruffio et al. (2018) method in approximate agreement with the scaling from Equation 4.11 (Ireland 2013) and recommended in JWST proposal materials. We adopt this as a good rule of thumb, but would also remind readers of the proposed observing strategies for AMIGO, emphasising that the photon loss arising from the current implementation is not directly accounted for in this equation and should be treated with care. This remains as a practical estimate of AMI performance though the methods explored in the work. We expect future calibrations and refinement of AMIGO to reclaim the photons lost by the first group, to improve performance across the board, and push beyond these estimates.

4.10 Conclusions and Future Work

The AMIGO model and pipeline provide a comprehensive re-commissioning of JWST/AMI by harnessing end-to-end differentiable models and the novel techniques that they enable. It has enabled a robust recovery of the unrealised potential of high-contrast interferometric imaging at the highest angular resolutions probed by JWST. This is the first pipeline in astronomy to fully integrate a differentiable physical forward model, including a full treatment of non-linear detector systematics and an embedded NN, and trained entirely from on-sky data. Through case studies of AB Dor AC and HD 206893, we demonstrate that AMIGO can recover faint companions in the high-contrast and high-angular resolution regime while preserving precision in both astrometry and photometry. Its performance approaches the theoretical photon-noise limits of the instrument, and surpasses the performance of available pipelines and approaches. This first version of the model is also only the start of its journey; many significant improvements can be made and are planned, placing it as a platform for both future high quality calibrations of JWST and improvements to existing and future AMI data and instruments.

Quantitative detection limits inferred from the AMIGO model provide a new characterisation of its true performance limits to date. We find that AMI can achieve contrasts of $\gtrsim 9$ magnitudes with astrometry up to the diffraction limit $\sim \lambda/D$. Detail is also recovered very close to an inner working angle of ~ 100 mas, but with unknown potential biases introduced through the non-linear detector effects. The contrast limits found match closely with the Equation 4.11, giving a detectable contrast of $\sim 10/\sqrt{N_{\text{photons}}}$. These values are suitable for inclusion in future proposals, offering guidance on expected sensitivity, inner working angle and spatial resolution.

Beyond its immediate utility for science, AMIGO showcases the broader impact that differ-

entiable forward models are poised to make on modelling in astronomical instrumentation and imaging as a whole. By treating systems in a physically principled end-to-end manner spanning astrophysics, optics, and detectors, AMIGO enables efficient and interpretable inference, rigorous uncertainty quantification, and coherent system calibration. This ‘pixels to planets’ philosophy – avoiding data reduction, and modelling data directly in their rawest possible form – represents a promising new direction for data analysis for other instruments affected by complex or non-linear systematics.

Concrete calibrations provided by this approach, namely the integral non-linearity of the ADC, are common mode to all observations made by NIRISS. We have already found identical signals in other non-AMI modes with the potential to induce biases in downstream scientific analysis. cursory examinations of NIRC*am* and NIRSpec data do not show the same signal. Nonetheless, its presence in NIRISS data warrants further consideration of its effects in broader JWST calibration efforts, including SOSS transmission spectroscopy observations.

Looking forwards, several avenues of improvement and wider development are clear. The architecture of the embedded NN is both compact and physically constrained, but remains limited by the availability of extensive calibration data. While the PSFs produced by AMI present highly structured and well known illuminance pattern, most flux remains concentrated in the core. The calibration data used for training fills peak pixels to approximately half their maximum depth, and only provides a handful of pixels that fill beyond much smaller fractions of their maximum, and are suitable for calibration of the strongest BFE regime. This makes generalisation to the dynamics of charge migration for resolved sources an unexplored space, where model performance could degrade. Our team’s Cycle 4 program GO 8330 plans to observe a bright, wide binary source at multiple primary and subpixel dither positions. These exposures will fill many pixels to half well depths while allowing the core to saturate. This should provide a dataset far richer in the complex non-linear dynamics of the BFE, while partially exploring the full pixel well depth in the saturated core. Further training of AMIGO on this data should provide a substantial improvement to its overall ability to directly model and understand the physics of this problem.

Provided a high fidelity EDM trained on this more informative dataset, it may become possible to develop a self-contained detector model able to *directly* invert the effects of the detector entirely by sampling input photon distributions. Such a ‘de-BFE’ approach has the potential to plug into a classic detector calibration pipeline that inverts all physical degradations, returning high quality super-resolved PSFs that accurately capture the statistical correlations induced by its measurements. Models trained on lab data during instrument pre-commissioning could be released in tandem with detector hardware itself as part of the supplier’s service.

The AMIGO model and this manuscript fully exclude all calibration and analysis in the F277W filter. This comes down to a lack of available high quality training data, along with a more complex set of physics associated with a larger spectral range. Despite its potentially unique benefits, the calibration dataset used for AMIGO did not observe in this filter as very few programs use it. Furthermore, the relative short spectral range and small band passes across the three primary filters F380M, F430M, and F480M allows for chromatic effects in the BFE to

be ignored. Furthermore, the implementation of the interferometric visibility model considers all sources to have negligible spectral variation across any filter. While these limitations are not fundamental, the present AMIGO model is neither designed nor calibrated to operate at the F277W filter and is expected to perform poorly. Future version of AMIGO may be expanded to operate at these wavelengths, although this would require dedicated calibrated data and change to the structure of various sub-modules within AMIGO.

The DISCO observables introduced in this work represent a new approach to information-preserving projections of interferometric data. Building on Fisher information and kernel phase and amplitudes, and accelerated through autodiff, it enables the exploration of the information theoretic limits of imaging systems. By restricting the basis vectors to the OTF through their eigenvalues and varying the number of resolution elements and eigenvectors used in the uv -plane, the subspace of recoverable images can be quantified and concrete informational limits on instrumental performance can be explored. These ideas parallel limits on point-source detection set by small-sample statistics near the diffraction limit (Mawet et al. 2014), recast in the Fourier domain. In this paper we have not delved into depth to establish precise limits for DISCO linearity, scaling laws, or calibration strategies; future work will explore these ideas concretely for AMI (Charles et al., in prep.). Given the excellent long-term stability and frequent monitoring of the wavefront using NIRC*am* (Lajoie et al. 2023), it may even be possible to do calibrator-free AMI imaging by trusting the NIRC*am*-measured OPD between the observation epoch and the wavefront map in the AMIGO base model and relying on the DISCOs to deal with residual wavefront error.

This more general approach may also be able to improve the performance of spectrally-dispersed wavefront sensing and kernel phase on IFU data (Martinache et al. 2016; N'Diaye et al. 2022; Chaushev et al. 2023). The NIRS*pec* Integral Field Unit (IFU) has been used to achieve extremely high contrast ($\sim 3 \times 10^{-6}$), for which Ruffio et al. (2024) depend on a forward model of the IFU PSF using WEBBPSF; a trained ∂ LUX model as in AMIGO will be more flexible and, we expect, enable us to push to deeper systematics floors.

An advantage of modelling visibilities in a latent basis which very finely samples the uv plane is that the interferometric field of view (FOV) (which corresponds to the Nyquist limit of this uv sampling) is large, potentially as large as the native 80×80 subarray FOV if the uv sampling is the FFT coordinate grid conjugate to the raw image. This is ideal for image reconstruction/deconvolution. An implementation in JAX of regularised maximum likelihood image reconstruction, similar to MPOL (Czekala et al. 2025), is likely to be straightforward, and better able to recover complex scenes than reconstructions limited to one visibility per baseline and the basic closure phases. It remains an open question whether this is better suited to a DISCO likelihood or direct to pixels. As an extension, it may be advantageous to apply score-based methods to very expressively build priors over images (Feng et al. 2023; Dia et al. 2025), and likelihoods over detector effects (Adam et al. 2025), applications which will only be possible with differentiable instrument models.

The methods, ideas, and philosophy laid out in this work are not limited to either JWST nor interferometric systems. Natural extensions to NIRC*am* and archival HST NIR camera and

multi-object spectrometers are straightforwardly possible. The tighter PSFs and redundant apertures provide significant but not insurmountable challenges. Various methods harnessing kernel phase or data driven empirical PSF subtraction remain effective and have been used on both JWST and HST data to reveal faint companions at high resolution (De Furio et al. 2023; Calissendorff et al. 2023). These observations provide an ideal testing ground for deeper exploration of the ideas and models presented in this work to other imaging modes and instruments. One set of low-hanging fruit may be the HST/NICMOS LT dwarfs, which were previously studied with kernel phase but without accurate aperture calibration (Pope et al. 2013; Martinache et al. 2020).

Finally, there is growing potential to apply the pixels-to-planets philosophy to coronagraphic observations. PSF subtraction using the conventional optical modelling software WEBBPSF can improve recovered image deconvolution, accounting for the spatially-varying PSF (Balmer et al. 2025). Replacing these nominal PSF models with on-sky calibrated predictions from a model like AMIGO with data-driven estimates of the coronagraph metrology, and directly accounting for and calibrating non-linear detector effects, could unlock deeper contrasts and higher precision observables. This is likely to be helpful for the Roman coronagraph’s wavefront ground-in-the-loop control (Bailey et al. 2023), and future work with HabWorlds (Academies 2023).

4.11 Code and Data Availability

All codes and data used to produce this work are publicly available and open source. The AMIGO model and pipeline are hosted [on GitHub](#), along with a series of [example notebooks](#) used in this paper’s results. All data are available from MAST under the appropriate proposal numbers; in this pipeline paper, we have relied only on publicly available data with no proprietary data included. We encourage researchers to adapt and apply AMIGO to their own interferometric datasets, and to contact the development team with questions and contributions.

Acknowledgments

We would like to thank Steph Sallum, Laurent Pueyo, Marshall Perrin, Eddie Bergeron, Joel Sanchez-Bermudez, Thomas Vandal, Loïc Albert, Laurence Perreault-Levasseur, Alexandre Adam, Noé Dia, Adam Taras, Lucinda Lilley, Hayden Greer, Shashank Dholakia, and Shishir Dholakia for their helpful discussions, advice, and contributions.

We acknowledge and pay respect to the traditional owners of the land on which the University of Sydney, Macquarie University, and the University of Queensland are situated, upon whose unceded, sovereign, ancestral lands we work. We pay respects to their Ancestors and descendants, who continue cultural and spiritual connections to Country. We acknowledge and respect the Lək’wəḡən (Songhees and Esquimalt) Peoples on whose territory the University of Victoria stands, and the Lək’wəḡən and WSÁNEĆ Peoples whose historical relationships with

the land continue to this day.

BP, PT, and SR have been supported by the Australian Research Council grant DP230101439 and BP by DE210101639; and LD and MC have been supported by the Australian Government Research Training Program (RTP) award. We are grateful to the Australian public for enabling this science. BP and SR would like to thank the Big Questions Institute for their philanthropic support. DB and DJ acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). DJ also acknowledges support from NRC Canada. AS and KV acknowledge support from the NSF, NASA, and the STScI Director’s Discretionary Fund. Development of ∂ LUX has been supported by the Breakthrough Foundation through their Toliman project as a part of the Breakthrough Watch initiative.

This work is based on observations made with the NASA/ESA/CSA James Webb Space Telescope. The NIRISS instrument was funded by the Canadian Space Agency and built in Canada by Honeywell. The AMIGO base model is constructed from observations associated with JWST program CAL 4481 (PI: A. Sivaramakrishnan). The data were obtained from the Mikulski Archive for Space Telescopes at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127 for JWST.

We would also like to thank the reviewer for their insightful and highly constructive comments and feedback. These have helped strengthen this paper and its interpretability to the wider community.

This research made use of NUMPY (Harris et al. 2020); Matplotlib (Hunter 2007); JAX (Bradbury et al. 2018); NUMPYRO (Phan et al. 2019b; Bingham et al. 2019); equinox (Kidger and Garcia 2021b); OPTAX (Hessel et al. 2020); and ChainConsumer (Hinton 2016).

Chapter 5

Conclusions & Future Work

5.1 Conclusions

In this thesis, I have explored new methods made possible by differentiable forwards modelling applied to data and instrumentation challenges in precision contemporary astrophysics. These techniques are part of a philosophy I have termed ‘pixels to planets’ that was developed to meet the stringent calibration requirements of exoplanetary science. This method forges all relevant physical processes — from photon generation to detection — into a single, differentiable computational graph. The resulting unified models support tasks spanning the full life-cycle of an instrument, from design and calibration through to data analysis. I have shown how this approach gives access to rigorous, statistically principled inference, and is computationally tractable even for complex high-dimensional problems. I have also explored the limitations of previously-standard inverse pipelines in the presence of degenerate, non-linear, and coupled processes; showing that the alternative differentiable forward models can address these problems directly. Finally, I have shown how hybrid modelling methods that embed machine learning into physical models can restore precision observables while also maintaining physical intuition and interpretation, in contrast to long-standing epistemological tensions between Machine Learning (ML) and the physical sciences (Hogg and Villar 2024). This thesis provides a new path towards the ambitious goals for next generation astrophysical measurement, seeking to push beyond the capabilities of modern instruments.

A major outcome of this research is the development of the open-source software ∂LUX , designed to make differentiable optical models broadly accessible to the astronomical community. This framework was developed entirely within the context of the problems addressed in this thesis and embraces differentiability as a core architectural principle. Chapter 2 describes how these models can unify calibration and scientific analysis within a single framework. I show that Automatic Differentiation (autodiff; Margossian 2019) enables efficient solutions to high-dimensional problems, even those with over a million parameters — allowing for simultaneous phase retrieval, pixel-level calibration, and parameter inference directly from science data. This provides a robust way to correct for time-dependent systematics and degradation effects, particularly relevant in space-based environments. Chapter 3 extends these capabilities to hardware design, showing how efficient computation of second- and higher-order derivatives provides di-

rect access to Fisher information and the Cramér-Rao Lower Bound (CRLB; Radhakrishna Rao 1945; Cramé, Haraldr 1947). This allows statistically grounded design optimisation without requiring the simplifications typically required for analytical tractability. Although both chapters focus on simulated data, they directly address real-world challenges faced by the TOLIMAN (Telescope for Orbital Locus Interferometric Monitoring of our Astronomical Neighbourhood; Tuthill et al. 2018) mission and actively contribute to its advancement.

The most consequential application of the methods developed in this thesis is their use in addressing performance limitations of the Aperture Masking Interferometer (AMI; Sivaramakrishnan et al. 2012; Soulain et al. 2020) mode on James Webb Space Telescope (JWST; Gardner et al. 2006). Since launch, AMI has not achieved its pre-flight performance expectations. This shortfall does not arise from its optical design, but from charge redistribution effects within the detector that subtly distort the Point Spread Function (PSF) (an issue that impacts all methods attempting precision of sensor data at the required levels of precision, not only AMI mode). Despite their subtlety, these effects alter the morphology of the image in ways that have evaded correction by conventional modular calibration pipelines. As a result, the performance of AMI has, in practice, hitherto hardly surpassed that of ground-based interferometric observations. Drawing on the modelling frameworks developed in the preceding chapters, Chapter 4 demonstrates how this problem can be solved using a fully differentiable end-to-end forwards models, despite the failure of currently dominant “data reduction” paradigms.

Chapter 4 demonstrates the effectiveness of this modelling framework through several concrete calibration results, many of which are broadly relevant beyond AMI mode. A previously unrecognised non-linearity was identified as an uncalibrated but commonly encountered component to the response curve of the analogue-to-digital converter. The corrected calibration presented has implications for all Near-Infrared Imager and Slitless Spectrograph (NIRISS; Doyon et al. 2023b) observations. A discrepancy in mask hole size, reported inconsistently across official documentation and various pipelines, was resolved through direct recovery of pupil geometry from observed data. The model also recovered a calibration of the non-linearly distorted mask metrology and showed that there exists a filter-dependent defocus in the optical path is sufficiently large to require a full Fresnel treatment of the wavefront. The high-dimensional inference techniques developed in Chapter 2 were applied to real data, enabling simultaneous phase retrieval, aperture calibration, and pixel-level response modelling within a unified inference framework. These results were made possible by adopting a hybrid modelling approach, in which a neural network was embedded within the detector model to capture the non-linear dynamics of the Brighter-Fatter Effect (BFE). The entire system was trained end-to-end directly utilising on-sky observations.

This calibrated, differentiable model then empowered a new approach to extracting interferometric observables through forward modelling. By leveraging the model’s differentiability, a latent representation of the complex visibilities was constructed, encoding both the physics of diffraction and the optical configuration of the instrument while accounting for defocus effects. Building on this, the concepts of kernel amplitudes and phases were extended to define a flexible null space that suppresses dominant instrumental phase modes while preserving the underlying astrophysical signal. These methods were applied to real observations in regimes

where conventional pipelines had previously failed to deliver scientifically useful results. The moderate-contrast AB Dor AC companion was recovered with high-precision astrometric and photometric information at moderate separation. The same framework was then employed, pushing recovery to instrumental limits, delivering both the B and c companions of HD 206893. The measured astrometry was consistent with prior high-precision ground-based predictions, while the inner companion was detected at a contrast greater than 9 magnitudes and a separation of approximately 100 mas — a distance about 30% smaller than the nominal (λ/D) diffraction limit. These results re-establish AMI as a scientifically competitive observing mode, an important outcome given the limited operational lifetime of JWST.

5.2 Future Work

The next generation of astronomical instrumentation will demand levels of precision that exceed those that current analytical frameworks and calibration strategies can reliably provide. Meeting such requirements will involve not only advances in hardware, but also in how we conceptualise and construct models of the measurement process itself. This thesis has explored how automatic differentiation offers a path forward, when treated not merely as a computational tool but as a foundational principle of model construction. By embedding the entire causal chain of observation — from photon emission to detector readout — into a single differentiable graph, we unify physical modelling and statistical inference into a coherent framework. This perspective reorients the modelling process: rather than building separate components for calibration, simulation, and analysis, we instead describe the physical system in full, and allow inference to operate end-to-end. The work presented here should therefore be viewed not simply as an extension of existing techniques, but as a step toward a modelling philosophy capable of addressing the complexity and precision requirements of future astronomical observations.

The hybrid modelling framework explored in this thesis opens new directions for addressing problems that have proven resistant to both traditional pipelines and unconstrained machine learning. By embedding learnable components within physically structured models, we can extend inference into domains where the systematics are too complex, non-linear, or poorly characterised for conventional methods to handle. This framework sets the stage for addressing calibration and inference challenges by combining physical structure with targeted flexibility. As the complexity of astronomical instruments increases, this approach offers a scalable and principled way to ensure that modelling and data inference keep pace.

