

Lights in Motion
Observing Nearby Planets with Imaging, Wavefront Sensing, Orbital Detection, and
Spectroscopy

by

William Thompson

B.Comp., Queen's University, 2016

B.Sc., Queen's University, 2018

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Physics and Astronomy

© William Thompson, 2023

University of Victoria

All rights reserved. This dissertation may not be reproduced in whole or in part, by photocopy or other means, without the permission of the author.

Lights in Motion
Observing Nearby Planets with Imaging, Wavefront Sensing, Orbital Detection, and
Spectroscopy

by

William Thompson

B.Comp., Queen's University, 2016

B.Sc., Queen's University, 2018

Supervisory Committee

Dr. C. Marois, Supervisor
(Department of Physics and Astronomy)

Dr. F. Herwig, Co-Supervisor
(Department of Physics and Astronomy)

Dr. C. Bradley, Committee Member
(Department of Mechanical Engineering)

Abstract

To place our solar system into a wider context, astronomers must study a broad sample of planets around Sun-like stars in detail. This will require a combination of indirect evidence and direct imaging, which is the focus of this dissertation. Directly imaging solar system analogues is a challenging endeavour that is to a large extent, limited by our instruments and analysis techniques. This dissertation describes how some of these challenges can be overcome from many directions. First, it presents a new analysis technique that re-evaluates how we treat the problem of analyzing direct imaging data called direct signal-to-noise optimization. This approach can provide a three to five times reduction of speckle noise close to the star when applied to angular differential imaging data. Second, it presents applications of an approach for combining images in the presence of orbital motion. This removes a sensitivity limit to direct imaging caused by orbital smearing. It results in near-ideal scaling of sensitivity with square root of the number of observations. Additionally, this technique is extended to arbitrarily combine direct and indirect sources of evidence for planets. Next, this dissertation demonstrates improved instrumentation that could increase the sensitivity of future instruments. It demonstrates the Fast Atmospheric Self-Coherent Camera Technique in a laboratory environment and presents a five hundred times reduction in quasi-static speckles. It then presents a concept for an imaging Fourier transform spectrograph that could combine a self-coherent camera with high resolution spectral information at a resolution of 5,000 to 20,000. It demonstrates such an imaging spectrograph in a laboratory environment and shows how spectro-coherent differential imaging can lead to an approximately forty times reduction in speckle noise. Lastly, it describes a speculative concept for a constellation of orbital retroreflector beacons that could one day lead to the imaging of Earth-like planets from the ground. The analysis techniques developed in this dissertation are applied to a deep, targeted survey of the HR 8799 planetary system. This results in tight limits on any additional outer planets and the detection of a fifth candidate planet at just 4 AU separation, which would be one of the closest separation planets ever directly detected. These results will change how future surveys and searches for planets are completed, and ultimately contribute to understanding the Earth's place in our local neighbourhood.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	iv
List of Tables	ix
List of Figures	x
Glossary	xvi
Acknowledgements	xx
Authorship	xxiii
Dedication	xxiv
1 Introduction	1
1.1 Planets and their Orbits	4
1.2 The Detection Parameter Space	9
1.3 Velocimetry	13
1.3.1 Doppler Radial Velocity	13
1.3.2 Astrometric Motion	15
1.4 Transit	16
1.5 Direct Imaging	17
1.5.1 Adaptive Optics	20
1.5.2 Coronagraphy	25
1.5.3 Quasi-Static Speckles	26
1.5.4 Spectroscopy	29

1.5.5	Differential Imaging	30
1.5.6	Self-Coherent Camera	33
1.6	Combining Exoplanet Detection Methods	37
1.7	Goals of this Dissertation	38
2	Improved Contrast in Images of Exoplanets using Direct SNR Optimization	41
2.1	Introduction	41
2.2	SNR Optimization	44
2.2.1	Formulation as a Constrained Least-Squares Problem	47
2.2.2	Reduction to Linear Least Squares in Limit of no Self-Subtraction	48
2.3	Multi-Target Image Optimization	49
2.3.1	Batching	51
2.3.2	Controlling Overfitting	52
2.4	Implementation	53
2.5	Demonstration	55
2.6	Discussion	63
2.7	Conclusion	64
3	Octofitter: Fast, Flexible, and Accurate Orbit Modelling to Detect Exoplanets	65
3.1	Introduction	65
3.2	Data Models	68
3.2.1	Relative Astrometry	69
3.2.2	Images	70
3.2.3	Interferometric Observables	72
3.2.4	Radial Velocity	73
3.2.5	Proper Motion Anomaly	74
3.3	Methods	76
3.3.1	Assessing Detections	76
3.3.2	Orbital Bases and Priors	78
3.3.3	Modelling Language	79
3.3.4	Numerical Methods	81
3.3.5	Kepler Solver	82
3.3.6	Analysis and Visualization	83
3.3.7	Simulation-Based Calibration	85
3.3.8	Other Packages	86

3.4	Results	88
3.4.1	Demonstration with Relative Astrometry	88
3.4.2	Demonstration with Relative Astrometry, RV, and PMA	92
3.4.3	Demonstration with Images	93
3.4.4	Demonstration with Aperture Masking Interferometry and Radial Velocity	95
3.5	Conclusion	96
4	Deep Orbital Search for Additional Planets in the HR 8799 System	99
4.1	Introduction	99
4.2	Observations and Processing	102
4.2.1	Observations	102
4.2.2	Archival data selection	105
4.2.3	ADI Reduction	105
4.2.4	Photometric calibration	108
4.3	Modelling	108
4.3.1	Detection and Limits	113
4.3.2	Sampling	115
4.3.3	Stability	116
4.4	Results	117
4.4.1	Recovery of Known Planets	117
4.4.2	Limits on additional outer planets	120
4.4.3	Evidence for a fifth inner planet	120
4.4.4	Contribution of the 2021 Epoch	124
4.4.5	Noise Distribution and Sample Size	127
4.4.6	Mass and Proper Motion Anomaly	128
4.4.7	Stability	131
4.4.8	Sensitivity to planets besides the 4-5 AU candidate	133
4.5	Conclusion	133
5	Performance of the Fast Atmospheric Self Coherent Camera and an Im- proved Measurement Algorithm	139
5.1	Introduction	139
5.2	Experimental Setup	141
5.3	VENOMS: A New Software Toolkit for Adaptive Optics Laboratories	142

5.4	SCC: Measurement Algorithm and Results	146
5.4.1	Calibration	146
5.4.2	Measurement	149
5.4.3	Laboratory Results	151
5.5	Post Processing with Coherent Differential Imaging	153
5.5.1	Fourier Transform and Model Based Algorithm	155
5.5.2	Self-Consistent Image Plane Algorithm	157
5.6	Conclusion	158
6	A Direct Imaging Fourier Transform Spectrograph	159
6.1	Motivation	159
6.2	Proposed Solution	162
6.3	Simulations	164
6.4	Implementation	168
6.5	Results	172
6.6	Outlook	177
7	Extremely Bright Orbital Guide Beacons for Extremely Large Telescopes	178
7.1	Introduction	178
7.2	Guide Beacons	179
7.3	Orbital Parameters	180
7.4	Brightness	183
7.5	Outlook	184
8	Conclusions	186
	Bibliography	190
A	Derivation of Cartesian and Celestial Position, Velocity, and Acceleration	216
A.1	Cartesian Coordinates	217
A.2	Celestial Coordinates	220
B	Thiele-Innes Elements	222
C	Derivation of the Parameter τ from Position Angle θ	224
D	Loading, Manipulating, and Visualizing Astronomical Images	227

D.1 Introduction	227
D.2 Background	228
D.3 Existing Tools	229
D.4 The AstroImages.jl package	229

List of Tables

Table 3.1	Relative astrometry input sample.	70
Table 3.2	Images input sample	72
Table 3.3	Visibilities input sample	73
Table 3.4	Radial velocity input sample	75
Table 4.1	Observations	102
Table 4.2	SNR in combined images	103
Table 4.3	Model priors	112
Table 4.4	Stable orbits	132

List of Figures

Figure 1.1	Schematic showing how a planet's orbit is specified by Campbell elements.	5
Figure 1.2	Image of planets orbiting HR 8799 with arrows indicating their orbital velocity.	6
Figure 1.3	Planet cooling tracks	8
Figure 1.4	Mass vs. semi-major axis distribution of all confirmed exoplanets.	10
Figure 1.5	The mass vs. semi-major axis distribution of confirmed planets broken down by our distance from the system.	11
Figure 1.6	The mass vs. semi-major axis distribution of confirmed planets broken down by spectral type.	12
Figure 1.7	Illustration of using velocimetry to discover planets	14
Figure 1.8	Colour-magnitude diagram of directly imaged planets and brown dwarfs.	18
Figure 1.9	Planet brightness vs. mass and wavelength.	20
Figure 1.10	Simulated Komolgrov phase screen	22
Figure 1.11	Conceptual schematic of an adaptive optics system followed by a coronagraph	23
Figure 1.12	A conceptual block diagram of a traditional adaptive optics system.	23
Figure 1.13	Simulated telescope PSFs with and without the effects of the atmosphere and adaptive optics system.	24
Figure 1.14	Coronagraph simulations	27
Figure 1.15	Simulations of ADI and SDI	31
Figure 1.16	Simulations of two types of modified FAST SCC coronagraph.	34
Figure 1.17	Conceptual schematic of an adaptive optics system with focal plane wavefront sensor	35
Figure 2.1	Examples of the region geometry used this chapter.	45

Figure 2.2	Image preparation for global optimization	50
Figure 2.3	LOCI compared with SNR optimization for three frames in an image sequence.	56
Figure 2.4	Comparison between coefficients chosen by LOCI and SNR optimization.	57
Figure 2.5	The benefits of multi-target SNR optimization on one batch of ten images	58
Figure 2.6	Comparison between LOCI, SNR optimization, and global SNR optimization on an L'-band ADI sequence taken with NIRC2 .	59
Figure 2.7	Contrast improvement between SNR optimization and LOCI. .	60
Figure 2.8	Histograms of residual noise.	61
Figure 2.9	Comparison of pixel distributions between algorithms	61
Figure 3.1	Conceptual schematic of three different ways to use Octofitter. Other possibilities, like combining images with Doppler or astrometric velocimetry or using a mix of relative astrometry and images without detections to constrain orbits, are also possible..	68
Figure 3.2	Posterior density of photometry versus semi-major axis for a simulated detection and non-detection. The solid lines mark the mean of the marginal photometry posterior, and the dashed lines mark $\pm 1\sigma$. We consider a planet detected when the SNR based on the photometry marginalized over all other parameters is greater than some chosen threshold, e.g. 5σ	78
Figure 3.3	Sample results of running the simulation-based calibration procedure on a model consisting of a single planet parameterized with Thiele-Innes elements (A, B, F, G , etc.) and position angle θ . Each count represents a model fit to a different simulated system. The horizontal band gives a 99% range around the expected value for a perfect sampling procedure. The observation epochs and uncertainties are taken from Table 3.1.	84
Figure 3.4	Orbit fitting posteriors visualized in the plane of the sky, compared between three packages	86

- Figure 3.5 Corner plot of key orbital parameters for the case shown in Figure 3.4 compared between three packages: `orvara`, `orbitize!`, and `Octofitter`. Semi-major axis is plotted on a log scale to reveal how the sampler behaviour differs in the long tail. Though we expect all packages to eventually converge to near-identical results, we find that there are small differences in the ω , a , and e marginal posteriors that persist even after running for multiple days. 87
- Figure 3.6 Comparison between packages of the average time until chains converged to a stationary distribution and the rate at which independent posterior samples are generated. The effective sample size (ESS) rate was measured separately using the bulk and tail methods of `MCMCDiagnostics.jl`. Note that the \hat{R} and ESS of the slowest variable for a given sampler are used as this is what ultimately limits sampling performance. These results are based on the astrometry presented in Table 3.1, and are expected to depend strongly on hardware, input data, and choice of priors. 90
- Figure 3.7 Sample plot output from `Octofitter` using data from the HD 91312 system. The top row visualizes the orbital posterior compared with velocity measurements of the host. The horizontal bars in the proper motion panels show the timespans over which the average velocity was measured. The middle row shows the posterior compared with astrometry measurements in the plane of the sky and in separation and position angle over time. The bottom row shows the orbital posterior in physical coordinates to compliment the astrometry plot. The rightmost panel shows a deprojected view of the system where orbits have been rotated face on and to place periastron at the bottom. The conventions used by `Octofitter` are described in Appendix A. 91

Figure 3.8	GPI sequences used for simulations in this section. Each image was normalized to have the same average contrast at 200 mas separation from the star (just outside the edge of the mask). The images displayed above contain a simulated planet orbiting CCW at an average of just SNR 1 per epoch, spaced one month apart. Given this data, the model recovers the simulated planet at SNR 7.	93
Figure 3.9	Octofitter’s ability to recover planets from simulations of circular, face-on orbits as a function of signal to ratio per epoch and number of epochs. Left: Recovered SNR. Cells above the yellow line would be detected with a 5σ threshold. The cell outlined in cyan corresponds to the data shown in Figure 3.2. Right: The same SNRs relative to an ideal \sqrt{N} improvement with the number of epochs. We find that the SNR grows as expected unless the final, combined SNR is below ≈ 5 . The recovered SNR falls off quickly below this value and levels off at approximately 1. This explains why the recovered SNR of 1 exceeds the injected value of 0.5 in the bottom left corner.	94
Figure 3.10	Joint modelling of AMI and RV data. A-C: Recovered χ^2 maps at each independent AMI epoch created using <i>Fouriever</i> (presented in Kammerer et al., 2023). D and E: joint radial velocity and AMI posteriors visualized as an RV time series (D) and orbit posterior in the plane of the sky (E).	98
Figure 4.1	Combined images of HR 8799 for each epoch.	104
Figure 4.2	5σ contrast limits at each of the five L' epochs on a log-log scale	106
Figure 4.3	Schematic showing how we model planets across epochs with orbital motion	110
Figure 4.4	Visualization of orbit posteriors for planets b, c, d, and e. . . .	118
Figure 4.5	Marginal posteriors of photometry and orbital elements compared between the four known planets and the inner planet model.	119
Figure 4.6	Marginal photometry vs. semi-major axis for planets b, c, d, and e.	121

Figure 4.7	Flux-ratio and semi-major axis marginal posteriors for an additional outer planet between planet b and the start of the outer debris disk.	122
Figure 4.8	Flux-ratio and semi-major axis posteriors for an additional inner planet.	122
Figure 4.9	Posterior propability of the candidate f at each of the five epochs.	123
Figure 4.10	Residuals after subtracting the candidate f model.	125
Figure 4.11	Comparison of the inner planet models against three different datasets	126
Figure 4.12	Extrapolation of emperical confidence level.	129
Figure 4.13	Impact of a fifth planet on the star’s astrometric motion.	130
Figure 4.14	Stability map for a fifth inner planet.	131
Figure 4.15	85th percentile upper limit of the marginal L band flux-ratio posterior for an additional fifth inner planet.	132
Figure 4.16	Position posterior density for the inner planet model at different epochs	136
Figure 4.17	Corner plot showing the posterior of the inner planet model applied to all epochs	137
Figure 4.18	Full combined image of HR 8799 at each epoch.	138
Figure 5.1	Screenshot of the VENOMS application showing closed loop operation.	144
Figure 5.2	Conceptual block diagram describing the architecture of the VENOMS real time lab control software.	145
Figure 5.3	Images illustrating the calibration of FAST at the NEW-EARTH lab’s VIPER bench.	148
Figure 5.4	Prepared slices from the interaction matrix for three different sinusoidal modes.	149
Figure 5.5	The SCC loop closing on the bench’s static aberations and local bench turbulence.	152
Figure 5.6	Half dark hole on the VIPER bench.	153
Figure 5.7	Converging to a dark hole despite saturated pixels.	154
Figure 5.8	Effect of detector pixel MTF on CDI.	155
Figure 5.9	Fourier transform based CDI subtraction.	156
Figure 5.10	Contrast improvement from CDI	156

Figure 5.11	Image-based CDI algorithm.	158
Figure 6.1	Schematic of an Imaging Fourier Transform Spectrograph (IFTS) using a flexural bearing and corner-cube design (Grandmont, 2006).	162
Figure 6.2	Template spectrum used for FTS simulations.	164
Figure 6.3	Simulated FTS interferogram.	166
Figure 6.4	Simulated FTS spectrum.	167
Figure 6.5	Cross-correlation of a planet model spectrum with a simulated FTS spectrum.	168
Figure 6.6	Simulation of a low resolution FTS scan	169
Figure 6.7	Image of the FTS beam-splitter and corner cubes	170
Figure 6.8	Image and schematic of the SPIDERS instrument	171
Figure 6.9	Recorded FTS interferogram	172
Figure 6.10	Images recorded from IFTS across I, Z, Y, J, and H bands.	173
Figure 6.11	Application of SDI to recorded IFTS images.	174
Figure 6.12	Contrast improvement from SDI using an IFTS data cube.	175
Figure 6.13	Comparison of CDI, SDI, and both CDI and SDI.	176
Figure 7.1	Guide beacon orbit observing window length.	181
Figure 7.2	Guide beacon path in sky.	182
Figure 7.3	Guide beacon constellation path in sky.	183
Figure 7.4	Guide beacon brightness.	184
Figure C.1	Comparison of the θ and τ orbital parameterizations.	226
Figure D.1	Schematic of a FITS file	229
Figure D.2	Example rederning of the Carina nebula.	230
Figure D.3	Example rendering of a velocity cube.	232

Glossary

aberration An error in an optical system.

ADI Angular Differential Imaging.

AO Adaptive Optics.

C-RED2 An infrared detector manufactured by First Light Imaging.

CAL2 Project to update the calibration unit in the Gemini Planet Imager from a shearing interferometer to a self-coherent camera.

CDI Coherent Differential Imaging.

contrast Brightness ratio between a bright and faint object. Equivalently, the faintest object that could be discerned at some separation from the bright object. Typically presented as 1 or 5 times the relative flux standard deviation measured in concentric annuli centred on the bright object.

coronagraph An optical spatial filter used to suppress a bright object, typically a star.

DAR Differential Atmospheric Refraction.

DM Deformable Mirror.

dynamical mass Mass determined directly rather than estimated using photometry or spectroscopy.

EDR3 Gaia's 3rd early data release catalogue.

EPDF Empirical Probability Density Function.

ESS Effective Sample Size.

FAST Fast Atmospheric Self-Coherent Camera Technique.

FFT Fast Fourier Transform.

FITS Flexible Image Transport System.

FoV Field of View.

FPF False Positive Fraction.

FPM Focal Plane Mask.

FTS Fourier Transform Spectrograph.

Gaia European Space Agency mission to map the three-dimensional position and velocity of trillions of point sources in the night sky.

GDPS Gemini Deep Planet Survey.

GPI Gemini Planet Imager.

HGCA Hipparcos-Gaia Catalogue of Accelerations.

Hipparcos Satellite mission to map the three-dimensional position and velocity of stars in the night sky. Predecessor to Gaia.

IDPS International Deep Planet Survey.

IFS Integral Field Spectrograph.

IFTS Imaging Fourier Transform Spectrograph.

IFU Integral Field Unit.

JWST James Webb Space Telescope.

KLIP Karhunen-Loeve Image Processing.

LBT Large Binocular Telescope.

LLOWFS Lyot-stop Low-Order Wavefront Sensor.

LOCI Locally Optimized Combination of Images.

MCMC Markov-Chain Monte-Carlo.

MEGNO Mean Exponential Growth factor of Nearby Orbits.

MJD Modified Julian Date.

MTF Modulation Transfer Function.

NCPA Non-Common Path Aberrations.

NEW-EARTH NRC Extreme Wavefront control for Exoplanet Adaptive optics Research Topics at Herzberg.

NIR Near Infrared.

NIRC2 Near Infrared Camera 2, an AO corrected infrared imager on the Nasmyth platform of the Keck II observatory.

NIRISS Near Infrared Imager and Slitless Spectrograph.

PCA Principle Component Analysis.

PDF Probability Density Function.

PDI Polarization Differential Imaging.

PMA Proper Motion Anomaly.

PSF Point Spread Function.

R Spectral resolution of an instrument, where $R = \frac{\lambda}{\Delta\lambda}$.

RDI Reference star Differential Imaging.

RMS Root-Mean-Square.

RV Radial Velocity.

SCC Self Coherent Camera.

SDI Spectral Differential Imaging.

SNAP Signal to Noise Analysis Pipeline.

SNR Signal to Noise Ratio (S/N).

SPIDERS Subaru Pathfinder Instrument for Detection of Exoplanets and Removal of Speckles.

SVD Singular Value Decomposition.

T-LOCI Template Locally Optimized Combination of Images.

VENOMS Versatile and Efficient New-earth-lab Operating and Monitoring Software.

VIPER Victoria Pathfinder for Exoplanet Research.

VLT The European Southern Observatory's Very Large Telescope.

VLT-SPHERE A dedicated direct imaging instrument at the VLT that supports imaging, spectroscopy, and polarimetry.

Acknowledgements

I offer my sincere thanks to

Dr. Christian Marois and **Dr. Falk Herwig** for their boundless support, guidance, and mentorship;

Dr. Quinn Konopacky for including me in many observing campaigns;

Dr. Olivier Lardière for his assistance and patience in and outside the lab;

Dr. Jean-Baptiste Ruffio for updating my priors on Bayesian statistics;

Dr. Garima Singh for her advice and mentorship;

the future **Dr. Adam Johnson** for his persistence and attention to detail on the SPIDERS project,

Dr. Jean-Pierre Véran, Glen Herriot, and Dr. Kathryn Jackson for lending their expertise in adaptive optics and control theory;

Dr. Jason Wang for his help fitting orbits (into my brain);

the **National Research Council** as a whole, for the many opportunities and resources I was provided in completing these projects;

and **my friends and collaborators** in Victoria and around the world.

I further thank my father **Simon Thompson** for his support and enthusiasm, and above all, my wife **Natalie Thompson** for being my guiding star. I would not be here today without her love and support.

In addition, I thank the National Science and Engineering Research Council, New Technologies for Canadian Observatories, the University of Victoria, and Alexander and Helen Stafford MacCarthy Muir. Their generous funding made this work possible.

Many observations presented in this dissertation were conducted using facilities on Maunakea. I wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

I further acknowledge and respect the Lekwungen peoples on whose traditional territory the university stands and the Songhees, Esquimalt and WSÁNEĆ peoples whose historical relationships with the land continue to this day.

Some of the data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California and the National Aeronautics and Space Administration.

The Observatory was made possible by the generous financial support of the W. M. Keck Foundation.

This research has made use of the Keck Observatory Archive (KOA), which is operated by the W. M. Keck Observatory and the NASA Exoplanet Science Institute (NExSci), under contract with the National Aeronautics and Space Administration.

Based in part on observations obtained at the international Gemini Observatory, a program of NSF's NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation on behalf of the Gemini Observatory partnership: the National Science Foundation (United States), National Research Council (Canada), Agencia Nacional de Investigación y Desarrollo (Chile), Ministerio de Ciencia, Tecnología e Innovación (Argentina), Ministério da Ciência, Tecnologia, Inovações e Comunicações (Brazil), and Korea Astronomy and Space Science Institute (Republic of Korea).

The Gemini Planet Imager project has been supported by Gemini Observatory, which is operated by AURA, Inc., under a cooperative agreement with the NSF on behalf of the Gemini partnership: the NSF (USA), the National Research Council (Canada), CONICYT (Chile), the Australian Research Council (Australia), MCTI (Brazil) and MINCYT (Argentina).

This research has made use of the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

This work has benefited from The UltracoolSheet maintained by Will Best, Trent Dupuy, Michael Liu, Rob Siverd, and Zhoujian Zhang, and developed from compilations by Dupuy, Liu (2012), Dupuy, Kraus (2013), Liu et al. (2016), Best et al. (2018), and Best et al. (2021).

This research used the Canadian Advanced Network For Astronomy Research (CANFAR) operated in partnership by the Canadian Astronomy Data Centre and The Digital Research Alliance of Canada with support from the National Research Council of Canada the Canadian Space Agency, CANARIE and the Canadian Foundation for Innovation.

This research was enabled in part by support provided by WestGrid, Compute Ontario, and Compute Canada.

Authorship

This dissertation includes content from five published or submitted papers and conference proceedings. Authorship of these works is described in the relevant chapters. In addition, the following people contributed this dissertation:

- **Christian Marois** and **Falk Herwig** provided input and comments on the entire text.
- **Ruobing Dong** provided comments on Chapter 1.
- The National Research Council NEW-EARTH Lab collaboration designed and constructed the VIPER bench and SPIDERS pathfinder instrument. Results from these experiments are presented in Chapters 5 and 6.
- The Imaging Fourier Transform Spectrograph described in Chapter 6 was enabled by an improved control circuit and alignment mechanism contributed by **Adam Johnson**.
- Appendix A and B contain mathematical derivations that accompany a submitted paper reproduced here as Chapter 4. These appendices were contributed to that paper by **Jensen Lawrence** and are reproduced here for completeness.

Dedication

In memory of my grandfathers Brian Thompson and J. David Raal.

Chapter 1

Introduction

Placing our planet and solar system into a broader context is one of the major goals of astronomy in the twenty-first century. To this end, we will need to study in detail a sample of planets that are analogous to those in our own system through light they absorb, emit, or reflect. This light contains the fingerprints of planets' dynamics, atmospheres, and even interior compositions, and could ultimately give us the information needed for a detailed understanding of planet demographics and formation mechanisms.

The central challenge that prevents us from accessing light from planets like those in our solar system is that from our distant perspective, planets appear to almost overlap with their vastly brighter stars. Given this high contrast and small separation, the shimmering mirage of the Earth's atmosphere and even tiny optical defects inside telescopes throw glare, or "speckles," across the field of view. This blinds us from planets, just as how when driving the glare from a setting sun can blind us from seeing a traffic light.

Consequently, astronomers must use a number of tricks to study planets outside our solar system. We can either attempt to find and study planets indirectly based on their effects on their host stars or attempt to suppress this glare and reveal the planet light directly. Indirect methods have been incredibly fruitful over the past two decades, and thousands of planets have now been confirmed. Nonetheless, very few resemble the planets in our own solar system and fewer still have been studied with spectroscopy. Finding and studying solar system analogues with direct methods is the current frontier in exoplanet science. At this time, we are limited entirely by the instruments and methods we have at our disposal.

The purpose of this dissertation is to describe how some of these challenges can be

overcome, especially as they relate to direct imaging. These will include a new analysis technique that re-evaluates how we treat the problem of analyzing direct imaging data, applications of an approach for combining images in the presence of orbital motion to increase our sensitivity, approaches for combining direct and indirect sources of evidence for planets, and applications of electro-optical techniques that will dramatically increase the sensitivity of future instruments. Finally, this dissertation presents the concept and results of a new spectrograph instrument that will, in combination with these techniques, provide information on the atmospheric composition of the planets with unprecedented detail. Demonstrations on laboratory data and real observations are provided throughout, culminating in the possible discovery of one of the closest separation planets ever directly detected.

These results will shape how future surveys and searches for planets are completed, and ultimately contribute to understanding the Earth's place in our local neighbourhood.

Agenda

The contents of this dissertation are organized as follows.

Chapter 1 provides background information on exoplanets, the ways in which they are currently studied, and context for previously and currently attempted techniques that aim to improve our sensitivity.

Chapter 2 presents a mathematical formalism the signal-to-noise ratio in combined, processed direct images, develops an analysis technique that optimizes this directly, and compares it with the current state of the art.

Chapter 3 develops a technique that compensates for the orbital motion of planets and merges direct imaging with other techniques for detecting planets.

Chapter 4 presents a large observing campaign of the multi-planet HR8799 system, applies the new data analysis and orbital detection techniques, and presents a candidate fifth inner planet.