The differentiable and deterministic forward models developed in this thesis also open the door to learning not just the model, but the inference procedure itself. In particular, methods inspired by amortised optimisation (Andrychowicz et al. 2016) and inference (e.g. Gershman and Goodman 2014; Díaz Baso et al. 2022) suggest that it is possible to train a Neural Network (NN) to act as a solver — mapping from the gradients of the forward model, conditioned on data, to the parameters that best reproduce the observation. This is made possible by the fact that the model’s predictive dynamics are fixed and governed by physical principles, making the relationship between gradient information and optimal solutions stable enough to be learned. Such an approach could significantly reduce the computational cost of inference,

particularly in settings where repeated optimisation over similar problems is required. While preliminary, this direction offers a promising route to integrating fast, learned solvers with physically interpretable models, enabling efficient inference without sacrificing transparency.

Accurate wavefront calibration remains a central challenge in both ground- and space-based high-contrast imaging, particularly in the presence of non-common path aberrations that cannot be sensed or corrected by traditional wavefront sensors. Focal-plane wavefront sensing offers a promising solution by inferring the phase directly from the science camera, avoiding differential errors between sensing and science paths. This approach has been demonstrated through the use of asymmetric pupil masks, which break the phase degeneracy and make focal-plane inference tractable (Martinache 2013; Martinache et al. 2016). However, most implementations rely on linearised models that fail in regimes where aberrations become large. Recent attempts to address this non-linear regime with machine learning have typically resorted to unconstrained neural networks, sacrificing physical interpretability and calibration transferability in the process (Andersen et al. 2019; Weinberger et al. 2024). The differentiable modelling framework developed in this thesis offers the best of both worlds: by solving for wavefront phase in a physically grounded model, we can operate in the non-linear regime without abandoning transparency. Computational speed for inference can be improved by taking advantage of sparsity; and further improved using amortised optimisation strategies, using a neural network trained to act as a solver — potentially allowing real-time wavefront retrieval from focal-plane data. These models can be calibrated on large archival datasets and deployed efficiently on GPUs, and critically, the model state used for wavefront sensing remains available to downstream science analysis, providing a high quality model for the inference required to deliver science observables.

The Aperture Masking Interferometry Generative Observations (AMIGO) framework introduced in Chapter 4 provides a foundation for several immediate extensions. While the current implementation has successfully recovered high-contrast companions, it simplifies this process by discarding the first group read from observations. This simplification has been useful for tractability, but it is not a fundamental limitation. With appropriate calibration data, the model can be extended to incorporate the first group and improve accuracy at higher pixel depths. A dedicated JWST calibration program has already been approved to provide this data, which will allow AMIGO to be refined and extended to brighter targets currently inaccessible due to increasing non-linear effects. These improvements will further increase the scientific yield of AMI and extend its observational bounds.

While AMIGO has been developed as a general-purpose forward modelling framework, currently its empirical validation has focused on point-source companions, which offer well-understood test cases and tractable benchmarks. A natural next step is to apply the model to observations of resolved or spatially extended sources, which introduce more complex illumination patterns at the detector. These scenarios are especially useful for testing the limits of the learned components within AMIGO, such as the ability of the embedded NNs to generalise beyond the training distribution defined by isolated sources. While this thesis does not explore this idea directly, it is the subject of ongoing work (Charles et al., in prep), and presents an opportunity to further assess the robustness of the forward model in more challenging inference settings.

Differentiable forward modelling enables, for the first time, the construction of a visibility basis that is directly tied to the physical response of the instrument. Rather than relying on fixed analytical constructs, visibilities are defined in terms of the specific interferometric modes to which the system is sensitive, shaped by diffraction, optical aberrations, and detector effects. This allows for a fully forward-modelled approach to interferometry, in which observables emerge from the model itself rather than being imposed externally under idealised assumptions (as is the case in closure phases or kernel phases for example). This basis provides a principled way to study how information is retained or lost under varying instrument configurations or imperfections. Moreover, with the instrumental response fully characterised, the associated null space can be constructed with greater flexibility, allowing specific aberrations or alignment modes to be explicitly suppressed or sensed. While demonstrated here in the context of AMI, this framework extends naturally to other interferometric and imaging systems where system-aware observables may offer improved robustness and diagnostic power.

Differentiable forward-modelling also offers a natural framework for image reconstruction and deconvolution. By incorporating the full instrumental response, reconstructions can be performed through a direct maximum likelihood, where both the scene and system parameters are jointly optimised. This allows for the nuisance instrumental effects to be accounted for explicitly during reconstruction, rather than corrected in preprocessing or approximated in the PSF. As explored in Chapter 4, recovered observables can also be projected into phase or instrumental null spaces to suppress modes dominated by systematics and improve robustness to residual miscalibration. These techniques can be directly applied to a number of archival JWST AMI datasets — including observations of GTO 1242 targets, the transition disks PDS 70, HD 100546, and the exozodiacal system ν Hor; and the AGN NGC 1068 and the colliding-wind binary WR 137, both observed as imaging standards — offering the potential for refined reconstructions and improved interpretation of spatially extended or complex astrophysical structures.

The degeneracy of image reconstruction problems can be further mitigated by integrating physically grounded forward models with learned generative priors. Recent advances in score-based diffusion models have made it possible to construct powerful data-driven priors that encode the distribution of plausible astrophysical scenes (Adam et al. 2022; Adam et al. 2023). The differentiable forward modelling framework developed in this thesis has already been directly integrated with these models, enabling joint posterior sampling over both the latent astrophysical scene and the instrumental parameters that affect the measurement process. Unlike two-stage approaches, this integration is fully differentiable and coherent: gradients flow from the data through the forward model and into the generative prior, allowing the system to exploit the strengths of both components. The forward model addresses the limitations of the generative prior, which typically lacks a physically realistic instrument model, while the prior helps regularise the ill-posed nature of deconvolution. This synergy enables reconstructions that are both physically constrained and statistically plausible. This approach can be further extended to incorporate learned score-based likelihoods, replacing hand-tuned noise models with data-driven approximations that better reflect the statistical structure of complex imaging systems and their systematics. Together, these techniques provide a robust foundation for

the analysis of archival datasets from Hubble Space Telescope (HST) and JWST, particularly in cases where calibration is limited and scene morphology is both rich and uncertain.

Many of the limitations that motivated the development of AMIGO arise from the Tele-dyne HAWAII-2RG (H2RG) detector hardware — which is common to most instruments on JWST — and so are not unique to AMI. Kernel phase observations on JWST, which apply interferometric techniques to full-pupil images, suffer from similar systematics that degrade their sensitivity (Kammerer et al. 2023; Calissendorff et al. 2023; De Furio et al. 2023). The modelling framework introduced in this thesis can be directly applied to kernel phase data, with minimal modification, using the same calibrations and inference tools. This offers a clear path to improving the robustness and scientific yield of kernel phase analysis, particularly in regimes where conventional calibration has proven insufficient. More broadly, these methods can be extended to other NIRISS imaging modes, enabling a unified approach to calibration and inference across the full suite of observations supported by the instrument.

The modelling approach developed for AMIGO is directly applicable to other instruments that share the same detector architecture, such as Near-Infrared Camera (NIRCam; Rieke et al. 2023). Like NIRISS, NIRCam exhibits charge migration and non-linear response effects, but its high-contrast coronagraphic modes introduce additional challenges. These include wavefront instabilities between reference and science exposures together with small misalignments of the various masks — both of which can degrade the fidelity of PSF subtraction. Differentiable forward models offer a way to model these effects coherently, enabling the joint calibration of optical and detector systematics within a unified framework. This extends the model-based subtraction approach taken by Balmer et al. (2025), by allowing gradient-based optimisation and uncertainty propagation to be integrated directly into the calibration process. In doing so, this framework provides a path to improved robustness and interpretability in coronagraphic analysis.

Future space missions such as the Roman space telescope (Bailey et al. 2023) and the Habitable Worlds Observatory (HWO or HabWorlds; Stark et al. 2025; Vaughan et al. 2023) represent a generational leap in precision instrumentation, as the first observatories to implement active wavefront control in orbit. Both instruments are designed to achieve unprecedented levels of wavefront stability, with HWO aiming for picometer-level control to support high-contrast imaging of Earth-like exoplanets. These goals reflect extraordinary engineering ambition, but rely heavily on conventional calibration architectures — including dedicated wavefront sensors that suffer from non-common path error. Even if the nominal wavefront stability targets are met, other instrumental effects such as mask and optical misalignments driven by thermal drifts, and detector-level systematics are likely to dominate the error budget when analysed using standard inverse-correction pipeline methods. Achieving the scientific precision demanded by these missions will require not only ambitious leaps in engineering, but equally ambitious approaches to calibration and analysis methodology. The end-to-end differentiable forward modelling paradigm developed in this thesis offers such an approach: providing the ability to encode and jointly marginalise over physical effects across the full optical and detector chain, and to recover science observables directly from raw measurements with physically meaningful uncertainty estimates. In this sense, the methods explored here are not only well matched to

the challenges posed by upcoming flagship missions — but they may be essential to fully realise their scientific potential.

The hybrid detector model developed in AMIGO also points toward a more general goal: the construction of unified forward models for HXRG-class detectors that can be applied across instruments and observing modes. With sufficient calibration data, such models could be used not only to forward-predict detector response, but to invert it — reconstructing high-fidelity photon distributions from raw images. This enables direct recovery of super-resolved PSFs (as sub-pixel information is encoded both thorough the sensitivity profile of each pixel and the dynamics of charge migration), removing reliance on static, empirical corrections. Realising this vision will require calibration datasets that span a broader range of illuminance patterns and dynamic ranges than are typically available on-sky. Laboratory systems, where illumination can be precisely controlled and systematically varied, offer a path to building such datasets. Beyond astronomy, these models may serve as tools for characterising detector hardware itself, providing insight into the underlying physics of charge migration, sensor nonlinearities, and their effect on science observables.

The instrument design framework developed in Chapter 3 employed Fisher Information to optimise hardware configurations for maximal sensitivity to scientific parameters, marginalised over the instrumental state. While effective, the current implementation operates locally — optimising performance around a single, fixed instrumental state. In practice, however, instruments are subject to a range of perturbations, including thermal drift, mechanical mis-alignments, and time-variable optical aberrations. A natural extension of this approach is to incorporate stochastic variation in the instrument state during optimisation, thereby marginalising over expected sources of instability. An example might be evaluating performance over an ensemble of atmospheric turbulence realisations, to find the best average or maximum instrument. Rather than identifying designs that are optimal only in an idealised or average configuration, this enables the discovery of configurations that retain high information content across a range of expected system perturbations. By embedding this variability directly into the optimisation process, the method may also become more robust to local minima in the loss topology and less dependent on initial conditions, improving both stability and generalisation.

While the Fisher Information approach used in Chapter 3 enables efficient design optimisation, it relies on the Laplace approximation and therefore on the assumption that the posterior distribution is well described by a Gaussian. For more complex optical systems such as coronagraphs, this assumption often breaks down. These systems exhibit strong non-linear parameter couplings and non-Gaussian posteriors that cannot be easily captured by local curvature alone. Variational Inference (VI; Jordan et al. 1999; Kucukelbir et al. 2017) offers a natural extension in this context, enabling posterior approximations that remain tractable but are better suited to the statistical structure of such systems. Moreover, the figures of merit in coronagraph design are often heuristic compromises between competing criteria such as off-axis throughput and on-axis suppression. These can be re-formulated more rigorously as Bayesian evidence ratios that quantify the relative support for models with and without an injected companion. This allows design choices to be evaluated directly in terms of detectability, and opens the door to statistically principled design workflows for non-linear high-contrast imaging systems.

This thesis has introduced and developed a modelling framework that unifies calibration, inference, and design within a single, physically grounded and differentiable architecture. In doing so, it departs from the modular pipeline paradigm that has traditionally structured astronomical data analysis, and proposes an alternative centred on coherence, interpretability, and statistical rigour. The methods presented here have been shown to recover lost performance in existing instruments, to enable new forms of inference, and to lay the groundwork for information-driven design of future hardware. More broadly, they offer a principled path forward as instruments grow more complex, systematics more subtle, and the scientific questions more demanding of the data obtained. It is my hope that the ideas developed here will continue to evolve — not as fixed solutions, but as tools that support a more unified, transparent and statistically principled approach to precision measurement in astronomy.

Bibliography

- Academies, National (Oct. 2023). *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. National Academies Press. ISBN: 9780309467346. DOI: [10.17226/26141](https://doi.org/10.17226/26141). URL: <http://dx.doi.org/10.17226/26141> (cit. on p. 125).
- Adam, Alexandre et al. (2022). *Posterior samples of source galaxies in strong gravitational lenses with score-based priors*. arXiv: [2211.03812](https://arxiv.org/abs/2211.03812) [astro-ph.IM]. URL: <https://arxiv.org/abs/2211.03812> (cit. on p. 131).
- Adam, Alexandre et al. (2023). *Echoes in the Noise: Posterior Samples of Faint Galaxy Surface Brightness Profiles with Score-Based Likelihoods and Priors*. arXiv: [2311.18002](https://arxiv.org/abs/2311.18002) [astro-ph.IM]. URL: <https://arxiv.org/abs/2311.18002> (cit. on p. 131).
- Adam, Alexandre et al. (May 2025). “Echoes in the Noise: Posterior Samples of Faint Galaxy Surface Brightness Profiles with Score-based Likelihoods and Priors”. In: 169.5, 254, p. 254. DOI: [10.3847/1538-3881/adb039](https://doi.org/10.3847/1538-3881/adb039). arXiv: [2311.18002](https://arxiv.org/abs/2311.18002) [astro-ph.IM] (cit. on p. 124).
- Aerts, C. (Jan. 2021). “Probing the interior physics of stars through asteroseismology”. In: *Rev. Mod. Phys.* 93 (1), p. 015001. DOI: [10.1103/RevModPhys.93.015001](https://doi.org/10.1103/RevModPhys.93.015001). URL: <https://link.aps.org/doi/10.1103/RevModPhys.93.015001> (cit. on p. 41).
- Aigrain, S., Parviainen, H., and Pope, B. J. S. (July 2016). “K2SC: flexible systematics correction and detrending of K2 light curves using Gaussian process regression”. In: 459.3, pp. 2408–2419. DOI: [10.1093/mnras/stw706](https://doi.org/10.1093/mnras/stw706). arXiv: [1603.09167](https://arxiv.org/abs/1603.09167) [astro-ph.SR] (cit. on p. 41).
- Aigrain, Suzanne and Foreman-Mackey, Daniel (Sept. 2022). “Gaussian Process regression for astronomical time-series”. In: *arXiv e-prints*, arXiv:2209.08940, arXiv:2209.08940. arXiv: [2209.08940](https://arxiv.org/abs/2209.08940) [astro-ph.IM] (cit. on p. 53).
- Airy, G. B. (Jan. 1835). “On the Diffraction of an Object-glass with Circular Aperture”. In: *Transactions of the Cambridge Philosophical Society* 5, p. 283 (cit. on p. 19).
- Akhare, Deepak, Luo, Tengfei, and Wang, Jian-Xun (2023). “Physics-integrated neural differentiable (PiNDiff) model for composites manufacturing”. In: *Computer Methods in Applied Mechanics and Engineering* 406, p. 115902. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2023.115902>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782523000257> (cit. on pp. 79, 95).
- Amari, Shun-ichi (Feb. 1998). “Natural Gradient Works Efficiently in Learning”. In: *Neural Computation* 10.2, pp. 251–276. ISSN: 0899-7667. DOI: [10.1162/089976698300017746](https://doi.org/10.1162/089976698300017746). eprint: <https://direct.mit.edu/neco/article-pdf/10/2/251/813415/089976698300017746.pdf>. URL: <https://doi.org/10.1162/089976698300017746> (cit. on p. 32).

- Andersen, Torben, Owner-Petersen, Mette, and Enmark, Anita (Sept. 2019). “Neural networks for image-based wavefront sensing for astronomy”. In: *Opt. Lett.* 44.18, pp. 4618–4621. DOI: [10.1364/OL.44.004618](https://doi.org/10.1364/OL.44.004618). URL: <https://opg.optica.org/ol/abstract.cfm?URI=ol-44-18-4618> (cit. on pp. 34, 130).
- Andrews, Sean M. et al. (Dec. 2018). “The Disk Substructures at High Angular Resolution Project (DSHARP). I. Motivation, Sample, Calibration, and Overview”. In: 869.2, L41, p. L41. DOI: [10.3847/2041-8213/aaf741](https://doi.org/10.3847/2041-8213/aaf741). arXiv: [1812.04040](https://arxiv.org/abs/1812.04040) [astro-ph.SR] (cit. on p. 6).
- Andrychowicz, Marcin et al. (2016). *Learning to learn by gradient descent by gradient descent*. arXiv: [1606.04474](https://arxiv.org/abs/1606.04474) [cs.NE]. URL: <https://arxiv.org/abs/1606.04474> (cit. on p. 129).
- Antilogus, P et al. (Mar. 2014). “The brighter-fatter effect and pixel correlations in CCD sensors”. In: *Journal of Instrumentation* 9.03, pp. C03048–C03048. ISSN: 1748-0221. DOI: [10.1088/1748-0221/9/03/c03048](https://doi.org/10.1088/1748-0221/9/03/c03048). URL: <http://dx.doi.org/10.1088/1748-0221/9/03/C03048> (cit. on pp. 15, 77).
- Argyriou, Ioannis et al. (Dec. 2023). “The brighter-fatter effect in the JWST MIRI Si:As IBC detectors: I. Observations, impact on science, and modeling”. In: *Astronomy & Astrophysics* 680, A96. ISSN: 1432-0746. DOI: [10.1051/0004-6361/202346490](https://doi.org/10.1051/0004-6361/202346490). URL: <http://dx.doi.org/10.1051/0004-6361/202346490> (cit. on pp. 15, 77, 97).
- Bailey, Vanessa P. et al. (Oct. 2023). *Nancy Grace Roman Space Telescope Coronagraph Instrument Overview and Status*. DOI: [10.1117/12.2679036](https://doi.org/10.1117/12.2679036). arXiv: [2309.08672](https://arxiv.org/abs/2309.08672) [astro-ph.IM]. URL: <https://arxiv.org/abs/2309.08672> (cit. on pp. 125, 132).
- Baldwin, J. E. et al. (Apr. 1986). “Closure phase in high-resolution optical imaging”. In: *Nature* 320.6063, pp. 595–597. DOI: [10.1038/320595a0](https://doi.org/10.1038/320595a0). URL: <https://doi.org/10.1038/320595a0> (cit. on pp. 9, 10, 24, 75).
- Balmer, William O. et al. (Mar. 2025). “JWST-TST High Contrast: Living on the Wedge, or, NIRCcam Bar Coronagraphy Reveals CO₂ in the HR 8799 and 51 Eri Exoplanets’ Atmospheres”. In: *The Astronomical Journal* 169.4, p. 209. ISSN: 1538-3881. DOI: [10.3847/1538-3881/adb1c6](https://doi.org/10.3847/1538-3881/adb1c6). URL: <https://dx.doi.org/10.3847/1538-3881/adb1c6> (cit. on pp. 125, 132).
- Bard, Yonathan (1974). *Nonlinear parameter estimation*. Tech. rep. (cit. on p. 47).
- Barnett, Alexander H et al. (Sept. 2020). “Geometry of the phase retrieval problem”. In: *Inverse Problems* 36.9, p. 094003. DOI: [10.1088/1361-6420/aba5ed](https://doi.org/10.1088/1361-6420/aba5ed). URL: <https://doi.org/10.1088/1361-6420/aba5ed> (cit. on p. 41).
- Barthelemy, J. -F. M. and Hall, L. E. (1995). “Automatic differentiation as a tool in engineering design”. In: *Structural optimization* 9.2, pp. 76–82. DOI: [10.1007/BF01758823](https://doi.org/10.1007/BF01758823). URL: <https://doi.org/10.1007/BF01758823> (cit. on p. 26).
- Benner, Peter, Gugercin, Serkan, and Willcox, Karen (2015). “A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems”. In: *SIAM Review* 57.4, pp. 483–531. DOI: [10.1137/130932715](https://doi.org/10.1137/130932715). eprint: <https://doi.org/10.1137/130932715>. URL: <https://doi.org/10.1137/130932715> (cit. on pp. 32, 90).

- Betancourt, Michael (Jan. 2017). “A Conceptual Introduction to Hamiltonian Monte Carlo”. In: *arXiv e-prints*, arXiv:1701.02434, arXiv:1701.02434. DOI: [10.48550/arXiv.1701.02434](https://doi.org/10.48550/arXiv.1701.02434). arXiv: [1701.02434](https://arxiv.org/abs/1701.02434) [stat.ME] (cit. on p. 31).
- Beuzit, J.-L. et al. (2019). “SPHERE: the exoplanet imager for the Very Large Telescope”. In: *A&A* 631, A155. DOI: [10.1051/0004-6361/201935251](https://doi.org/10.1051/0004-6361/201935251). URL: <https://doi.org/10.1051/0004-6361/201935251> (cit. on p. 7).
- Bezanson, Jeff et al. (Sept. 2012). “Julia: A Fast Dynamic Language for Technical Computing”. In: *arXiv e-prints*, arXiv:1209.5145, arXiv:1209.5145. arXiv: [1209.5145](https://arxiv.org/abs/1209.5145) [cs.PL] (cit. on pp. 42, 57).
- Bhandari, Naren et al. (Jan. 2021). “Fisher Matrix Stability”. In: *arXiv e-prints*, arXiv:2101.00298, arXiv:2101.00298. DOI: [10.48550/arXiv.2101.00298](https://doi.org/10.48550/arXiv.2101.00298). arXiv: [2101.00298](https://arxiv.org/abs/2101.00298) [astro-ph.CO] (cit. on pp. 32, 58).
- Bingham, Eli et al. (2019). “Pyro: Deep Universal Probabilistic Programming”. In: *J. Mach. Learn. Res.* 20, 28:1–28:6. URL: <http://jmlr.org/papers/v20/18-403.html> (cit. on pp. 111, 126).
- Björck, Å. (1994). “Numerics of Gram-Schmidt orthogonalization”. In: *Linear Algebra and its Applications* 197-198, pp. 297–316. ISSN: 0024-3795. DOI: [https://doi.org/10.1016/0024-3795\(94\)90493-6](https://doi.org/10.1016/0024-3795(94)90493-6). URL: <https://www.sciencedirect.com/science/article/pii/0024379594904936> (cit. on p. 64).
- Blakely, Dori et al. (May 2022). “Two Rings and a Marginally Resolved, 5 au Disk around LkCa 15 Identified via Near-infrared Sparse Aperture Masking Interferometry”. In: 931.1, 3, p. 3. DOI: [10.3847/1538-4357/ac6586](https://doi.org/10.3847/1538-4357/ac6586). arXiv: [2204.07055](https://arxiv.org/abs/2204.07055) [astro-ph.EP] (cit. on p. 76).
- Blakely, Dori et al. (2024). *The James Webb Interferometer: Space-based interferometric detections of PDS 70 b and c at 4.8 μm*. arXiv: [2404.13032](https://arxiv.org/abs/2404.13032) [astro-ph.EP]. URL: <https://arxiv.org/abs/2404.13032> (cit. on pp. 77, 109).
- Blakely, Dori et al. (Mar. 2025). “A Tentative Detection of a Point Source in the Disk g Gap of HD 100546 with VLT/SPHERE-IRDIS Sparse Aperture Masking Interferometry”. In: 169.3, 152, p. 152. DOI: [10.3847/1538-3881/ada944](https://doi.org/10.3847/1538-3881/ada944). arXiv: [2502.07759](https://arxiv.org/abs/2502.07759) [astro-ph.SR] (cit. on p. 76).
- Böker, T. et al. (Mar. 2023). “In-orbit Performance of the Near-infrared Spectrograph NIRSpec on the James Webb Space Telescope”. In: *Publications of the Astronomical Society of the Pacific* 135.1045, p. 038001. DOI: [10.1088/1538-3873/acb846](https://doi.org/10.1088/1538-3873/acb846). URL: <https://dx.doi.org/10.1088/1538-3873/acb846> (cit. on p. 12).
- Bonneau, D. et al. (Oct. 2012). *SearchCal: The JMMC Evolutive Search Calibrator Tool*. ASCL, record ascl:1210.012 (cit. on p. 81).
- Born, Max et al. (1999). *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. 7th ed. Cambridge University Press (cit. on pp. 18, 87).
- Borucki, William J. et al. (Feb. 2010). “Kepler Planet-Detection Mission: Introduction and First Results”. In: *Science* 327.5968, p. 977. ISSN: 0036-8075. DOI: [10.1126/science.1185402](https://doi.org/10.1126/science.1185402). eprint: <https://science.sciencemag.org/content/327/5968/977.full.pdf>. URL: <https://science.sciencemag.org/content/327/5968/977> (cit. on pp. 3, 40, 41).

- Bowler, Brendan P. (Oct. 2016). “Imaging Extrasolar Giant Planets”. In: 128.968, p. 102001. DOI: [10.1088/1538-3873/128/968/102001](https://doi.org/10.1088/1538-3873/128/968/102001). arXiv: [1605.02731](https://arxiv.org/abs/1605.02731) [astro-ph.EP]. URL: <https://dx.doi.org/10.1088/1538-3873/128/968/102001> (cit. on pp. 4, 6, 40).
- Box, George E. P. (Dec. 1976). “Science and Statistics”. In: *Journal of the American Statistical Association* 71.356, pp. 791–799. ISSN: 1537-274X. DOI: [10.1080/01621459.1976.10480949](https://doi.org/10.1080/01621459.1976.10480949). URL: <http://dx.doi.org/10.1080/01621459.1976.10480949> (cit. on p. 32).
- Brabandere, Bert De et al. (2016). *Dynamic Filter Networks*. arXiv: [1605.09673](https://arxiv.org/abs/1605.09673) [cs.LG]. URL: <https://arxiv.org/abs/1605.09673> (cit. on p. 98).
- Bracewell, R. N. and MacPhie, R. H. (Apr. 1979). “Searching for nonsolar planets”. In: 38.1, pp. 136–147. DOI: [10.1016/0019-1035\(79\)90093-9](https://doi.org/10.1016/0019-1035(79)90093-9) (cit. on p. 7).
- Bradbury, James et al. (2018). *JAX: composable transformations of Python+NumPy programs*. Version 0.1.55. URL: <http://github.com/google/jax> (cit. on pp. 26, 40, 42, 54, 57, 72, 78, 126).
- Broeg, C. et al. (2013). “CHEOPS: A transit photometry mission for ESA’s small mission programme”. In: *EPJ Web of Conferences* 47. Ed. by Roberto Saglia, p. 03005. ISSN: 2100-014X. DOI: [10.1051/epjconf/20134703005](https://doi.org/10.1051/epjconf/20134703005). URL: <http://dx.doi.org/10.1051/epjconf/20134703005> (cit. on p. 41).
- Brooks, Steve et al. (May 2011). *Handbook of Markov Chain Monte Carlo*. Ed. by Steve Brooks et al. 1st. Boca Raton, FL: Chapman and Hall/CRC. ISBN: 9780429138508. DOI: [10.1201/b10905](https://doi.org/10.1201/b10905). URL: <http://dx.doi.org/10.1201/b10905> (cit. on p. 31).
- Broyden, C. G. (Mar. 1970). “The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations”. In: *IMA Journal of Applied Mathematics* 6.1, pp. 76–90. ISSN: 0272-4960. DOI: [10.1093/imamat/6.1.76](https://doi.org/10.1093/imamat/6.1.76). eprint: <https://academic.oup.com/imamat/article-pdf/6/1/76/2233756/6-1-76.pdf>. URL: <https://doi.org/10.1093/imamat/6.1.76> (cit. on pp. 31, 107).
- Burtscher, L. et al. (2013). “A diversity of dusty AGN tori - Data release for the VLTI/MIDI AGN Large Program and first results for 23 galaxies”. In: *A&A* 558, A149. DOI: [10.1051/0004-6361/201321890](https://doi.org/10.1051/0004-6361/201321890). URL: <https://doi.org/10.1051/0004-6361/201321890> (cit. on p. 5).
- Bushouse, H. et al. (2025). *JWST Calibration Pipeline (1.18.1)*. DOI: [10.5281/zenodo.15632984](https://doi.org/10.5281/zenodo.15632984). URL: <https://doi.org/10.5281/zenodo.15632984> (cit. on p. 13).
- Calissendorff, Per et al. (Apr. 2023). “JWST/NIRCam Discovery of the First Y+Y Brown Dwarf Binary: WISE J033605.05-014350.4”. In: 947.2, L30, p. L30. ISSN: 2041-8213. DOI: [10.3847/2041-8213/acc86d](https://doi.org/10.3847/2041-8213/acc86d). arXiv: [2303.16923](https://arxiv.org/abs/2303.16923) [astro-ph.SR]. URL: <http://dx.doi.org/10.3847/2041-8213/acc86d> (cit. on pp. 125, 132).
- Campagne, Jean-Eric et al. (Apr. 2023). “JAX-COSMO: An End-to-End Differentiable and GPU Accelerated Cosmology Library”. In: *The Open Journal of Astrophysics* 6, 15, p. 15. DOI: [10.21105/astro.2302.05163](https://doi.org/10.21105/astro.2302.05163). arXiv: [2302.05163](https://arxiv.org/abs/2302.05163) [astro-ph.CO] (cit. on p. 78).
- Candes, Emmanuel J. et al. (Sept. 2011). “Phase Retrieval via Matrix Completion”. In: *arXiv e-prints*, arXiv:1109.0573, arXiv:1109.0573. arXiv: [1109.0573](https://arxiv.org/abs/1109.0573) [cs.IT] (cit. on p. 41).

- Cantalloube, F. et al. (Jan. 2021). “Exoplanet Imaging Data Challenge: benchmarking the various image processing methods for exoplanet detection”. In: *arXiv e-prints*, arXiv:2101.05080, arXiv:2101.05080. arXiv: [2101.05080](https://arxiv.org/abs/2101.05080) [[astro-ph.IM](#)] (cit. on p. 53).
- Chaloner, Kathryn and Verdinelli, Isabella (1995). “Bayesian Experimental Design: A Review”. In: *Statistical Science* 10.3, pp. 273–304. ISSN: 08834237. URL: <http://www.jstor.org/stable/2246015> (visited on 05/17/2023) (cit. on p. 58).
- Charbonneau, David et al. (Dec. 1999). “Detection of Planetary Transits Across a Sun-like Star”. In: *The Astrophysical Journal* 529.1, p. L45. DOI: [10.1086/312457](https://doi.org/10.1086/312457). URL: <https://dx.doi.org/10.1086/312457> (cit. on p. 3).
- Chaushev, Alexander et al. (Apr. 2023). “Spectrally dispersed kernel phase interferometry with SCEXAO/CHARIS: proof of concept and calibration strategies”. In: 9, 028004, p. 028004. DOI: [10.1117/1.JATIS.9.2.028004](https://doi.org/10.1117/1.JATIS.9.2.028004). arXiv: [2305.17065](https://arxiv.org/abs/2305.17065) [[astro-ph.IM](#)] (cit. on p. 124).
- Cheetham, Anthony C. et al. (Dec. 2012). “Fizeau interferometric cophasing of segmented mirrors”. In: *Optics Express* 20.28, p. 29457. DOI: [10.1364/OE.20.029457](https://doi.org/10.1364/OE.20.029457) (cit. on pp. 75, 87).
- Chen, Ricky T. Q. et al. (2019). *Neural Ordinary Differential Equations*. arXiv: [1806.07366](https://arxiv.org/abs/1806.07366) [[cs.LG](#)]. URL: <https://arxiv.org/abs/1806.07366> (cit. on p. 36).
- Cittert, P.H van (1934). “Die Wahrscheinliche Schwingungsverteilung in Einer von Einer Lichtquelle Direkt Oder Mittels Einer Linse Beleuchteten Ebene”. In: *Physica* 1.1, pp. 201–210. ISSN: 0031-8914. DOI: [https://doi.org/10.1016/S0031-8914\(34\)90026-4](https://doi.org/10.1016/S0031-8914(34)90026-4). URL: <https://www.sciencedirect.com/science/article/pii/S0031891434900264> (cit. on p. 22).
- Codona, J. L. et al. (June 2006). “A high-contrast coronagraph for the MMT using phase apodization: design and observations at 5 microns and 2 λ/D radius”. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by Ian S. McLean and Masanori Iye. Vol. 6269. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. International Society for Optics and Photonics. SPIE, 62691N, 62691N. DOI: [10.1117/12.672727](https://doi.org/10.1117/12.672727). URL: <https://doi.org/10.1117/12.672727> (cit. on p. 57).
- Coe, Dan (June 2009). “Fisher Matrices and Confidence Ellipses: A Quick-Start Guide and Software”. In: *arXiv e-prints*, arXiv:0906.4123, arXiv:0906.4123. arXiv: [0906.4123](https://arxiv.org/abs/0906.4123) [[astro-ph.IM](#)] (cit. on pp. 53, 60).
- Collaboration, Euclid et al. (Jan. 2022). “Euclid preparation. XVI. Exploring the ultra-low surface brightness Universe with Euclid/VIS”. In: 657, A92, A92. DOI: [10.1051/0004-6361/202141935](https://doi.org/10.1051/0004-6361/202141935). arXiv: [2108.10321](https://arxiv.org/abs/2108.10321) [[astro-ph.IM](#)] (cit. on p. 42).
- Costa, Joaquim F. Pinto da (2011). “Weighted Correlation”. In: *International Encyclopedia of Statistical Science*. Ed. by Miodrag Lovric. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 1653–1655. ISBN: 978-3-642-04898-2. DOI: [10.1007/978-3-642-04898-2_612](https://doi.org/10.1007/978-3-642-04898-2_612). URL: https://doi.org/10.1007/978-3-642-04898-2_612 (cit. on p. 51).
- Coulton, William R. and Wandelt, Benjamin D. (May 2023). “How to estimate Fisher matrices from simulations”. In: *arXiv e-prints*, arXiv:2305.08994, arXiv:2305.08994. arXiv: [2305.08994](https://arxiv.org/abs/2305.08994) [[stat.ME](#)] (cit. on p. 58).
- Cramé, Haraldr (1947). “Mathematical methods of statistics”. In: (cit. on pp. 58, 128).