Chapter 5 demonstrates the self-coherent camera in a laboratory environment, develops simplified and improved techniques for measuring the electric field from self-coherent camera images, and presents results from laboratory testing.

Chapter 6 describes a concept, implementation, and results of a novel imaging spectrograph designed to fit within the constraints of self-coherent camera system. It further presents a concept and laboratory results of multi-spectral coherent differential imaging.

Chapter 7 presents an exploratory concept for orbital retroreflectors, which could present a path for directly imaging Earth-like planets with current 8 m class observatories.

Chapter 8 concludes with a discussion placing these contributions into the wider state of the field. It also suggests paths for future work on these concepts.

1.1 Planets and their Orbits

There is significant contention over what should, or should not be called a planet. The one part of this definition that is widely agreed upon is that a planet orbits a star.¹ Accordingly, a planet's orbital parameters are some of its most defining attributes. For one, the scale of a planet's orbit controls how much light it receives from the star and thus its surface equilibrium temperature. For our purposes, the orbit also ties together all the ways we have of detecting planets. Each exoplanet detection method is sensitive to a different overlapping set of orbital parameters.

A two body orbit obeys Kepler's laws as set out in [Kepler \(1609\)](#) and [Kepler \(1618\)](#). Keplerian orbits are typically represented by the seven Campbell elements:

- a , the semi-major axis;
- e , the eccentricity;
- M , the total mass (star and planet);
- i , the inclination;
- ω , the argument of periastron;
- Ω , the longitude of the ascending node;
- t_{peri} , the epoch of periastron passage.

One can think of these as corresponding to six degrees of freedom for describing a planet's three-dimensional position and velocity and one degree of freedom to describe the strength of the star's gravitational pull. Other orbital parameterizations will be discussed in [Chapter 3](#).

The parameters a , e , and M are perhaps the most fundamental. a and M combine to describe the orbital period of the planet ($P^2 \propto a^3$; Kepler's third law). a and e combine to describe the range of separations the planet orbits from the star, with a setting the overall scale of the orbit and e controlling how circular or elliptical it is. The angular parameters i , ω , and Ω describe the orientation of the orbit relative to our perspective. Finally, t_{peri} specifies the location of the planet along its orbit by specifying a time when it made its closest approach to the star (periastron). The location and speed of

¹or once orbited a star, in the case of rogue planets that were ejected from a solar system. Even this broad definition leaves some objects out.




Figure 1.1 Schematic of a planet’s orbit (green) around a star showing how it is specified by Campbell orbital elements. The blue circle shows the inclined plane of the orbit compared to the non-inclined grey circle. This plane is specified by inclination and longitude of the ascending node. The green ellipse shows how an eccentric orbit deviates from a circular inclined orbit.

the planets motion along this orbit is then given by Kepler’s second law and must be computed numerically in the general case (see Chapter 3).

The star’s luminosity, related to its mass M ,² determines the equilibrium temperature of a planet at a given orbital separation which in turn drives chemical processes on the surface. Planets further from their stars receive less light, and are therefore cooler once mature, which can be plainly seen in our own solar system (compare the ice giants Neptune and Uranus with scorching hot Mercury).

In fact, this falling equilibrium temperature with increasing separation affects planets even while they form. In disks of gas and dust that surround young stars, the

²Main-sequence stars have predictable luminosity as a function of mass or spectral type (the “mass-luminosity relation”).

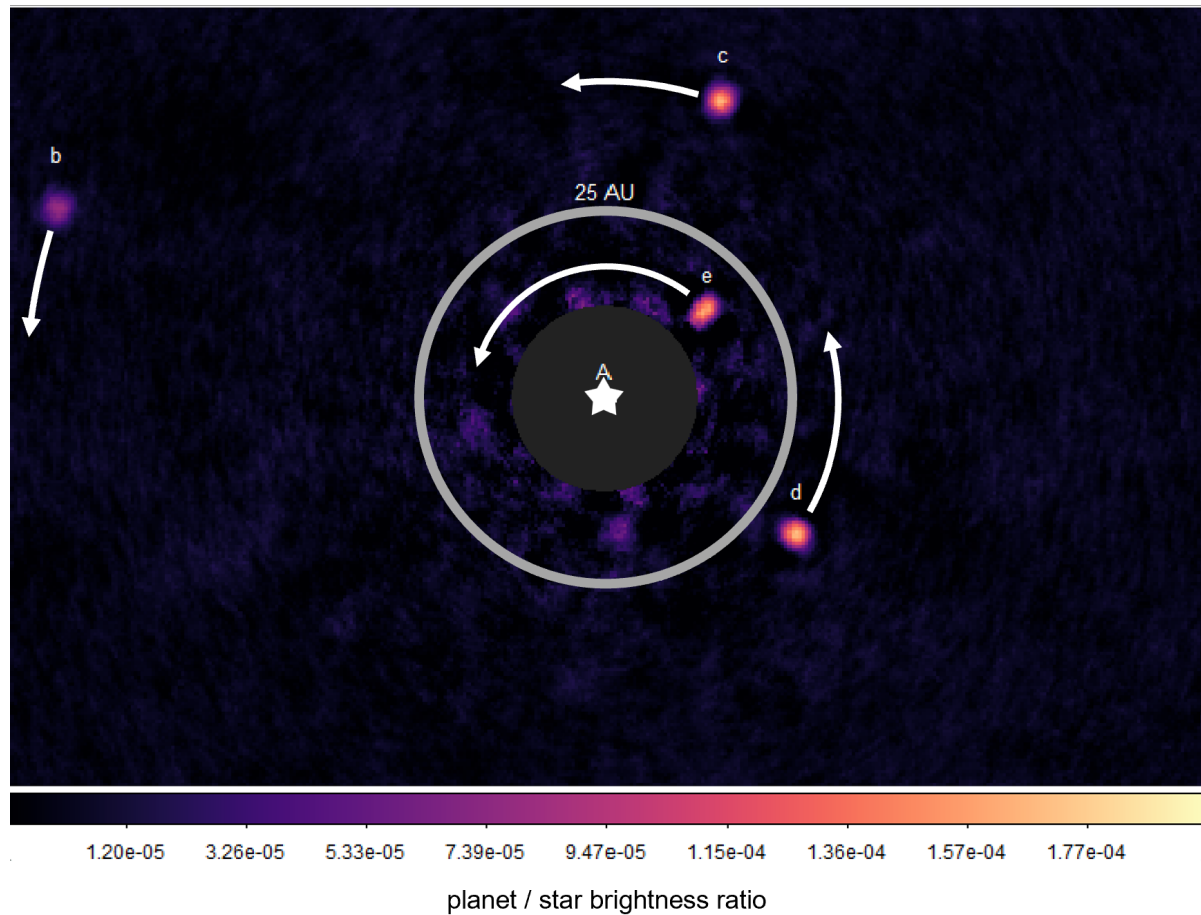


Figure 1.2 Image of four planets orbiting the star HR 8799 (discovered in [Marois et al., 2008a, 2010b](#)) taken from [Thompson et al. \(2023\)](#), with a 25 AU circle shown for scale. The length of the arrows indicate how far each planet would orbit in fifteen years. Note that the inclination and eccentricity of the real orbits are neglected for this simplified picture.

dropping temperature causes volatile molecules like water to freeze out on to dust particles at so-called “snow lines”. This creates regions with differing compositions. When planets eventually form from this stratified material, they carry abundance signatures like the C/O ratio from the locations of their births.

This formation process is believed to proceed through two separate pathways, both described in [Armitage \(2020\)](#). In massive disks, the material’s own self-gravity may lead to instability, fragmentation, and then direct collapse. This is referred to as the gravitational instability method. Because it is thought to occur quickly, material would not have time to radiate away much heat. For this reason, planets formed via gravitational instability would start with greater initial entropy and so are called “hot

start” planets.

Less massive disks would not trigger this process, and are instead believed to instead form planets more gradually. Dynamics within disks, especially near snow lines, create regions with higher densities of icy dust. Dust continues to orbit the star at Keplerian velocities, while gas, supported by its own pressure, orbits slightly more slowly. This creates a headwind that causes icy dust to spiral inwards until it reaches a region of higher pressure where it then coagulates into larger bodies and eventually planetesimal cores. Once large enough, these rocky and icy cores pull in remaining gas and dust to form protoplanets and complete the formation process. This is called the core accretion method. “Cold start” planets formed in this way would have sufficient time to radiate away more heat, and therefore begin start with lower initial entropy.

Figure 1.3 shows theoretical models of planet temperature over time for a range of planet masses (a “cooling track”). The difference between hot and cold start models is most apparent when planets are less than 200 Myr old.

Once planets form, their gravity may perturb each other’s orbits. This can lead to orbital migration, which may be a chaotic process. In our own solar system, Jupiter and Saturn are believed to have migrated significantly before arriving at their present locations (according to both the Grand Tack hypothesis ([Raymond, Morbidelli, 2014](#)) or the Nice model ([Desch, 2007](#))). Many multi-body systems have planets that orbit with mean-motion resonances which are strong indicators of previous migration (e.g. [Goździewski, Migaszewski, 2014](#)).

Emerging evidence suggests that dynamical properties like eccentricity can be connected with their formation scenario ([Bowler et al., 2020](#); [Nagpal et al., 2023](#); [Bowler et al., 2023](#), [Do Ó et al., in prep.](#)) Accordingly, though a planet’s orbital separation has a strong effect on its current conditions, it cannot be assumed to have formed in its present location.




Figure 1.3 Cooling tracks for a range of planet masses using the models of Spiegel, Burrows (2012). Young hot start planets are the easiest to detect, and so make up the majority of all direct detections. Planet data from the UltracoolSheet and NASA Exoplanet Archive.

1.2 The Detection Parameter Space

Figure 1.4 presents all confirmed planets listed by the NASA Exoplanet Archive as of May 2023 with their masses, semi-major axes, and method of detection. I describe the most relevant detection methods shortly (velocimetry, imaging, and transit), but start by making a few general observations. First, the majority of detected planets lie above a diagonal line in the mass–semi-major axis parameter space that excludes most solar system planets. Below this line, we simply do not yet have the sensitivity to detect such planets. The exceptions are mostly single microlensing events which cannot in general be followed up with additional observations and are not given further consideration in this dissertation. Second, as we saw in Figure 1.3 the majority of directly studied planets are consistent with the hot-start model. This is likely a selection effect, as young hot start planets would be much brighter in the near-infrared while they are young.

Figure 1.5 and 1.6 show these same detections disaggregated by their stars' spectral type and distance from Earth. These figures show that the vast majority of planets have been detected around distant, low-mass stars. This distance makes most of these targets both challenging for spectroscopic follow-up and not representative of our own solar system.

These point to significant gaps in our knowledge. We do not have the sensitivity to answer if planets like Jupiter and Saturn are common, let alone the ice giants and terrestrial planets of our solar system. Estimates for these occurrence rates have been extrapolated from the detected population and have not been measured directly (Zink, Hansen, 2019). Indirect surveys may be able to answer if planets with Jupiter and Saturn's mass and semi-major axes are common, but not if they formed in similar ways or are otherwise typical.

To tie a given object back to its formation roots, we require a much fuller picture. This includes a full characterization of the planets orbit, dynamical mass, radius, bolometric luminosity, and composition.³ As we will see, this level of characterization will require direct spectroscopic measurements of planets in combination with indirect evidence.

³It will in general only be possible to characterize planets' atmospheric compositions; however, some insights into their bulk compositions may be possible through the study of polluted white dwarfs (Bonsor et al., 2023).

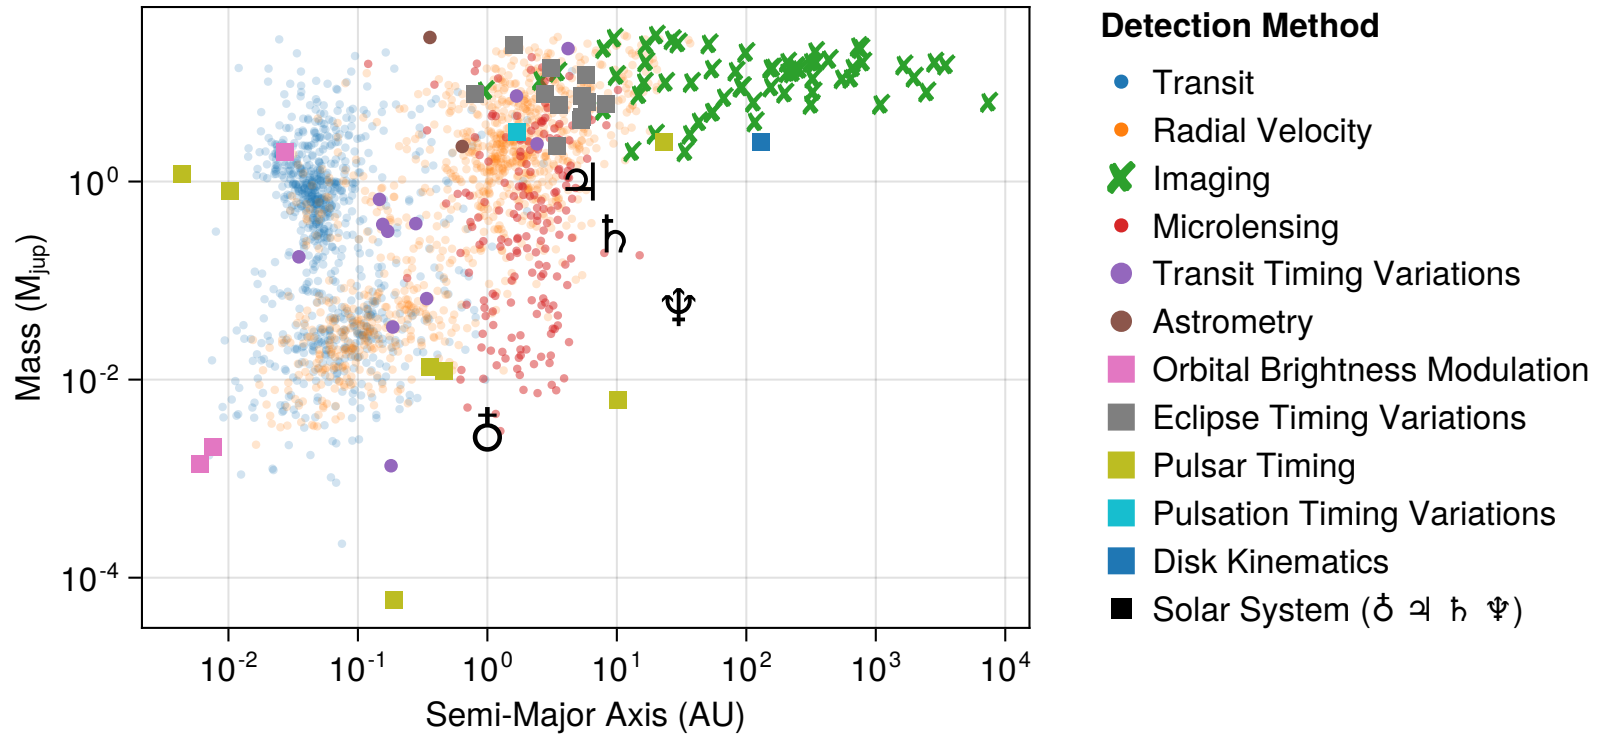


Figure 1.4 The mass vs. orbital semi-major axis distribution of all confirmed planets recorded by the NASA exoplanet archive (accessed 2023-05-05). To a large extent, this plot reveals the strong biases present in different detection techniques rather than the underlying population of planets. For most objects presented, the mass is estimated from another variable like transit depth, luminosity, etc. For radial velocity planets, $m \sin(i)$ is used.

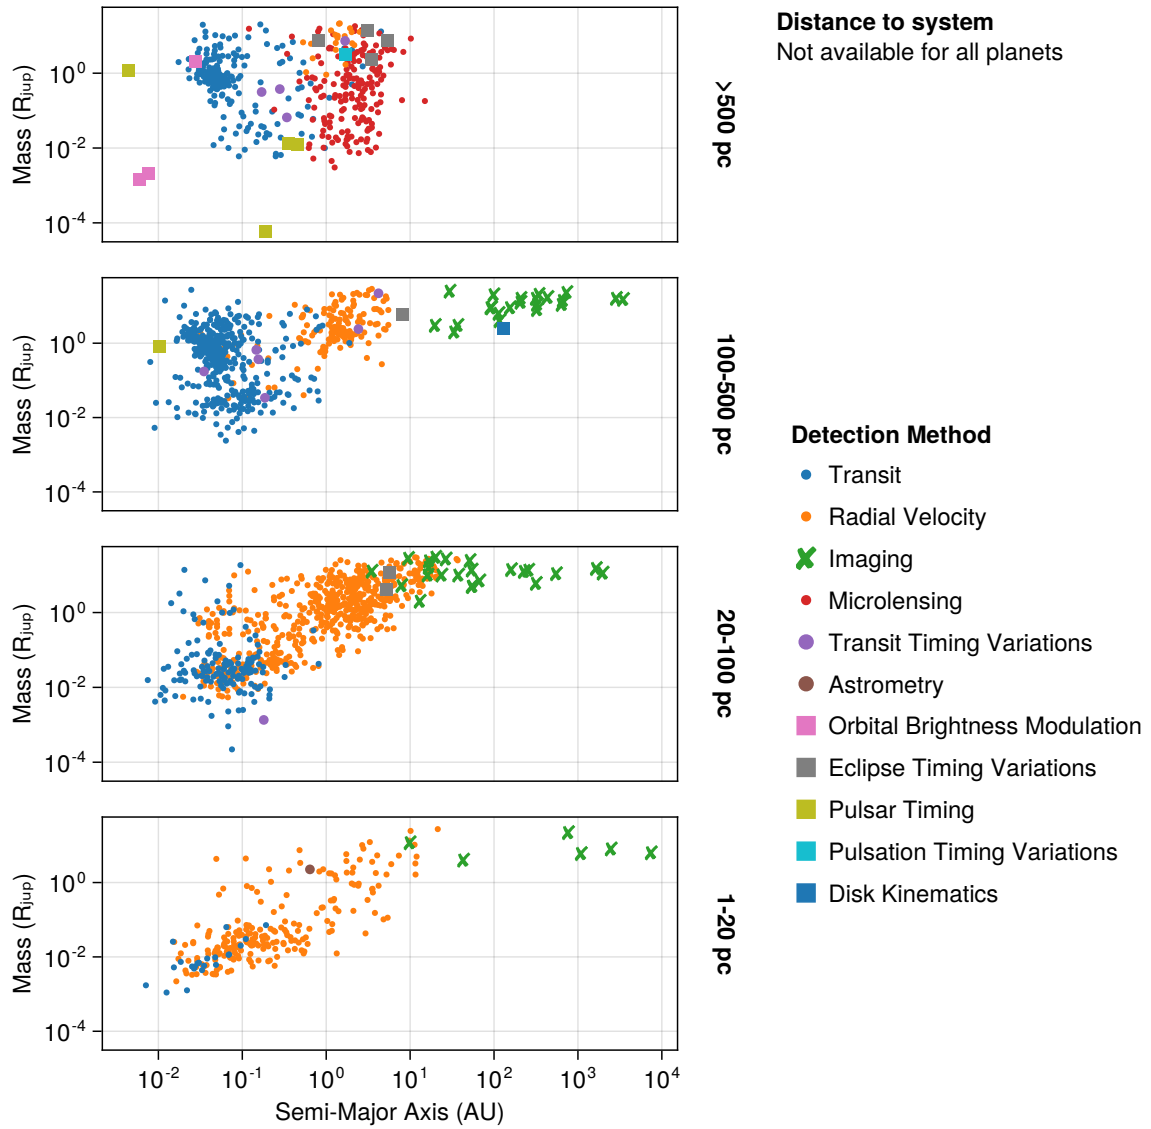


Figure 1.5 The mass vs. semi-major axis distribution of confirmed planets broken down by our distance from the system.

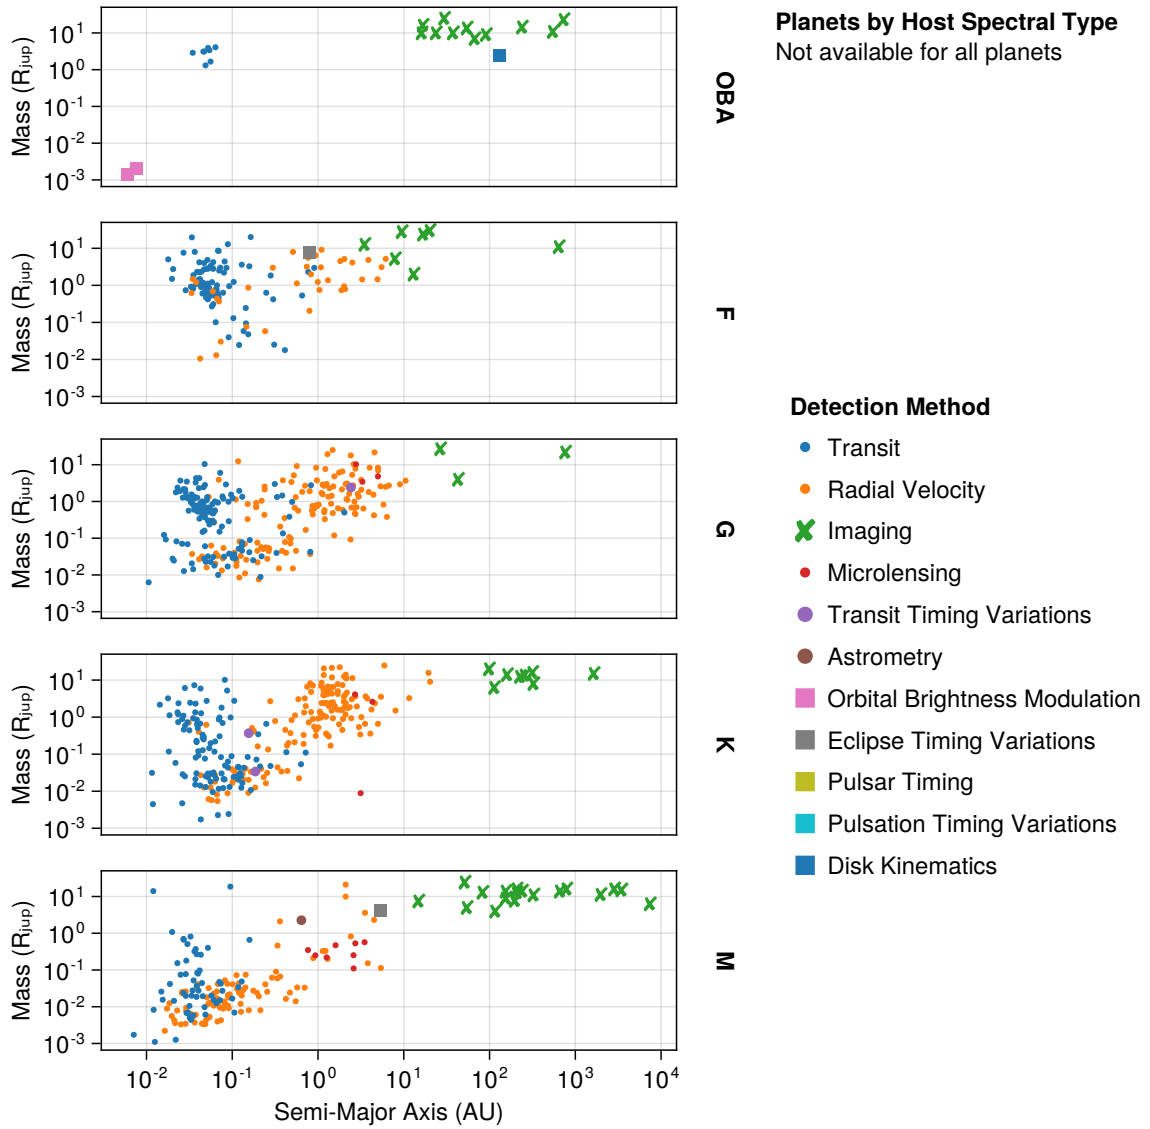


Figure 1.6 The mass vs. semi-major axis distribution of confirmed planets broken down by spectral type.

1.3 Velocimetry

Velocimetry is a way to infer the presence and mass of planets based on their dynamical effects on the motion of their host star. As a planet orbits its star, the star completes a proportional and opposite orbit around the system's centre of mass. From our perspective on Earth, the orientation of planets' orbits are randomly distributed in three-dimensional space—that is, other planetary systems are not aligned with our own solar system nor the plane of the Milky Way. Stars and their accompanying planets have a net centre of mass motion through the galaxy. Fortunately, we can appropriate this motion as constant for most all stars and search for the effects of planets as a perturbation on top of this steady velocity (Perryman et al., 2014). There are two ways we can measure this perturbation, depending on the direction of motion.

1.3.1 Doppler Radial Velocity

First, we can measure the velocity of the host star towards or away from us by monitoring the Doppler shift of its spectrum. This is most famously known as the radial velocity, or RV method. This method was pioneered by the Canadian astronomers Bruce Campbell, Gordon Walker and Stephenson Yang (Brennan, 2018) and was responsible for the earliest generally accepted detection of an exoplanet around a Sun-like star, 51 Pegasi b (Mayor, Queloz, 1995).⁴ The RV method is responsible for an impressive $1/5$ of all confirmed planets to date (NASA Exoplanet Archive, accessed 2023-05-05).

The radial velocity method is most sensitive to planets that orbit close to edge-on and is entirely insensitive to planets that orbit face on. It is also more sensitive to more massive planets and to planets that orbit closer to their stars. This latter effect has two sources. First, the same planet placed closer to the star creates a larger radial velocity signal that is easier to detect in a given time window. Second, planets that orbit further from their star have longer orbital periods according to Kepler's third law, $P^2 \propto a^3$, where P is the orbital period and a is the orbital semi-major axis. This means that very long observation campaigns are required to distinguish the perturbation of a planet from the steady velocity of the system.

This requires longer and longer observation time series to detect velocity shifts. A practical limit appears to be on the order of one decade, which is roughly the replacement lifecycle of telescope instruments (with many exceptions of course). Observations from

⁴Planets had previously been detected using the pulsar timing (Bisnovatyi-Kogan, 1993) and astrometry methods (J.H., 1988) to be discussed shortly, but these proved somewhat controversial.