- Currie, Thayne et al. (Apr. 2018). “Laboratory and On-sky Validation of the Shaped Pupil Coronagraph’s Sensitivity to Low-order Aberrations With Active Wavefront Control”. In: 130.986, p. 044505. DOI: [10.1088/1538-3873/aaab41](https://doi.org/10.1088/1538-3873/aaab41). arXiv: [1801.09760](https://arxiv.org/abs/1801.09760) [astro-ph.IM] (cit. on p. 72).
- Currie, Thayne et al. (2023). “Direct imaging and astrometric detection of a gas giant planet orbiting an accelerating star”. In: *Science* 380.6641, pp. 198–203. DOI: [10.1126/science.abo6192](https://doi.org/10.1126/science.abo6192). eprint: <https://www.science.org/doi/pdf/10.1126/science.abo6192>. URL: <https://www.science.org/doi/abs/10.1126/science.abo6192> (cit. on p. 6).
- Cybenko, G. (1989). “Approximation by superpositions of a sigmoidal function”. In: *Mathematics of Control, Signals and Systems* 2.4, pp. 303–314. DOI: [10.1007/BF02551274](https://doi.org/10.1007/BF02551274). URL: <https://doi.org/10.1007/BF02551274> (cit. on p. 31).
- Czekala, Ian et al. (Jan. 2025). “Million Points of Light (MPoL): a PyTorch library for radio interferometric imaging and inference”. In: *arXiv e-prints*, arXiv:2502.00100, arXiv:2502.00100. DOI: [10.48550/arXiv.2502.00100](https://doi.org/10.48550/arXiv.2502.00100). arXiv: [2502.00100](https://arxiv.org/abs/2502.00100) [astro-ph.IM] (cit. on p. 124).
- D d’Assignies, William et al. (May 2023). “Cosmological Fisher forecasts for next-generation spectroscopic surveys”. In: 521.3, pp. 3648–3662. DOI: [10.1093/mnras/stad611](https://doi.org/10.1093/mnras/stad611). arXiv: [2301.02289](https://arxiv.org/abs/2301.02289) [astro-ph.CO] (cit. on p. 58).
- De Furio, Matthew et al. (May 2023). “JWST Observations of the Enigmatic Y-Dwarf WISE 1828+2650. I. Limits to a Binary Companion”. In: 948.2, 92, p. 92. DOI: [10.3847/1538-4357/acbf1e](https://doi.org/10.3847/1538-4357/acbf1e). arXiv: [2302.12723](https://arxiv.org/abs/2302.12723) [astro-ph.SR] (cit. on pp. 125, 132).
- Desdoigts, Louis, Pope, Benjamin, and Tuthill, Peter (Aug. 2022). “Optical design, analysis, and calibration using ∂ Lux”. In: *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*. Ed. by Laura E. Coyle, Marshall D. Perrin, and Shuji Matsuura. SPIE. DOI: [10.1117/12.2629774](https://doi.org/10.1117/12.2629774). URL: <https://doi.org/10.1117/12.2629774> (cit. on p. 41).
- Desdoigts, Louis et al. (June 2023). “Differentiable optics with ∂ Lux: I—deep calibration of flat field and phase retrieval with automatic differentiation”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 9.02. ISSN: 2329-4124. DOI: [10.1117/1.jatis.9.2.028007](https://doi.org/10.1117/1.jatis.9.2.028007). URL: <https://doi.org/10.1117/1.jatis.9.2.028007> (cit. on pp. 39, 58, 78, 84).
- Desdoigts, Louis et al. (2025). *Differentiable Optics with d Lux II: Optical Design Maximising Fisher Information*. arXiv: [2406.08704](https://arxiv.org/abs/2406.08704) [astro-ph.IM]. URL: <https://arxiv.org/abs/2406.08704> (cit. on pp. 56, 78).
- Dholakia, Shashank, Dholakia, Shishir, and Pope, Benjamin J. S. (Oct. 2024). “A General, Differentiable Transit Model for Ellipsoidal Occulters: Derivation, Application, and Forecast of Planetary Oblateness and Obliquity Constraints with JWST”. In: *arXiv e-prints*, arXiv:2410.03449, arXiv:2410.03449. DOI: [10.48550/arXiv.2410.03449](https://doi.org/10.48550/arXiv.2410.03449). arXiv: [2410.03449](https://arxiv.org/abs/2410.03449) [astro-ph.EP] (cit. on p. 78).
- Dia, Noé et al. (Jan. 2025). “IRIS: A Bayesian Approach for Image Reconstruction in Radio Interferometry with expressive Score-Based priors”. In: *arXiv e-prints*, arXiv:2501.02473, arXiv:2501.02473. DOI: [10.48550/arXiv.2501.02473](https://doi.org/10.48550/arXiv.2501.02473). arXiv: [2501.02473](https://arxiv.org/abs/2501.02473) [astro-ph.IM] (cit. on p. 124).

- Díaz Baso, C. J., Asensio Ramos, A., and de la Cruz Rodríguez, J. (Mar. 2022). “Bayesian Stokes inversion with normalizing flows”. In: 659, A165, A165. DOI: [10.1051/0004-6361/202142018](https://doi.org/10.1051/0004-6361/202142018). arXiv: [2108.07089](https://arxiv.org/abs/2108.07089) [astro-ph.SR] (cit. on p. 129).
- Doyon, Rene et al. (2023a). “The Near Infrared Imager and Slitless Spectrograph for the James Webb Space Telescope – I. Instrument Overview and in-Flight Performance”. In: arXiv: [2306.03277](https://arxiv.org/abs/2306.03277) [astro-ph.IM]. URL: <https://arxiv.org/abs/2306.03277> (cit. on p. 75).
- Doyon, René et al. (2012). “The JWST Fine Guidance Sensor (FGS) and Near-Infrared Imager and Slitless Spectrograph (NIRISS)”. In: *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*. Ed. by Mark C. Clampin et al. Vol. 8442. International Society for Optics and Photonics. SPIE, 84422R. DOI: [10.1117/12.926578](https://doi.org/10.1117/12.926578). URL: <https://doi.org/10.1117/12.926578> (cit. on p. 75).
- Doyon, René et al. (Sept. 2023b). “The Near Infrared Imager and Slitless Spectrograph for the James Webb Space Telescope. I. Instrument Overview and In-flight Performance”. In: *Publications of the Astronomical Society of the Pacific* 135.1051, p. 098001. DOI: [10.1088/1538-3873/acd41b](https://doi.org/10.1088/1538-3873/acd41b). URL: <https://dx.doi.org/10.1088/1538-3873/acd41b> (cit. on pp. 12, 73, 128).
- Duane, Simon et al. (1987). “Hybrid Monte Carlo”. In: *Physics Letters B* 195.2, pp. 216–222. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X). URL: <https://www.sciencedirect.com/science/article/pii/037026938791197X> (cit. on p. 31).
- Dube, Brandon (2019). “prysm: A Python optics module”. In: *Journal of Open Source Software* 4.37, p. 1352. DOI: [10.21105/joss.01352](https://doi.org/10.21105/joss.01352). URL: <https://doi.org/10.21105/joss.01352> (cit. on pp. 57, 84).
- Ewart, Paul (2019). *Optics*. 2053-2571. Morgan & Claypool Publishers. ISBN: 978-1-64327-676-2. DOI: [10.1088/2053-2571/ab2231](https://doi.org/10.1088/2053-2571/ab2231). URL: <https://dx.doi.org/10.1088/2053-2571/ab2231> (cit. on p. 17).
- Fedorov, V.V. (1972). *Theory of Optimal Experiments*. Cellular Neurobiology. Academic Press. ISBN: 9780122507502. URL: <https://books.google.com.au/books?id=v6vTAvqGny4C> (cit. on p. 58).
- Feng, Berthy T. et al. (Apr. 2023). “Score-Based Diffusion Models as Principled Priors for Inverse Imaging”. In: *arXiv e-prints*, arXiv:2304.11751, arXiv:2304.11751. DOI: [10.48550/arXiv.2304.11751](https://doi.org/10.48550/arXiv.2304.11751). arXiv: [2304.11751](https://arxiv.org/abs/2304.11751) [cs.CV] (cit. on p. 124).
- Fienup, J. R. et al. (Apr. 1993). “Hubble Space Telescope characterized by using phase-retrieval algorithms”. In: 32.10, pp. 1747–1767. DOI: [10.1364/AO.32.001747](https://doi.org/10.1364/AO.32.001747) (cit. on p. 41).
- Fisher, R. A. (1925). “Theory of Statistical Estimation”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 22.5, pp. 700–725. DOI: [10.1017/S0305004100009580](https://doi.org/10.1017/S0305004100009580) (cit. on pp. 32, 59).
- Fizeau, Hippolyte (1868). “Prix Bordin: Rapport sur le concours de l’année 1867”. In: *Comptes Rendus de l’Académie des Sciences* 66, p. 932 (cit. on p. 9).
- Fletcher, R. (Jan. 1970). “A new approach to variable metric algorithms”. In: *The Computer Journal* 13.3, pp. 317–322. ISSN: 0010-4620. DOI: [10.1093/comjnl/13.3.317](https://doi.org/10.1093/comjnl/13.3.317). eprint: <https://academic.oup.com/comjnl/article-pdf/13/3/317/988678/130317.pdf>. URL: <https://doi.org/10.1093/comjnl/13.3.317> (cit. on pp. 31, 107).

- Follette, Katherine B (2023a). *An Introduction to High Contrast Differential Imaging of Exoplanets and Disks*. arXiv: 2308.01354 [astro-ph.IM]. URL: <https://arxiv.org/abs/2308.01354> (cit. on p. 75).
- (Sept. 2023b). “An Introduction to High Contrast Differential Imaging of Exoplanets and Disks”. In: *Publications of the Astronomical Society of the Pacific* 135.1051, p. 093001. DOI: 10.1088/1538-3873/aceb31. URL: <https://dx.doi.org/10.1088/1538-3873/aceb31> (cit. on pp. 7, 8).
- Foreman-Mackey, Daniel (2023). *dfm/tinygp: The tiniest of Gaussian Process libraries*. DOI: 10.5281/ZENODO.7646759. URL: <https://zenodo.org/record/7646759> (cit. on p. 78).
- Frater, R. H. et al. (Jan. 1987). “High Resolution Interferometric Imaging Using a Large Optical Telescope”. In: *Image Detection and Quality*. Ed. by Lucien F. Guyot. Vol. 702. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 255. DOI: 10.1117/12.966771 (cit. on p. 10).
- Fulton, Benjamin J. and Petigura, Erik A. (Nov. 2018). “The California-Kepler Survey. VII. Precise Planet Radii Leveraging Gaia DR2 Reveal the Stellar Mass Dependence of the Planet Radius Gap”. In: *The Astronomical Journal* 156.6, p. 264. DOI: 10.3847/1538-3881/aae828. URL: <https://dx.doi.org/10.3847/1538-3881/aae828> (cit. on p. 4).
- Gaia Collaboration et al. (2016). “The Gaia mission★”. In: *A&A* 595, A1. DOI: 10.1051/0004-6361/201629272. URL: <https://doi.org/10.1051/0004-6361/201629272> (cit. on p. 5).
- Gallenne, A. et al. (July 2015). “Robust high-contrast companion detection from interferometric observations. The CANDID algorithm and an application to six binary Cepheids”. In: 579, A68, A68. DOI: 10.1051/0004-6361/201525917. arXiv: 1505.02715 [astro-ph.SR] (cit. on p. 109).
- Gardner, Jonathan P. et al. (Apr. 2006). “The James Webb Space Telescope”. In: *Space Science Reviews* 123.4, pp. 485–606. DOI: 10.1007/s11214-006-8315-7. arXiv: astro-ph/0606175 [astro-ph]. URL: <https://doi.org/10.1007/s11214-006-8315-7> (cit. on pp. 1, 38, 53, 73, 75, 128).
- Gardner, Jonathan P. et al. (June 2023). “The James Webb Space Telescope Mission”. In: 135.1048, 068001, p. 068001. DOI: 10.1088/1538-3873/acd1b5. arXiv: 2304.04869 [astro-ph.IM]. URL: <https://dx.doi.org/10.1088/1538-3873/acd1b5> (cit. on p. 75).
- Gaudi, B. Scott et al. (Jan. 2020). “The Habitable Exoplanet Observatory (HabEx) Mission Concept Study Final Report”. In: *arXiv e-prints*, arXiv:2001.06683, arXiv:2001.06683. DOI: 10.48550/arXiv.2001.06683. arXiv: 2001.06683 [astro-ph.IM] (cit. on p. 7).
- Gerchberg, R. W. and Saxton, W. O. (1972). “A Practical Algorithm for the Determination of Phase from Image and Diffraction Plane Pictures”. In: *Optik* 35.2, pp. 237–246. ISSN: 0030-4026 (cit. on p. 41).
- Gershman, Samuel J. and Goodman, Noah D. (2014). “Amortized Inference in Probabilistic Reasoning”. In: *Cognitive Science* 36. URL: <https://api.semanticscholar.org/CorpusID:924780> (cit. on p. 129).
- Gilmozzi, R. and Spyromilio, J. (Mar. 2007). “The European Extremely Large Telescope (E-ELT)”. In: *The Messenger* 127, p. 11 (cit. on p. 7).

- Girard, Julien H. et al. (Aug. 2022). “JWST/NIRCam coronagraphy: commissioning and first on-sky results”. In: *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*. Ed. by Laura E. Coyle, Shuji Matsuura, and Marshall D. Perrin. Vol. 12180. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 121803Q, 121803Q. DOI: [10.1117/12.2629636](https://doi.org/10.1117/12.2629636). arXiv: [2208.00998](https://arxiv.org/abs/2208.00998) [astro-ph.IM] (cit. on p. 40).
- Goldfarb, Donald (1970). “A Family of Variable-Metric Methods Derived by Variational Means”. In: *Mathematics of Computation* 24, pp. 23–26. DOI: <https://doi.org/10.1090/S0025-5718-1970-0258249-6> (cit. on pp. 31, 107).
- Goodman, Joseph W. (2005). *Introduction to Fourier Optics*. 3rd. Roberts and Company Publishers (cit. on pp. 18, 22, 23, 88).
- Goudfrooij, Paul et al. (Jan. 2024). “An Algorithm to Mitigate Charge Migration Effects in Data from the Near Infrared Imager and Slitless Spectrograph on the James Webb Space Telescope*”. In: *Publications of the Astronomical Society of the Pacific* 136.1, p. 014503. ISSN: 1538-3873. DOI: [10.1088/1538-3873/ad1c98](https://doi.org/10.1088/1538-3873/ad1c98). URL: <http://dx.doi.org/10.1088/1538-3873/ad1c98> (cit. on pp. 15, 77).
- GRAVITY Collaboration et al. (2018). “Detection of the gravitational redshift in the orbit of the star S2 near the Galactic centre massive black hole”. In: *A&A* 615, p. L15. DOI: [10.1051/0004-6361/201833718](https://doi.org/10.1051/0004-6361/201833718). URL: <https://doi.org/10.1051/0004-6361/201833718> (cit. on p. 5).
- GRAVITY Collaboration et al. (Apr. 2020). “Detection of the Schwarzschild precession in the orbit of the star S2 near the Galactic centre massive black hole”. In: 636, L5, p. L5. DOI: [10.1051/0004-6361/202037813](https://doi.org/10.1051/0004-6361/202037813). arXiv: [2004.07187](https://arxiv.org/abs/2004.07187) [astro-ph.GA] (cit. on p. 5).
- Greenbaum, Alexandra Z. et al. (Jan. 2015). “An Image-plane Algorithm for JWST’s Non-redundant Aperture Mask Data”. In: 798.2, 68, p. 68. DOI: [10.1088/0004-637X/798/2/68](https://doi.org/10.1088/0004-637X/798/2/68). arXiv: [1411.3446](https://arxiv.org/abs/1411.3446) [astro-ph.IM] (cit. on pp. 13, 76).
- Greenbaum, Alexandra Z. et al. (June 2019). “Performance of the Gemini Planet Imager Non-redundant Mask and Spectroscopy of Two Close-separation Binaries: HR 2690 and HD 142527”. In: 157.6, 249, p. 249. DOI: [10.3847/1538-3881/ab17db](https://doi.org/10.3847/1538-3881/ab17db). arXiv: [1904.09006](https://arxiv.org/abs/1904.09006) [astro-ph.IM] (cit. on p. 76).
- Griewank, A. and Walther, A. (2008). *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation, Second Edition*. Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104). ISBN: 9780898717761. URL: <https://books.google.com.au/books?id=x0iiLaRxcbEC> (cit. on p. 78).
- Griewank, Andreas (1988). “On automatic differentiation”. In: URL: <https://api.semanticscholar.org/CorpusID:264691130> (cit. on p. 26).
- Gully-Santiago, Michael and Morley, Caroline V. (Dec. 2022). “An Interpretable Machine-learning Framework for Modeling High-resolution Spectroscopic Data”. In: 941.2, 200, p. 200. DOI: [10.3847/1538-4357/aca0a2](https://doi.org/10.3847/1538-4357/aca0a2). arXiv: [2210.01827](https://arxiv.org/abs/2210.01827) [astro-ph.IM] (cit. on p. 78).
- Guyon, O. et al. (Nov. 2006). “Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs”. In: 167.1, pp. 81–99. ISSN: 1538-4365. DOI: [10.1086/507630](https://doi.org/10.1086/507630). arXiv:

- [astro-ph/0608506](https://doi.org/10.1086/507630) [astro-ph]. URL: <http://dx.doi.org/10.1086/507630> (cit. on pp. 6, 9, 57, 75).
- Guyon, Olivier (Aug. 2005). “Limits of Adaptive Optics for High-Contrast Imaging”. In: *The Astrophysical Journal* 629.1, p. 592. DOI: [10.1086/431209](https://doi.org/10.1086/431209). URL: <https://dx.doi.org/10.1086/431209> (cit. on p. 6).
- Guyon, Olivier et al. (June 2012). “High-precision Astrometry with a Diffractive Pupil Telescope”. In: *The Astrophysical Journal Supplement Series* 200.2, 11, p. 11. DOI: [10.1088/0067-0049/200/2/11](https://doi.org/10.1088/0067-0049/200/2/11). arXiv: [1304.0370](https://arxiv.org/abs/1304.0370) [astro-ph.IM] (cit. on pp. 53, 57, 61).
- Guyon, Olivier et al. (Apr. 2013). “Simultaneous Exoplanet Characterization and Deep Wide-field Imaging with a Diffractive Pupil Telescope”. In: 767.1, 11, p. 11. DOI: [10.1088/0004-637X/767/1/11](https://doi.org/10.1088/0004-637X/767/1/11). arXiv: [1304.0373](https://arxiv.org/abs/1304.0373) [astro-ph.IM] (cit. on p. 53).
- Guyonnet, A. et al. (Feb. 2015). “Evidence for self-interaction of charge distribution in charge-coupled devices”. In: *Astronomy & Astrophysics* 575, A41. ISSN: 1432-0746. DOI: [10.1051/0004-6361/201424897](https://doi.org/10.1051/0004-6361/201424897). URL: <http://dx.doi.org/10.1051/0004-6361/201424897> (cit. on pp. 15, 77).
- Haffert, S. Y. et al. (2019). “Two accreting protoplanets around the young star PDS 70”. In: *Nature Astronomy* 3.8, pp. 749–754. DOI: [10.1038/s41550-019-0780-5](https://doi.org/10.1038/s41550-019-0780-5). URL: <https://doi.org/10.1038/s41550-019-0780-5> (cit. on p. 6).
- Hanbury Brown, R. (Nov. 1956). “A Test of a New Type of Stellar Interferometer on Sirius”. In: 178.4541, pp. 1046–1048. DOI: [10.1038/1781046a0](https://doi.org/10.1038/1781046a0) (cit. on p. 9).
- Hanbury Brown, R. and Twiss, R.Q. (July 1954). “LXXIV. A new type of interferometer for use in radio astronomy”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 45.366, pp. 663–682. ISSN: 1941-5990. DOI: [10.1080/14786440708520475](https://doi.org/10.1080/14786440708520475). URL: <http://dx.doi.org/10.1080/14786440708520475> (cit. on p. 9).
- Hanbury Brown, R. et al. (June 1974). “The effects of limb darkening on measurements of angular size with an intensity interferometer”. In: 167, pp. 475–484. DOI: [10.1093/mnras/167.3.475](https://doi.org/10.1093/mnras/167.3.475) (cit. on p. 9).
- Haniff, Chris (2007). “An introduction to the theory of interferometry”. In: *New Astronomy Reviews* 51.8. Observation and Data Reduction with the VLT Interferometer, pp. 565–575. ISSN: 1387-6473. DOI: <https://doi.org/10.1016/j.newar.2007.06.002>. URL: <https://www.sciencedirect.com/science/article/pii/S1387647307000619> (cit. on p. 24).
- Hara, Nathan C. and Ford, Eric B. (2023). “Statistical Methods for Exoplanet Detection with Radial Velocities”. In: *Annual Review of Statistics and Its Application* 10. Volume 10, 2023, pp. 623–649. ISSN: 2326-831X. DOI: <https://doi.org/10.1146/annurev-statistics-033021-012225>. URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-statistics-033021-012225> (cit. on p. 3).
- Harris, Charles R. et al. (2020). “Array programming with NumPy”. In: *Nature* 585.7825, pp. 357–362 (cit. on pp. 42, 54, 72, 126).
- Hastings, W. K. (Apr. 1970). “Monte Carlo sampling methods using Markov chains and their applications”. In: *Biometrika* 57.1, pp. 97–109. ISSN: 0006-3444. DOI: [10.1093/biomet/57.1.97](https://doi.org/10.1093/biomet/57.1.97). eprint: <https://academic.oup.com/biomet/article-pdf/57/1/97/23940249/57-1-97.pdf>. URL: <https://doi.org/10.1093/biomet/57.1.97> (cit. on p. 31).

- Hattori, Soichiro et al. (Mar. 2024). *exoplanet-dev/jaxoplanet: Astronomical time series analysis with JAX*. Version v0.0.2. DOI: [10.5281/zenodo.10736936](https://doi.org/10.5281/zenodo.10736936). URL: <https://doi.org/10.5281/zenodo.10736936> (cit. on p. 78).
- Hecht, Eugene (2002). *Optics*. 4th ed. San Francisco: Addison-Wesley. ISBN: 9780321188786 (cit. on p. 88).
- Hedges, Christina et al. (Sept. 2021). “Linearized Field Deblending: Point-spread Function Photometry for Impatient Astronomers”. In: 162.3, 107, p. 107. DOI: [10.3847/1538-3881/ac0825](https://doi.org/10.3847/1538-3881/ac0825). arXiv: [2106.08411](https://arxiv.org/abs/2106.08411) [astro-ph.IM] (cit. on p. 53).
- Hessel, Matteo et al. (2020). *Optax: composable gradient transformation and optimisation, in JAX!* Version 0.0.1. URL: <http://github.com/deepmind/optax> (cit. on pp. 43, 46, 54, 72, 126).
- Higgins, Michael E. and Bell, Keaton J. (Apr. 2022). “Localizing Sources of Variability in Crowded TESS Photometry”. In: *arXiv e-prints*, arXiv:2204.06020, arXiv:2204.06020. arXiv: [2204.06020](https://arxiv.org/abs/2204.06020) [astro-ph.IM] (cit. on p. 53).
- Hinkley, S. et al. (Mar. 2023). “Direct discovery of the inner exoplanet in the HD 206893 system. Evidence for deuterium burning in a planetary-mass companion”. In: 671, L5, p. L5. DOI: [10.1051/0004-6361/202244727](https://doi.org/10.1051/0004-6361/202244727). arXiv: [2208.04867](https://arxiv.org/abs/2208.04867) [astro-ph.EP] (cit. on pp. 114, 116, 121).
- Hinton, S. R. (Aug. 2016). “ChainConsumer”. In: *The Journal of Open Source Software* 1, 00045, p. 00045. DOI: [10.21105/joss.00045](https://doi.org/10.21105/joss.00045) (cit. on pp. 72, 126).
- Hinz, P. M. et al. (Aug. 2016). “Overview of LBTI: a multipurpose facility for high spatial resolution observations”. In: *Optical and Infrared Interferometry and Imaging V*. Ed. by Fabien Malbet, Michelle J. Creech-Eakman, and Peter G. Tuthill. Vol. 9907. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 990704, p. 990704. DOI: [10.1117/12.2233795](https://doi.org/10.1117/12.2233795) (cit. on p. 7).
- Hoffman, Matt et al. (2013). *Stochastic Variational Inference*. arXiv: [1206.7051](https://arxiv.org/abs/1206.7051) [stat.ML] (cit. on p. 72).
- Hogg, David W. and Villar, Soledad (2024). *Is machine learning good or bad for the natural sciences?* arXiv: [2405.18095](https://arxiv.org/abs/2405.18095) [stat.ML]. URL: <https://arxiv.org/abs/2405.18095> (cit. on pp. 34, 79, 127).
- Hornik, Kurt (1991). “Approximation capabilities of multilayer feedforward networks”. In: *Neural Networks* 4.2, pp. 251–257. ISSN: 0893-6080. DOI: [https://doi.org/10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T). URL: <https://www.sciencedirect.com/science/article/pii/089360809190009T> (cit. on p. 31).
- Horta, Danny et al. (June 2025). “Lux: A Generative, Multioutput, Latent-variable Model for Astronomical Data with Noisy Labels”. In: 169.6, 314, p. 314. DOI: [10.3847/1538-3881/adcc20](https://doi.org/10.3847/1538-3881/adcc20). arXiv: [2502.01745](https://arxiv.org/abs/2502.01745) [astro-ph.IM] (cit. on p. 78).
- Howell, Steve B. et al. (Apr. 2014). “The K2 Mission: Characterization and Early Results”. In: *Publications of the Astronomical Society of the Pacific* 126.938, pp. 398–408. ISSN: 1538-3873. DOI: [10.1086/676406](https://doi.org/10.1086/676406). URL: <http://dx.doi.org/10.1086/676406> (cit. on p. 41).
- Huijser, David, Goodman, Jesse, and Brewer, Brendon J. (Feb. 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. DOI: [10.1111/anzs.12358](https://doi.org/10.1111/anzs.12358). URL: <https://doi.org/10.1111/anzs.12358> (cit. on pp. 31, 42, 57).
- Hunter, J. D. (2007). “Matplotlib: A 2D graphics environment”. In: *Computing In Science & Engineering* 9.3, pp. 90–95 (cit. on pp. 54, 72, 126).
- Huygens, Christiaan (1690). *Treatise on Light*. Original work published 1690. Macmillan (cit. on pp. 17, 20).
- Ireland, M. J. (Aug. 2013). “Phase errors in diffraction-limited imaging: contrast limits for sparse aperture masking”. In: 433.2, pp. 1718–1728. DOI: [10.1093/mnras/stt859](https://doi.org/10.1093/mnras/stt859). arXiv: [1301.6205](https://arxiv.org/abs/1301.6205) [astro-ph.IM] (cit. on pp. 9, 10, 94, 107–109, 122).
- Isella, Andrea et al. (Dec. 2016). “Ringed Structures of the HD 163296 Protoplanetary Disk Revealed by ALMA”. In: *Phys. Rev. Lett.* 117 (25), p. 251101. DOI: [10.1103/PhysRevLett.117.251101](https://doi.org/10.1103/PhysRevLett.117.251101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.251101> (cit. on p. 6).
- Işil, Çağatay, Oktem, Figen S., and Koç, Aykut (July 2019). “Deep iterative reconstruction for phase retrieval”. In: *Appl. Opt.* 58.20, pp. 5422–5431. DOI: [10.1364/AO.58.005422](https://doi.org/10.1364/AO.58.005422). URL: <http://ao.osa.org/abstract.cfm?URI=ao-58-20-5422> (cit. on p. 41).
- Jackiewicz, Jason (Mar. 2021). “Solar-Like Oscillators in the Kepler Era: A Review”. In: *Frontiers in Astronomy and Space Sciences* 7. DOI: [10.3389/fspas.2020.595017](https://doi.org/10.3389/fspas.2020.595017). URL: <https://doi.org/10.3389/fspas.2020.595017> (cit. on p. 41).
- Jefferys, W. H. et al. (Nov. 1985). “Prospects for astrometry with the Hubble Space Telescope”. In: *Celestial Mechanics* 37.3, pp. 299–305. DOI: [10.1007/BF02285054](https://doi.org/10.1007/BF02285054) (cit. on p. 75).
- Jennison, R. C. (Jan. 1958). “A phase sensitive interferometer technique for the measurement of the Fourier transforms of spatial brightness distributions of small angular extent”. In: 118, p. 276. DOI: [10.1093/mnras/118.3.276](https://doi.org/10.1093/mnras/118.3.276) (cit. on pp. 9, 24, 75, 92).
- Johns, Matt (Apr. 2008). “The Giant Magellan Telescope (GMT)”. In: *Extremely Large Telescopes: Which Wavelengths? Retirement Symposium for Arne Ardeberg*. Ed. by Torben E. Andersen. Vol. 6986. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 698603, p. 698603. DOI: [10.1117/12.801254](https://doi.org/10.1117/12.801254) (cit. on p. 7).
- Jordan, Michael I. et al. (1999). “An Introduction to Variational Methods for Graphical Models”. In: *Machine Learning* 37.2, pp. 183–233. DOI: [10.1023/A:1007665907178](https://doi.org/10.1023/A:1007665907178). URL: <https://doi.org/10.1023/A:1007665907178> (cit. on p. 133).
- Jovanovic, N. et al. (Sept. 2015). “The Subaru Coronagraphic Extreme Adaptive Optics System: Enabling High-Contrast Imaging on Solar-System Scales”. In: *Publications of the Astronomical Society of the Pacific* 127.955, p. 890. DOI: [10.1086/682989](https://doi.org/10.1086/682989). URL: <https://dx.doi.org/10.1086/682989> (cit. on p. 7).
- Jurling, Alden S. and Fienup, James R. (July 2014). “Applications of algorithmic differentiation to phase retrieval algorithms”. In: *J. Opt. Soc. Am. A* 31.7, pp. 1348–1359. DOI: [10.1364/JOSAA.31.001348](https://doi.org/10.1364/JOSAA.31.001348). URL: <http://josaa.osa.org/abstract.cfm?URI=josaa-31-7-1348> (cit. on pp. 41, 58).
- Kammerer, Jens et al. (June 2019). “Kernel phase imaging with VLT/NACO: high-contrast detection of new candidate low-mass stellar companions at the diffraction limit”. In: 486.1,

- pp. 639–654. DOI: [10.1093/mnras/stz882](https://doi.org/10.1093/mnras/stz882). arXiv: [1903.11252](https://arxiv.org/abs/1903.11252) [astro-ph.EP] (cit. on p. 42).
- Kammerer, Jens et al. (Oct. 2022). “The Near Infrared Imager and Slitless Spectrograph for JWST – V. Kernel Phase Imaging and Data Analysis”. In: *arXiv e-prints*, arXiv:2210.17528, arXiv:2210.17528. arXiv: [2210.17528](https://arxiv.org/abs/2210.17528) [astro-ph.IM] (cit. on p. 40).
- Kammerer, Jens et al. (Feb. 2023). “The Near Infrared Imager and Slitless Spectrograph for JWST. V. Kernel Phase Imaging and Data Analysis”. In: *Publications of the Astronomical Society of the Pacific* 135.1043, 014502, p. 014502. DOI: [10.1088/1538-3873/ac9a74](https://doi.org/10.1088/1538-3873/ac9a74). arXiv: [2210.17528](https://arxiv.org/abs/2210.17528) [astro-ph.IM]. URL: <https://dx.doi.org/10.1088/1538-3873/ac9a74> (cit. on pp. 76, 132).
- Karniadakis, George Em et al. (2021). “Physics-informed machine learning”. In: *Nature Reviews Physics* 3.6, pp. 422–440. DOI: [10.1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5). URL: <https://doi.org/10.1038/s42254-021-00314-5> (cit. on pp. 79, 95).
- Kass, Robert, Tierney, Luke, and Kadane, Joseph (1991). *Laplace’s method in Bayesian analysis*. DOI: [10.1090/conm/115/07](https://doi.org/10.1090/conm/115/07). URL: <https://doi.org/10.1090/conm/115/07> (cit. on pp. 32, 58, 102).
- Kenworthy, Matthew A. and Haffert, Sebastiaan Y. (May 2025). “High-Contrast Coronagraphy”. In: *Annual Review of Astronomy and Astrophysics*. ISSN: 1545-4282. DOI: [10.1146/annurev-astro-021225-022840](https://doi.org/10.1146/annurev-astro-021225-022840). URL: <http://dx.doi.org/10.1146/annurev-astro-021225-022840> (cit. on pp. 4, 7, 8).
- Keppler, M. et al. (2018). “Discovery of a planetary-mass companion within the gap of the transition disk around PDS 70”. In: *A&A* 617, A44. DOI: [10.1051/0004-6361/201832957](https://doi.org/10.1051/0004-6361/201832957). URL: <https://doi.org/10.1051/0004-6361/201832957> (cit. on p. 6).
- Kidger, Patrick (2021). “On Neural Differential Equations”. PhD thesis. University of Oxford (cit. on p. 32).
- (2022a). *On Neural Differential Equations*. arXiv: [2202.02435](https://arxiv.org/abs/2202.02435) [cs.LG]. URL: <https://arxiv.org/abs/2202.02435> (cit. on p. 36).
 - (2022b). *On Neural Differential Equations*. arXiv: [2202.02435](https://arxiv.org/abs/2202.02435) [cs.LG]. URL: <https://arxiv.org/abs/2202.02435> (cit. on p. 95).
- Kidger, Patrick and Garcia, Cristian (2021a). “Equinox: neural networks in JAX via callable PyTrees and filtered transformations”. In: *Differentiable Programming workshop at Neural Information Processing Systems 2021* (cit. on pp. 42, 54, 72).
- (2021b). *Equinox: neural networks in JAX via callable PyTrees and filtered transformations*. arXiv: [2111.00254](https://arxiv.org/abs/2111.00254) [cs.LG]. URL: <https://arxiv.org/abs/2111.00254> (cit. on pp. 78, 126).
- Kidger, Patrick et al. (2020). *Neural Controlled Differential Equations for Irregular Time Series*. arXiv: [2005.08926](https://arxiv.org/abs/2005.08926) [cs.LG]. URL: <https://arxiv.org/abs/2005.08926> (cit. on p. 36).
- Kingma, Diederik P. and Ba, Jimmy (Dec. 2014). “Adam: A Method for Stochastic Optimization”. In: *arXiv e-prints*, arXiv:1412.6980, arXiv:1412.6980. arXiv: [1412.6980](https://arxiv.org/abs/1412.6980) [cs.LG] (cit. on pp. 46, 103).
- (2017). *Adam: A Method for Stochastic Optimization*. arXiv: [1412.6980](https://arxiv.org/abs/1412.6980) [cs.LG] (cit. on p. 65).

- Kirchhoff, G. (1883). “Zur Theorie der Lichtstrahlen”. In: *Annalen der Physik* 254.4, pp. 663–695. DOI: <https://doi.org/10.1002/andp.18832540409>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.18832540409>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.18832540409> (cit. on pp. 17, 20).
- Kochkov, Dmitrii et al. (July 2024). “Neural general circulation models for weather and climate”. In: *Nature* 632.8027, pp. 1060–1066. ISSN: 1476-4687. DOI: [10.1038/s41586-024-07744-y](https://doi.org/10.1038/s41586-024-07744-y). URL: <http://dx.doi.org/10.1038/s41586-024-07744-y> (cit. on p. 96).
- Kraus, Adam L. et al. (July 2016). “The Impact of Stellar Multiplicity on Planetary Systems. I. The Ruinous Influence of Close Binary Companions”. In: 152.1, 8, p. 8. DOI: [10.3847/0004-6256/152/1/8](https://doi.org/10.3847/0004-6256/152/1/8). arXiv: [1604.05744](https://arxiv.org/abs/1604.05744) [astro-ph.EP] (cit. on pp. 5, 8).
- Krist, John E., Hook, Richard N., and Stoehr, Felix (2011). “20 years of Hubble Space Telescope optical modeling using Tiny Tim”. In: *Optical Modeling and Performance Predictions V*. Ed. by Mark A. Kahan. Vol. 8127. International Society for Optics and Photonics. SPIE, 81270J. DOI: [10.1117/12.892762](https://doi.org/10.1117/12.892762). URL: <https://doi.org/10.1117/12.892762> (cit. on p. 88).
- Kucukelbir, Alp et al. (2016). *Automatic Differentiation Variational Inference*. arXiv: [1603.00788](https://arxiv.org/abs/1603.00788) [stat.ML] (cit. on p. 72).
- (2017). “Automatic Differentiation Variational Inference”. In: *Journal of Machine Learning Research* 18.14, pp. 1–45. URL: <http://jmlr.org/papers/v18/16-107.html> (cit. on p. 133).
- Labeyrie, A. (May 1970). “Attainment of Diffraction Limited Resolution in Large Telescopes by Fourier Analysing Speckle Patterns in Star Images”. In: 6, p. 85 (cit. on p. 9).
- (Mar. 1975). “Interference fringes obtained on Vega with two optical telescopes.” In: 196, pp. L71–L75. DOI: [10.1086/181747](https://doi.org/10.1086/181747) (cit. on pp. 19, 24).
- Lafrenière, David et al. (May 2007). “A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging”. In: 660.1, pp. 770–780. DOI: [10.1086/513180](https://doi.org/10.1086/513180). arXiv: [astro-ph/0702697](https://arxiv.org/abs/astro-ph/0702697) [astro-ph] (cit. on p. 8).
- Lajoie, Charles-Philippe et al. (July 2023). *OTE Science Performance Memo 2 - A Year of Wavefront Sensing with JWST in Flight: Cycle 1 Telescope Monitoring & Maintenance Summary*. Technical Report JWST-STScI-008497. DOI: [10.48550/arXiv.2307.11179](https://doi.org/10.48550/arXiv.2307.11179). arXiv: [2307.11179](https://arxiv.org/abs/2307.11179) [astro-ph.IM] (cit. on p. 124).
- Lakshminarayanan, Vasudevan and Fleck, Andre (Apr. 2011). “Zernike polynomials: A guide”. In: *JModOptic* 58, pp. 1678–1678. DOI: [10.1080/09500340.2011.633763](https://doi.org/10.1080/09500340.2011.633763) (cit. on p. 87).
- Lau, Ryan M. et al. (2023). “A First Look with JWST Aperture Masking Interferometry (AMI): Resolving Circumstellar Dust around the Wolf-Rayet Binary WR 137 beyond the Rayleigh Limit”. In: arXiv: [2311.15948](https://arxiv.org/abs/2311.15948) [astro-ph.SR]. URL: <https://arxiv.org/abs/2311.15948> (cit. on pp. 76, 77).
- Lavin, Alexander et al. (2022). *Simulation Intelligence: Towards a New Generation of Scientific Methods*. arXiv: [2112.03235](https://arxiv.org/abs/2112.03235) [cs.AI] (cit. on p. 58).
- Lawson, Peter R., ed. (Jan. 2000). *Principles of Long Baseline Stellar Interferometry* (cit. on p. 8).