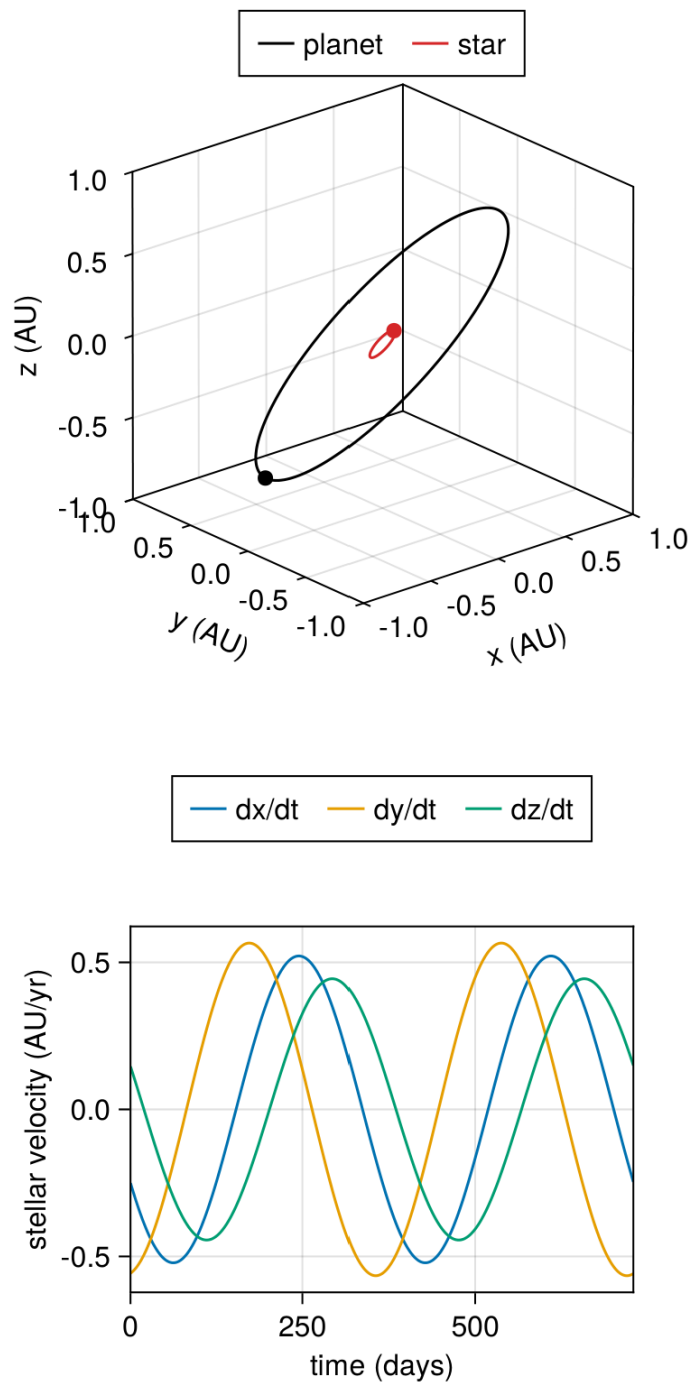


Figure 1.7 Planets perturb the position and velocity of their host stars over time. These signals can reveal the presence of a planet. $\partial z/dt$ can be measured with Doppler radial velocity while $\partial x/dt$ and $\partial y/dt$ can be measured with the astrometry method.

multiple successive instruments can be stitched together to access longer period planets (e.g. [Chilcote et al., 2021](#)), but this comes at the expense of reduced constraining power as instrument-specific RV zero points must be modelled as free parameters. Properties of the host star also affect the sensitivity of the radial velocity method. Stellar activity presents a noise limit to radial velocity searches. Young stars still settling onto the main sequence are typically quite variable, making this method most sensitive to quiet, mature stars.

The RV method is now just sensitive enough to detect planets with mass, separation, and host spectral type analogous to Jupiter (Figures [1.4](#) and [1.6](#)); analogues to all other solar system planets remain out of reach.

1.3.2 Astrometric Motion

The second way we can measure the velocity of the host star is to monitor its position in the plane of the sky over time. This is arguably the oldest technique for detecting planets, as similar ideas were used since antiquity to identify the planets in our own solar system (“planet” derives from the Greek *planētēs*, meaning “wanderer”). For exoplanets though, this technique has only recently become practical on a large scale, and the terminology has not been fully standardized. It is referred to as either “absolute astrometry”, “proper motion anomaly”(PMA), “astrometric acceleration”, or simply the “astrometry” method. The main way that the astrometry method is used today is through a combination of the Gaia ([GAIA Collaboration et al., 2021](#)) and Hipparcos ([van Leeuwen, 2007](#)) catalogues. These two satellites recorded the positions of stars in the sky periodically over the courses of their missions. These measurements were made for over one hundred thousand stars between 1989–1993 by Hipparcos, and hundreds of millions of stars between 2014–present by Gaia. The complete position measurement catalogue of Hipparcos is available, however at this time only certain processed and reduced data products are available from the Gaia mission, with data ending in 2017. For most stars, these include the average position and a single fitted straight-line velocity vector for the star during the 2014–2017 phase of the mission.

Despite only having access to these reduced data products from Gaia, the two catalogues can be cross-matched ([Brandt, T. D., 2021](#)) to calculate a total of three velocity vectors measured between 1989–1993, 2014–2017, and 1993–2016. The long baseline between the two mission epochs is approximately equal to the long term motion of the system’s centre of mass. Any discrepancy between the long term trend (“proper

motion”) and the velocity during the two mission epochs may then reveal an orbiting planet, hence the term “proper motion anomaly.”

In contrast to radial velocity, the astrometry technique is most sensitive to face on orbits. Since the astrometry method measures deflections in two directions, there is no orbital orientation where its sensitivity falls to zero. Only at most a reduction by a factor of $1/\sqrt{2}$ in sensitivity is lost for worst case, edge-on viewing angles. The astrometry method also is less sensitive to stellar activity than the radial velocity method (see for example [Kaplan-Lipkin et al., 2022](#), and the references therein).

By itself, cross-matched Gaia-Hipparcos proper motion anomaly data is not yet able to confirm the presence of planetary mass objects. Instead, it is most useful for guiding target selection for other methods (e.g. [Currie et al., 2021](#); [De Rosa et al., 2023](#)) and for determining planet dynamical masses which is not possible using any other single method. This is poised to change dramatically when the full Gaia ([GAIA Collaboration et al., 2021](#)) catalog is released. This catalogue will eventually contain epoch astrometry for millions of stars over the course of its planned five-year mission, and is ultimately expected to reveal on the order of 21,000 planets of mass $1M_{\text{jup}}$ or greater within 500 pc ([Perryman et al., 2014](#)).

Still, it would be possible to exceed Gaia’s sensitivity using targeted observations of individual stars. This has inspired the Toliman space telescope which hopes to perform a targeted astrometry search for planets around the α Cen system ([Tuthill et al., 2018](#)).

1.4 Transit

Like velocimetry, the transit method infers the presence of planets through effects on their host stars. Once every orbital period, planets that orbit nearly edge-on ($i \approx 90^\circ$) will pass in front of their stars and block some fraction of the starlight seen from the Earth. This fraction, called the transit depth, gives the ratio between the planet’s cross-sectional area and that of its star. By assuming that both bodies are spherical, this also gives the ratio between their radii. The time between repeated transits gives the orbital period. With knowledge of the star itself, these can be turned into estimates of the planets radius and semi-major axis.

Individual planets are unlikely to transit their stars, and this probability falls with increasing orbital separation. Thankfully, the advent of the satellites CoRoT ([Auvergne et al., 2009](#)), then Kepler ([Borucki et al., 2010](#)), and now TESS ([Ricker et al., 2014](#)) made

it possible to observe thousands of stars simultaneously with exquisit photometric accuracy. This made the transit method very effective for dragnet-style searches, and ultimately led to a massive increase in the number of detected exoplanets. That said, transit searches are poorly suited for studying very nearby systems since the relatively few nearby stars are unlikely to have planets that orbit with this lucky orientation.

Once transiting planets are found, they can in some circumstances be studied using transmittance spectroscopy. Instead of simply monitoring the photometry of the star over time, a spectrograph is used to split the light across many spectral channels and the transit depth is monitored separately in each. This provides the absorption spectrum of the planet's upper atmosphere. In this way, JWST may be able to detect chemical signatures associated with life on rocky planets around some M-dwarfs ([Wunderlich et al., 2019](#)). However, since transit depths are small for Earth sized planets around F, G, and K stars and because, statistically most stars with transmitting planets are far away ([Figure 1.5](#)), this means that many transits have to be averaged together to build up sufficient signal. This is further made difficult by the long orbital periods of solar system-like planets, resulting in infrequent observing windows. Indeed, even with the most optimistic assumptions about planet occurrence rates and instrument development, it will be infeasible to detect O₂ on Earth like planets using ground-based transit spectroscopy ([Hardegree-Ullman et al., 2023](#)) even with upcoming 30 m class observatories.

1.5 Direct Imaging

Direct imaging, also known as high contrast imaging, is a collection of techniques for producing actual images of planets that orbit nearby stars. In theory this could be either infrared light emitted thermally from the planet, light from the star reflected by the planet, or auroral emission from the planet's ionosphere. Direct imaging grew out of efforts to image brown dwarf companions in the 1990s, spurred on by advances to adaptive optics systems (to be discussed shortly). To date, only relatively bright thermal emission from young planets has been detected. These planets remain self-luminous as they cool from their formation tens or hundreds of million years ago.

This light provides a wealth of information about a planet. First, it provides dynamical information that fully describes its orbit including semi-major axis, inclination, and eccentricity, with the only exception being a \pm ambiguity in the planets motion towards

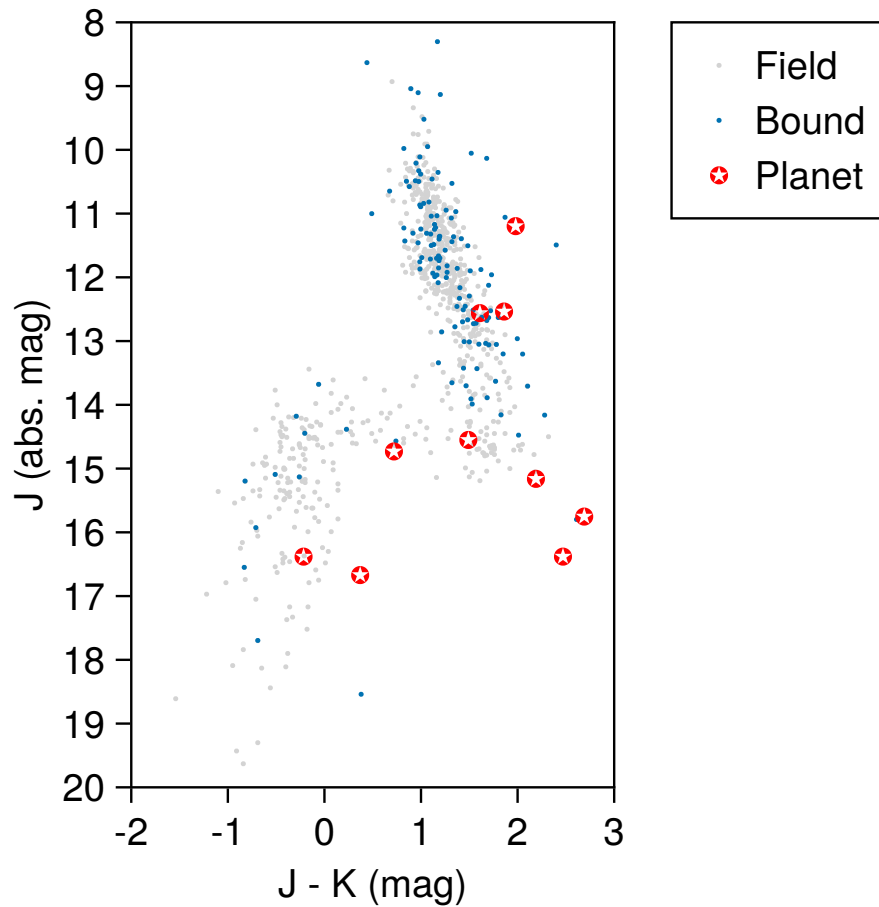


Figure 1.8 Colour-magnitude diagram comparing field brown dwarfs with bound brown dwarfs and high mass planets. Data from [Best et al. \(2020\)](#).

or away from the observer. Full information about a planet's orbit allows us to make inferences about its formation history, especially when multiple planets are present. In these cases, orbital resonances, eccentricity, and mutual inclinations all probe the migration history of the system.

For the planets themselves, multi-band photometry can be used to estimate the planet's mass, temperature, radius, bolometric luminosity, and composition (e.g. if the atmosphere has suspended silicates). For example, Figure 1.8 uses multi-band photometry to show how directly detected planets appear to fall into two populations, the L- and T-dwarfs. Medium resolution spectroscopy and model-based retrieval studies allow for detailed atmospheric characterization and determination of physical properties like mass, surface gravity, atmospheric structure, and $v \sin(i)$ (related to

planetary rotation rate). See for example [Konopacky et al. \(2013\)](#), [Wang et al. \(2020\)](#), and [Wang et al. \(2021\)](#).

In theory, accurate time-series spectro-photometry will allow for the detection of surface features in planets' atmospheres like hot spots and any other large scale structure. Longer term, it would also allow for Doppler-mapping of planet surfaces ([Cowan, Fujii, 2017](#)), radial velocity- and transit-based exomoon searches, and measurements of planet interior structures via seismology.⁵

For all of these results to be possible, one must first detect a planet. Even if a planet has already been discovered using an indirect method, its brightness, colour, and position cannot be fully known in advance. In the context of direct imaging, a detection means that one is sure that light in a particular location comes from the planet and not the star, with 5σ considered the usual minimum to claim an initial detection. To define the sensitivity of an instrument or particular resulting image, the term "contrast" is used. Contrast refers to the brightness ratio between a star and the faintest detectable planet and is typically reported as a function of separation. In most cases, the main challenge comes from achieving sufficient contrast between star and planet rather than sufficient total signal from the planet itself, though this may change as instruments improve and we begin to target fainter planets.

Direct imaging observations are currently conducted in the Near Infrared (NIR) between 1 and 5 micron. This is due to a balance of several factors. Adaptive optics systems (to be discussed shortly) provide the best quality image correction in the NIR. Further into the IR, however, and the resolution of our telescopes begin to suffer and the thermal background increases dramatically. JWST is the exception, and can observe planets at wavelengths as red as 25 micron thanks to its vantage point in space at the L2 Lagrange point, sun shield, and cooled detector. In the NIR, the most easily detectable objects are young planets 5-100 Myr old. These objects remain warm (e.g. on the order of 1500 K) from formation and their peak emission is redder than their host stars, reducing contrast requirements. [Figure 1.9](#) shows how planet photometry depends on mass and the choice of filter.

In rare occasions, it is also possible to observe accreting protoplanets at optical wavelengths thanks to their bright $H\alpha$ line emission ([Hashimoto et al., 2020](#)).

⁵Analogous to asteroseismology and kronoseismology (([Hedman, Nicholson, 2013](#)) and subsequent works)




Figure 1.9 Brightness of a 30 Myr old hot start planet as function of mass and wavelength. Interpolated from the COND0 models of [Baraffe et al. \(2003\)](#).

1.5.1 Adaptive Optics

The first step to directly imaging planets from the ground is to use an Adaptive Optics (AO) system. Adaptive optics systems were first deployed in astronomy in the 1990s and were key enabling technologies for the first brown dwarf and giant exoplanet surveys like the Gemini Deep Planet Survey (GDPS; [Lafrenière et al., 2007a](#), not to be confused with the Gemini Planet Imager Exoplanet Survey) and the international International Deep Planet Survey (IDPS; [Galicher et al., 2016](#)). Adaptive optics systems are responsible for measuring phase errors in an optical system and performing a real time correction. This is necessary from the ground primarily because of the Earth's shimmering atmosphere. Due to natural convection, the atmosphere is composed of pockets of different densities and therefore different optical indices of refraction. Since

these pockets arise from turbulence, they follow a Kolmogorov power law distribution with an exponent of roughly $-\frac{5}{3}$. That said, the actual turbulent evolution of the Earth's atmosphere is relatively slow. The dominant evolving source of error for optical and infrared observatories is in fact the wind blowing this convection pattern, or "phase screen," across the telescope's pupil (the "frozen flow" hypothesis or assumption). An example Kolmogorov phase screen is visible in Figure 1.10. These phase errors averaged over time lead to a "seeing"-limited Point Spread Function (PSF) as shown later in Figure 1.13. For these reasons, direct imaging is not possible from a seeing-limited instrument without prohibitively large integration times. In fact, even space observatories are set to adopt adaptive optics to achieve greater optical stability.

Uncorrected optical errors, called "aberrations", create a pattern of random "speckles" across images, falling in intensity with increasing separation from a star. Individual speckles come from optical errors at a particular spatial scale, and so typically each have their own associated temporal behaviour. Speckles very close to the star are caused by spatially large, lower order aberrations, while speckles further away are caused by increasingly fine spatial scale errors. Speckles evolve on a range of timescales from milliseconds to hours and days. These are typically described as either static errors if they last for days or more (indicating a miscalibration in the optical system), atmospheric speckles if they last seconds or less, and quasi-static speckles if they evolve on a timescale in between these two. The overall spatial distribution is usually described as following a modified Rician distribution (Soummer et al., 2007), though the exact pattern varies considerably by instrument and observing conditions.

To compensate for evolving phase errors, adaptive optics systems integrate three key components (see schematic in Figure 1.11). The first is a device to measure phase errors across the pupil known as a wavefront sensor. Wavefront sensors come in many forms, the most common of which are the Shack-Hartmann and pyramid wavefront sensors (Ragazzoni, 1996). The second is an active optical component that is capable of adjusting the phase of incoming light. These are most commonly Deformable Mirrors (DMs) that contain a reflective surface that bends according to computer control. Other active elements can be simple tip/tilt mirrors, active mirror segments, and spatial light modulators. The final component is a real time computing system that interprets the measurements of the wavefront sensor or sensors and sends an appropriate command to the active optical elements in order to correct the measured error.

Though digitally controllable, all active optical components are inherently analogue devices. DMs in particular suffer from non-linearities and hysteresis effects. For this




Figure 1.10 Simulated Kolmogorov phase screen representing atmospheric turbulence, or “seeing.” In the frozen-flow hypothesis, wind blows a static phase screen across the telescope pupil. The resulting star image is the complex Fourier transform of the telescope pupil as amplitude and this turbulence as phase.

reason, adaptive optics systems are typically used in a closed-loop configuration. That is, the wavefront sensor is placed after the correction mechanisms such that it only measures deviations from a perfect wavefront. The adaptive optics controller can then take incremental steps to improve the wavefront rather than having to perfectly correct each error without feedback. Closed loop operation greatly improves the quality of the AO correction in most situations. Closed loop control and the control theory that describes it is a rich and deeply studied field, with a range of mathematical tools available for analyzing, designing, and improving the behaviour of control loops such as these (See [Nise, 2020](#), for an introductory text).

The simplest widely used control scheme for adaptive optics is a basic linear integrator. In this scheme, the command to the DM is updated at each time step by adding the current error multiplied by a gain parameter between 0 and 1. Integrating controllers are used in the work described in Chapter 5.

To enable direct imaging from the ground, adaptive optics systems must reach near-diffraction limited image quality. That is, the actual point spread function of the telescope and atmosphere must approach a perfect state limited only by the physical




Figure 1.11 Conceptual schematic of an adaptive optics system followed by a basic Lyot coronagraph system. The adaptive optics system measures and corrects errors in the incoming wavefront, but is blind to errors that occur after the wavefront sensor. The focal plane mask and Lyot stop combine to create a spatial filter that suppresses the light of the star while allowing any off-axis sources like planets to pass through.

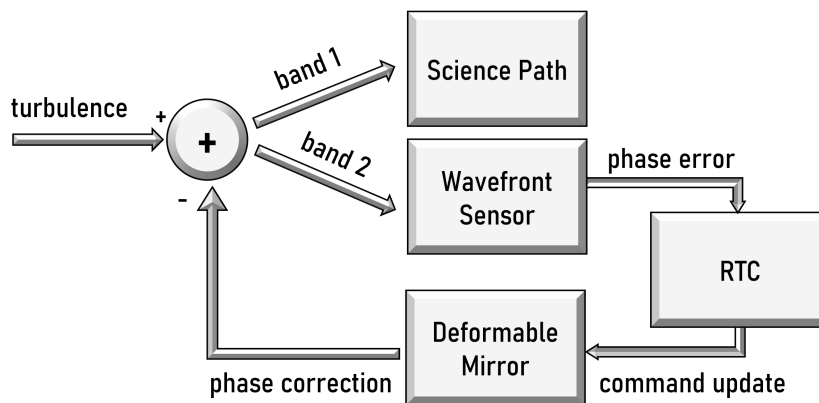


Figure 1.12 A conceptual block diagram of a traditional adaptive optics system.

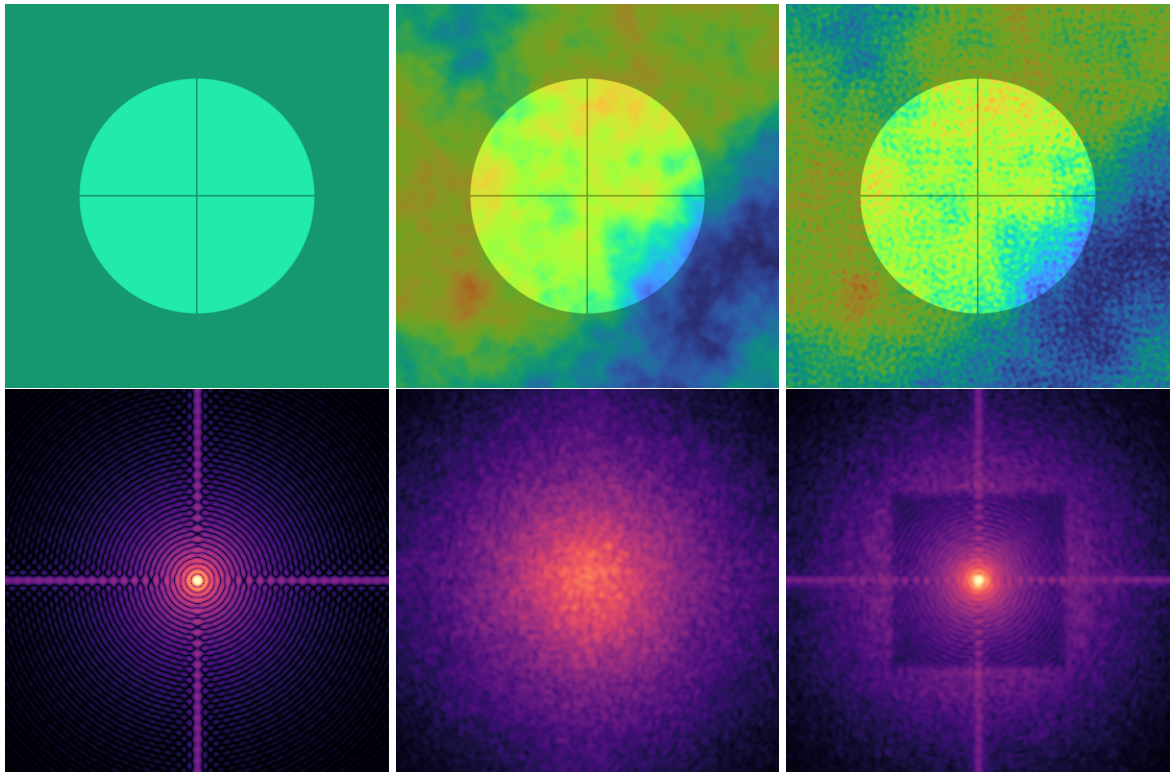


Figure 1.13 Simulated telescope PSFs with and without the effects of the atmosphere and adaptive optics system. The top row shows an instantaneous phase screen moving across a telescope pupil while the bottom row shows the resulting time-averaged PSF. **Left:** a simulated diffraction limited PSF. **Centre:** a simulated seeing limited PSF. **Right:** a simulated very well-corrected AO PSF. The core approaches the diffraction limited case but is surrounded by a halo of residual atmospheric speckles that lie beyond the DM's control region. The phase panel at the top right has a stretched linear colour scale to show the high spatial frequency residuals left by the AO system.

limitations of the finite telescope aperture. It should be noted however, that this is only necessary at spatial scales where one wishes to search for a planet. Deformable mirrors possess finite numbers of physical actuators that deform the optical surface. Accordingly, the Shannon-Nyquist sampling theorem dictates that they can only correct phase errors that are as finely spaced as twice the projected spacing between actuators. Transformed from the pupil plane to the focal plane, this finest correctable feature corresponds to a maximum corrected separation from the star. As a result, the adaptive optics systems used by latest generation planet imaging instruments create a “dark hole”—a small region immediately around stars that is well corrected, and therefore devoid of speckles (Figure 1.13, third panel).

1.5.2 Coronagraphy

After wavefront correction from an adaptive optics system, the next component used to enable direct imaging is a coronagraph. Coronagraphs, first developed by Lyot (1930) to study the sun’s corona, are used to suppress light from stars. Mathematically, coronagraphs are a form of spatial filter applied at one location in the image plane.

To use a coronagraph, an intermediate focal plane and pupil plane are introduced between the telescope and detector. The system is adjusted to direct the star onto a Focal Plane Mask (FPM) which is responsible for reflecting or diffracting the starlight to a different spatial scale at the subsequent pupil plane. This typically forms a ring of bright light near the outside edge of the telescope pupil. Next, a pupil stop is used to suppress this ring of starlight (a “Lyot” stop, named after its inventor) at the expense of slightly reduced overall angular resolution.

Since Lyot’s original work, a wide variety of coronagraphs have been developed with increasingly sophisticated methods and materials. For the work described in this dissertation, there are two relevant extensions. The first is an apodized Lyot coronagraph (Soummer, 2004). This coronagraph design adds an element in a preceding pupil plane that shapes, or “apodizes” the instrument’s PSF before it propagates to the focal plane mask. This approach suppresses diffraction rings at the expense of a broader PSF core, and is the analog equivalent of applying a window function to digital data prior to taking the Fourier transform (see Chapter 6 an application to spectroscopy). The second extension is an optical vortex coronagraph (Foo et al., 2005). This coronagraph design replaces the occulting disk focal plane mask with a spiral phase ramp. Away from the centre, the ramp is slow and has no effect on off-axis light, but at the core, it creates a

null in the electric field diffracting light out to the edge of the telescope pupil. Simulated examples of an apodized Lyot coronagraph and a vortex coronagraph are shown in Figure 1.14.

Besides suppressing starlight, coronagraphs also have some undesired impact on light from planets, known as throughput loss. The throughput of a coronagraph depends on its design and must be traded against its ability to suppress starlight. It is also a function of separation and falls impractically low at a location known as its inner working angle (IWA, typically defined as 50% planet throughput). By analogy, the furthest separation that is well corrected by an AO system may be referred to as the outer working angle (OWA).

Coronagraphs are used in almost all direct imaging instruments; however, there may be situations where coronagraphs degrade performance at very close separations from the star. In these situations, non-coronagraphic imaging may have better speckle stability which therefore enables better post-processing, as will be discussed in Chapter 4. For these reasons, many of the Keck/NIRC2 (PI: K. Matthews) observations presented later in this dissertation were taken without a coronagraph.

1.5.3 Quasi-Static Speckles

Unfortunately, coronagraphs by themselves (even with modern designs) failed to live up to expectations in real world applications. The issue was not simply imperfect correction of atmospheric seeing by adaptive optics systems. Such atmospheric residual speckles should rapidly average into a smooth halo that, while contributing photon noise, is easy to distinguish from a planetary companion. Instead, residual speckle patterns remained and showed little improvement with increasing integration time. These were termed “quasi-static” speckles since they changed too slowly to average into a smooth halo and too quickly to be measured once and digitally subtracted (Marois et al., 2003; Soummer et al., 2007; Marois et al., 2008b).

These quasi-static speckles were found to have a few origins, mostly relating to internal instrument challenges and interactions between an instrument and the atmosphere. The first category of quasi-static speckles are caused by non-common path errors. Adaptive optics systems operate by splitting light from the star into two paths: a wavefront sensing channel and a science channel. Any differences in alignment or optical surfaces after this beam-splitter cause the AO and science wavefronts to diverge. Then, a perfect correction of the wavefront as observed by the AO wavefront sensor




Figure 1.14 Simulations of an apodized Lyot coronagraph (top) and a charge seven vortex coronagraph (bottom). From left to right: telescope pupil, focal plane immediately before the focal plane mask, focal plane mask, Lyot stop, and resulting PSF. The top half of each panel shows the logarithm of the electric field intensity in grey and the bottom half shows the phase from 0 to 2π as a hue. The display of the final PSFs at the right are brightened by a factor of one million.

results in residual phase errors in the science channel. At the science focal plane, these phase errors cause speckles. Since they result from a kind of optical miscalibration they have a very different temporal evolution than atmospheric speckles. They change over seconds to hour timescales as the instrument and observatory flex, shift, and expand due to telescope tracking, changing gravity vectors, and uneven thermal expansion.

A second category are chromatic errors. Since the two channels typically receive different wavelengths of light,⁶ any error seen by the AO system that depends on wavelength results in an imperfect correction in the science channel.