- LeCun, Y., Bengio, Y., and Hinton, G. (May 2015). “Deep learning”. In: *Nature* 521.7553, pp. 436–444. DOI: [10.1038/nature14539](https://doi.org/10.1038/nature14539). URL: <https://doi.org/10.1038/nature14539> (cit. on pp. 42, 57, 78).
- Levenberg, Kenneth (1944). “A METHOD FOR THE SOLUTION OF CERTAIN NON-LINEAR PROBLEMS IN LEAST SQUARES”. In: *Quarterly of Applied Mathematics* 2.2, pp. 164–168. ISSN: 0033569X, 15524485. URL: <http://www.jstor.org/stable/43633451> (visited on 06/18/2025) (cit. on p. 31).
- Liaudat, Tobias et al. (Nov. 2021). “Rethinking the modeling of the instrumental response of telescopes with a differentiable optical model”. In: *arXiv e-prints*, arXiv:2111.12541, arXiv:2111.12541. arXiv: [2111.12541](https://arxiv.org/abs/2111.12541) [[astro-ph.IM](https://arxiv.org/abs/2111.12541)] (cit. on pp. 43, 58).
- (Mar. 2022). “Rethinking data-driven point spread function modeling with a differentiable optical model”. In: *arXiv e-prints*, arXiv:2203.04908, arXiv:2203.04908. arXiv: [2203.04908](https://arxiv.org/abs/2203.04908) [[astro-ph.IM](https://arxiv.org/abs/2203.04908)] (cit. on p. 43).
- (Mar. 2023a). “Rethinking data-driven point spread function modeling with a differentiable optical model”. In: *Inverse Problems* 39.3, 035008, p. 035008. ISSN: 1361-6420. DOI: [10.1088/1361-6420/acb664](https://doi.org/10.1088/1361-6420/acb664). arXiv: [2203.04908](https://arxiv.org/abs/2203.04908) [[astro-ph.IM](https://arxiv.org/abs/2203.04908)]. URL: <http://dx.doi.org/10.1088/1361-6420/acb664> (cit. on pp. 35, 58, 71, 78).
- Liaudat, Tobías I., Starck, Jean-Luc, and Kilbinger, Martin (Oct. 2023b). “Point spread function modelling for astronomical telescopes: a review focused on weak gravitational lensing studies”. In: *Frontiers in Astronomy and Space Sciences* 10. ISSN: 2296-987X. DOI: [10.3389/fspas.2023.1158213](https://doi.org/10.3389/fspas.2023.1158213). URL: <http://dx.doi.org/10.3389/fspas.2023.1158213> (cit. on p. 35).
- Lin, Henry W., Tegmark, Max, and Rolnick, David (July 2017). “Why Does Deep and Cheap Learning Work So Well?” In: *Journal of Statistical Physics* 168.6, pp. 1223–1247. ISSN: 1572-9613. DOI: [10.1007/s10955-017-1836-5](https://doi.org/10.1007/s10955-017-1836-5). URL: <http://dx.doi.org/10.1007/s10955-017-1836-5> (cit. on p. 27).
- Ling, Julia, Kurzwaski, Andrew, and Templeton, Jeremy (2016). “Reynolds averaged turbulence modelling using deep neural networks with embedded invariance”. In: *Journal of Fluid Mechanics* 807, pp. 155–166. DOI: [10.1017/jfm.2016.615](https://doi.org/10.1017/jfm.2016.615) (cit. on p. 36).
- Liu, Haonan and Bunn, Emory F. (Jan. 2016). “Fisher matrix optimization of cosmic microwave background interferometers”. In: 93.2, 023512, p. 023512. DOI: [10.1103/PhysRevD.93.023512](https://doi.org/10.1103/PhysRevD.93.023512). arXiv: [1511.03635](https://arxiv.org/abs/1511.03635) [[astro-ph.CO](https://arxiv.org/abs/1511.03635)] (cit. on p. 58).
- Liu, Xuewen et al. (June 2020). “Wavefront prediction using artificial neural networks for open-loop adaptive optics”. In: *Monthly Notices of the Royal Astronomical Society* 496.1, pp. 456–464. ISSN: 1365-2966. DOI: [10.1093/mnras/staa1558](https://doi.org/10.1093/mnras/staa1558). URL: <http://dx.doi.org/10.1093/mnras/staa1558> (cit. on p. 34).
- Lucas, Miles et al. (Nov. 2024). “Visible-light High-contrast Imaging and Polarimetry with SCEXAO/VAMPIRES”. In: 136.11, 114504, p. 114504. DOI: [10.1088/1538-3873/ad89af](https://doi.org/10.1088/1538-3873/ad89af). arXiv: [2410.12199](https://arxiv.org/abs/2410.12199) [[astro-ph.IM](https://arxiv.org/abs/2410.12199)] (cit. on p. 76).
- Luger, Rodrigo et al. (Oct. 2016). “EVEREST: Pixel Level Decorrelation of K2 Light Curves”. In: *The Astronomical Journal* 152.4, p. 100. ISSN: 1538-3881. DOI: [10.3847/0004-6256/152/4/100](https://doi.org/10.3847/0004-6256/152/4/100). URL: <http://dx.doi.org/10.3847/0004-6256/152/4/100> (cit. on p. 41).

- Lyot, Bernard and Marshall, Roy K. (Aug. 1933). “The Study of the Solar Corona without an Eclipse”. In: 27, p. 225 (cit. on p. 7).
- Macintosh, Bruce et al. (2014). “First light of the Gemini Planet Imager”. In: *Proceedings of the National Academy of Sciences* 111.35, pp. 12661–12666. DOI: [10.1073/pnas.1304215111](https://doi.org/10.1073/pnas.1304215111). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1304215111>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1304215111> (cit. on p. 7).
- MacKay, David J. C. (2002). *Information Theory, Inference & Learning Algorithms*. USA: Cambridge University Press. ISBN: 0521642981 (cit. on pp. 58, 102).
- Margossian, Charles C. (Nov. 2018). “A Review of automatic differentiation and its efficient implementation”. In: *arXiv e-prints*, arXiv:1811.05031, arXiv:1811.05031. arXiv: [1811.05031](https://arxiv.org/abs/1811.05031) [cs.MS] (cit. on pp. 40, 57).
- (Mar. 2019). “A review of automatic differentiation and its efficient implementation”. In: *WIREs Data Mining and Knowledge Discovery* 9.4. ISSN: 1942-4795. DOI: [10.1002/widm.1305](https://doi.org/10.1002/widm.1305). URL: <http://dx.doi.org/10.1002/WIDM.1305> (cit. on pp. 25, 78, 127).
- Marois, Christian et al. (Apr. 2006). “Angular Differential Imaging: A Powerful High-Contrast Imaging Technique”. In: 641.1, pp. 556–564. DOI: [10.1086/500401](https://doi.org/10.1086/500401). arXiv: [astro-ph/0512335](https://arxiv.org/abs/astro-ph/0512335) [astro-ph] (cit. on p. 8).
- Marois, Christian et al. (2008). “Direct Imaging of Multiple Planets Orbiting the Star HR 8799”. In: *Science* 322.5906, pp. 1348–1352. DOI: [10.1126/science.1166585](https://doi.org/10.1126/science.1166585). eprint: <https://www.science.org/doi/pdf/10.1126/science.1166585>. URL: <https://www.science.org/doi/abs/10.1126/science.1166585> (cit. on pp. 4, 6).
- Marquardt, Donald W. (1963). “An Algorithm for Least-Squares Estimation of Nonlinear Parameters”. In: *Journal of the Society for Industrial and Applied Mathematics* 11.2, pp. 431–441. DOI: [10.1137/0111030](https://doi.org/10.1137/0111030). eprint: <https://doi.org/10.1137/0111030>. URL: <https://doi.org/10.1137/0111030> (cit. on p. 31).
- Martens, James (2020). *New insights and perspectives on the natural gradient method*. arXiv: [1412.1193](https://arxiv.org/abs/1412.1193) [cs.LG]. URL: <https://arxiv.org/abs/1412.1193> (cit. on pp. 32, 102).
- Martín Abadi et al. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. URL: <http://tensorflow.org/> (cit. on pp. 42, 43, 57).
- Martinache, Frantz (Nov. 2010). “Kernel Phase in Fizeau Interferometry”. In: 724.1, pp. 464–469. DOI: [10.1088/0004-637X/724/1/464](https://doi.org/10.1088/0004-637X/724/1/464). arXiv: [1009.3933](https://arxiv.org/abs/1009.3933) [astro-ph.IM]. URL: <https://dx.doi.org/10.1088/0004-637X/724/1/464> (cit. on pp. 9, 24, 40, 76, 91, 93).
- (Apr. 2013). “The Asymmetric Pupil Fourier Wavefront Sensor”. In: *Publications of the Astronomical Society of the Pacific* 125.926, p. 422. ISSN: 1538-3873. DOI: [10.1086/670670](https://doi.org/10.1086/670670). arXiv: [1303.6678](https://arxiv.org/abs/1303.6678) [astro-ph.IM]. URL: <http://dx.doi.org/10.1086/670670> (cit. on pp. 44, 91, 130).
- Martinache, Frantz and Ireland, Michael J. (Nov. 2018). “Kernel-nulling for a robust direct interferometric detection of extrasolar planets”. In: 619, A87, A87. DOI: [10.1051/0004-6361/201832847](https://doi.org/10.1051/0004-6361/201832847). arXiv: [1802.06252](https://arxiv.org/abs/1802.06252) [astro-ph.IM] (cit. on p. 9).
- Martinache, Frantz, Jovanovic, Nemanja, and Guyon, Olivier (Sept. 2016). “Closed-loop focal plane wavefront control with the SCExAO instrument”. In: SPIE Conference Series 593,

- A33. Ed. by Fabien Malbet, Michelle J. Creech-Eakman, and Peter G. Tuthill, A33. DOI: [10.1051/0004-6361/201628496](https://doi.org/10.1051/0004-6361/201628496). arXiv: [1604.08787](https://arxiv.org/abs/1604.08787) [astro-ph.IM] (cit. on pp. 124, 130).
- Martinache, Frantz et al. (Apr. 2020). “Kernel-phase analysis: Aperture modeling prescriptions that minimize calibration errors”. In: 636, A72, A72. DOI: [10.1051/0004-6361/201936981](https://doi.org/10.1051/0004-6361/201936981). arXiv: [2003.02032](https://arxiv.org/abs/2003.02032) [astro-ph.IM] (cit. on pp. 44, 95, 125).
- Mawet, D. et al. (Jan. 2010). “The Vector Vortex Coronagraph: Laboratory Results and First Light at Palomar Observatory”. In: 709.1, pp. 53–57. DOI: [10.1088/0004-637x/709/1/53](https://doi.org/10.1088/0004-637x/709/1/53). arXiv: [0912.2287](https://arxiv.org/abs/0912.2287) [astro-ph.IM] (cit. on p. 6).
- Mawet, D. et al. (Aug. 2012). “Direct imaging of extra-solar planets in star forming regions. Lessons learned from a false positive around IM Lupi”. In: 544, A131, A131. DOI: [10.1051/0004-6361/201219662](https://doi.org/10.1051/0004-6361/201219662). arXiv: [1207.6017](https://arxiv.org/abs/1207.6017) [astro-ph.EP] (cit. on p. 5).
- Mawet, D. et al. (Aug. 2014). “Fundamental Limitations of High Contrast Imaging Set by Small Sample Statistics”. In: 792.2, p. 97. ISSN: 1538-4357. DOI: [10.1088/0004-637x/792/2/97](https://doi.org/10.1088/0004-637x/792/2/97). URL: <http://dx.doi.org/10.1088/0004-637x/792/2/97> (cit. on pp. 91, 124).
- Mayor, Michel (2024). “Plurality of Worlds”. In: *Annual Review of Astronomy and Astrophysics* 62. Volume 62, 2024, pp. 1–20. ISSN: 1545-4282. DOI: <https://doi.org/10.1146/annurev-astro-071422-101131>. URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-astro-071422-101131> (cit. on p. 3).
- Mayor, Michel and Queloz, Didier (1995). “A Jupiter-mass companion to a solar-type star”. In: *Nature* 378.6555, pp. 355–359. DOI: [10.1038/378355a0](https://doi.org/10.1038/378355a0). URL: <https://doi.org/10.1038/378355a0> (cit. on p. 3).
- McDougall, Hugh G., Davis, Tamara M., and Pope, Benjamin J. S. (May 2025). “LITMUS: Bayesian Lag Recovery in Reverberation Mapping with Fast Differentiable Models”. In: *arXiv e-prints*, arXiv:2505.09832, arXiv:2505.09832. DOI: [10.48550/arXiv.2505.09832](https://doi.org/10.48550/arXiv.2505.09832). arXiv: [2505.09832](https://arxiv.org/abs/2505.09832) [astro-ph.GA] (cit. on p. 78).
- Metropolis, Nicholas et al. (June 1953). “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6, pp. 1087–1092. ISSN: 0021-9606. DOI: [10.1063/1.1699114](https://doi.org/10.1063/1.1699114). eprint: https://pubs.aip.org/aip/jcp/article-pdf/21/6/1087/18802390/1087\1\1_online.pdf. URL: <https://doi.org/10.1063/1.1699114> (cit. on pp. 31, 57).
- Metzler, Christopher et al. (Oct. 2018). “prDeep: Robust Phase Retrieval with a Flexible Deep Network”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by Jennifer Dy and Andreas Krause. Vol. 80. Proceedings of Machine Learning Research. Stockholmsmässan, Stockholm Sweden: PMLR, pp. 3501–3510. URL: <http://proceedings.mlr.press/v80/metzler18a.html> (cit. on p. 41).
- Michelson, A. A. (Dec. 1891). “Measurement of Jupiter’s Satellites by Interference”. In: 45.1155, pp. 160–161. DOI: [10.1038/045160a0](https://doi.org/10.1038/045160a0) (cit. on p. 9).
- Michelson, A. A. and Pease, F. G. (May 1921). “Measurement of the Diameter of α Orionis with the Interferometer.” In: 53, pp. 249–259. DOI: [10.1086/142603](https://doi.org/10.1086/142603) (cit. on pp. 9, 22).
- Milli, J. et al. (July 2018). “Low wind effect on VLT/SPHERE: impact, mitigation strategy, and results”. In: *Adaptive Optics Systems VI*. Ed. by Laird M. Close, Laura Schreiber, and Dirk Schmidt. Vol. 10703. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference

- Series, 107032A, 107032A. DOI: [10.1117/12.2311499](https://doi.org/10.1117/12.2311499). arXiv: [1806.05370](https://arxiv.org/abs/1806.05370) [[astro-ph.IM](#)] (cit. on p. 72).
- Monnier, J. D., Tuthill, P. G., and Danchi, W. C. (Oct. 1999). “Pinwheel Nebula around WR 98a”. In: *The Astrophysical Journal* 525.2, p. L97. ISSN: 0004-637X. DOI: [10.1086/312352](https://doi.org/10.1086/312352). URL: <https://dx.doi.org/10.1086/312352> (cit. on pp. 10, 75).
- Monnier, J. D. et al. (Feb. 2007). “The Keck Aperture-masking Experiment: Near-Infrared Sizes of Dusty Wolf-Rayet Stars”. In: 655.2, pp. 1033–1045. ISSN: 1538-4357. DOI: [10.1086/509873](https://doi.org/10.1086/509873). arXiv: [astro-ph/0610437](https://arxiv.org/abs/astro-ph/0610437) [[astro-ph](#)]. URL: <http://dx.doi.org/10.1086/509873> (cit. on pp. 5, 8).
- Monnier, John D (Apr. 2003). “Optical interferometry in astronomy”. In: *Reports on Progress in Physics* 66.5, pp. 789–857. ISSN: 1361-6633. DOI: [10.1088/0034-4885/66/5/203](https://doi.org/10.1088/0034-4885/66/5/203). URL: <http://dx.doi.org/10.1088/0034-4885/66/5/203> (cit. on pp. 5, 8, 9, 75).
- Mordasini, C. et al. (2012). “Characterization of exoplanets from their formation - I. Models of combined planet formation and evolution”. In: *A&A* 547, A111. DOI: [10.1051/0004-6361/201118457](https://doi.org/10.1051/0004-6361/201118457). URL: <https://doi.org/10.1051/0004-6361/201118457> (cit. on p. 6).
- Morse, Philip M. and Feshbach, Herman (1953). *Methods of Theoretical Physics*. Vol. 1. McGraw-Hill (cit. on pp. 18, 87).
- N’Diaye, Mamadou et al. (Aug. 2022). “Statistical tests with multi-wavelength Kernel-phase analysis for the detection and characterization of planetary companions”. In: *Optical and Infrared Interferometry and Imaging VIII*. Ed. by Antoine Mérand, Stephanie Sallum, and Joel Sanchez-Bermudez. Vol. 12183. SPIE Conference Series, 121831M, p. 121831M. DOI: [10.1117/12.2629961](https://doi.org/10.1117/12.2629961). arXiv: [2208.04594](https://arxiv.org/abs/2208.04594) [[astro-ph.IM](#)] (cit. on p. 124).
- Nardiello, D. et al. (Dec. 2019). “A PSF-based Approach to TESS High quality data Of Stellar clusters (PATHOS) - I. Search for exoplanets and variable stars in the field of 47 Tuc”. In: 490.3, pp. 3806–3823. DOI: [10.1093/mnras/stz2878](https://doi.org/10.1093/mnras/stz2878). arXiv: [1910.03592](https://arxiv.org/abs/1910.03592) [[astro-ph.SR](#)] (cit. on p. 53).
- Nielsen, Eric L. et al. (June 2019). “The Gemini Planet Imager Exoplanet Survey: Giant Planet and Brown Dwarf Demographics from 10 to 100 au”. In: *The Astronomical Journal* 158.1, p. 13. DOI: [10.3847/1538-3881/ab16e9](https://doi.org/10.3847/1538-3881/ab16e9). URL: <https://dx.doi.org/10.3847/1538-3881/ab16e9> (cit. on p. 6).
- Nishizaki, Yohei et al. (2020). “Analysis of non-iterative phase retrieval based on machine learning”. In: *Optical Review* 27.1, pp. 136–141. DOI: [10.1007/s10043-019-00574-8](https://doi.org/10.1007/s10043-019-00574-8). URL: <https://doi.org/10.1007/s10043-019-00574-8> (cit. on p. 41).
- Norris, Ray P. (2016). “Dawes Review 5: Australian Aboriginal Astronomy and Navigation”. In: *Publications of the Astronomical Society of Australia* 33. ISSN: 1448-6083. DOI: [10.1017/pasa.2016.25](https://doi.org/10.1017/pasa.2016.25). URL: <http://dx.doi.org/10.1017/pasa.2016.25> (cit. on p. 1).
- Page, Josue and Favaro, Paolo (May 2020). “Learning to Model and Calibrate Optics via a Differentiable Wave Optics Simulator”. In: *arXiv e-prints*, arXiv:2005.08562, arXiv:2005.08562. DOI: [10.1109/ICIP40778.2020.9190870](https://doi.org/10.1109/ICIP40778.2020.9190870). arXiv: [2005.08562](https://arxiv.org/abs/2005.08562) [[eess.IV](#)] (cit. on pp. 43, 78).
- Paszke, Adam et al. (Dec. 2019). “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *arXiv e-prints*, arXiv:1912.01703, arXiv:1912.01703. arXiv: [1912](https://arxiv.org/abs/1912.01703).