A third category are amplitude errors. Just like phase errors, non-uniform illumination of the telescope pupil creates an imperfect PSF marred by speckles. For fully static features like the secondary mirror support structure, these are typically referred to as “diffraction spikes”, but these spikes are just one extreme of a general class of amplitude error. Besides the instrument itself, the atmosphere can be a source of amplitude errors due to the Talbot effect (Marois et al., 2006c). Phase errors from atmospheric turbulence at elevations that are not conjugated to the DM propagate into an amplitude error. Amplitude errors are typically ignored by adaptive optics systems with a single deformable mirror because they are mostly small compared to phase errors in general applications; because unlike phase errors, they cannot be simultaneously corrected on both sides of the telescope’s PSF; and because even if one side of the telescope’s PSF is sacrificed, the resulting correction is chromatic.

Quasi-static speckles are challenging to remove since, essentially by definition, they change too quickly to calibrate at the start of an observing sequence, but too slowly to average into a smooth halo. Today, quasi-static speckles are suppressed using a variety of observing strategies and digital post-processing techniques that will be described in Section 1.5.5. In addition, work is now progressing on suppressing quasi-static speckles optically (Bottom et al., 2016; N’Diaye et al., 2016; Galicher et al., 2019; Vigan et al., 2019; Potier et al., 2022).

Instruments like the Gemini Planet Imager (GPI; Macintosh et al., 2014) partially mitigate non-common path errors by using super-polished optics, where each optical element before the FPM is super-polished to a nanometer RMS smooth surface. Unfortunately there is a practical limit to how smooth optics can be polished, and optical components can also be expected to shift and flex. To go beyond this, one strategy is to develop methods for calibrating the two optical paths on the fly by adding an additional wavefront sensor near or at the science detector, and in particular, after the

⁶This is not a strict requirement; a grey beam-splitter could be used in place of a dichroic.

coronagraph. It should be noted that we cannot simply move the AO wavefront sensor to this location as the coronagraph works to suppress that vast majority of the starlight, leaving insufficient signal for the AO system to respond to the atmosphere.

Techniques for calibrating the science and AO wavefront sensor paths include temporal modulation approaches like Electric Field Conjugation (EFC; [Give'on et al., 2007](#)), pupil plane interferometers like the Zernike wavefront sensor ([N'Diaye et al., 2016](#)) or GPI's original shearing interferometer ([Wallace et al., 2010](#)), and finally, focal plane interferometers like the Self-Coherent Camera (SCC; [Baudoz et al., 2005](#)), which will be explored in this dissertation.

1.5.4 Spectroscopy

One of the promises of direct imaging is that it enables the detailed characterization of individual planets in cases where transit spectroscopy is not viable. A very coarse way to access spectral information is to take multiple observations using different filters, either one after another or simultaneously using multi-band imagers like TRIDENT ([Marois et al., 2005](#)) and IRDIS ([Beuzit et al., 2019](#)). To access finer grained information, one can use an Integral Field Spectrograph (IFS) that captures images with light decomposed into dozens or hundreds of wavelength slices, like a prism. For known planets with accurately measured positions, long-slit and fibre-fed spectrographs can unlock detailed spectral features.

Spectroscopy has been used with direct imaging to measure many features. These include the effective temperature of planets (different from the equilibrium temperature) and atomic (Na, K) and chemical (CO, FeH) compositions (For a review, see [Currie et al., 2022](#)). Measurements of broad absorption features like the CH₄ (methane) bands can trace the L-T dwarf transition, while derived properties like the C/O ratio may be able to trace back a planets formation pathway ([Hoch et al., 2022](#)). Higher resolution spectroscopy can also provide insights into planets' surface gravity, rotation rate, and orbital velocity ([Wang et al., 2021](#)).

High resolution spectroscopy also promises to increase our sensitivity to faint planets. This is accomplished by cross-correlating a template spectrum against a measured spectrum that contains both planet and starlight ([Wang et al., 2017](#)). A challenge with this approach is that either the location of the planet must be known in advance with relatively high accuracy, or a bundle of many fibres and spectrographs must be used at great expense, and still resulting in fibre-coupling losses. A promising middle ground

between current low-resolution IFSs and the concepts of high-dispersion coronagraphy is presented in Chapter 6.

1.5.5 Differential Imaging

To help suppress quasi-static speckles and reveal planets, a suite of observing and post-processing techniques are used. Three widely used differential imaging techniques that are relevant to this dissertation are Reference star Differential Imaging (RDI; recently demonstrated in [Wahhaj et al., 2021](#)), Angular Differential Imaging (ADI; [Marois et al., 2006a](#)), and Spectral Differential Imaging (SDI; [Walker et al., 1999](#); [Racine et al., 1999](#); [Marois et al., 2000](#)). They are differential in the sense that the intensity of the planet signal is modulated in a different and controlled way than light from the star. One then attempts to model the stellar point spread function in a given target image using a combination of other stellar point spread function reference images. If this model is well-constructed, subtracting it from the target image removes some quasi-static speckles without removing the planet.

The simplest case is reference differential imaging. In reference differential imaging, one captures images of another star or stars before and/or after observing the target. Ideally, the reference star or stars are chosen such that they appear near the target system in the sky, are of similar brightness, and have similar spectra in the band of interest. These considerations ensure that the reference PSFs are as similar as possible to the PSF in the target images. RDI has proven to have only a small benefit for ground based instruments as the conditions are insufficiently stable to preserve the quasi-static speckle pattern across sequences. Space based telescopes including Hubble and Webb on the other hand are far more stable, and accordingly RDI with large archives of reference images are one of the main strategies used for direct imaging in space.

The most powerful observing technique for detecting planets is arguably Angular Differential Imaging (ADI), also known as roll subtraction in the space telescope community. In ADI, one rotates the entire observatory about its axis during the observing sequence. For large ground based observatories, this would of course be mechanically impractical, but thankfully is accomplished naturally by the rotation of the Earth over the course of a night.⁷ For instruments like the Gemini Planet Imager (GPI) at prime or Cassegrain focus, the rotation angle of the telescope pupil remains

⁷This is the case for telescopes on Altitude-Azimuth mounts. Equatorial mounts counteract this rotation making them unsuitable for ADI but are rare on telescopes greater than 1-2 m in diameter.




Figure 1.15 Simulations of angular differential imaging (ADI, top) and spectral differential imaging (SDI, bottom).

fixed on the detector while astrophysical sources like planets appear to rotate around the star (see example in Figure 1.15). For instruments like Keck’s NIRC2 that occupy a Nasmyth focus, the pupil-derotator is adjusted to accomplish the same goal i.e., fix the rotation of telescope pupil in the instrument making physical sources appear to rotate. ADI is very effective because it uses images captured close in time without opening and closing an AO loop,⁸ and without changing wavelengths. That said, ADI is less effective at suppressing quasi-static speckles that are close to the star. This is because for a given rotation angle, the physical signal of a planet that is close in moves less far relative to the speckle pattern compared to a planet that is further away. Effectively, larger separations from the star benefit from self-reference images that are taken close in time for a given displacement requirement.

The third case considered in this dissertation is Spectral Differential Imaging (SDI). In SDI, one captures images of the same star in multiple wavelengths simultaneously. This can be done either with dual or multi-band imaging (e.g. TRIDENT and IRDIS) or with an integral field spectrograph. There are two separate effects that can be used to

⁸Stopping and restarting AO systems can scramble residual speckles making for unsuitable reference images. This can be observed with Keck/NIRC2 when dithering a target across a detector.

our benefit. The first, which does not require any assumptions about the planet, is the optical magnification effect that occurs between images at different wavelengths. The resolution of the instrument is measured in λ/D . For a fixed diameter D observatory, longer wavelength λ observations have reduced resolution. Another way to express this is to say that the PSF is magnified at longer, redder wavelengths. As a result, the speckle pattern in an image scales radially from the star proportional to the ratio of the wavelengths in question. Any planets, however, will remain at a fixed apparent separation from the star. This outwards radial motion of the speckles compared to the star provides a differential signal used in SDI. The second closely related effect is that the spectrum of the planet may differ significantly from the star. This is especially true if the observation wavelengths are chosen to maximize this effect. One then assumes a known spectrum for the planet which makes it possible to treat the brightening and darkening of the planet with wavelength as another source of intensity modulation. A benefit of SDI compared to ADI is that the reference images are taken simultaneously which prevents any time evolution. In practice, SDI is often less effective than ADI. This is because images taken in different wavelengths often differ due to chromatic wavefront errors within the instrument (Gerard et al., 2019a). This is caused even by camera detectors themselves since the magnification effect used by SDI changes the spatial scale of speckles relative to detector pixels (the “pixel MTF” effect described in Marois (2004) and further discussed in Chapter 5). Unlike time evolution used by ADI, which is random, chromatic evolution of the wavefront is largely a systematic effect which may not average away in longer sequences. For further discussion and a potential solution, see Chapter 6.

Both ADI and SDI are less effective at close separations from stars since there is less apparent displacement for the same magnification ratio or rotation angle, though the spectral diversity effect in SDI remains unaffected.

Besides RDI, ADI, and SDI, Polarization Differential Imaging (PDI; Baba, Murakami, 2003) is also used to study debris disks and may have some niche applications to detecting the ocean glint from habitable exoplanets (Lustig-Yaeger et al., 2018). Two additional types of differential imaging, orbital differential imaging and coherent differential imaging, are somewhat distinct from these methods and will be introduced later in this dissertation.

Once these differential imaging observations are made, the next step is to post-process the data into cleaned images. All widely used post-processing techniques create a model stellar PSF for each target image by finding the linear combination of reference

images that best match the target image. An in-depth discussion of these methods, their limitations, and an improved approach is presented in Chapter 2.

1.5.6 Self-Coherent Camera

The Self Coherent Camera (SCC) is a concept first proposed by [Baudoz et al. \(2005\)](#). The basic design of an SCC is a common path interferometer. When the light of the star hits the coronagraph, it is diffracted out to high spatial frequencies. This light is then blocked in the pupil plane by a Lyot stop, suppressing the light from the star. Usually, this light is absorbed or reflected away by the Lyot stop. In an SCC, however, a small part of this blocked light is re-introduced from off-axis. This light is spatially filtered through a pinhole to produce a stable reference beam. In this design, the two beams follow an identical path after the Focal Plane Mask (FPM) and so the reference beam is perfectly coherent with any starlight that leaks around the coronagraph. When both beams arrive at the focal plane of the instrument, they coherently interfere and residual starlight becomes fringed. These fringes encode the complete complex electric field of the residual starlight compared to the stable reference beam. Crucially, any light from other objects that did not hit the focal plane mask, such as light from a planet, remains unfringed. This is because over any reasonable bandpass, light from the star is incoherent with light from other sources.⁹ Once the electric field of the residual starlight is measured via the fringes, one can use a deformable mirror to actively correct the incoming wavefront and post-processing to digitally remove speckles (Figure 1.17).

It should also be noted that, for current 10-m class observatories, the disks of dwarf stars are unresolved. This means that light from across the stellar disk is perceived as coming from a single idealized point source. Future extremely large telescopes may begin to marginally resolve the disk of some nearby stars which would act to slightly reduce the fringe visibility.

A second option that can be used separately or in conjugation with active correction is to use this information in a post-processing step to digitally subtract any light that is fringed. This process is called Coherent Differential Imaging ([Guyon, 2004](#); [Baudoz et al., 2006](#); [Bordé, Traub, 2006](#); [Give'On et al., 2007](#); [Serabyn et al., 2011](#); [Sauvage et al., 2012](#); [Gerard et al., 2018a](#); [Gerard, 2020](#); [Potier, et. al., 2022](#)). CDI is powerful because it creates reference images using simultaneous information, like SDI, and that are in

⁹The term “coherence” is somewhat ambiguous. In theory, any two sources would create an interference pattern with a sufficiently narrow bandwidth or a sufficiently short integration time. What matters in practice is that planet and starlight are incoherent over a detector integration timescale.

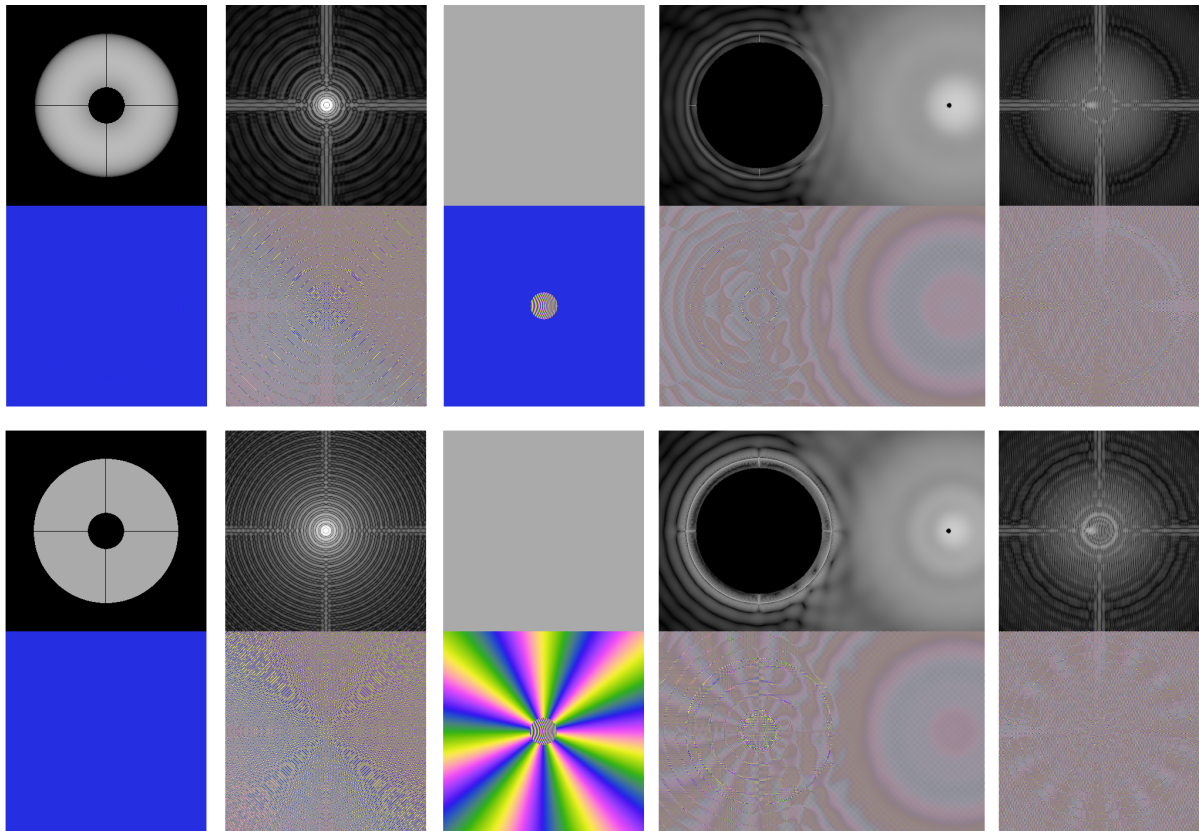


Figure 1.16 Simulation of modified FAST SCC coronagraphs based on an apodized Lyot coronagraph (top) and vortex coronagraph (bottom). Simulations of an apodized Lyot coronagraph (top) and vortex coronagraph (bottom). From left to right: telescope pupil, focal plane immediately before the focal plane mask, focal plane mask, Lyot stop, and resulting science PSF. The top half of each panel shows the logarithm of the electric field intensity in grey and the bottom half shows the phase from 0 to 2π as a hue. The display of the final PSFs at the right are brightened by a factor of one million.

the same wavelength, like ADI. No rotation or scaling is needed, which means that CDI does not drop in performance close to the star. There is also no negative impact on extended structure which would simplify the analysis of debris disks compared to ADI and SDI. That said, like SDI, CDI is still affected by the detector pixel MTF effect (see Chapter 5) and correlated residuals.

The self-coherent camera has been demonstrated on-sky by [Galicher et al. \(2019\)](#) with a $5\times$ improvement in contrast. However, with traditional coronagraphs the SCC fringes are very dim. If the residual wavefront evolves significantly before the fringes can be measured, then very little contrast gain can be achieved.




Figure 1.17 Conceptual schematic of an adaptive optics system and basic Lyot coronagraph system, followed by a focal plane wavefront sensor like the SCC. Additional information is measured at the science focal plane and used to adjust the AO response to remove NCPAs (green).

For this reason, [Gerard et al. \(2018a\)](#) proposed a new class of modified coronagraph designs that significantly boost the intensity of the reference beam and thus the fringes. They do this by literally tilting the light from the star off to the side and directing it to the reference beam pinhole. This decreases the time it takes to measure fringes from seconds or minutes, down to just tens or hundreds of milliseconds—enough time to measure and correct for many aberrations before they change. Simulations of modified FAST SCC coronagraphs are shown in [Figure 1.16](#).

If this modified SCC design were proven in realistic laboratory tests and deployed in a direct imaging instrument, sensitivity may increase by multiple orders of magnitude. In theory, [Gerard et al. \(2018a\)](#) and [Gerard et al. \(2019b\)](#) indicate that contrast could increase arbitrarily proportional to the square root of the integration time at least

until other contrast limits are reached. These other limits include atmospheric effects such as speckles that change too quickly to measure before they change, instrumental effects like incoherent ghosts caused by internal reflections, and astrophysical limits like exozodiacal light (Mennesson et al., 2018).

The key elements missing from deploying a FAST SCC for science operation include:

- developing effective calibration routines;
- mitigating real world sources of noise like saturated data, hot and dark pixels, and evolving detector bias current;
- demonstrating closed loop operation;
- developing strategies to prevent windup while still rejecting static speckles,
- testing that performance meets theoretical expectations;
- integrating the SCC as part of a larger adaptive optics system;
- and ensuring robust operation in changing atmospheric conditions (like passing clouds).

If these issues were addressed, the FAST SCC could live up to its theoretical potential: enabling routine access to Jupiter mass planets as close as 5 AU from their stars using current 8 m class facilities.

Building upon the FAST SCC, an open question is how technology can be extended to improve planet characterization in addition to detection sensitivity. The SCC and its FAST variant operate in single narrow bandwidths of 1-2% whereas deformable mirrors apply phase offsets to all wavelengths at once. Since the atmosphere and instrument are at least somewhat chromatic, a perfect correction across a wide bandwidth is not possible. Still, if one assumes that phase errors change slowly with wavelength, one may expect that the correction quality is adequate across a broader bandwidth of at least 10-20%. Accordingly, if we could combine the SCC with an integral field spectrograph, then we could use the improved starlight suppression of the SCC and CDI to also better characterize planets.

A major challenge to combining an SCC with an integral field spectrograph is that the SCC encodes phase offsets as fringes which require greater spatial sampling. Resolving these fringe structures across multiple wavelengths would increase pixel count requirements well beyond existing direct imaging instruments.

Solutions to all of these challenges will be presented in Chapters 5 and 6.

1.6 Combining Exoplanet Detection Methods

Different exoplanet detection techniques can be combined to learn more about individual planets. The two most common techniques to be combined are the radial velocity and transit methods. Typically, a planet is detected via the transit method and then the star's radial velocity is monitored over time. Both methods directly measure the period of the planet's orbit, so it is obvious when a perturbation to both signals is caused by the same planet. Additionally, since the planet is known to transit, the inclination of the orbit is 90 degrees. This resolves the $m \sin(i)$ ambiguity from the radial velocity data allowing the planet's dynamical mass to be measured. Since the planet's radius is known from its transit depth, one can then calculate its bulk density. At this point, large catalogues of transit and RV planets with densities have been compiled, including for example, the Transiting M-dwarf Planets catalogue (Trifonov et al., 2021).

Other combinations of exoplanet detection methods are less developed in the literature. One example is using velocimetry (RV and/or PMA) to guide direct imaging surveys towards stars that exhibit significant proper motion anomaly (Currie et al., 2021) or long-term RV trends (Ryu et al., 2016). Since large and wide-separation planets detectable by direct imaging are relatively rare (Nielsen et al., 2019; Fulton et al., 2021), this target selection guidance can greatly improve survey yields. This does come at the cost of not being able to determine planet population statistics from any found objects, but future Gaia astrometry (Sozzetti et al., 2015) and Nancy Grace Roman microlensing surveys (Johnson et al., 2020) should complete the basic picture of planet occurrence by mass and semi-major axis.

Once detected, direct imaging and velocimetry can be combined to accurately measure the dynamical mass of planets as was done using RV by Lacour et al. (2021) and using PMA by Currie et al. (2023). This is possible since relative astrometry extracted from direct images fully determines the planet's orbit up to a \pm ambiguity in Ω . Velocity signals, which happen to resolve this ambiguity, can directly determine to the planet's dynamical mass.

These approaches combine measurements that are independently extracted from each dataset. This requires an independent robust detection by each method and accordingly, only a subset of all detected planets can be studied in this way. Another more general approach would be to extract composite properties like dynamical mass directly from all available data. This would be possible since all detection methods constrain planets' orbits in ways that overlap with each other—that is, they are not

measurements of independent properties. The mathematical relationship between different kinds of data is not straightforward and is in general highly non-linear, but can nonetheless be connected through joint orbit models.

Joint modelling of direct imaging, radial velocity, and proper motion anomaly has been proposed and carried on the ϵ Eridani system (Mawet et al., 2019; Llop-Sayson et al., 2021). In these works, the authors constrained the orbit and photometry of the planet ϵ Eri b though in this case though were not yet able to constrain the planet's photometry above zero.

An adjacent problem is that of orbital motion. Thanks to Kepler's second law, planets closer to their stars move much faster than those that are further away. As direct imaging improves in sensitivity and begins to probe separations inwards of 5 AU, planets move fast enough that it can become difficult to detect them before they smear out across observations. Some efforts like those of Nowak et al. (2018) have improved this situation by developing tools to "de-rotate" or "de-orbit" sequences of images. Developing a process for finding and rigorously evaluating planets in sequences of images in combination with indirect methods would allow them to be broadly applied. Studying how well they work to detect new planets would allow future surveys to be designed around multi-epoch and joint detection, ultimately improving survey yield.

Such strategies will be essential for realizing the potential of the Nancy Grace Roman Coronagraph Instrument and the proposed Habitable Worlds Observatory, and ultimately the detection of Earth like planets and other solar system analogues.

1.7 Goals of this Dissertation

The goals of this dissertation are to present and demonstrate complimentary ways of increasing our sensitivity to solar system analogue planets around Sun-like stars. The focus is on direct imaging and ways in which direct imaging can be combined with other kinds of data. Concrete and practical techniques are presented and applied, with a final chapter describing more speculative ideas. On the other hand, details of what chemical and dynamical features we should look for on these planets in order to understand their formation, composition, and demographics, lie outside the scope of this work. Since we are currently so starved for directly detected planets, it is presumed that for now, increasing our sensitivity and size of this sample is a worthwhile goal in and of itself.

To reach these goals, we begin with Chapter 2. This chapter concerns the methods by which one combines the raw images taken of stellar systems into order to create a composite that removes, as much as possible, the glare from the star. I present an algorithm that improves on current techniques and show that it can provide up to a $5\times$ improvement in contrast close to the star. This will allow for the detection of fainter, lower mass, and closer in companions using the same datasets and observing techniques widely used today.

Chapter 3 then presents approaches for improving our sensitivity by combining direct imaging with velocimetry, and removing the sensitivity limit caused by orbital motion. I demonstrate this approach on real and simulated data from different instruments and methods, including HD 91312, simulated JWST/NIRISS aperture masking interferometry observations, radial velocity curves, and grids of images from the Gemini Planet Imager. I show that Octofitter can reliably recover faint planets in long sequences of images with arbitrary orbital motion. This tool will enable broad application of multi-epoch and multi-method exoplanet detection, which could improve how future targeted ground- and space-based surveys are performed.

After developing these methods, Chapter 4 then presents a large observing campaign of the HR 8799 planetary system and analyzes it with these improved techniques. This system is of particular interest since it hosts four massive planets orbiting 15 and 80 AU and because studies have opened the possibility of additional planets, both interior to and exterior to the known system. I present a deep L-band imaging campaign using NIRC2 at Keck comprising 14 observing sequences and re-reduce archival data for a total of 16.75 hours, one of the largest uniform datasets of a single direct imaging target. Using the techniques of the previous chapters for detecting planets in images while compensating for plausible orbital motion, I then present deep limits on the existence of additional planets in the HR 8799 system. The final combination shows a tentative candidate, consistent with $4 - 7 M_{\text{jup}}$ at $4 - 5$ AU, detected with an equivalent false alarm probability better than 3σ .

After this point, the dissertation focuses on instrumentation techniques and concepts that could improve our sensitivity in the future. Chapter 5 begins by describing efforts at the National Research Council's NEW-EARTH lab to develop and demonstrate the Fast Atmospheric SCC (Self Coherent Camera) Technique. This approach to wavefront sensing at the science focal plane aims to dynamically suppress speckles as soon as they are detected by the science camera. I develop the software infrastructure and control loops needed to run the SCC and present the results of our testing. I further present an

improved and simplified algorithm for performing active wavefront control with the SCC. This work is an important step on the way to the upcoming Subaru Pathfinder Instrument for Detection of Exoplanets and Removal of Speckles (SPIDERS) and the Gemini Planet Imager's CAL2 upgrade.

After presenting these results from the SCC, Chapter 6 then proposes that an Imaging Fourier Transform Spectrograph (IFTS) that would work well in combination with an SCC. This would allow for post-processing based speckle suppression across wavelengths at high spectral resolution. I present numerical simulations and present results from a laboratory demonstration. The chapter finishes by demonstrating a net speckle suppression of approximately 40× across a wide bandwidth IZYJH band dark hole.

Finally, Chapter 7 presents a speculative instrument concept for a constellation of orbital retroreflector beacons. The techniques described up until this point all focus on the removal of quasi-static speckles which will allow for the detection of young sub-Jupiter mass planets, but to eventually image analogues to the solar system's terrestrial planets from the ground we will need to also improve our suppression of rapidly changing atmospheric speckles. For this endeavour, natural and laser guide stars will simply be too dim. For these proposed orbital beacons, I present eccentric orbits that would place guide beacons so that they appear near a target from for one to three hours every 3-10 nights. I discuss these orbital parameters, the brightness of the beacons, and other considerations for high contrast imaging.

Chapter 2

Improved Contrast in Images of Exoplanets using Direct SNR Optimization

This chapter is based on material published in the *Astronomical Journal*, Volume 161, Issue 5, pp. 263 (2021). Authors: William Thompson and Christian Marois. *My contribution to this work was as the lead author. I wrote the full text. Christian Marois contributed to discussions and provided comments on the text. The data were acquired as part of the HR8799 observing campaign described in Chapter 4.*

2.1 Introduction

Direct imaging of exoplanets offers astronomers a wealth of information – from initial detection and astrometry, to detailed orbital and spectroscopic characterization. Due to the challenges of detecting faint companions close to their stars, direct imaging is currently limited in most cases to large, self luminous planets on wide orbits. A major goal in the field of direct imaging is thus to improve the planet-to-star contrast ratio to image fainter planets, closer to their stars. Evidence from radial-velocity and direct imaging surveys suggest that the occurrence rate of giant planets peaks near 2 – 3 AU and increases with lower mass (Fernandes et al., 2019; Nielsen et al., 2020). Improving contrast close to stars is therefore a clear path towards more direct detections.

Despite significant progress over the past decade, the dominant source of noise in direct imaging today remains quasi-static speckle noise. This noise represents

aberrations in the Point Spread Function (PSF) of the atmosphere, telescope, and instrument that are not corrected by the adaptive optics system, and have lifetimes on the order of seconds to tens of minutes. These quasi-static speckles are sufficiently correlated between images that simply averaging data over longer and longer integrations does not significantly improve the final contrast (Walker et al., 1999; Marois et al., 2003). Since speckle fields can be filtered until only the component that appears planet-like remains, these residual, long living speckles appear identical to planets and can be much brighter.