- 01703 [cs.LG]. URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf> (cit. on pp. 26, 42, 57, 78).
- Perrin, Marshall D. et al. (Sept. 2012). “Simulating point spread functions for the James Webb Space Telescope with WebbPSF”. In: *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*. Ed. by Mark C. Clampin et al. Vol. 8442. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 84423D, p. 84423D. DOI: [10.1117/12.925230](https://doi.org/10.1117/12.925230) (cit. on pp. 42, 57, 79, 84).
- Perrin, Marshall D. et al. (Aug. 2014). “Updated point spread function simulations for JWST with WebbPSF”. In: *Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave*. Ed. by Jr. Oschmann Jacobus M. et al. Vol. 9143. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 91433X, p. 91433X. DOI: [10.1117/12.2056689](https://doi.org/10.1117/12.2056689) (cit. on pp. 42, 79, 84).
- Perryman, Michael (2018). *The Exoplanet Handbook*. 2nd ed. Cambridge University Press (cit. on p. 4).
- Phan, Du, Pradhan, Neeraj, and Jankowiak, Martin (Dec. 2019a). “Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro”. In: *arXiv e-prints*, arXiv:1912.11554, arXiv:1912.11554. arXiv: [1912.11554](https://arxiv.org/abs/1912.11554) [stat.ML] (cit. on pp. 43, 72).
- (2019b). “Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro”. In: *arXiv preprint arXiv:1912.11554* (cit. on pp. 111, 126).
- Pinder, Thomas and Dodd, Daniel (2022). “GPJax: A Gaussian Process Framework in JAX”. In: *Journal of Open Source Software* 7.75, p. 4455. DOI: [10.21105/joss.04455](https://doi.org/10.21105/joss.04455). URL: <https://doi.org/10.21105/joss.04455> (cit. on p. 53).
- Pollack, James B. et al. (1996). “Formation of the Giant Planets by Concurrent Accretion of Solids and Gas”. In: *Icarus* 124.1, pp. 62–85. ISSN: 0019-1035. DOI: <https://doi.org/10.1006/icar.1996.0190>. URL: <https://www.sciencedirect.com/science/article/pii/S0019103596901906> (cit. on p. 6).
- Pope, B (2017). “Observing bright stars and their planets from the Earth and from space”. PhD thesis. University of Oxford (cit. on p. 23).
- Pope, Benjamin, Martinache, Frantz, and Tuthill, Peter (Apr. 2013). “DANCING IN THE DARK: NEW BROWN DWARF BINARIES FROM KERNEL PHASE INTERFEROMETRY”. In: *The Astrophysical Journal* 767.2, p. 110. ISSN: 1538-4357. DOI: [10.1088/0004-637x/767/2/110](https://doi.org/10.1088/0004-637x/767/2/110). URL: <http://dx.doi.org/10.1088/0004-637X/767/2/110> (cit. on pp. 9, 95, 125).
- Pope, Benjamin et al. (May 2014). “A demonstration of wavefront sensing and mirror phasing from the image domain”. In: 440.1, pp. 125–133. DOI: [10.1093/mnras/stu218](https://doi.org/10.1093/mnras/stu218). arXiv: [1401.7566](https://arxiv.org/abs/1401.7566) [astro-ph.IM] (cit. on pp. 87, 91).
- Pope, Benjamin J. S. (Sept. 2016). “Kernel phase and kernel amplitude in Fizeau imaging”. In: *Monthly Notices of the Royal Astronomical Society* 463.4, pp. 3573–3581. ISSN: 1365-2966. DOI: [10.1093/mnras/stw2215](https://doi.org/10.1093/mnras/stw2215). URL: <http://dx.doi.org/10.1093/mnras/stw2215> (cit. on pp. 9, 25, 90, 93).
- Pope, Benjamin J. S. et al. (Nov. 2019). “The K2 Bright Star Survey. I. Methodology and Data Release”. In: *The Astrophysical Journal Supplement Series* 245.1, p. 8. ISSN: 1538-4365. DOI:

- [10.3847/1538-4365/ab3d29](https://doi.org/10.3847/1538-4365/ab3d29). URL: <http://dx.doi.org/10.3847/1538-4365/ab3d29> (cit. on p. 41).
- Pope, Benjamin J. S. et al. (Jan. 2021). “Kernel Phase and Coronagraphy with Automatic Differentiation”. In: *The Astrophysical Journal* 907.1, p. 40. ISSN: 1538-4357. DOI: [10.3847/1538-4357/abcb00](https://doi.org/10.3847/1538-4357/abcb00). URL: <https://doi.org/10.3847/1538-4357/abcb00> (cit. on pp. 42, 93, 94).
- Por, E. H. et al. (2018). “High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator”. In: *Adaptive Optics Systems VI*. Vol. 10703. Proc. SPIE. DOI: [10.1117/12.2314407](https://doi.org/10.1117/12.2314407). URL: <https://doi.org/10.1117/12.2314407> (cit. on p. 57).
- Por, Emiel H. (2017). “Optimal design of apodizing phase plate coronagraphs”. In: *Techniques and Instrumentation for Detection of Exoplanets VIII*. Ed. by Stuart Shaklan. Vol. 10400. International Society for Optics and Photonics. SPIE, pp. 236–247. DOI: [10.1117/12.2274219](https://doi.org/10.1117/12.2274219). URL: <https://doi.org/10.1117/12.2274219> (cit. on p. 57).
- Quanz, S. P. et al. (Aug. 2022). “Large Interferometer For Exoplanets (LIFE). I. Improved exoplanet detection yield estimates for a large mid-infrared space-interferometer mission”. In: 664, A21, A21. DOI: [10.1051/0004-6361/202140366](https://doi.org/10.1051/0004-6361/202140366). arXiv: [2101.07500](https://arxiv.org/abs/2101.07500) [[astro-ph.EP](https://arxiv.org/abs/2101.07500)] (cit. on p. 7).
- Rader, Jason, Lyons, Terry, and Kidger, Patrick (2024). “Optimistix: modular optimisation in JAX and Equinox”. In: *arXiv:2402.09983* (cit. on p. 107).
- Radhakrishna Rao, C. (1945). “Information and the accuracy attainable in the estimation of statistical parameters”. In: *Bull. Calcutta Math. Soc.* 37, pp. 81–91. ISSN: 0008-0659 (cit. on pp. 58, 128).
- Raissi, M., Perdikaris, P., and Karniadakis, G.E. (2019). “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378, pp. 686–707. ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2018.10.045>. URL: <https://www.sciencedirect.com/science/article/pii/S0021999118307125> (cit. on pp. 36, 95).
- Rauscher, Bernard J. et al. (July 2007). “Detectors for the James Webb Space Telescope Near-Infrared Spectrograph. I. Readout Mode, Noise Model, and Calibration Considerations”. In: 119.857, pp. 768–786. ISSN: 1538-3873. DOI: [10.1086/520887](https://doi.org/10.1086/520887). arXiv: [0706.2344](https://arxiv.org/abs/0706.2344) [[astro-ph](https://arxiv.org/abs/0706.2344)]. URL: <http://dx.doi.org/10.1086/520887> (cit. on pp. 5, 14, 82).
- Rauscher, Bernard J. et al. (Sept. 2017). “Improved Reference Sampling and Subtraction: A Technique for Reducing the Read Noise of Near-infrared Detector Systems”. In: 129.980, p. 105003. ISSN: 1538-3873. DOI: [10.1088/1538-3873/aa83fd](https://doi.org/10.1088/1538-3873/aa83fd). URL: <http://dx.doi.org/10.1088/1538-3873/aa83fd> (cit. on p. 101).
- Ray, Shrishmoy et al. (Feb. 2023). “Detecting planetary mass companions near the water frost-line using JWST interferometry”. In: 519.2, pp. 2718–2735. DOI: [10.1093/mnras/stac3425](https://doi.org/10.1093/mnras/stac3425). arXiv: [2211.09830](https://arxiv.org/abs/2211.09830) [[astro-ph.EP](https://arxiv.org/abs/2211.09830)] (cit. on p. 75).
- Ray, Shrishmoy et al. (2025). *The JWST Early Release Science Program for Direct Observations of Exoplanetary Systems III: Aperture Masking Interferometric Observations of the star HIP 65426 at 3.8 um*. arXiv: [2310.11508](https://arxiv.org/abs/2310.11508) [[astro-ph.EP](https://arxiv.org/abs/2310.11508)]. URL: <https://arxiv.org/abs/2310.11508> (cit. on pp. 13, 15, 77).

- Rayleigh, Lord (Feb. 1880). “Investigations in optics, with special reference to the spectro-scope”. In: 40, p. 254. DOI: [10.1093/mnras/40.4.254](https://doi.org/10.1093/mnras/40.4.254) (cit. on p. 19).
- Readhead, A. C. S. et al. (1980). “Mapping radio sources with uncalibrated visibility data”. In: *Nature* 285.5761, pp. 137–140. DOI: [10.1038/285137a0](https://doi.org/10.1038/285137a0). URL: <https://doi.org/10.1038/285137a0> (cit. on p. 75).
- Ricker, George R. et al. (Jan. 2015). “Transiting Exoplanet Survey Satellite (TESS)”. In: *Journal of Astronomical Telescopes, Instruments, and Systems* 1.1, 014003, p. 014003. DOI: [10.1117/1.JATIS.1.1.014003](https://doi.org/10.1117/1.JATIS.1.1.014003). URL: <https://doi.org/10.1117/1.JATIS.1.1.014003> (cit. on pp. 3, 40, 41).
- Rieke, Marcia J. et al. (Feb. 2023). “Performance of NIRCcam on JWST in Flight”. In: *Publications of the Astronomical Society of the Pacific* 135.1044, p. 028001. DOI: [10.1088/1538-3873/acac53](https://dx.doi.org/10.1088/1538-3873/acac53). URL: <https://dx.doi.org/10.1088/1538-3873/acac53> (cit. on pp. 12, 132).
- Rigby, Jane et al. (Apr. 2023). “The Science Performance of JWST as Characterized in Commissioning”. In: *Publications of the Astronomical Society of the Pacific* 135.1046, p. 048001. ISSN: 1538-3873. DOI: [10.1088/1538-3873/acb293](https://dx.doi.org/10.1088/1538-3873/acb293). URL: <http://dx.doi.org/10.1088/1538-3873/acb293> (cit. on pp. 11, 34, 87).
- Robertson, J. et al. (Jan. 1991). “High-resolution Imaging by Optical Aperture Synthesis: First Results from the MAPPIT Project”. In: 9, p. 162. DOI: [10.1017/S1323358000025418](https://doi.org/10.1017/S1323358000025418) (cit. on p. 75).
- Roddier, F. (Jan. 1981). “The effects of atmospheric turbulence in optical astronomy”. In: *Progress in Optics* 19, pp. 281–376. DOI: [10.1016/S0079-6638\(08\)70204-X](https://doi.org/10.1016/S0079-6638(08)70204-X) (cit. on p. 23).
- Rowlands, Neil, Midwinter, Calvin, and Warner, Gerry (2018). “The impact of the brighter-fatter effect on the performance of the JWST fine guidance sensor”. In: *High Energy, Optical, and Infrared Detectors for Astronomy VIII*. Ed. by Andrew D. Holland and James Beletic. Vol. 10709. International Society for Optics and Photonics. SPIE, 107091K. DOI: [10.1117/12.2312338](https://doi.org/10.1117/12.2312338). URL: <https://doi.org/10.1117/12.2312338> (cit. on pp. 15, 77).
- Royal Swedish Academy of Sciences, The (2024). *The Nobel Prize in Physics 2024 – Press release*. URL: <https://www.nobelprize.org/prizes/physics/2024/press-release> (cit. on p. 26).
- Ruder, Sebastian (Sept. 2016). “An overview of gradient descent optimization algorithms”. In: *arXiv e-prints*, arXiv:1609.04747, arXiv:1609.04747. DOI: [10.48550/arXiv.1609.04747](https://doi.org/10.48550/arXiv.1609.04747). arXiv: [1609.04747 \[cs.LG\]](https://arxiv.org/abs/1609.04747) (cit. on p. 78).
- Ruffio, Jean-Baptiste et al. (Oct. 2018). “A Bayesian Framework for Exoplanet Direct Detection and Non-detection”. In: 156.5, p. 196. DOI: [10.3847/1538-3881/aade95](https://doi.org/10.3847/1538-3881/aade95). URL: <https://dx.doi.org/10.3847/1538-3881/aade95> (cit. on pp. 112, 114, 116, 120, 122).
- Ruffio, Jean-Baptiste et al. (Aug. 2024). “JWST-TST High Contrast: Achieving Direct Spectroscopy of Faint Substellar Companions Next to Bright Stars with the NIRSpec Integral Field Unit”. In: 168.2, 73, p. 73. DOI: [10.3847/1538-3881/ad5281](https://doi.org/10.3847/1538-3881/ad5281). arXiv: [2310.09902 \[astro-ph.EP\]](https://arxiv.org/abs/2310.09902) (cit. on p. 124).

- Rumelhart, David E., Hinton, Geoffrey E., and Williams, Ronald J. (1986). “Learning representations by back-propagating errors”. In: *Nature* 323.6088, pp. 533–536. DOI: [10.1038/323533a0](https://doi.org/10.1038/323533a0). URL: <https://doi.org/10.1038/323533a0> (cit. on p. 78).
- Rustankulov, Zafar et al. (Mar. 2022). “Analysis of a JWST NIRSpec Lab Time Series: Characterizing Systematics, Recovering Exoplanet Transit Spectroscopy, and Constraining a Noise Floor”. In: *The Astrophysical Journal Letters* 928.1, p. L7. DOI: [10.3847/2041-8213/ac5b6f](https://doi.org/10.3847/2041-8213/ac5b6f). URL: <https://doi.org/10.3847/2041-8213/ac5b6f> (cit. on p. 53).
- Ryan, EG et al. (2016). “Fully Bayesian optimal experimental design: A review”. In: *International Statistical Review* 84, pp. 128–154 (cit. on p. 58).
- Sallum, Steph et al. (Aug. 2023). “Systematic Multiepoch Monitoring of LkCa 15: Dynamic Dust Structures on Solar System Scales”. In: 953.1, 55, p. 55. DOI: [10.3847/1538-4357/ace16c](https://doi.org/10.3847/1538-4357/ace16c). arXiv: [2306.15069](https://arxiv.org/abs/2306.15069) [astro-ph.EP] (cit. on p. 76).
- Sallum, Steph et al. (2024a). *The JWST Early Release Science Program for Direct Observations of Exoplanetary Systems IV: NIRISS Aperture Masking Interferometry Performance and Lessons Learned*. arXiv: [2310.11499](https://arxiv.org/abs/2310.11499) [astro-ph.EP]. URL: <https://arxiv.org/abs/2310.11499> (cit. on pp. 13, 15).
- (2024b). “The JWST Early Release Science Program for Direct Observations of Exoplanetary Systems IV: NIRISS Aperture Masking Interferometry Performance and Lessons Learned”. In: arXiv: [2310.11499](https://arxiv.org/abs/2310.11499) [astro-ph.EP]. URL: <https://arxiv.org/abs/2310.11499> (cit. on pp. 77, 121).
- Sallum, Stephanie, Ray, Shrishmoy, and Hinkley, Sasha (2022). “SAMpy: a Fourier-plane pipeline for JWST/NIRISS aperture masking interferometry (and more!)” In: *Optical and Infrared Interferometry and Imaging VIII*. Ed. by Antoine Mérand, Stephanie Sallum, and Joel Sanchez-Bermudez. Vol. 12183. International Society for Optics and Photonics. SPIE, p. 121832M. DOI: [10.1117/12.2630401](https://doi.org/10.1117/12.2630401). URL: <https://doi.org/10.1117/12.2630401> (cit. on pp. 13, 76).
- Schütt, Kristof T. et al. (2017). “Quantum-chemical insights from deep tensor neural networks”. In: *Nature Communications* 8.1, p. 13890. DOI: [10.1038/ncomms13890](https://doi.org/10.1038/ncomms13890). URL: <https://doi.org/10.1038/ncomms13890> (cit. on p. 36).
- Shanno, David F. (1970). “Conditioning of Quasi-Newton Methods for Function Minimization”. In: *Mathematics of Computation* 24, pp. 647–656. DOI: <https://doi.org/10.1090/S0025-5718-1970-0274029-X> (cit. on pp. 31, 107).
- Shechtman, Y., Beck, A., and Eldar, Y. C. (2014). “GESPAR: Efficient Phase Retrieval of Sparse Signals”. In: *IEEE Transactions on Signal Processing* 62.4, pp. 928–938. DOI: [10.1109/TSP.2013.2297687](https://doi.org/10.1109/TSP.2013.2297687) (cit. on p. 41).
- Siskind, Jeffrey Mark and Pearlmutter, Barak A. (2008). “Nesting forward-mode AD in a functional framework”. In: *Higher-Order and Symbolic Computation* 21.4, pp. 361–376. DOI: [10.1007/s10990-008-9037-1](https://doi.org/10.1007/s10990-008-9037-1). URL: <https://doi.org/10.1007/s10990-008-9037-1> (cit. on p. 78).
- Sitzmann, Vincent et al. (July 2018a). “End-to-End Optimization of Optics and Image Processing for Achromatic Extended Depth of Field and Super-Resolution Imaging”. In: *ACM*