For quasi-static speckles, most observing strategies depend on differential imaging. These are a range of techniques, most notably angular differential imaging (ADI; Marois et al., 2006a) and spectral differential imaging (SDI; Walker et al., 1999; Racine et al., 1999; Marois et al., 2000) that produce a geometric offset that varies from image to image between the stellar PSF and any planets. These images taken close in time and / or wavelength can then serve to build a model of the stellar PSF that ideally is representative of the speckle noise at each moment.

Most current algorithms based on LOCI (Lafrenière et al., 2007b) reduce differential imaging sequences by considering one image at a time—the “target image”—and select a set of “reference images” from the other images of the sequence. Then, they find the linear combination of reference images that, when subtracted from the target image, create a “processed image” with minimal residual noise. The images are usually divided into small annular “subtraction regions” so that the model is specialized to the local noise distribution in the target image. Finally, they transform all of the processed images to align the signals of any planets and then stack them to produce the final output.

There are three challenges with this approach that must be given careful consideration: over-subtraction, overfitting, and self-subtraction. First, over-subtraction occurs when an algorithm fits the signal of the planet in the target image using some combination of speckle noise in the reference images and reduces planet throughput. This can be mitigated by running the algorithm on a separate “optimization region” paired with each subtraction region, and separating the two by a small buffer to prevent the linear least-squares solution from incorporating any planets into its model. When the subtraction and optimization regions are sufficiently close together, the speckle noise is typically correlated and the model created using the optimization region is effective at removing speckles in the subtraction region.

Second, when the subtraction and optimization regions are separated, overfitting

occurs when the model becomes too specialized to the optimization region and fails to generalize to the subtraction region, producing a model that does not remove speckle noise effectively. There are several ways to control overfitting, but the most common strategies all serve to reduce the dimensionality of the model. This can be done by only selecting the reference images that are the most correlated with the target image (LOCI; Lafrenière et al., 2007b; Marois et al., 2010a) or by linearly transforming the images into an orthogonal basis and only keeping the eigenimages with the highest variance (KLIP/PCA; Soummer et al., 2012). The size of the optimization regions and the best number of reference images/eigenimages to include are usually correlated, with larger optimization regions being less specialized to the noise near the subtraction region, but allowing the use of more reference images before overfitting occurs.

Finally, and perhaps most importantly, planet self-subtraction occurs when the signal of a planet is suppressed by reference images contaminated with that same planet. Linear methods have a strong tendency towards selecting the most contaminated reference images in order to minimize the noise, since these are the images that are the most correlated with the target image. Traditionally, these reference images are rejected by applying a threshold to the local planet displacement between images (e.g. due to varying parallactic over time), the expected amount of flux contamination via forward modelling (T-LOCI; Marois et al., 2010a), or other similar criteria. This limits self-subtraction at the expense of increased noise, so this parameter (sometimes called the aggressiveness) and other hyper-parameters are varied to find the combination that maximizes the SNR of injected, forward modelled planets (Lafrenière et al., 2007b; Marois et al., 2010a; Soummer et al., 2012; Meshkat et al., 2014, and others). Some amount of self-subtraction is expected from these methods, so planet throughput must be calibrated by injecting forward modelled planets into the raw images.

Rejecting some of the most correlated references, followed by minimizing the noise, and then correcting the planet throughput might be sub-optimal. One improvement to this was considered in Pueyo et al. (2012) in the Damped-LOCI algorithm. In addition to limiting spectral cross-talk when extracting spectra from SDI sequences, Damped-LOCI aims to improve planet SNR by simultaneously minimizing noise in the optimization region and maximizing the variance in the subtraction region, with an additional tunable hyper-parameter to balance the trade off. One possible issue with this approach is that by attempting to boost the variance in the subtraction region, speckle noise under the planet might be amplified.

To directly achieve an optimal planet signal to noise ratio, we present a new

algorithm in this paper that re-frames the problem from minimizing the noise to directly maximizing the signal to noise ratio (SNR) of any point sources using the same forward modelling facilities implemented in most pipelines for throughput correction. This relatively straightforward approach allows us to use any reference image regardless of its proximity to or flux contamination with the target image. This removes the need to reject the most correlated reference images, and removes the need to correct throughput loss due to self-subtraction. Finally, this formulation can be generalized to treat multiple target images simultaneously so that planet SNR in the overall stack of processed images is optimized, further suppressing correlated noise. This approach is a general algorithm for maximizing the signal to noise ratio when combining measurements with correlated noise, and a known variation in signal intensity.

We begin by describing the new SNR optimization algorithm. We show how it relates to linear methods like LOCI in a limiting case, and why it should be expected to meet or exceed the performance of linear methods even when hyper-parameters of those methods are tuned to maximize SNR. Next, we show how the algorithm can be generalized to optimize a grid of multiple target and reference images simultaneously and describe how this can further reduce correlated noise in some sequences. Finally, we analyze the performance of the SNR optimization algorithm on a sample ADI sequence.

2.2 SNR Optimization

To achieve an optimal planet signal to noise ratio, we propose directly optimizing an expression for SNR in each region of the sequence, instead of only minimizing the noise. For a given location in the sequence, if we know the relative photometry of a planet in each input image by forward modeling, we can write a formula for the planet SNR for any linear combination of images.

First, we divide the images into pairs of optimization regions, R_O , and subtraction regions, R_S , just as in previous methods (Figure 2.1). For each pair of regions, we consider a forward modelled planet centred in R_S and noise evaluated in R_O . We use the noise in the optimization region to find a solution that we then apply to the subtraction region. This makes the assumption that the noise in the two regions are correlated, but is necessary to avoid finding solutions that fit the signals of any planets using noise from the other images (called planet over-subtraction). In previous works,

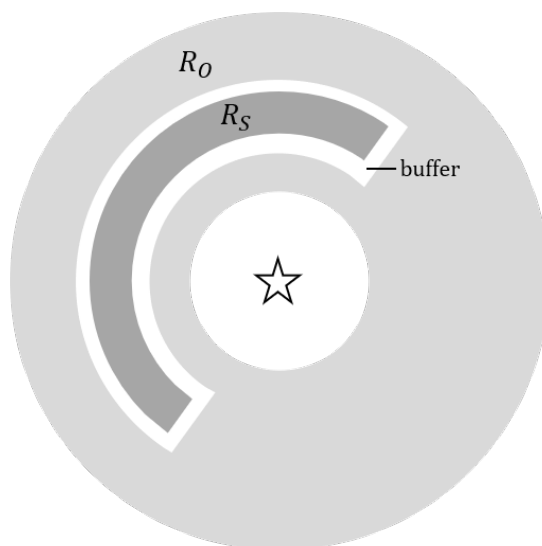


Figure 2.1 Examples of the region geometry used this chapter. The annulus sector R_S is the subtraction region and R_O is the optimization region. The optimization region is chosen to be slightly thinner on the side closer to the star since the brighter speckles in that area could bias the model towards an overly aggressive solution. The white area between the two regions is a buffer to prevent any planets from leaking into the optimization region and causing over-subtraction. The geometry of the regions can be chosen differently as long as a sufficient buffer is preserved between them to prevent over-subtraction.

authors have considered regions that partially overlap (beginning with [Lafrenière et al., 2007b](#)); however, the methods presented here are aggressive enough that a buffer larger than approximately $0.4\lambda/D$ should be kept between the two regions to prevent over-subtraction. Though we use the noise in the optimization region, we are free to model photometry for a planet located at the centre of the subtraction region in the target image.

The signal of a linear combination of images is then simply the linear combination of the forward modelled relative photometries in each image:

$$\text{signal} = \mathbf{p} \cdot \mathbf{c}, \quad (2.1)$$

where \mathbf{p} is a vector of forward modelled relative planet photometry, and \mathbf{c} is a vector of coefficients for the linear combination of images. The forward model should be calculated using either satellite spots ([Marois et al., 2006b](#); [Sivaramakrishnan, Oppenheimer, 2006](#)) on instruments where they are present, or an unocculted and unsaturated PSF taken before or after the sequence. This forward modelling is already implemented in

most direct imaging reduction pipelines for the purpose of throughput correction (e.g. [Marois et al., 2010a](#); [Ruffio et al., 2017](#)).

By assuming that the noise is centred around zero, we can write the noise as the root-mean-square (RMS) of the pixel values in the same linear combination of images

$$\text{RMS} = \sqrt{\sum_j^M \left(\sum_i^N O_{ij} c_i \right)^2 / M} \quad (2.2)$$

where i ranges from 1 to N , the number of images; j ranges from 1 to M , the number of pixels; and O is a matrix of N images by M pixels. The RMS is chosen for simplicity and should be equivalent to the standard deviation if a radial profile is pre-subtracted, or the images are high-pass filtered. Note that this does not assume the underlying noise distribution is Gaussian, since we are not using the RMS to calculate confidence intervals. Minimizing the RMS of any reasonable distribution should minimize the noise in the images. It is possible that this metric could be tweaked to, for example, penalize spatially correlated noise, but this is outside the scope of this chapter.

By combining these expressions, we arrive at an equation for the signal to noise ratio of the modelled planet for any given linear combination of images:

$$\text{SNR} = \frac{\text{signal}}{\text{RMS}} = \frac{\sum_i^N p_i c_i}{\sqrt{\sum_j^M \left(\sum_i^N O_{ij} c_i \right)^2 / M}} \quad (2.3)$$

By noting that maximizing the SNR is equivalent to maximizing an expression proportional to the SNR^2 , we can see that this optimization problem is a quadratically constrained quadratic problem:

$$\text{argmax}_{\mathbf{c}} \text{SNR} = \text{argmax}_{\mathbf{c}} \text{SNR}^2 \quad (2.4)$$

$$= \text{argmax}_{\mathbf{c}} \frac{\left(\sum_i^N p_i c_i \right)^2}{\sum_j^M \left(\sum_i^N O_{ij} c_i \right)^2} \quad (2.5)$$

The values of the coefficient vector \mathbf{c} can then be chosen such that the planet SNR is maximized. Applying that vector of coefficients to the subtraction region produces the linear combination of images that gives close to the optimal planet SNR for that region, limited only by the degree of correlation between the chosen optimization and subtraction regions.

Since the SNR ratio is preserved if the \mathbf{c} vector is scaled by a constant, the problem is under constrained. We therefore apply the additional constraint that the planet-throughput is kept at unity, or

$$\mathbf{p} \cdot \mathbf{c} = 1 \quad (2.6)$$

Then, if the forward modelling is done with a correctly scaled image of the same star, the resulting processed image is then naturally in units of relative contrast. This assumes that the planet forward model accurately reflects changes in, for example, atmospheric variations, or that these effects are small.

2.2.1 Formulation as a Constrained Least-Squares Problem

In order for this formulation to be useful, it must be possible to efficiently optimize the coefficients to find a global optimum. With the application of the constraint in Equation (2.6), that is $\mathbf{c} \cdot \mathbf{p} = 1$, we can express this optimization problem as a constrained system of linear equations as follows:

$$\operatorname{argmax}_{\mathbf{c}} \text{SNR} = \operatorname{argmax}_{\mathbf{c}} \frac{1}{\sqrt{\sum_j^M (\sum_i^N O_{ij} c_i)^2}} \quad (2.7)$$

$$= \operatorname{argmin}_{\mathbf{c}} \sqrt{\sum_j^M (\sum_i^N O_{ij} c_i)^2} \quad (2.8)$$

$$= \operatorname{argmin}_{\mathbf{c}} |\mathbf{O}\mathbf{c}| \quad (2.9)$$

$$\text{subject to } \mathbf{c} \cdot \mathbf{p} = 1 \quad (2.10)$$

This finds the linear combination of reference images that minimizes the noise in the optimization zone, subject to the constraint that the throughput is preserved at unity. Without the constraint, the solution would have all $\mathbf{c} = 0$ and the output image would be entirely empty. With this formulation, the photometry vector \mathbf{p} decides which is the “target” image. The target image has its own coefficient, just like a reference. Since we chose to model the planet photometry with a planet centred in the subtraction zone of the target image, the index of the target image should correspond to the maximum value of \mathbf{p} , and likely (though not necessarily) the highest value in \mathbf{c} as well.

This formulation is a valid expression in the methodology of Disciplined Convex Programming (DCP; [Grant et al., 2006](#)). Many programming languages have libraries

that can solve such expressions efficiently to a global optimum.

In some cases, such as when using Reference star Differential Imaging (RDI) to “star-hop” between two stars, images with significant flux contamination and images with no contamination are intermixed. In this case, the constraints can be treated as sparse since only images with significant planet flux contribution must enter into the constraint term.

2.2.2 Reduction to Linear Least Squares in Limit of no Self-Subtraction

To show how this method relates to previous algorithms, we show that SNR optimization reduces to a linear problem if we neglect planet contamination in the reference images. This situation would occur if a target image was reduced using only reference images taken from another star (RDI) or if all neighbouring reference images were rejected (e.g. in LOCI based algorithms). Taking the target image to be $i = 1$, then $p_1 = 1$, all other $p_i = 0$, and $c_1 = 1$ due to the throughput constraint. Equation (2.5) then reduces to:

$$\operatorname{argmax}_{\mathbf{c}} \frac{1}{\sum_{j=1}^M (\sum_{i=1}^N O_{ij} c_i)^2} \quad (2.11)$$

$$= \operatorname{argmin}_{\mathbf{c}} \sum_{j=1}^M \left(O_{1j} + \sum_{i=2}^N O_{ij} c_i \right)^2 \quad (2.12)$$

Continuing now as in [Lafrenière et al. \(2007b\)](#), the minimum occurs when all of the partial derivatives with respect to \mathbf{c} are zero:

$$\frac{\delta}{\delta c_k} \sum_{j=1}^M \left(O_{1j} + \sum_{i=2}^N O_{ij} c_i \right)^2 = 0, \forall k \quad (2.13)$$

$$(2.14)$$

$$\sum_{j=1}^M O_{kj} O_{1j} = \sum_{i=2}^N -c_i \left(\sum_{j=1}^M O_{ij} O_{kj} \right), \forall k \quad (2.15)$$

Which is the familiar system of linear equations of the form $\mathbf{Ax} = \mathbf{b}$ where

$$A_{ik} = \sum_{j=1}^M O_{ij} O_{kj}, \quad x_k = -c_k, \quad \text{and} \quad b_k = \sum_{j=1}^M O_{kj} O_{1j}. \quad (2.16)$$

This is exactly equation 7 of [Lafrenière et al. \(2007b\)](#) apart from notational differences and the sign convention on c . We can therefore consider SNR optimization to be a generalization of LOCI for reference images that contain overlapping flux from the planet. Since this occurs in almost every ADI or SDI reduction, we can expect the resulting SNR to always be higher when using SNR optimization for the same set of references—ignoring secondary effects such as how well the solution generalizes from the optimization region to the subtraction region. Note that SNR optimization could be used directly on the subtraction regions analogously to some LOCI pipelines; however, in this chapter we maintain a buffer to prevent over-subtraction. Example target images reduced with both LOCI and SNR optimization are presented in Section 2.5.

This also shows that we cannot further improve RDI sequences since the two formulations would reduce to the same system of linear equations. If a sequence also contains field of view rotation or multiple wavelengths in addition to the reference images, this algorithm can still offer improved SNR by better handling the images that do have planet flux contamination.

2.3 Multi-Target Image Optimization

In Section 2.2, we described the process of applying SNR optimization to a single target image with a set of reference images. The results of each image must then be rotated and/or scaled to re-align the signals of any planets and then stacked, just as in previous approaches. A limitation of this single-target optimization is the assumption that solutions giving the highest SNR for an individual target image also lead to the highest SNR in the stacked image.

This is a reasonable assumption to first order, but is often violated at close separations. For example, there are often bright, randomly fluctuating speckles in addition to a floor of dimmer, more consistent speckles. When considering each target individually, the SNR is maximized by suppressing the bright, fluctuating speckles, perhaps at the expense of the dimmer ones; however, in the final stack, it may be the dimmer persistent speckles that are correlated between images that contribute the most noise.

It is therefore worth considering ways to avoid this assumption. Since the SNR equations above do not distinguish between target images and reference images except by how the forward modelled photometry is calculated, it is possible to combine multiple target images into a single optimization problem. The output of the optimization

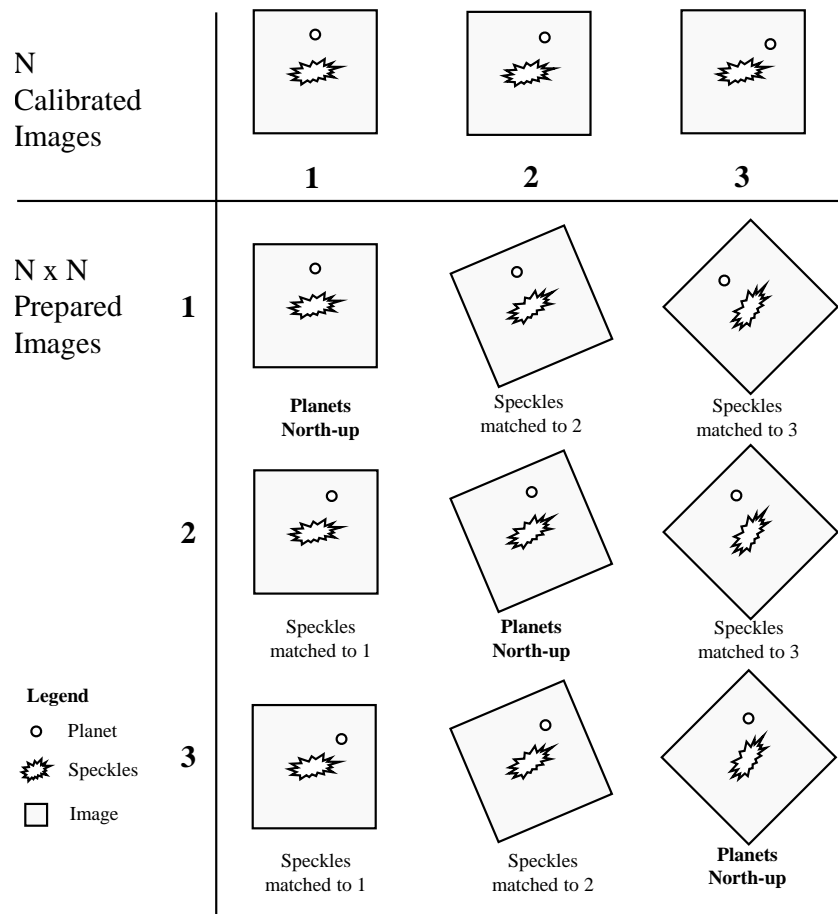


Figure 2.2 Preparing a grid of N images for multi-target SNR optimization. In this example of a simplified 3-image ADI sequence, a distorted star shape represents the star’s PSF and a small circle represents a planet directly to the North of the star. We begin by transforming each target such that any companion would appear in their physical, e.g. North-up, locations (the diagonal elements). Next, for each transformed target, we apply the same transformation to each corresponding reference (the off-diagonals). This aligns the speckles in the references with the speckles in that transformed target. We then follow the usual approach of dividing the field into annuli or annulus sectors for local optimization. An analogous matrix is prepared for the relative photometry using forward modelling for each region. In previous approaches, one would effectively reduce each column of this grid independently and then stack the results. By preparing this grid of transformed images as a single problem, the optimizer becomes aware of noise that is correlated between target images, and we achieve a higher point source SNR in the final stacked image.

algorithm then becomes a single, larger set of coefficients that describes both how each target image is combined with its matching references, and how those resulting images are combined into the final processed image.

There are three steps to preparing this problem. First, transform the target images to align any planet signals (usually by rotating them North-up). Then, transform copies of each reference image so that its speckles are aligned with each transformed target image. This is best visualized as a grid of reference images with the target images placed along the diagonal (see Figure 2.2). Finally, generate subtraction and optimization regions that are in physical coordinates (e.g. North-up), rather than in the coordinates of the detector.

At this point, the algorithm can continue as usual, with a few adjustments to how the forward modelled photometry is calculated to account for those coordinate transformations. If all the target images are combined into a single problem, then there is no need to stack the results using a median or other statistic. Applying the coefficients to the subtraction region produces a single output image.

At first glance, this appears to increase the size of the model from N images or eigenimages to N^2 images; however, the total number of coefficients to calculate is the same. Looking again at Figure 2.2, standard linear methods or SNR optimization would calculate a coefficient for each image, on N -image column at a time, and then stack the resulting processed images. If an arithmetic mean is used to stack the residuals, it's clear that the output consists of a linear combination of $N \times N$ images in either case.

By directly optimizing the SNR of that combination, we no longer assume that each target image is independent, and that the brightest noise, regardless of its temporal evolution, contributes the most to the final result. Secondly, we also use information in the optimization region to guide how much weight should be given to each column of the grid, rather than giving each target image an equal weight. A comparison of the performance of single-target versus multi-target SNR optimization is presented in Section 2.5.

2.3.1 Batching

Multi-target SNR optimization is useful for improving noise that is correlated between target images, and that is correlated between the subtraction and optimization region, but that does not dominate until multiple reduced images are stacked. On the other hand, a traditional stacking technique like a median is more effective at rejecting large

outliers that are not spatially and temporally correlated since it operates inside the subtraction region.

These two techniques can be balanced by dividing the grid into groups of columns called batches. Optimization can proceed one batch at a time to produce a smaller set of output images. The residuals can then be stacked analogously to per-target optimization, though without needing to first rotate or otherwise transform the results. By grouping correlated targets together and then stacking the results, we retain most of the correlated noise rejection and the outlier rejection of a median. We can then see that multi-target SNR optimization is a spectrum, ranging from considering only a single target at a time to an entire sequence. For the same contrast, smaller batches have an additional benefit of reducing the computational cost of the reduction.

In principle, a correlation analysis could be used to group target images into batches, but a simple method is to group targets according to the time they were captured since images captured close in time are likely to be the most correlated. The ideal size of a batch depends on both the sequence and instrument. It should be sufficiently large that target images outside the batch are no longer highly correlated, but of sufficient number that outliers can be rejected with a median. This parameter can be optimized for each sequence or estimated by first reducing the sequence using single-target optimization and estimating the residual speckle lifetime.

Finally, the target batches do not need to be disjoint. It is reasonable to have, for example, a batch size of fifteen targets that advances five target images per batch. This slightly improves robustness against outliers at the expense of increased computation time. An example of batched of multi-target SNR optimization is presented in [Section 2.5](#).

2.3.2 Controlling Overfitting

In the context of differential imaging, overfitting occurs when a model is given too much flexibility and insufficient data. For example, this can occur if hundreds of reference images are used with a small optimization region to reduce one target image. The large set of references may be combined in such a way that random read noise, photon noise, or a transient speckle in the optimization region is minimized, but the solution does not effectively generalize to the noise in the subtraction region. This limits the performance of an algorithm but should not bias the signal of any planets as long as sufficient buffers are maintained.

Overfitting can be limited in several ways. First, one can restrict the set of reference images to only include those most correlated to the target (Marois et al., 2010a). Second, one can perform an SVD (Marois et al., 2010a) or PCA analysis (Soummer et al., 2012), and similarly restrict the set of eigenimages. Third, one can increase the size of the optimization regions until there is sufficient data to train the model (though if the optimization regions are too large, the model may not be sufficiently specialized to the local noise distribution). Another option is to introduce a regularization term to the model to penalize overly complex combinations of coefficients. Lastly, one average multiple reductions with variations in their parameters to reduce the impact of overfitting (sometimes referred to as “bagging”).

The challenge of overfitting is exacerbated with multi-target SNR optimization. When many target images are combined into a single problem, the size of the reference set must grow quadratically. Unfortunately, SNR optimization is a non-linear process that does not admit a simple PCA solution. Such a solution would need to decompose the system of quadratic equations into non-linear “SNR modes.” Instead, we approach the problem of overfitting by combining four of the above techniques. First, we eliminate poorly correlated reference images from the grid, reducing the model size. Next, we add an L2 regularization parameter into the objective. Finally, we divide the sequence of target images into overlapping batches, and average their results. Each of these hyper-parameters can be optimized in turn to arrive at a balance that is ideal for any given sequence. This tuning process can be relatively efficient when using an optimizer that supports warm-starting from a previous solution.

When tuning hyper-parameters, it is best to use a stand-in such as a separate validation region, or set of prepared images with rotation reversed. By tuning the algorithm on this different but statistically similar data, we ensure that the hyper-parameters are not biased towards combinations that, by chance, reduce the flux of any planets and diminish overall throughput.

2.4 Implementation

The full algorithm for multi-target image SNR optimization with regularization is presented succinctly in Algorithm 1. This takes both the number of most correlated references to include and the regularization parameter α to limit overfitting. It also takes the number of images in a batch to control the trade off between outlier rejection

Algorithm 1: Multi-Target Image SNR Optimization for ADI.

```

procedure SNROPT(images, regions,  $N_{\text{batch}}$ ,  $N_{\text{refs}}$ ,  $\alpha$ )
  for  $i \in 1 : N_{\text{images}}$  do                                 $\triangleright$  Prepare images
     $\theta_{\text{targ}} \leftarrow \text{PA}(\text{images}[i])$  for  $j \in 1 : N_{\text{images}}$  do
       $\text{prepared}[i, j] \leftarrow \text{rotate}(\text{images}[j], \theta_{\text{targ}})$ 
    end
  end
  for  $i \in 1 : N_{\text{batch}} : N_{\text{images}}$  do                     $\triangleright$  Iterate through batches
    for  $k \in 1 : N_{\text{batch}}$  do                                 $\triangleright$  For each target in batch, select references
       $\text{selected}[k] \leftarrow N_{\text{refs}}$  most correlated references in column k of batch
    end
    foreach  $(R_S, R_O) \in \text{regions}$  do                     $\triangleright$  Iterate through each pair of regions
       $\mathbf{p} \leftarrow \text{model photometry of selected in } R_S \text{ for planet centred in targets}$ 
       $\mathbf{O} \leftarrow R_O$  regions of selected images
       $\mathbf{c} \leftarrow \text{argmin} |\mathbf{O}\mathbf{c} + \alpha \sum_i^N c_i^2|$  subject to  $\mathbf{c} \cdot \mathbf{p} = 1$   $\triangleright$ 
      Optimize SNR in optimization re-
      gion
       $\mathbf{S} \leftarrow R_S$  regions of selected images
       $\text{out}_i[R_S] \leftarrow \mathbf{S}\mathbf{c}$                              $\triangleright$  Apply coefficients to subtraction
                                                                region
    end
  end
  Measure contrast curves of for each  $\text{out}_i$  and take the weighted median
end procedure

```

The hyper-parameters N_{batch} , N_{refs} , and α should be optimized against either a separate validation region R_v or by reversing the angle θ_{targ} . This ensures values are not selected that, by chance, hurt the throughput of any planets in R_S .

($N_{\text{batch}} = 1$) and reducing correlated noise ($N_{\text{batch}} = N_{\text{images}}$). These hyper-parameters should be tuned for each sequence. The image preparation steps on lines 4 and 5 can be generalized for any type of differential imaging. We do not include hyper-parameter optimization steps for the geometry of the subtraction and optimization regions, as these are physically motivated, relatively insensitive, and already discussed in previous works (Lafrenière et al., 2007b; Currie et al., 2012)

To apply the algorithm, we built a new () Implemented in Julia (Bezanson et al., 2012), it can process sequences with RDI, ADI, SDI from Keck NIRC2 (PI: K. Matthews), VLT-SPHERE (Beuzit et al., 2019), GPI (Macintosh et al., 2014), and LBT LMIRCam (Skrutskie

et al., 2010). Single-image and multi-target image SNR optimization is supported along with LOCI and T-LOCI-style algorithms for the sake of comparison. The LOCI style algorithms include optimization steps for the number of reference images and rejection threshold, making it roughly comparable to other “optimized” implementations like Meshkat et al. (2014).