- Trans. Graph.* 37.4, p. 114. ISSN: 0730-0301. DOI: [10.1145/3197517.3201333](https://doi.org/10.1145/3197517.3201333). URL: <https://doi.org/10.1145/3197517.3201333> (cit. on pp. 36, 43).
- Sitzmann, Vincent et al. (2018b). “End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging”. In: *ACM Transactions on Graphics (TOG)* 37.4, p. 114 (cit. on p. 78).
- Sivaramakrishnan, Anand et al. (2012). “Non-redundant Aperture Masking Interferometry (AMI) and segment phasing with JWST-NIRISS”. In: *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*. Ed. by Mark C. Clampin et al. Vol. 8442. International Society for Optics and Photonics. SPIE, 84422S. DOI: [10.1117/12.925565](https://doi.org/10.1117/12.925565). URL: <https://doi.org/10.1117/12.925565> (cit. on pp. 1, 7, 10, 13, 38, 73, 75, 128).
- Sivaramakrishnan, Anand et al. (Oct. 2022). “The Near Infrared Imager and Slitless Spectrograph for the James Webb Space Telescope – IV. Aperture Masking Interferometry”. In: *arXiv e-prints*, arXiv:2210.17434, arXiv:2210.17434. arXiv: [2210.17434](https://arxiv.org/abs/2210.17434) [[astro-ph.IM](https://arxiv.org/abs/2210.17434)] (cit. on pp. 13, 15, 40, 42).
- Sivaramakrishnan, Anand et al. (Jan. 2023). “The Near Infrared Imager and Slitless Spectrograph for the James Webb Space Telescope. IV. Aperture Masking Interferometry”. In: 135.1043, p. 015003. ISSN: 1538-3873. DOI: [10.1088/1538-3873/acaebd](https://doi.org/10.1088/1538-3873/acaebd). URL: <http://dx.doi.org/10.1088/1538-3873/acaebd> (cit. on pp. 75, 77, 109, 110).
- Soulain, A. et al. (Dec. 2020). “The James Webb Space Telescope aperture masking interferometer”. In: *Optical and Infrared Interferometry and Imaging VII*. Ed. by Peter G. Tuthill, Antoine Mérand, and Stephanie Sallum. Vol. 11446. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. International Society for Optics and Photonics. SPIE, 1144611, p. 1144611. DOI: [10.1117/12.2560804](https://doi.org/10.1117/12.2560804). arXiv: [2201.01524](https://arxiv.org/abs/2201.01524) [[astro-ph.IM](https://arxiv.org/abs/2201.01524)]. URL: <https://doi.org/10.1117/12.2560804> (cit. on pp. 1, 10, 12, 13, 38, 73, 75, 76, 114, 128).
- Soummer, R. et al. (Jan. 2007). “Fast computation of Lyot-style coronagraph propagation”. In: *Optics Express* 15.24, p. 15935. DOI: [10.1364/OE.15.015935](https://doi.org/10.1364/OE.15.015935). arXiv: [0711.0368](https://arxiv.org/abs/0711.0368) [[astro-ph](https://arxiv.org/abs/0711.0368)] (cit. on pp. 44, 50).
- Soummer, Rémi, Pueyo, Laurent, and Larkin, James (Aug. 2012). “DETECTION AND CHARACTERIZATION OF EXOPLANETS AND DISKS USING PROJECTIONS ON KARHUNEN-LOÈVE EIGENIMAGES”. In: *The Astrophysical Journal Letters* 755.2, p. L28. DOI: [10.1088/2041-8205/755/2/L28](https://doi.org/10.1088/2041-8205/755/2/L28). URL: <https://dx.doi.org/10.1088/2041-8205/755/2/L28> (cit. on p. 8).
- Sozzetti, Alessandro (Oct. 2005). “Astrometric Methods and Instrumentation to Identify and Characterize Extrasolar Planets: A Review”. In: 117.836, pp. 1021–1048. DOI: [10.1086/444487](https://doi.org/10.1086/444487). arXiv: [astro-ph/0507115](https://arxiv.org/abs/astro-ph/0507115) [[astro-ph](https://arxiv.org/abs/astro-ph/0507115)] (cit. on p. 57).
- Sparks, William B. and Ford, Holland C. (Oct. 2002). “Imaging Spectroscopy for Extrasolar Planet Detection”. In: 578.1, pp. 543–564. DOI: [10.1086/342401](https://doi.org/10.1086/342401). arXiv: [astro-ph/0209078](https://arxiv.org/abs/astro-ph/0209078) [[astro-ph](https://arxiv.org/abs/astro-ph/0209078)] (cit. on p. 8).
- Spergel, D. et al. (2015). *Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report*. arXiv: [1503.03757](https://arxiv.org/abs/1503.03757) [[astro-ph.IM](https://arxiv.org/abs/1503.03757)]. URL: <https://arxiv.org/abs/1503.03757> (cit. on p. 7).

- Stark, Christopher C. et al. (2025). *Paths to Robust Exoplanet Science Yield Margin for the Habitable Worlds Observatory*. arXiv: [2405.19418](https://arxiv.org/abs/2405.19418) [astro-ph.EP]. URL: <https://arxiv.org/abs/2405.19418> (cit. on p. 132).
- Stefansson, Gudmundur et al. (Oct. 2017). “Toward Space-like Photometric Precision from the Ground with Beam-shaping Diffusers”. In: 848.1, 9, p. 9. DOI: [10.3847/1538-4357/aa88aa](https://doi.org/10.3847/1538-4357/aa88aa). arXiv: [1710.01790](https://arxiv.org/abs/1710.01790) [astro-ph.IM] (cit. on p. 44).
- Stengel, Karen et al. (2020). “Adversarial super-resolution of climatological wind and solar data”. In: *Proceedings of the National Academy of Sciences* 117.29, pp. 16805–16815. DOI: [10.1073/pnas.1918964117](https://doi.org/10.1073/pnas.1918964117). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1918964117>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1918964117> (cit. on p. 36).
- Stone, Connor J. et al. (Nov. 2023). “ASTROPHOT: fitting everything everywhere all at once in astronomical images”. In: 525.4, pp. 6377–6393. DOI: [10.1093/mnras/stad2477](https://doi.org/10.1093/mnras/stad2477). arXiv: [2308.01957](https://arxiv.org/abs/2308.01957) [astro-ph.IM] (cit. on p. 35).
- Sturm, E. et al. (2018). “Spatially resolved rotation of the broad-line region of a quasar at sub-parsec scale”. In: *Nature* 563.7733, pp. 657–660. DOI: [10.1038/s41586-018-0731-9](https://doi.org/10.1038/s41586-018-0731-9). URL: <https://doi.org/10.1038/s41586-018-0731-9> (cit. on p. 5).
- Team, Theano Development (May 2016). “Theano: A Python framework for fast computation of mathematical expressions”. In: *arXiv e-prints* abs/1605.02688. URL: <http://arxiv.org/abs/1605.02688> (cit. on p. 42).
- Traub, W. A. and Oppenheimer, B. R. (2010). “Direct Imaging of Exoplanets”. In: *Exoplanets*. Ed. by S. Seager, pp. 111–156 (cit. on p. 4).
- Tuthill, P. G., Monnier, J. D., and Danchi, W. C. (Jan. 1999). “A Menagerie of Stars: New Images from the Diffraction Limit”. In: *Working on the Fringe: Optical and IR Interferometry from Ground and Space*. Ed. by Stephen Unwin and Robert Stachnik. Vol. 194. Astronomical Society of the Pacific Conference Series, p. 188 (cit. on p. 10).
- Tuthill, P. G. et al. (Nov. 2000a). “Smoke Signals from IRC +10216. I. Milliarcsecond Proper Motions of the Dust”. In: *The Astrophysical Journal* 543.1, p. 284. DOI: [10.1086/317082](https://doi.org/10.1086/317082). URL: <https://dx.doi.org/10.1086/317082> (cit. on pp. 5, 7, 8, 10).
- Tuthill, P. G. et al. (Oct. 2002). “Imaging the Disk around the Luminous Young Star LkH α 101 with Infrared Interferometry”. In: *The Astrophysical Journal* 577.2, p. 826. DOI: [10.1086/342235](https://doi.org/10.1086/342235). URL: <https://dx.doi.org/10.1086/342235> (cit. on p. 10).
- Tuthill, P.G. et al. (Apr. 2000b). “Michelson Interferometry with the Keck I Telescope”. In: 112.770, pp. 555–565. ISSN: 1538-3873. DOI: [10.1086/316550](https://doi.org/10.1086/316550). URL: <http://dx.doi.org/10.1086/316550> (cit. on p. 75).
- Tuthill, Peter et al. (2018). “The TOLIMAN space telescope”. In: *Optical and Infrared Interferometry and Imaging VI*. Ed. by Michelle J. Creech-Eakman, Peter G. Tuthill, and Antoine Mérand. Vol. 10701. International Society for Optics and Photonics. SPIE, pp. 432–441. DOI: [10.1117/12.2313269](https://doi.org/10.1117/12.2313269). URL: <https://doi.org/10.1117/12.2313269> (cit. on pp. 1, 15, 16, 39, 40, 44, 56–58, 62, 128).
- Twiss, R. Q., Carter, A. W. L., and Little, A. G. (Aug. 1960). “Brightness distribution over some strong radio sources at 1427 Mc/s”. In: *The Observatory* 80, pp. 153–159 (cit. on p. 75).

- Unger, Eric and Hall, Laura (1994). “The use of automatic differentiation in an aircraft design problem”. In: *5th Symposium on Multidisciplinary Analysis and Optimization*. DOI: [10.2514/6.1994-4260](https://doi.org/10.2514/6.1994-4260). eprint: <https://arc.aiaa.org/doi/pdf/10.2514/6.1994-4260>. URL: <https://arc.aiaa.org/doi/abs/10.2514/6.1994-4260> (cit. on p. 26).
- Vanderburg, Andrew and Johnson, John Asher (Oct. 2014). “A Technique for Extracting Highly Precise Photometry for the Two-Wheeled Kepler Mission”. In: *Publications of the Astronomical Society of the Pacific* 126.944, pp. 948–958. ISSN: 1538-3873. DOI: [10.1086/678764](https://doi.org/10.1086/678764). URL: <http://dx.doi.org/10.1086/678764> (cit. on p. 41).
- Vaughan, Sophia R et al. (Aug. 2023). “Chasing rainbows and ocean glints: Inner working angle constraints for the Habitable Worlds Observatory”. In: *Monthly Notices of the Royal Astronomical Society* 524.4, pp. 5477–5485. ISSN: 0035-8711. DOI: [10.1093/mnras/stad2127](https://doi.org/10.1093/mnras/stad2127). eprint: <https://academic.oup.com/mnras/article-pdf/524/4/5477/51031631/stad2127.pdf>. URL: <https://doi.org/10.1093/mnras/stad2127> (cit. on p. 132).
- Vides, Christina L. et al. (Dec. 2023). “High-angular-resolution Imaging of the V892 Tau Binary System: A New Circumprimary Disk Detection and Updated Orbital Constraints”. In: 958.2, 123, p. 123. DOI: [10.3847/1538-4357/acfda6](https://doi.org/10.3847/1538-4357/acfda6). arXiv: [2310.02241](https://arxiv.org/abs/2310.02241) [astro-ph.SR] (cit. on p. 76).
- Vigan, A. et al. (2021). “The SPHERE infrared survey for exoplanets (SHINE) - III. The demographics of young giant exoplanets below 300 au with SPHERE”. In: *A&A* 651, A72. DOI: [10.1051/0004-6361/202038107](https://doi.org/10.1051/0004-6361/202038107). URL: <https://doi.org/10.1051/0004-6361/202038107> (cit. on p. 6).
- Wagner, Kevin, Apai, Dániel, and Kratter, Kaitlin M. (May 2019). “On the Mass Function, Multiplicity, and Origins of Wide-orbit Giant Planets”. In: 877.1, p. 46. DOI: [10.3847/1538-4357/ab1904](https://doi.org/10.3847/1538-4357/ab1904). URL: <https://dx.doi.org/10.3847/1538-4357/ab1904> (cit. on p. 75).
- Wang, Dun et al. (June 2016). “A Causal, Data-driven Approach to Modeling the Kepler Data”. In: *Publications of the Astronomical Society of the Pacific* 128.967, p. 094503. ISSN: 1538-3873. DOI: [10.1088/1538-3873/128/967/094503](https://doi.org/10.1088/1538-3873/128/967/094503). URL: <http://dx.doi.org/10.1088/1538-3873/128/967/094503> (cit. on p. 41).
- Weinberger, Camilo et al. (2024). “Transformer neural networks for closed-loop adaptive optics using nonmodulated pyramid wavefront sensors”. In: *A&A* 687, A202. DOI: [10.1051/0004-6361/202349118](https://doi.org/10.1051/0004-6361/202349118). URL: <https://doi.org/10.1051/0004-6361/202349118> (cit. on pp. 34, 130).
- Wengert, R. E. (Aug. 1964). “A simple automatic derivative evaluation program”. In: *Commun. ACM* 7.8, pp. 463–464. ISSN: 0001-0782. DOI: [10.1145/355586.364791](https://doi.org/10.1145/355586.364791). URL: <https://doi.org/10.1145/355586.364791> (cit. on p. 26).
- White, T. R. et al. (Aug. 2017). “Beyond the Kepler/K2 bright limit: variability in the seven brightest members of the Pleiades”. In: *Monthly Notices of the Royal Astronomical Society* 471.3, pp. 2882–2901. ISSN: 1365-2966. DOI: [10.1093/mnras/stx1050](https://doi.org/10.1093/mnras/stx1050). URL: <http://dx.doi.org/10.1093/mnras/stx1050> (cit. on p. 41).
- Willard, Jared et al. (Nov. 2022). “Integrating Scientific Knowledge with Machine Learning for Engineering and Environmental Systems”. In: *ACM Comput. Surv.* 55.4. ISSN: 0360-0300. DOI: [10.1145/3514228](https://doi.org/10.1145/3514228). URL: <https://doi.org/10.1145/3514228> (cit. on pp. 79, 95).

- Wong, Alison et al. (Aug. 2021). “Phase retrieval and design with automatic differentiation: tutorial”. In: *Journal of the Optical Society of America B* 38.9, p. 2465. ISSN: 1520-8540. DOI: [10.1364/josab.432723](https://doi.org/10.1364/josab.432723). arXiv: [2107.00952](https://arxiv.org/abs/2107.00952) [astro-ph.IM]. URL: <http://dx.doi.org/10.1364/JOSAB.432723> (cit. on pp. 41, 42, 53, 58, 63, 78).
- Wright, Gillian S. et al. (May 2023). “The Mid-infrared Instrument for JWST and Its In-flight Performance”. In: *Publications of the Astronomical Society of the Pacific* 135.1046, p. 048003. DOI: [10.1088/1538-3873/acbe66](https://doi.org/10.1088/1538-3873/acbe66). URL: <https://dx.doi.org/10.1088/1538-3873/acbe66> (cit. on p. 12).
- Young, Thomas (1804). “I. The Bakerian Lecture. Experiments and calculations relative to physical optics”. In: *Philosophical Transactions of the Royal Society of London* 94, pp. 1–16. DOI: [10.1098/rstl.1804.0001](https://doi.org/10.1098/rstl.1804.0001). eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rstl.1804.0001>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rstl.1804.0001> (cit. on p. 20).
- Ypma, Tjalling J. (1995). “Historical Development of the Newton–Raphson Method”. In: *SIAM Review* 37.4, pp. 531–551. DOI: [10.1137/1037125](https://doi.org/10.1137/1037125). eprint: <https://doi.org/10.1137/1037125>. URL: <https://doi.org/10.1137/1037125> (cit. on p. 31).
- Zellem, Robert T. et al. (Aug. 2022). “Nancy Grace Roman Space Telescope coronagraph instrument observation calibration plan”. In: *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*. Ed. by Laura E. Coyle, Shuji Matsuura, and Marshall D. Perrin. Vol. 12180. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 121801Z, 121801Z. DOI: [10.1117/12.2627567](https://doi.org/10.1117/12.2627567). arXiv: [2202.05923](https://arxiv.org/abs/2202.05923) [astro-ph.IM] (cit. on p. 53).
- Zernike, F. (Aug. 1938). “The concept of degree of coherence and its application to optical problems”. In: *Physica* 5.8, pp. 785–795. DOI: [10.1016/S0031-8914\(38\)80203-2](https://doi.org/10.1016/S0031-8914(38)80203-2) (cit. on p. 22).
- Zhu, Wei and Dong, Subo (Sept. 2021). “Exoplanet Statistics and Theoretical Implications”. In: *Annual Review of Astronomy and Astrophysics* 59.1, pp. 291–336. DOI: [10.1146/annurev-astro-112420-020055](https://doi.org/10.1146/annurev-astro-112420-020055). URL: <https://doi.org/10.1146/annurev-astro-112420-020055> (cit. on p. 41).
- Zhu, Wei et al. (Oct. 2017). “Extracting Microlensing Signals from K2 Campaign 9”. In: 129.980, p. 104501. DOI: [10.1088/1538-3873/aa7dd7](https://doi.org/10.1088/1538-3873/aa7dd7). arXiv: [1704.08692](https://arxiv.org/abs/1704.08692) [astro-ph.IM] (cit. on p. 53).