To abstract RDI, ADI, and SDI as a single process, we used the library `CoordinateTransforms.jl` to store arbitrary transformations between speckle-aligned coordinates and planet-aligned coordinates. For multi-target image SNR optimization, we combined these transformation objects with `WarpedViews` from the `ImageTransformations.jl` package. This allows us to avoid storing duplicates of each reference image transformed to match each target. When a given pixel of a reference image is required, the transformation and interpolation are performed on the fly. This reduces the number of prepared reference images stored in memory from N^2 to simply N , and removes much of the memory overhead for multi-target SNR optimization.

Forward modelling is also necessary for the SNR optimization algorithm. To this end, the pipeline uses the transformation objects to compute a model for a planet in each subtraction region for each second of the sequence to account for the rotational smearing that takes place during ADI integrations. These abstractions allow the pipeline to straightforwardly optimize the stacked SNR of a sequence containing multiple wavelengths and rotation angles without knowing the details of the transformations applied to the images.

For the SNR optimization itself, we implemented two modules. The first uses either Newton’s method or BFGS (Fletcher, 1987; Mogensen, Riseth, 2018) to directly optimize the SNR using Equation (2.5) and its partial derivatives. This performs very well on smaller problems up to a few hundred reference images. On larger problem sizes, `Convex.jl` (Udell et al., 2014) with the `COSMO.jl` (Garstka et al., 2019) backend is used with Equation (2.10). This outperforms the direct solution, and is particularly efficient on multi-target image SNR optimization problems which may have Hessians with very high condition numbers ($> 10^7$).

2.5 Demonstration

We evaluated the performance of the SNAP pipeline on an ADI sequence taken of HR8799 (Marois et al., 2008a) with NIRC2 at Keck on 2020 November 17 (PI: Q.




Figure 2.3 Three examples of target images reduced with both LOCI (top) and with SNR optimization (bottom). Each example was taken from a different part of the sequence with different FoV, rotation rates. In all three examples, we see that SNR optimization is very effective at suppressing speckles from the star. Images are displayed on the same color scale.

Konopacky). The sequence consists of 90 separate 40-second integrations taken at L' band while the system transited the meridian. We took the sequence in good seeing without a coronagraph, causing the central core of the star to saturate. We acquired separate unsaturated images at the beginning and end of the sequence.

We reduced the sequence three times using the same code, changing only whether the coefficients were selected by our baseline implementation of LOCI, by single-target image SNR optimization, or by multi-target image SNR optimization. This reduces the likelihood that any unrelated differences between implementations affect the results. For all three algorithms, hyper-parameters such as the number of reference images and the regularization were independently optimized to select their best values for this sequence. That said, general details such as the geometry of the optimization and subtraction regions were chosen to give the best performance using SNR optimization—it's possible that these, or other choices are not ideal values for LOCI style algorithms, and that comparable results could be achieved given sufficient effort. It is challenging to compare different algorithms across sequences and instruments in a completely

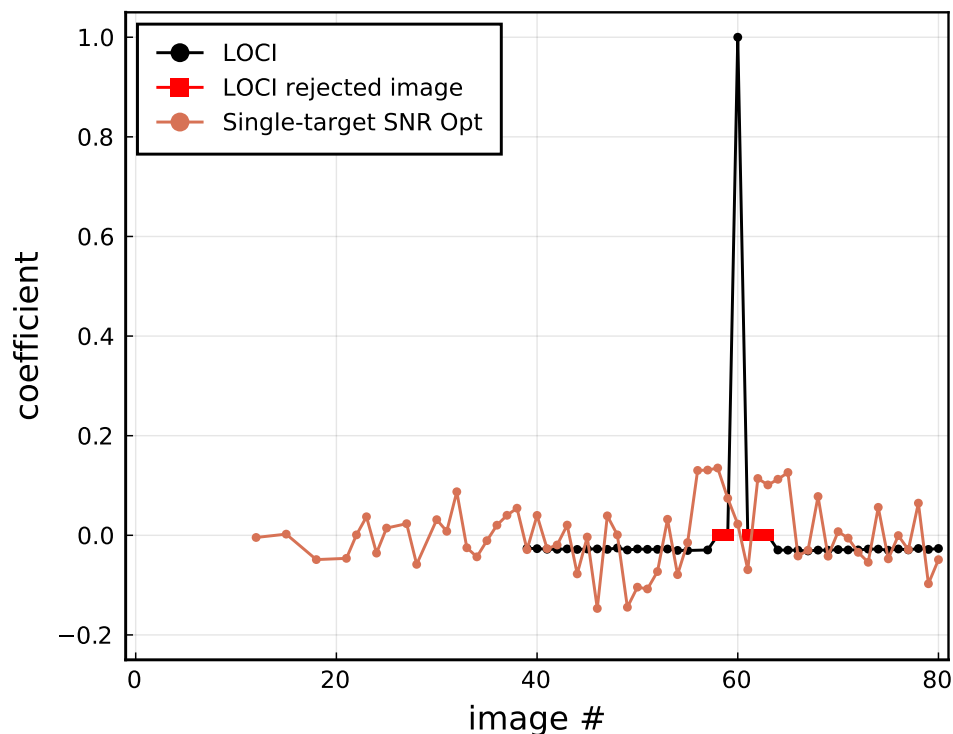


Figure 2.4 The coefficients chosen by LOCI and by SNR optimization for the images shown in the third column of Figure 2.3 and a region close to planet HR8799e. The black line shows the coefficients chosen by the LOCI algorithm for the target image # 60. The red squares indicate images that were rejected for the LOCI reduction due to insufficient displacement from the target. The orange line, on the other hand, shows the coefficients that give the highest SNR. Note that in SNR optimization, the target image is treated just like a reference image with a coefficient.

fair way, so this should be considered only as an illustrative example for how SNR optimization performs.

For this sequence we used subtraction regions that are annulus sectors $1.4\lambda/D$ thick and 180° wide, and optimization regions that surround the subtraction regions with $2\lambda/D$ on the outside, and $1.2\lambda/D$ on the inner side. A buffer of $0.4\lambda/D$ separates the two to prevent over-subtraction. See Figure 2.1 for a schematic.

To tune the hyper-parameters of SNR optimization, we increased the number of included reference images and the L2 regularization parameter until the forward modelled SNR stopped improving. For LOCI, the rejection window and the number of included reference images were similarly adjusted. We tuned the hyper-parameters using the backwards-rotated set of images to avoid unintentionally favouring solutions

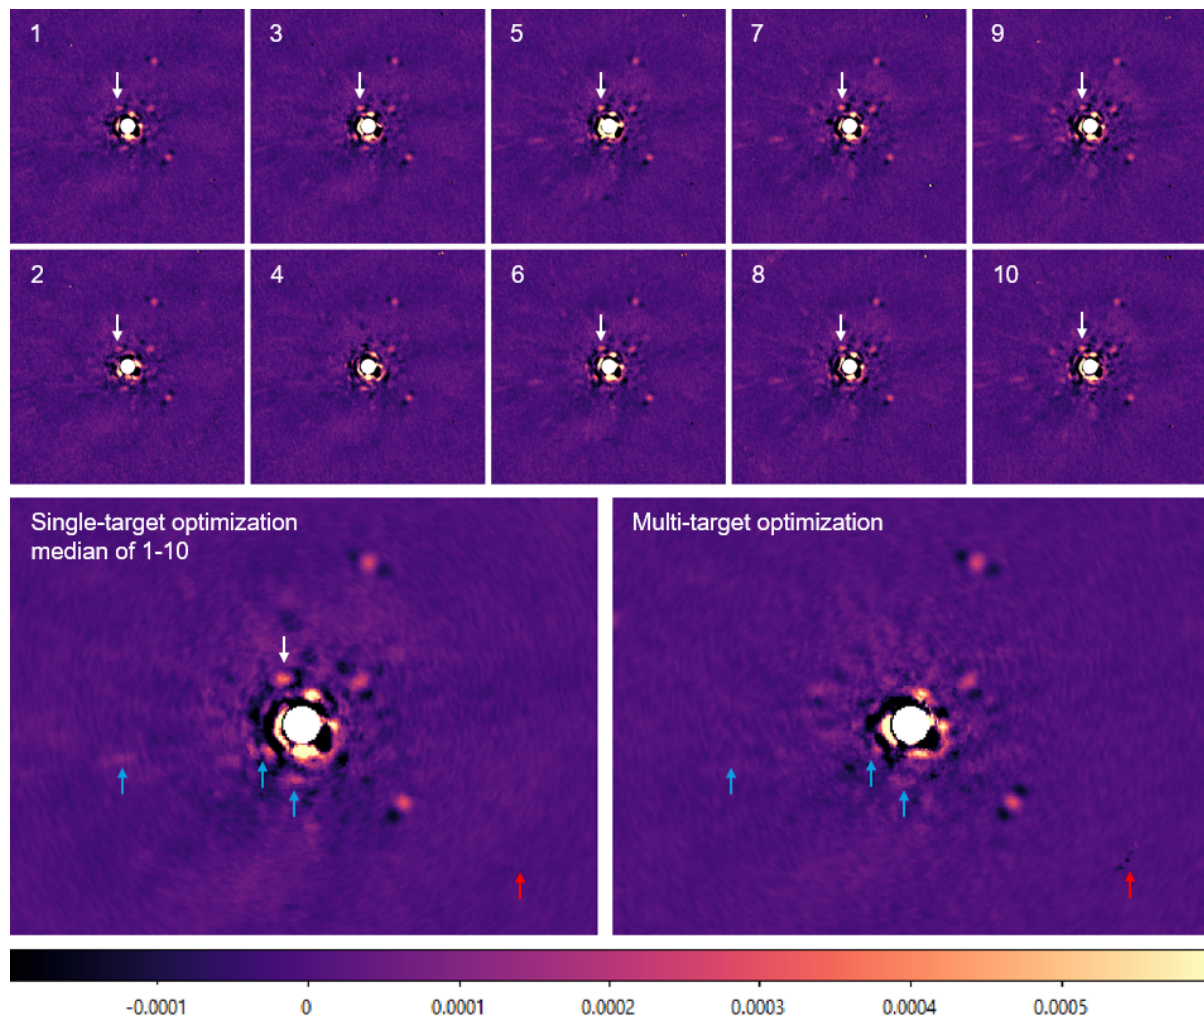


Figure 2.5 The benefits of multi-target SNR optimization on one batch of ten images. **Top:** Target images reduced individually with SNR optimization. A white arrow points to an example of noise that is correlated between nine out of the ten images. **Bottom left:** Median stack of the ten images at the top. The highly correlated residuals are not rejected by the median. **Bottom right:** The same images reduced simultaneously with multi-target image SNR optimization. By considering the ten targets simultaneously, we prevent the buildup of correlated noise between targets. The blue arrows highlight locations where noise that was correlated between the ten target images was suppressed by multi-target SNR optimization. The red arrow shows a counter example where multi-target SNR optimization fails to reject a bad pixel. Taking the median of multiple batches is an effective way to suppress these cosmetics. Images are in units of brightness relative to the star.

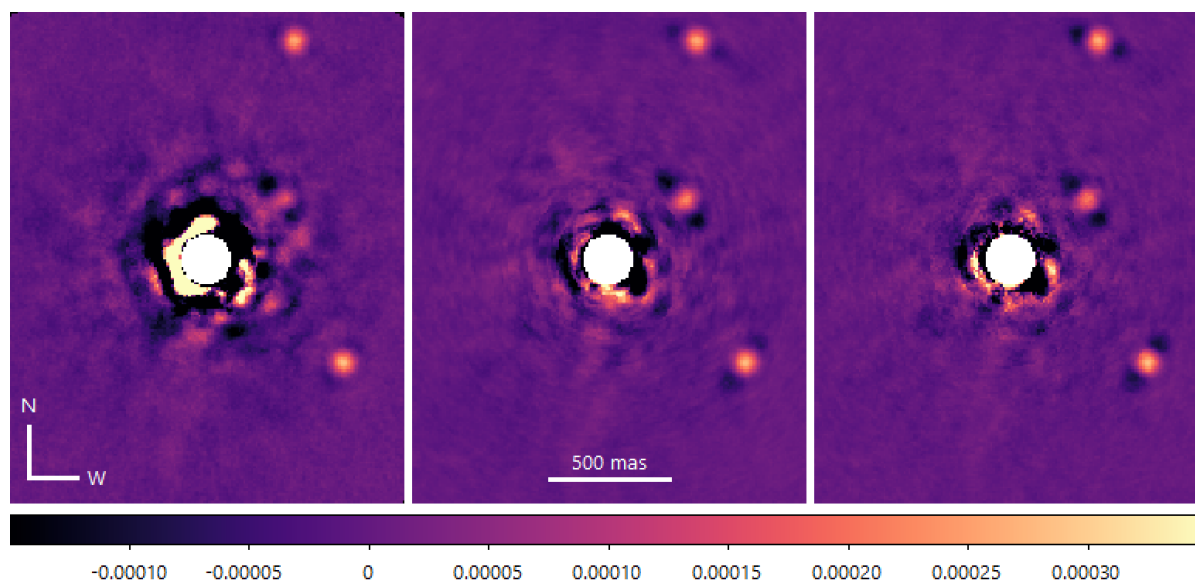


Figure 2.6 A non-coronagraphic L'-band ADI sequence taken of HR8799 at Keck, reduced with LOCI (left), single-target image SNR optimization (middle), and multi-target image SNR optimization (right). While HR8799e is near the level of the noise in the LOCI reduction, the new algorithm recovers it robustly and even recovers usable data from behind the first Airy ring of the stellar PSF. The central $1.2 \lambda/D$ region marked in white is saturated. The outputs of the three algorithms are each combined in a contrast weighted median, which tends to favor images with the highest FoV rotation rate—hence the uncharacteristically faint dark wings in the LOCI reduction. Much of the central 500 mas in the left panel is outside the plotted color scale.

that bias planet photometry and risk losing planet throughput. In both cases, the reference images were ranked against each target by their correlation to the subtraction region. Finally, the LOCI images were throughput corrected using forward modelling in order to match the results of SNR optimization. This may lead it to look qualitatively worse than similar images in the literature where throughput correction might be neglected but reflects the true ability of the algorithm to recover planets.

We begin by examining three individual target images, # 9, 31, and 60, reduced with LOCI and with single-target image SNR optimization (Figure 2.3). The new algorithm significantly outperforms the baseline reduction and reveals the innermost known planet (Marois et al., 2010b). In Figure 2.4, we compare the coefficients chosen by both algorithms for a subtraction region near planet HR8799e. For the sake of comparison, we adopt the convention that the target image has a coefficient held at 1 for the LOCI reduction. Reference images that lie within the LOCI rejection window are considered as having coefficients held at 0.

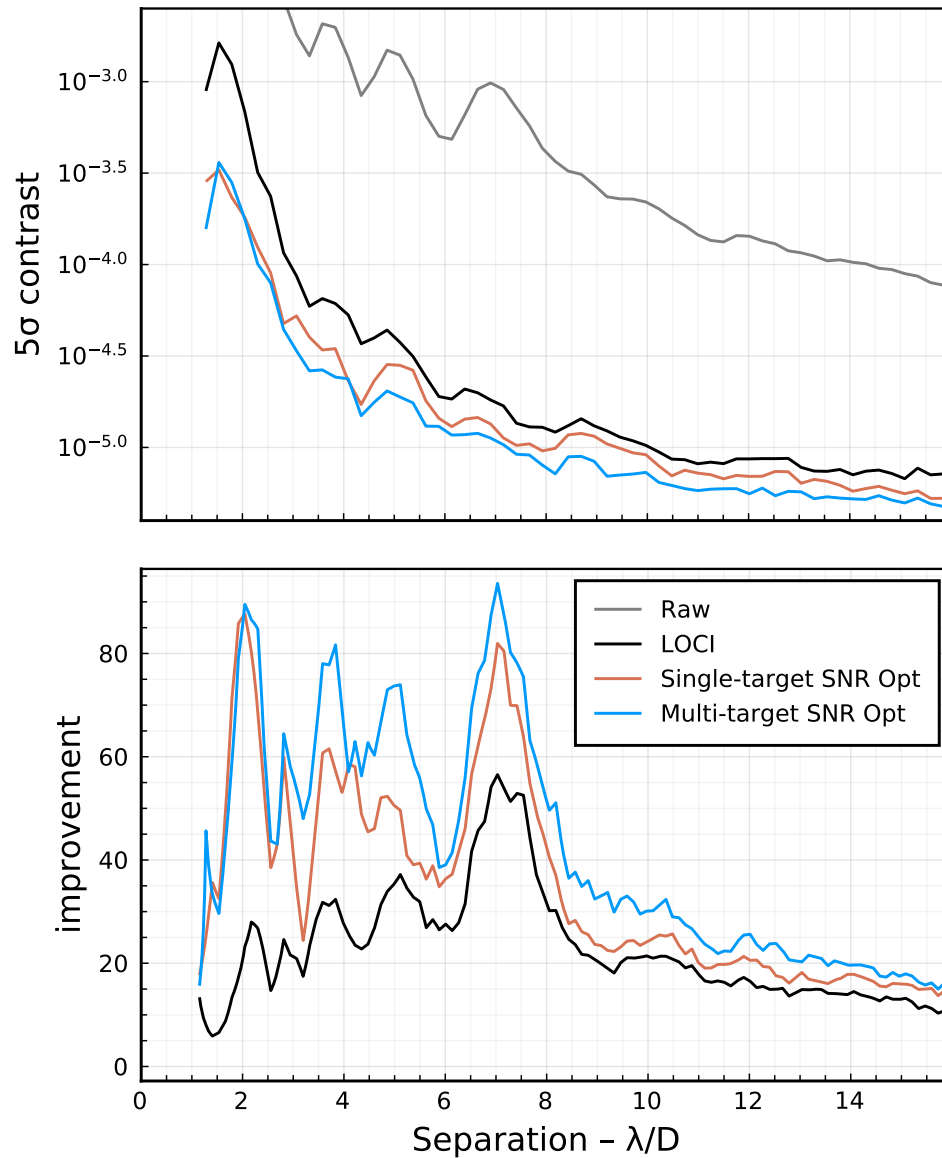


Figure 2.7 **Top:** planet-to-star contrast compared for the different reductions, and one raw image. **Bottom:** Contrast improvement compared to the raw image. Multi-target SNR optimization outperforms LOCI by three to five times at small separations from the star, and achieves a consistent 10-20% improvement at wide separations.

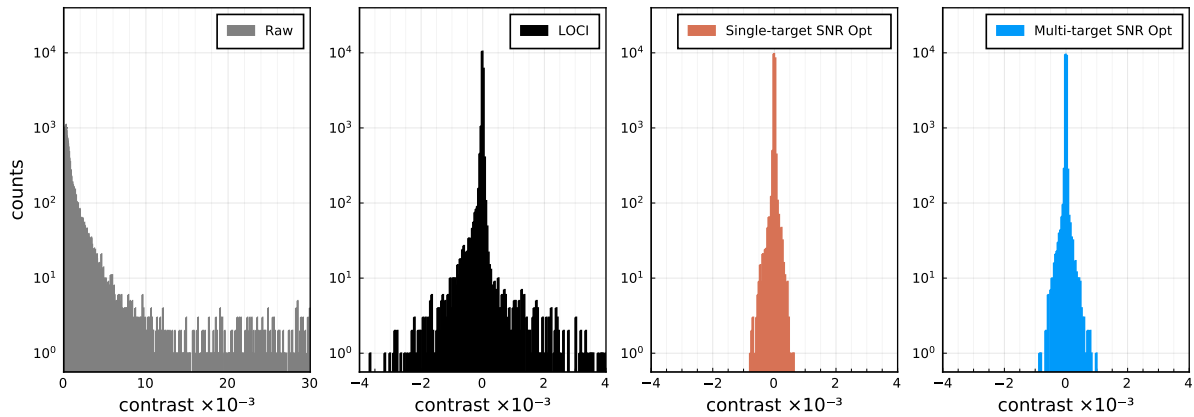


Figure 2.8 Histograms of residual noise in the inner $3 \lambda/D$ region around the star for one raw image and the reduced images presented in Figure 2.6. In addition to the lower overall noise, SNR optimization produces residuals with a more symmetric distribution on this sequence.

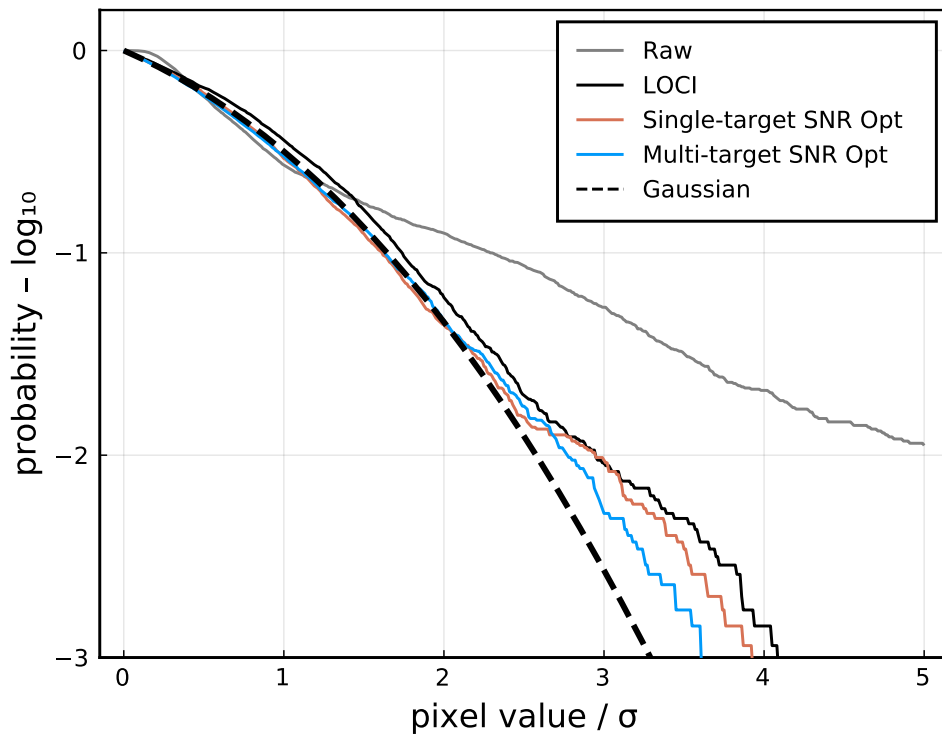


Figure 2.9 Comparison of the pixel distributions between the different algorithms and one raw image. The values shown are $1 -$ the cumulative distribution of pixel values taken from an annulus between 500 and 600 milliarseconds, normalized to one standard deviation and integrated symmetrically from zero. The algorithm presented in this chapter produces residuals with a distribution that more closely resembles a Gaussian.

The difference in the chosen coefficients for the linear combination is striking. The images that were rejected by LOCI for having too much overlap are used to a significant extent by SNR optimization. Surprisingly, one of the directly adjacent reference images was subtracted from the target, despite significant planet overlap. This may indicate that SNR optimization is compensating for a poor quality image using its neighbours. This, and other surprising deviations from the coefficients chosen by LOCI shows that there is significant room for the new algorithm to outperform existing techniques.

Next, we compare ten target images reduced with single- and multi-target SNR optimization. Figure 2.5 shows the ten images each reduced independently with SNR optimization. Many areas of the reduced images are highly correlated from image to image. The information necessary to remove these speckles must therefore be present in the reference images; however, it was not used. This could be because removing those speckles would have been a poor trade off for the SNR of a single reduced image. When stacked, these residual correlated speckles remain, limiting the SNR of the final image. When the ten target images are considered simultaneously, however, the objective function is sensitive to the final combined SNR, and therefore finds a better overall solution. This shows that reducing correlated target images in batches prevents the build up of correlated residuals in the final stack.

Finally, we examine the full reduction of the sequence with LOCI and SNR optimization. We adopt a batch size of ten target images, and advance the window five targets at a time (i.e. producing 16 output images). The processed images from LOCI and SNR optimization were stacked using a contrast-weighted median where the contrast was evaluated using a matching backwards-rotated reduction to prevent planet signals from biasing the weights. The final outputs of the two algorithms are shown in Figure 2.6. The images are cropped to the inner region of the system to highlight the improvement close to the star.

Planets b, c and d (Marois et al., 2008a) are recovered in both reductions, while planet e (Marois et al., 2010b) is only recovered robustly from this sequence using the new algorithm. Though the performance of both types of algorithm are comparable at wider separations, regions closer to the star are dramatically improved. We hypothesize that the improved performance is due to its more optimal treatment of the reference images that overlap with the target image, and its use of multi-target SNR optimization to limit the build up of correlated noise.

The shapes of the planet PSFs and negative side-lobes are quite different between the two reductions, indicating that SNR optimization chooses very different coefficients.

Unlike LOCI, SNR optimization is free to add adjacent reference images, or even improve the planet signal using the first Airy ring of the planet from a reference image. The net effect at small separations from the star is typically a broader planet PSF than LOCI.

Figure 2.7 shows the contrast improvement of each algorithm compared to a raw image. We find that at wider separations, the contrast was improved by roughly 20%, while at smaller separations, the contrast was improved by nearly 5 times.

Figure 2.8 shows histograms of the pixel values of a region close to the star for the three fully reduced images. The spread of pixel values is much narrower in the SNR optimization reductions, and is also more symmetrical with fewer outliers. Figure 2.9 shows how, in addition to producing images with lower residual noise, SNR optimization can also produce residuals whose distribution more closely resembles a Gaussian. This reduces the effects of non-Gaussian noise on detection, astrometry, and photometry confidence intervals.

2.6 Discussion

Now, we discuss a few considerations when applying this technique and directions that warrant further study.

First, there may be sequences where LOCI outperforms single-target image SNR optimization because a more aggressive subtraction, while sub-optimal for each processed image, improves the final stacked image. The core of the SNR optimization algorithm is parameter-free, so it is not possible to tune the aggressiveness to replicate this effect. Instead, multi-target SNR optimization should be used so that an optimal solution for the final stacked image is found. Compared to using LOCI with a tuned aggressiveness, this does rely more heavily on the assumption that the subtraction and optimization regions are well correlated.

Second, this technique does not permit a straightforward application of PCA methods which are useful for rejecting noise on some sequences. A potential future direction would be to transform the images into an approximately orthogonal basis for the SNR function. In the meantime, an archive of images collected from reference stars or images where the planet PSF does not overlap with the target image could still be used to generate an orthogonal basis using an algorithm like KLIP. These eigenimages could then be used in combination with the overlapping reference images

to combine the strengths of both algorithms. The use of a regularization parameter as described in Section 2.3.2 should achieve a similar result by favouring solutions with lower complexity.

Third, this method may distort a planet’s PSF, shifting its apparent centroid and affecting its photometry. Just as with previous techniques, astrometry and photometry can be extracted robustly using forward modelling (described in [Marois et al., 2010a](#); [Galicher, Marois, 2011](#)). Ultimately, it could be combined with a full forward-model matched filter [Ruffio et al. \(2017\)](#) to further improve the contrast. In general many additional steps and constraints have been shown to improve the performance of linear methods, and will likely be equally applicable to SNR optimization.

Finally, SNR optimization has a higher computational complexity than solving a system of linear equations. When combined with multi-target SNR optimization on a large sequence, the computational cost becomes significant. Our Julia-based SNAP implementation is able to reduce a single wavelength ADI sequence in roughly two hours on a 2015 Mac Pro workstation, or in just a few minutes when restricting the space explored by hyper-parameter tuning. Large SDI sequences can push the computational requirements higher such that a compute cluster becomes more appropriate.

2.7 Conclusion

We have presented a new algorithm for processing high contrast images based the direct optimization of the non-linear SNR function. We showed that this approach improves on previous techniques that find linear least squares solutions that minimize the noise, even if the parameters like aggressiveness are optimized. The new algorithm no longer requires us to reject reference images taken too close in time or wavelength. We also showed how this new formulation allows us to reduce multiple target images simultaneously. This allows the optimizer to work across the entire sequence, taking into account the temporal coherence of noise in addition to its spatial properties, and delivering optimized full-stack SNR. Finally, we presented our implementation in SNAP, the Signal to Noise Analysis Pipeline, and demonstrated a significant improvement to contrast in the challenging region close to the star. We expect this algorithm will improve the contrast of ground based facility-class instruments and future flagship space missions. This will enable the detection of fainter, lower mass, and closer in companions, or achieve the same sensitivity with less telescope time.

Chapter 3

Octofitter: Fast, Flexible, and Accurate Orbit Modelling to Detect Exoplanets

This chapter is based on material accepted for publication in the *Astronomical Journal*. This work was completed chronologically after Chapter 4 but is presented first as provides a more thorough exposition of the methods used in that work. Authors: William Thompson, Jensen Lawrence, Dori Blakely, Christian Marois, Jason Wang, Mosè Giordano, Timothy Brandt, Doug Johnstone, Jean-Baptiste Ruffio, S. Mark Ammons, Katie A. Crotts, Clarissa R. Do Ó, Eileen C. Gonzales, and Malena Rice. *My contribution to this work was as the lead author. The likelihood function for radial velocity and software implementation of simulation based calibration was contributed in part by Jensen Lawrence with my guidance. The expositive derivation of orbital positions, velocities, and accelerations was also contributed by Jensen Lawrence. The likelihood function for interferometric visibilities was contributed by Dori Blakely. The software implementation of one of the Kepler solvers was contributed by Mosè Giordano. The remaining authors provided comments or were contributors to the Gemini Planet Imager Exoplanet Survey.*

3.1 Introduction

Instruments for directly studying exoplanets are steadily improving in sensitivity. Current facilities are now accessing planets less than 10 AU from their stars. Below these separations, orbital motion can become significant over mere months. This will

be especially true for facilities with high angular resolving power thanks to their larger apertures and/or shorter operating wavelengths. This is an advantage for those who wish to determine the orbits of already-known planets but presents a significant challenge when searching for new companions.

When planets move from observation to observation, naive image stacking causes their signals to blur out and fade away. Reaching a necessary integration in a single epoch hits practical scheduling constraints and eventually physical limitations—for a sufficiently faint planet, it would not be possible to detect a significant number of photons before it moves by a full resolution element. These constraints apply equally to images as they do to integral field units and interferometers, including aperture masking interferometry (AMI) on JWST (Sivaramakrishnan et al., 2023), VLTI-GRAVITY (GRAVITY Collaboration et al., 2017), and even ALMA (Wootten, Thompson, 2009) in the case of accreting protoplanets.

A number of projects have sought to solve this challenge for image data only by compensating for orbital motion between epochs. These include a search for planets around Sirius B (Skemer, Close, 2011), K-Stacker (Nowak et al., 2018; Le Coroller et al., 2020), PACOME (Dallant et al., submitted), the search for planets around ϵ Eri (Mawet et al., 2019; Llop-Sayson et al., 2021), and the search for additional HR 8799 planets by Thompson et al. (2023). These have now led to promising evidence for α Cen AB b (Le Coroller et al., 2022) and HR 8799 f (Thompson et al., 2023). Moving the analysis of direct imaging data into the orbital domain enables a further extension: joint models of both images or interferometric observables with indirect exoplanet detection techniques, including radial velocity, astrometric motion, and transit (not directly considered in this paper). These have previously been explored in Mawet et al. (2019) and Llop-Sayson et al. (2021). This opens many possible scenarios. In addition to combining images to search for planets despite orbital motion, this allows one to freely combine Doppler or astrometric velocimetry with image data. This can then be used to constrain orbits using images with tentative detections or non-detections, improve photometry accuracy or limits, better constrain a planet's mass, and/or detect planets where any individual kind of data fails to reach significance (e.g. Llop-Sayson et al., 2021).

These scenarios are possible since all exoplanet detection methods provide orbital constraints that at least partially overlap. Imaging, RV, and transit all provide the orbital period (P); proper motion anomaly (PMA) and RV constrain eccentricity (e), the argument of periaapsis (ω), and either mass (m) or $m \sin(i)$ where i is the orbital inclination; and multi-epoch imaging/interferometry constrains all orbital parameters

up to a \pm ambiguity on the longitude of the ascending node (Ω). These connections could, in principle, allow information from all methods to flow into a single orbit model and, ultimately, the detection of a new planet.

To apply these ideas broadly, the community will need a tool capable of modelling all different types of exoplanet data. The `orbitize!` (Blunt et al., 2020) and `orvara` (Brandt, T. D. et al., 2021) packages come close: they support Bayesian modelling of relative astrometry, RV, and PMA. We needed a publicly available and generally applicable package that goes further to directly model image and interferometer data with or without independent detections at each epoch.

It is, generally, challenging to accurately compute orbital posteriors since the traditional Campbell orbital elements ($a, e, i, \omega, \Omega, t_{\text{peri}}$) possess complex co-dependencies and degeneracies (e.g. when $e = 0$ or $i = 0$). This task becomes even more challenging when working with short orbital arcs because they lack the constraining power to independently determine each Campbell element, meaning orbit posteriors are typically complex and very sensitive to their priors (e.g. O’Neil et al., 2019). Introducing image and interferometer data to the model exacerbates issues further, as they produce multi-modal posteriors that are challenging to traverse. Any inaccuracies in the calculation of an orbit posterior can lead to errors in mass and/or photometry and a spurious detection.

These challenges motivate the development of our new orbit- and data-modelling framework, “Octofitter.”¹⁰ Named for the eight types of data through which it aims to grasp new planets, Octofitter is designed from the ground up to be a flexible platform for modelling and experimentation while providing very high computational performance. These advances are thanks to our implementation of a pure-Julia (Bezanson et al., 2012), fully differentiable, and non-allocating modelling language, and our use of the higher-order No-U-Turn Hamiltonian Monte Carlo sampler (Xu et al., 2020). We further present the use of Simulation-Based Calibration (SBC; Talts et al., 2020) to confirm the accuracy of orbital posteriors, which is made practical by Octofitter’s speed. Figure 3.1 shows a schematic of a few of the ways Octofitter can be applied to exoplanet data.

As a publicly available code with these advances, Octofitter will enable the wide application of multi-epoch, multi-instrument, and multi-method exoplanet detection and modelling. These approaches have the ability to improve how direct surveys are completed and improve the yield of the upcoming Nancy Grace Roman Space Telescope

¹⁰Octofitter is publicly available on the Julia General registry, and extensive usage examples are provided at <https://sefffal.github.io/Octofitter.jl/dev/>


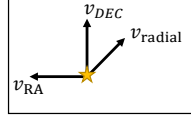
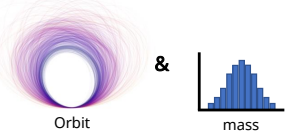
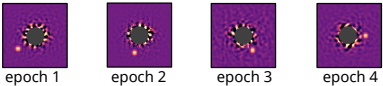
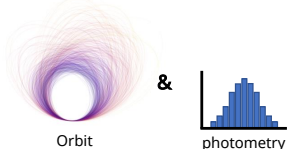

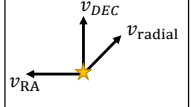
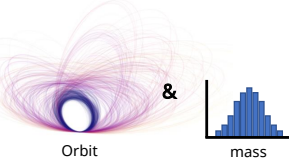
Input Data Examples	Computed Posterior	Detection
Relative Astrometry  & R.V. and/or P.M.A. 		
Images 		<input checked="" type="checkbox"/> Detection
Interferometric Observables  & R.V. and/or P.M.A. 		<input checked="" type="checkbox"/> Detection

Figure 3.1 Conceptual schematic of three different ways to use Octofitter. Other possibilities, like combining images with Doppler or astrometric velocimetry or using a mix of relative astrometry and images without detections to constrain orbits, are also possible..

Coronagraphic Instrument (Kasdin et al., 2020), the Habitable Worlds Observatory (HWO), and future facilities for extremely large telescopes.

3.2 Data Models

This section describes how the eight kinds of exoplanet data can be modelled by Octofitter. These types of data are relative astrometry, images, interferometric observables, star and planet radial velocity (absolute and relative), and proper motion anomaly.

The Octofitter framework is structured around three concepts: likelihood functions for observations, system models to tie observations to parameters, and generative functions to create synthetic observations. These break the problem down into orthogonal components that can be freely combined to solve a wide range of orbit modelling problems.

Each kind of observation in Octofitter is supported by its own Julia data type. Every

observation type is a wrapper for a data table of observations with a preset list of required columns. The observational data, in turn, can be provided to its associated type directly in the code, loaded from a local CSV or Arrow file, or obtained from a remote SQL database. For each observation type, a method of the `Octofitter.In_like` function is provided that computes the likelihood of the data it contains given a specific set of parameters. We now describe the observation types and likelihood functions included in Octofitter.

3.2.1 Relative Astrometry

The position measured between a directly imaged planet and its host star is one of the most fundamental measurements gathered from direct observations. It can be extracted from any image where the planet is robustly detected in a single epoch. Relative astrometry can be expressed either as separation in milliarcseconds and position angle in degrees, or as offsets in milliarcseconds in the R.A.–Dec. tangent plane.

Relative astrometry measurements can be provided to Octofitter using `AstrometryLikelihood`, which accepts a table with columns for `epoch`, the date of the measurement in units of modified Julian days; `sep`, the separation between primary and secondary in mas; `pa`, the position angle of the secondary measured East from North; `σ_pa`, `σ_sep`, and `cor` to specify the measurement uncertainties and correlation, respectively. Alternatively, `ra` (relative), `dec` (relative), `σ_ra`, and `σ_dec` can be substituted if preferred. A sample relative astrometry input table is presented in Table 3.1.

Based on these quantities, the likelihood function is simply taken as the Gaussian likelihood that the residual between the model and measurement for each parameter would be seen, given the provided measurement uncertainty. For relative positions measured along the R.A. and Dec. axes, this is

$$\begin{aligned} \log \mathcal{L}_{\text{astrom}} = & -\frac{1}{2} \frac{(\tilde{\Delta RA} - \Delta RA)^2}{\sigma_{\tilde{\Delta RA}}^2} - \log \sqrt{2\pi\sigma_{\tilde{\Delta RA}}^2} \\ & -\frac{1}{2} \frac{(\tilde{\Delta DEC} - \Delta DEC)^2}{\sigma_{\tilde{\Delta DEC}}^2} - \log \sqrt{2\pi\sigma_{\tilde{\Delta DEC}}^2}, \end{aligned} \quad (3.1)$$

where ΔRA and ΔDEC are the separation from the star along the R.A. and Dec. axes, $\sigma_{\tilde{\Delta RA}}$ and $\sigma_{\tilde{\Delta DEC}}$ are the corresponding uncertainties, and where the tilde distinguishes measured quantities from values calculated from model parameters. A more complex

Table 3.1. Relative astrometry input sample.

epoch (mjd)	pa (°)	sep (mas)	σ_{sep} (mas)	σ_{pa} (°)	cor
58849.0	224.93	615.2	30.0	0.8	0.0
58879.0	228.53	606.4	30.0	0.8	0.0
58909.0	229.66	663.3	30.0	0.8	0.0
58939.0	232.48	635.9	30.0	0.8	0.0
58969.0	233.25	610.3	30.0	0.8	0.0
58999.0	235.22	669.7	30.0	0.8	0.0
59029.0	236.67	666.8	30.0	0.8	0.0
59059.0	237.85	654.3	30.0	0.8	0.0
59089.0	239.68	713.6	30.0	0.8	0.0
59215.0	245.67	747.3	30.0	0.8	0.0
59245.0	245.54	736.4	30.0	0.8	0.0
59275.0	247.94	709.7	30.0	0.8	0.0
59305.0	248.13	791.4	30.0	0.8	0.0
59335.0	251.17	777.8	30.0	0.8	0.0
59365.0	251.43	773.8	30.0	0.8	0.0
59395.0	250.44	866.5	30.0	0.8	0.0
59425.0	253.49	789.5	30.0	0.8	0.0
59455.0	253.98	839.1	30.0	0.8	0.0
59945.0	268.72	986.3	30.0	0.8	0.0
59975.0	268.89	941.6	30.0	0.8	0.0
60005.0	269.96	959.2	30.0	0.8	0.0
60035.0	270.0	928.9	30.0	0.8	0.0
60065.0	270.67	952.9	30.0	0.8	0.0
60095.0	272.02	977.5	30.0	0.8	0.0
60125.0	270.99	950.1	30.0	0.8	0.0
60155.0	272.15	953.9	30.0	0.8	0.0
60185.0	274.33	985.2	30.0	0.8	0.0

expression is used when the correlations between uncertainties are non-zero.

3.2.2 Images

Directly modelling point sources in images is one of the key features of Octofitter. This is accomplished with the same approach described in [Ruffio et al. \(2018\)](#) and applied in [Mawet et al. \(2019\)](#), [Llop-Sayson et al. \(2021\)](#), and [Thompson et al. \(2023\)](#). In the case where a planet is robustly detected in each image and the uncertainties are well-approximated by Gaussian uncertainties, it is mathematically equivalent to extracting relative astrometry and photometry and then modelling those measurements. Where this approach goes beyond the two-step process, however, is when a planet is not detectable in a single epoch or when one wants to constrain the orbit of a planet based on a non-detection. The first case may arise when searching for planets that are too faint to detect before they exhibit orbital motion. The second case may arise any time a monitored planet passes too close to its star to detect. In this case, it's not

possible to extract astrometry. By directly modelling the images, large swathes of the orbital parameter space may still be ruled out. This can improve constraints on orbital parameters and improve predictions of the planet’s future location.

In Octofitter, image data can be modelled across any number of photometric bands, instruments, and epochs. Data can be provided using the `ImageLikelihood` observation type which accepts a data table with the columns `epoch`, `band`, `platescale`, `image`, and `contrast`. A sample input table is presented in Table 3.2. The platescale and parallax system parameters (which may be fixed or fit to the data) are used to map an orbit to a pixel location in the image. The `contrast` entry for each image allows one to pass an arbitrary function of position that gives 1σ contrast. If not provided, Octofitter calculates a contrast curve automatically from the image itself. This is sufficient for images with no clear detection, but should be avoided if the image contains a bright planet as the contrast curve will overestimate the noise at the planet’s separation. Extended emission from disks is not currently considered.

Given this data, the likelihood function used by Octofitter for each image by Octofitter is that of [Ruffio et al. \(2018\)](#):

$$\log \mathcal{L}_{\text{img}} = \frac{1}{2\sigma_{B,x,i}^2} (f_B^2 - 2f_B \tilde{f}_{B,x,i}), \quad (3.2)$$

where f_B is the model flux parameter for the photometric band B , x is the position in the image determined from orbital parameters, i is the epoch, σ is the uncertainty, and \tilde{f}_B is the measured flux at that same location.

This likelihood function assumes that flux is constant from epoch to epoch, but could easily be adapted to use an orbital phase function for planets imaged in reflected light (e.g. [Pogorelyuk et al., 2022](#)). Taken to the extreme, the flux can be fit independently at each epoch (e.g. [Nowak et al., 2018](#)) at the expense of reduced constraining power. For both images and relative astrometry, we do not currently consider uncertainty in the instrument’s North-angle or platescale. These might need to be considered when combining data from multiple instruments, in which case they could be added straightforwardly. We also note that this likelihood function assumes Gaussian distributed noise in each image. Where this is not a good assumption, standard techniques for inflating uncertainties could be used before applying Octofitter. Furthermore, spatial correlations between pixels are not currently handled. Finally, this likelihood function is only specified up to a constant and is not suitable for techniques like nested sampling.

Table 3.2. Images input sample

epoch (mjd)	band (symbol)	platescale (mas/px)	image (matrix)	contrast (function)
59976.0	:H	10.2
59976.0	:J	10.2
60576.0	:H	10.2

3.2.3 Interferometric Observables

Just as with imaging, combining interferometric observations across multiple epochs using orbital modelling removes the need to detect a companion in a single observation. We implement a model assuming an unresolved point source primary and N unresolved point source companions. The complex visibilities of this model are given by

$$V_{\text{bin}} = \frac{1 + \sum_{i=1}^N f_i \exp\left(-2\pi i\left(\Delta\text{RA}_i u + \Delta\text{DEC}_i v\right)\right)}{1 + \sum_{i=1}^N f_i}, \quad (3.3)$$

where f_i is the (companion/primary) contrast of the i th companion, ΔRA_i and ΔDEC_i are the right ascension and declination of the i th companion, and u and v are the Fourier domain coordinates, with magnitudes given by the interferometer baseline lengths divided by the observing wavelength (e.g., [Berger, 2003](#); [Kammerer et al., 2023](#)). A sample input table is presented in [Table 3.3](#).

The squared visibilities are calculated from the squared moduli of the complex visibilities, and the closure phases are calculated by summing the phases of the three complex visibilities calculated from triangles of stations in the interferometer. We construct likelihood functions for the squared visibilities and the closure phases separately, assuming Gaussian noise statistics and diagonal covariance matrices. We also assume that there is at least a moderate contrast between the primary and any companions such that we can neglect phase wrapping in the closure phase ([Thiébaud Éric, Young, 2017](#)). These likelihood functions are given by

$$\log \mathcal{L}_{\text{CP}} = -\frac{1}{2} \sum_i \frac{(\text{CP}_i - \tilde{\text{CP}}_i)^2}{\sigma_{\text{CP},i}^2} - \frac{1}{2} \log 2\pi\sigma_{\text{CP},i}^2 \quad (3.4)$$

Table 3.3. Visibilities input sample

epoch (mjd)	band (symbol)	pa (°)	sep (mas)	contrast
60096.0	:F480M	-92.7	180.5	0.00036
60171.0	:F480M	-61.5	159.1	0.00036
60462.0	:F480M	56.9	213.9	0.00036

and

$$\log \mathcal{L}_{V^2} = -\frac{1}{2} \sum_i \frac{(V_i^2 - \tilde{V}_i^2)^2}{\sigma_{V^2,i}^2} - \frac{1}{2} \log 2\pi\sigma_{V^2,i}^2, \quad (3.5)$$

where CP_i is the i th closure phase, V_i^2 is the i th squared visibility and $\sigma_{CP,i}$, and $\sigma_{V^2,i}$ are the uncertainties in the i th squared visibility and closure phase, respectively.

3.2.4 Radial Velocity

Similarly to other packages, Octofitter allows one to model radial velocity data in combination with other observation types. If combined with relative astrometry, proper motion anomaly, or image data, this allows one to directly model the dynamical mass of a planet.

RV data can be specified for the star or the planet using `RadialVelocityLikelihood`, which accepts a table with columns for epoch in MJD, `rv` in m/s, `sigma_rv`, the uncertainty on `rv` in the same units, and `inst_idx`. `inst_idx` is an integer between 1 and 4 used to specify which instrument the measurement corresponds to. The zero point `rv0_i` and jitter `jitter_i` must be specified as variables in the model where `i` corresponds to the instrument index.

A sample input table is presented in Table 3.4. The input format is the same for both cases. Depending on how the zero point is modelled, it is possible to use either relative or absolute RVs. The zero point can also modelled as an arbitrary function of other variables, allowing one to fit, for example, linear trends.

When combining radial velocity data with direct imaging modelling, it is possible (though not required) to connect the planet's photometry with its dynamical mass by using a user-supplied model. This model can either map a mass and/or system age to photometry in each band, or vice-versa. Connecting these two variables may be useful in cases such as when the orbit is determined by radial velocity up to the inclination

degeneracy but has not yet been detected with direct imaging.

In a similar manner to relative astrometry, we define a radial velocity likelihood function that allows us to fit orbital parameters to radial velocity data. This function is

$$\log \mathcal{L}_{\text{rv}} = -\frac{1}{2} \sum_i \frac{(v_{r,i} - \tilde{v}_{r,i})^2}{\sigma_{v_{r,i}}^2} - \log 2\pi\sigma_{v_{r,i}}^2, \quad (3.6)$$

where $v_{r,i}$ is the measured radial velocity, $\tilde{v}_{r,i}$ is the maximum likelihood estimate of the radial velocity, and $\sigma_{v_{r,i}}$ is the uncertainty in the radial velocity, all at the epoch i .

Octofitter supports multiple instruments, each with their own RV zero-point and jitter term and all with independently selectable priors. This is accomplished via the `inst_idx` column which associates each RV measurement to a particular instrument, zero point, and jitter. As a convenience, Octofitter includes a helper function to load RV curves from the publicly available HARPS RV Bank (Trifonov et al., 2020). Accessing this data will prompt the user to accept its license and will automatically fetch the RV curves.

3.2.5 Proper Motion Anomaly

Many systems that are candidates for direct imaging have their positions and proper motions measured accurately by both the Gaia (GAIA Collaboration et al., 2021) and Hipparcos (van Leeuwen, 2007) missions. The Hipparcos-Gaia Catalog of Accelerations (HGCA Brandt, T. D., 2021) cross-calibrates measurements from these two satellites and inflates uncertainties to cover most known systematics. This results in projected velocity measurements of the system’s photocentre at the Hipparcos epoch, the Gaia epoch, and between the two.

In Octofitter, proper motion anomaly must be loaded directly from the HGCA by specifying the host star’s Gaia identifier (the DR3 version at the time of writing). As with the radial velocity loaders, this will prompt the user to accept a license for the catalog and automatically download the HGCA.

In order to connect proper motion anomaly with the orbital image modelling described above, we define a likelihood function based on the HGCA data and the same orbital parameters:

Table 3.4. Radial velocity input sample

epoch (mjd)	rv (m/s)	σ_{rv} (m/s)	inst_idx
58849.0	47.2342	5.0	1
58879.0	14.4347	5.0	1
58909.0	12.9645	5.0	1
58939.0	26.2633	5.0	1
58969.0	-8.27905	5.0	1
58999.0	4.94685	5.0	1
59029.0	0.863664	5.0	1
59059.0	10.9524	5.0	1
59089.0	3.92389	5.0	1
59215.0	21.0509	5.0	1
59245.0	29.2009	5.0	1
59275.0	16.77	5.0	1
59305.0	19.0421	5.0	1
59335.0	52.9403	5.0	2
59365.0	21.8173	5.0	2
59395.0	55.0851	5.0	2
59425.0	20.444	5.0	2
59455.0	0.915145	5.0	2
59945.0	21.6418	5.0	2
59975.0	-0.54618	5.0	2
60005.0	11.4391	5.0	2
60035.0	15.3402	5.0	2
60065.0	17.7835	5.0	2
60095.0	5.35458	5.0	2
60125.0	60.3179	5.0	2
60155.0	43.5441	5.0	2
60185.0	11.4066	5.0	2

$$\log \mathcal{L}_{\text{pma}} = -\frac{1}{2} \frac{(v_{\text{RA}} - \tilde{v}_{\text{RA}})^2}{\sigma_{\tilde{v}_{\text{RA}}}^2} - \log 2\pi\sigma_{\tilde{v}_{\text{RA}}}^2 - \frac{1}{2} \frac{(v_{\text{DEC}} - \tilde{v}_{\text{DEC}})^2}{\sigma_{\tilde{v}_{\text{DEC}}}^2} - \log 2\pi\sigma_{\tilde{v}_{\text{DEC}}}^2, \quad (3.7)$$

where v_{RA} and \tilde{v}_{DEC} are the R.A. and Dec. proper motion at each epoch and σ_{RA} and $\sigma_{\tilde{v}_{\text{DEC}}}$ are the associated uncertainties. The input table format for this observation type matches Table 4 of [Brandt, T. D. \(2021\)](#).

Compared with [Brandt, T. D. et al. \(2021\)](#) and [Brandt, G. M. et al. \(2021\)](#) we adopt a simpler model of the Gaia and Hipparcos data. Rather than use the exact epochs each mission scanned the star (or estimate using the Gaia Observation Forecast Tool¹¹), we presume that the position and proper motion of the star were measured

¹¹<https://gaia.esac.esa.int/gost/>

independently 25 times over the duration of each mission. For orbits with periods longer than the observation windows of Gaia (≈ 668 days) and Hipparcos (≈ 1227 days), using these 25 epochs approximates the smearing effect caused by the moving photocenter during observations. Future work could refine this model further following their more sophisticated approach by re-implementing the calculations of HTOF (Brandt, G. M. et al., 2021) in Julia. Once it is available, it should also be possible to use the full intermediate GAIA astrometry catalogue in Octofitter.

3.3 Methods

3.3.1 Assessing Detections

This section describes how we choose to interpret the results of a model fit with Octofitter, though it should be noted that the modelling code makes no such prescriptions. For evaluating upper limits and detections from image or visibility data alone, we follow the conventions of Ruffio et al. (2018).

As set out in Ruffio et al. (2018), we calculate upper photometry or mass limits by finding where the cumulative distribution function equals some threshold e.g. 97.7%. That is, $\text{CDF}(f_{\text{lim}}) = 97.7\%$. This can be stated as an overall value for the system to say, for example, that we believe with 97.7% confidence that there are no planets present at all with photometry above some f_{lim} . It can also be calculated over small ranges of orbital parameters, in order to, for example, calculate a flux upper limit as a function of semi-major axis. This value will be driven by the brightest speckles close to the star, so a much more useful metric is to provide f_{lim} as a function of orbital parameters.

This model does presuppose that a planet exists and follows a Keplerian orbit around the star; however, as long as the photometry prior is broad, the model is free to drive the photometry towards zero. We can therefore use the photometry posterior to calculate a signal to noise ratio (SNR) analogous to the typical metric used in direct imaging. This approach has the benefit of being familiar to those with experience evaluating direct images for detections.

The procedure we adopt for assessing detections is to

1. sample from the posterior,
2. marginalize over all orbital parameters,

3. calculate the SNR as the mean of the marginal photometry posterior divided by its standard deviation,
4. and then compare this SNR with a pre-selected threshold based on a tolerable false positive rate, e.g. 5σ .

The results of this procedure are visualized in Figure 3.2. This gives a single SNR value for existence of any planet in that dataset consistent with the user's choice of priors. This is somewhat different than the typical SNR detection thresholds used because this SNR is calculated by marginalizing over all credible orbits, or equivalently, if applied to a single image, marginalizing over all credible positions. By contrast, the standard 5σ threshold is applied to the position in an image with the highest SNR, rather than a weighted combination of all positions.

In general, SNRs obtained using this approach assume that

1. the noise is approximately Gaussian distributed,
2. the noise is uncorrelated along orbital tracks between images,
3. the images contain only a single planet,
4. and that planet closely follows a Keplerian orbit.

Deviations from assumptions (1-2) may increase the false positive rate, while deviations from (3-4) may lower the recovered photometry and SNR.

Multi-planet systems could be accommodated easily by introducing additional planets to the model and/or restricting the priors to include only a subset of the parameter space.

For models that combine direct and indirect data, the relationship between mass and photometry variables may be complex, and this simple scheme may not be appropriate. In this case, it may be better to calculate a Bayes factor between planet and no-planet models. Such a Bayes factor can be treated as a direct measurement of our belief that a planet exists and follows a Keplerian orbit. This can be carried out in a limited fashion in Octofitter by calculating the Savage-Dickey density ratio of the mass or photometry marginal posterior (Dickey (1971); Koop (2003); for a pedagogical text, see Wagenmakers et al. (2010)). This is more flexible since it does not require a uniform prior—any prior that includes zero mass or photometry would suffice—and because it does not presume that the marginal posterior is Gaussian distributed.

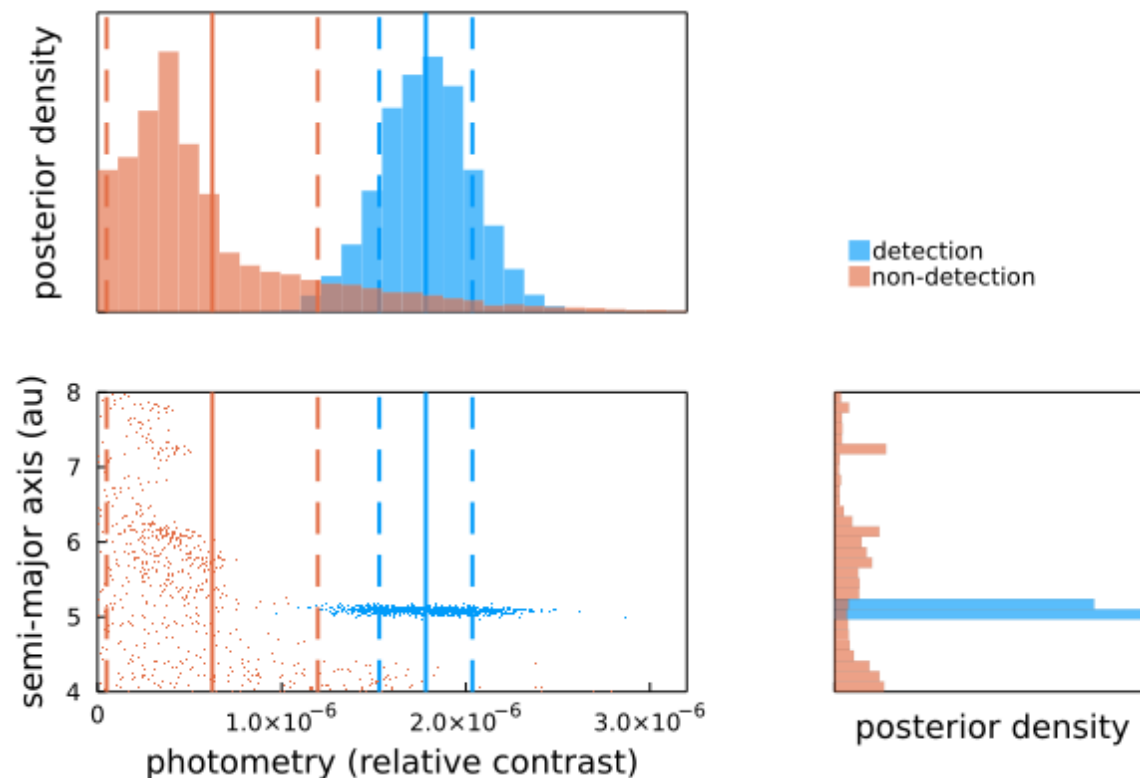


Figure 3.2 Posterior density of photometry versus semi-major axis for a simulated detection and non-detection. The solid lines mark the mean of the marginal photometry posterior, and the dashed lines mark $\pm 1\sigma$. We consider a planet detected when the SNR based on the photometry marginalized over all other parameters is greater than some chosen threshold, e.g. 5σ .

3.3.2 Orbital Bases and Priors

Octofitter supports a wide range of different orbital bases for use in different situations. These include traditional Campbell elements (a, e, i, ω, Ω), Thiele-Innes elements (e, A, B, F, G , Hartkopf et al., 1989; Wright, Howard, 2009; O’Neil et al., 2019), Cartesian elements (x, y, z, v_x, v_y, v_z , Ferrer-Chávez et al., 2021a), and a reduced basis set for modelling radial velocity only ($a, e, m \sin(i), \omega$). Users can specify priors using arbitrary distributions from `Distributions.jl`¹² and functions of those distributions.

For the analyses presented in this paper, we adopt either Campbell elements or Thiele-Innes elements with the following priors and modifications. When using Campbell elements, we adopt a log-uniform prior on semi-major axis, a uniform prior

¹²<https://juliastats.org/Distributions.jl>

on eccentricity, a sine prior on inclination, a Gaussian prior on host mass, and uniform priors on the remaining angular parameters. Note Octofitter reports the argument of periastron of the planet (and not the star) as ω and adopts $+z$ increasing away from the observer. These conventions match those used by `orvara` and `orbitize!` (Householder, Weiss, 2023). The full conventions used in Octofitter are described in Appendix A. When using Thiele-Innes elements, we adopt a log uniform prior on a “scale” parameter and multiply this with standard Gaussian priors on the constants A, B, F, and G. This maintains a log-uniform prior on semi-major axis and sine prior on inclination. For both cases, we replace the parameter τ , which gives the position of a planet along its orbit (Blunt et al., 2020), with θ , the position angle at the average epoch of the observations. This improves convergence relative to sampling from the τ parameter directly since θ is directly constrained by relative astrometry and is not sensitive to other orbital parameters. A derivation of τ from position angle θ is presented in Appendix C.

3.3.3 Modelling Language

To specify the structure of a system model, Octofitter provides a domain-specific modelling language. This simple language allows for the parameterizations and observations associated with each planet and the host star to be independently specified. For example, we could attach separate radial velocity measurements to each object in a system. Orbital parameters can be drawn from arbitrary prior distributions, fixed to particular values, or computed from combinations of other system parameters.

This modelling language makes it convenient to work with simple systems, like fitting the orbit of one planet to relative astrometry measurements, as well as more complex multi-planet systems.

The following is an example of a planet model using traditional Campbell orbital parameters, user-specified priors, and relative astrometry data:

```
table = Table(CSV.read("astrom.csv"))
astrom = AstrometryLikelihood(table)
@planet b Visual{KepOrbit} begin
  a ~ LogUniform(2.5, 25)
  i ~ Sine()
  e ~ truncated(
    Normal(0.2, 0.2),
    lower=0,
    upper=1
  )
end
```

```

)
 $\Omega$  ~ UniformCircular()
 $\omega$  ~ UniformCircular()
 $\tau$  ~ UniformCircular(1.0)
end astrom

```

A similar syntax is used to specify stellar properties:

```

@system HD1234 begin
  plx ~ truncated(
    Normal(41.123, 0.012), lower=0)
  M ~ truncated(
    Normal(1.5, 0.05), lower=0)
  age_Myr ~ Uniform(30, 300)
end b

```

The planet and system definition blocks contain pairs of variable names and values which can be constants, prior distributions, or arbitrary functions of other variables. Variables drawn from priors are specified by `~` whereas variables defined as constants or calculated as a function of other variables are defined with `=`. In this example, the convenience function `UniformCircular` creates two independent variables with standard normal priors and computes the angle between them using the arctangent. This, in effect, creates a circular domain where samples can smoothly wrap past 0 and 2π . Users are free to create their own parameterizations and likelihoods and combine them arbitrarily. The only requirements are that they are differentiable, smooth, and return a finite value at all points in the domain given by the priors.

This modelling approach, by being declarative rather than imperative, as in `exoplanet.py` (Foreman-Mackey et al., 2021), allows us to transform and evaluate the model in several ways. One key restriction is that each prior is proper, meaning it is a true probability distribution that integrates to unity. This is in contrast to some of the defaults used by `orvara`, which adopts fully scale-independent, but improper (in the statistical sense of the word) log-uniform priors. Octofitter requires proper priors to support tasks that require sampling directly from the prior distributions, such as simulation-based calibration, as will be discussed in Section 3.3.7.

3.3.4 Numerical Methods

From a model definition, Octofitter can generate efficient machine code using runtime-generated functions and Julia’s just-in-time compiler. This code generation step allows Octofitter to support a rich variety of observation types without paying any runtime overhead for features that aren’t used.

For the purposes of sampling from the posterior, Octofitter begins by remapping all variables from their possibly limited support (for example, eccentricity constrained between 0 and 1) into new variables defined across the entire real line. This makes it so that by construction, invalid parameter values like negative masses or semi-major axes are not possible to construct and will not be hit by a sampler. This transformation is performed by the Julia package `Bijectors.jl`¹³.

Next, a log-prior function is created for the model that simply evaluates the log-probability density of each prior distribution given a set of parameters. To preserve the user-specified prior distributions in place of the automatic bijections, a correction is applied based on the Jacobian of the transformation.

Similarly, a log-likelihood function is created based on the provided model and observations. Various constants are pre-calculated and reused between orbit solutions.

To enable the use of higher-order samplers, including Hamiltonian Monte Carlo, forward mode automatic differentiation is used to differentiate through the generated log-prior and log-likelihood functions (Revels et al., 2016). This provides the gradient of the log-posterior density, in addition to the value itself, without the overhead of calculating finite differences. The overhead of calculating both the log-posterior density and its gradient using forward mode automatic differentiation can be as low as just 2× compared to 10.2× for finite differences.

Special care was taken to remove all dynamic memory allocations from the generated log-density and gradient functions to prevent overhead from the Julia garbage collector.

Octofitter implements the Julia `LogDensityProblems` interface so that user models can be sampled from a wide variety of Julia-based MCMC sampler packages, including `AdvancedMH`, `AdvancedHMC`, and `MCMCTempering`. This allows users to select the best sampler for their particular problem and to compare results against samplers used by other popular packages.

The No-U-Turn Sampler (NUTS) variant of Hamiltonian Monte Carlo (Hoffman, Gelman, 2014) provided by `AdvancedHMC` is the default used by Octofitter. It allows

¹³<https://turinglang.org/Bijectors.jl/>

the code to explore complex posterior distributions with many fewer log-posterior density evaluations by simulating physical trajectories across the posterior landscape. The performance of NUTS in Octofitter will be discussed in Section 3.4.1.

When using the default NUTS sampler, we use a dense mass matrix, a jittered leapfrog integrator, a multinomial trajectory sampler, and allow the user to specify a maximum tree depth. We initialize the sampler by drawing 250,000 samples from the priors and selecting the value with the highest posterior density as the starting point. We also initialize the diagonal elements of the mass matrix using the interquartile range of the priors. We found this procedure was more robust than trying to determine the maximum *a posteriori* value with an optimizer for multi-modal posteriors like those found when modelling images. In particular, this procedure is less likely than an optimization pass to get stuck in a local optimum or other pathological position. A downside of this approach is that it may be inefficient if the priors are narrow and far from the posterior. For this case, the user can supply a starting point manually as in other tools. Next, we adapt the mass matrix and step size according to AdvancedHMC's implementation of the Stan windowed adaptation strategy. Finally, sampling proceeds until a preset number of accepted proposals are found—typically on the order of 1,000 to 10,000.

3.3.5 Kepler Solver

To sample from planet models, one must map from the orbital parameters specified in the system model to simulated observations, such as a planet's position over time in the plane of the sky or the radial velocity of the host star. In all cases, this requires solving Kepler's equation at every observation epoch and parameter draw. Kepler's equation (Eq. A.3) connects eccentricity and mean anomaly, a pseudo-angle somewhat analogous to the amount a planet has moved around its orbit, into eccentric anomaly, which can be used to find the actual location of the planet in its orbit.

Kepler's equation is transcendental and can't be solved analytically outside of special cases. The traditional approach to solving the equation is to use an iterative procedure like Newton's method with an initial guess chosen carefully from the mean anomaly and eccentricity. Octofitter's strategy is to wrap many different pluggable Kepler solvers that can be useful in different scenarios and use a fast non-iterative method as a default.

A non-exhaustive list of solvers supported by Octofitter are Markley (Markley, 1995), Goat (Philcox et al., 2021), and Newton. Newton's method and many other

available root finding algorithms are provided by the `Roots.jl`¹⁴ julia package.

Since these solvers are all implemented in pure Julia, there is no overhead from calling between Python and a C/Cython¹⁵ library and full performance is achieved with or without vectorization. Additionally, the solvers support changing the numerical precision between 16-bit, 32-bit, 64-bit, and arbitrary precision. Currently, no particular effort has been made to exploit hardware SIMD vectorization across epochs.

The `Markley` algorithm, the default choice, converges to nearly machine precision for all bound orbits when used with 64-bit floating point values. The `Newton` method can be combined with arbitrary precision floating point numbers to achieve arbitrarily tight tolerances if needed for niche applications. The default `Markley 1995` based method executes in just 90 ns on 64-bit floating point values for any valid eccentricity and mean anomaly (benchmarked on a Skylake Intel Xeon processor). This is slightly slower than the state-of-the-art results reported by [Brandt et al. \(2021b\)](#); however, their results of ≈ 60 ns are only achieved when the solver is vectorized over many epochs.

Applying automatic differentiation through an iterative Kepler equation solver would lead to poor performance. Even though Kepler's equation has no closed solution, the derivatives of eccentric anomaly with respect to mean anomaly and eccentricity can be found analytically using implicit differentiation (Eqs. [A.5](#) and [A.6](#)); that is, if we have already solved Kepler's equation for eccentric anomaly, we can calculate its gradient inexpensively. We supply these equations as a manual rule to the automatic differentiation library.

3.3.6 Analysis and Visualization

Once sampling is complete, Octofitter supports the user in testing the convergence of their chain or chains. The sampling results are returned as `Chains` objects from `MCMCChains.jl`. This table-like structure includes entries for each accepted MCMC proposal as well as metadata about the sampling process, such as the compute time used. All variables are returned, including those that were fixed to a constant or calculated deterministically from combinations of other variables. An automatic summary is output that includes plausible intervals, expectation values, effective sample sizes (ESS), and \hat{R} statistics for each parameter. The user can test the convergence of their chain or chains by assessing those ESS and \hat{R} values, creating a trace plot, and/or by calculating figures

¹⁴<https://juliamath.github.io/Roots.jl/>

¹⁵<https://cython.org/>

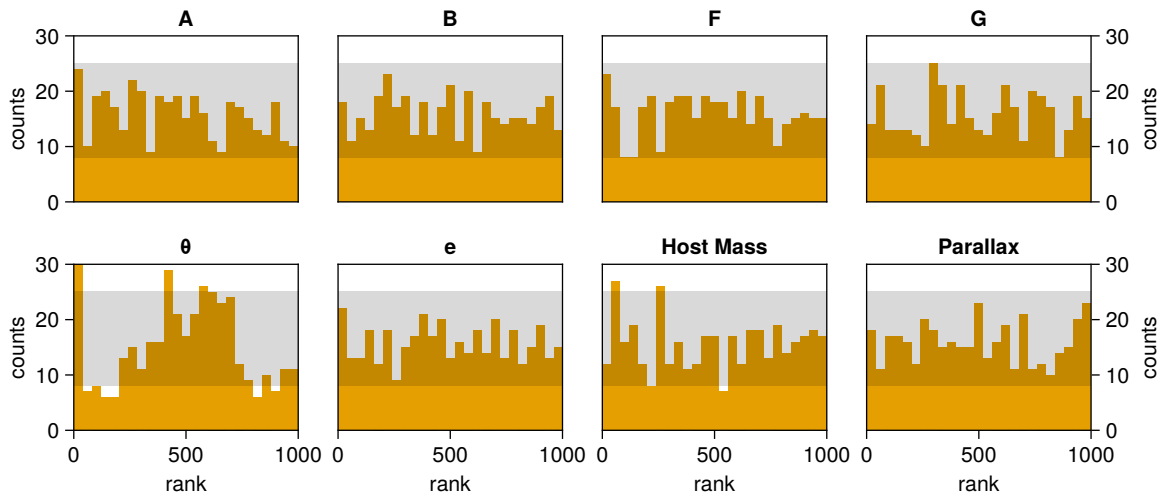


Figure 3.3 Sample results of running the simulation-based calibration procedure on a model consisting of a single planet parameterized with Thiele-Innes elements (A , B , F , G , etc.) and position angle θ . Each count represents a model fit to a different simulated system. The horizontal band gives a 99% range around the expected value for a perfect sampling procedure. The observation epochs and uncertainties are taken from Table 3.1.

like the Gelman, Rubin, and Brooks diagnostic (Gelman, Rubin, 1992; Brooks, Gelman, 1998) using tools provided by `MCMCChains.jl` and `MCMCDiagnosticTools.jl`.

Octofitter can be used to visualize orbit fits in several ways. The orbits represented by chains can be plotted in the plane of the sky, in physical system coordinates (i.e. AU), or as a time series (e.g. for Doppler or astrometric velocimetry).

When visualizing an orbital posterior, a common challenge is ensuring enough data points are used to create a smooth arc. This becomes especially challenging with eccentric orbits or ensembles of orbits with widely varying periods. Tracing out orbits in equal time steps will lead to most points clustering closer to apoastron, where the planet is moving the most slowly. The strategy used by Octofitter is to trace orbits in equal steps of eccentric anomaly, as suggested in Berry, Healy (2002). This places points in regions of greater curvature, creating smooth arcs with fewer points.

Octofitter includes functions to generate plots that visualize the posterior and compare it to the input data. This includes astrometry, separation, position angle, radial velocity, and proper motion anomaly. It also includes a function for visualizing orbits in spatial coordinates (units of AU) in one, two, or three dimensions. Examples of these plots are shown throughout the text and documented with the online tutorials.

3.3.7 Simulation-Based Calibration

Simulation-based calibration (SBC; [Cook et al., 2006](#); [Talts et al., 2020](#)) is a technique that allows one to verify that the output of a Bayesian modelling procedure is unbiased. This includes any part of the computation, from model specification to the sampling procedure after the choice of priors. Verifying that this choice of priors is reasonable is the domain of other procedures like prior and posterior predictive tests.

A conceptually related procedure was carried out in [Ferrer-Chávez et al. \(2021b\)](#) in which the authors systematically tested the `orbitize!` package, parameterization, and default priors for biases. Our contribution here is to present a procedure that is well explored in the statistics literature, is specific to a given set of observation epochs and measurement uncertainties, and can be applied in an automated fashion.

The calibration procedure requires a generative model—that is, a way to take a given set of parameters and create a simulated observation. A familiar example from direct imaging is the injection of fake point sources to test image processing algorithms. To follow this procedure, one repeatedly draws a set of parameters from the priors, creates simulated observations based on those parameters, samples from the resulting posterior, and then compares the true parameter values (originally sampled from the priors) with the resulting posterior. By doing this many times, one creates a histogram of rank statistics that can reveal many sources of biases present in the model and sampling process. To be clear, this does not evaluate how the choice of priors impacts the posterior.

We implemented support for performing simulation-based calibration automatically in Octofitter. In this procedure, Octofitter takes a given model’s observational data, discards the actual measurements, and keeps the epochs and uncertainties associated with each measurement. Octofitter then repeatedly creates simulated data at each epoch by drawing from the priors and performs SBC on this simulated data. SBC should, in general, be applied to each model to confirm that it is working as expected.

Figure 3.3 shows the results of the simulation-based calibration procedure applied to a model of a planet parameterized by Thiele-Innes elements and position angle at the average epoch. For the most part, these histograms do not reveal any systematic biases in the Octofitter sampling procedure. One exception is the “bump” in the position angle histogram. This shape indicates that Octofitter is under-confident and the sampled posterior distribution is wider than the true posterior. By contrast, seemingly small tweaks to the choice of priors can result in histograms with strong biases. For example,

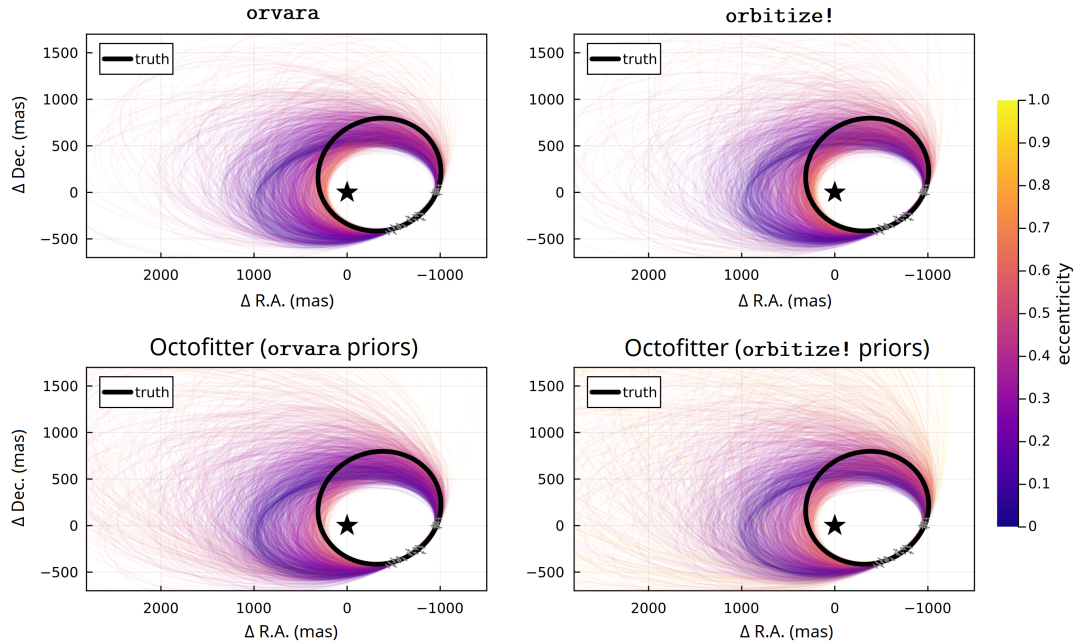


Figure 3.4 Orbit fitting posteriors visualized in the plane of the sky, compared between three packages: *orvara*, *orbitize!*, and *Octofitter*.

drawing the Thiele-Innes constants (A, B) and (F, G) from log-uniform prior distributions rather than Gaussian distributions centred around a single scale, itself drawn from a log-uniform prior, results in noticeable issues with the SBC histograms. A guide to interpreting the results of the SBC procedure is available in [Talts et al. \(2020\)](#).

These tests illustrate the value of performing the simulation-based calibration procedure for each new model and data combination a user wishes to use. Given the complexity of orbit models and the difficulty of sampling from them, we do not expect our sampling procedure to be entirely unbiased. Rather, we hope that by creating an easy way to diagnose these biases, users of *Octofitter* can interpret their results with an appropriate level of skepticism in accordance with the level of bias detected.

3.3.8 Other Packages

Fitting observations of planets and binary stars to orbits has been widely addressed in the literature, dating back to Kepler’s seminal work. More recently, a variety of software packages have been released following both frequentist and Bayesian approaches. Some of these packages include *EXOFAST* ([Eastman et al., 2019](#)), *rvfit* ([Iglesias-Marzoa et al., 2015](#)), *radvel* ([Fulton et al., 2018](#)), and the Python *exoplanet* package ([Foreman-Mackey](#)

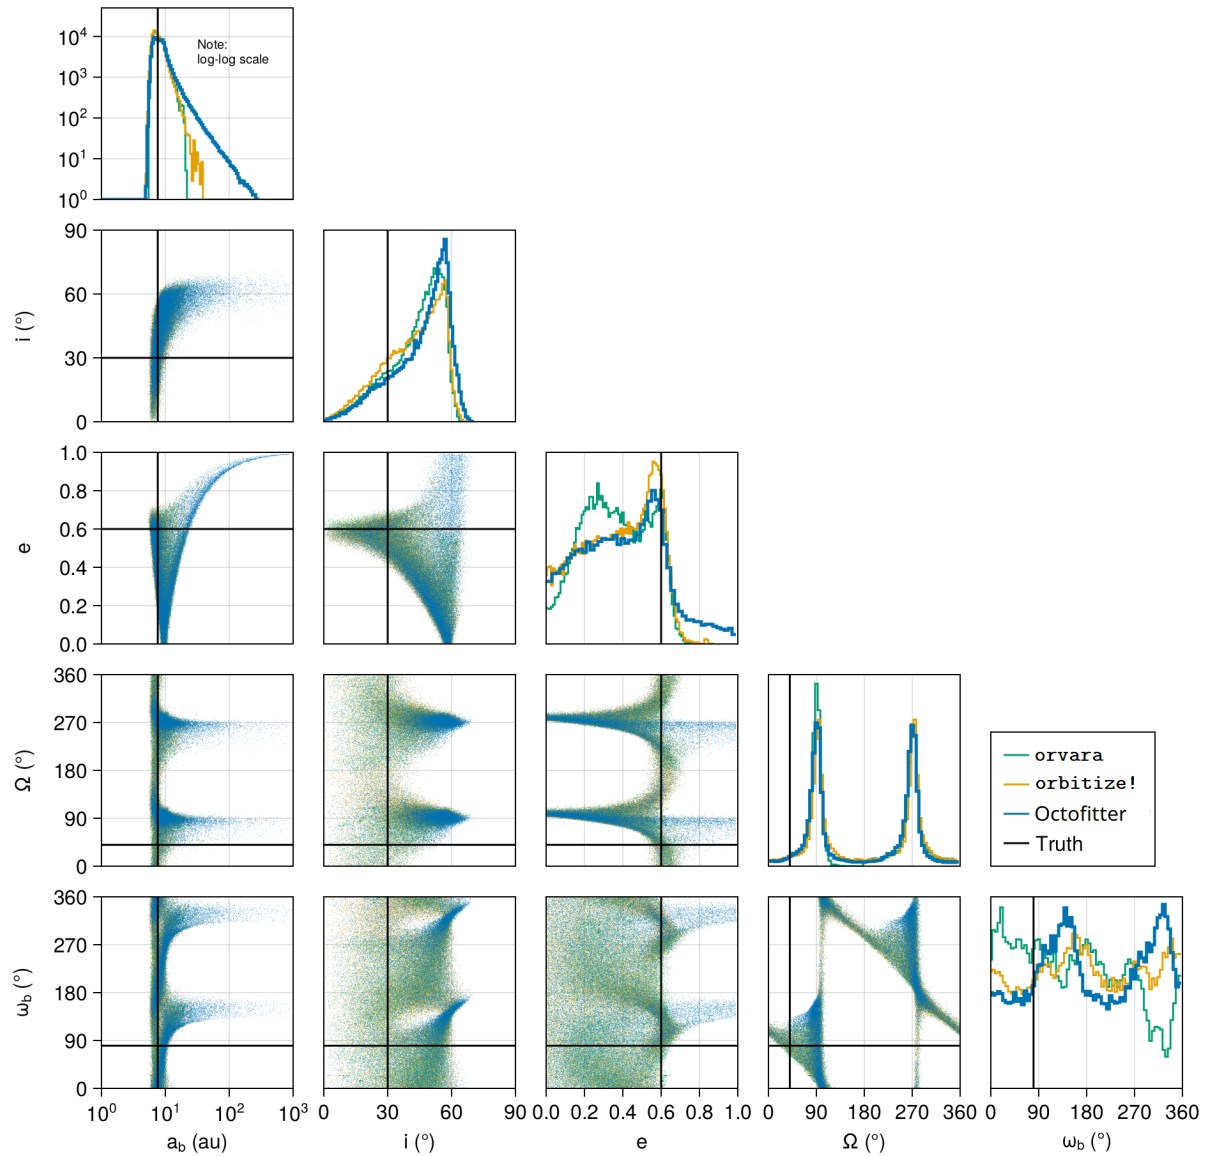


Figure 3.5 Corner plot of key orbital parameters for the case shown in Figure 3.4 compared between three packages: *orvara*, *orbitize!*, and Octofitter. Semi-major axis is plotted on a log scale to reveal how the sampler behaviour differs in the long tail. Though we expect all packages to eventually converge to near-identical results, we find that there are small differences in the ω , a , and e marginal posteriors that persist even after running for multiple days.

et al., 2021). Most relevant to image- and visibility-based modelling are packages for fitting relative astrometry, radial velocity, and proper motion anomaly, like `orbitize!` (Blunt et al., 2020), `orvara` (Brandt, T. D. et al., 2021), and `Efit` (presented in Meyer et al., 2012), as well as image searching codes like `K-Stacker` (Nowak et al., 2018) and `PACOME` (Dallant et al, submitted).

These tools have employed a variety of methods for approximating posteriors, including affine invariant (Foreman-Mackey et al., 2013) and parallel tempered affine invariant Markov chain Monte Carlo (Vousden et al., 2016) in `orvara` and `orbitize!`, Hamiltonian Monte Carlo (Hoffman, Gelman, 2014) in `exoplanet`, and nested sampling (Skilling, 2004) in `Efit`. With the exception of `exoplanet`, the majority of these codes have used first-order samplers. That is to say, they rely only on evaluating the posterior density and do not calculate or make use of gradient information.

3.4 Results

3.4.1 Demonstration with Relative Astrometry

We begin our demonstration of Octofitter by fitting orbits to simulated relative astrometry measurements, which is one of the simplest use cases for the package. We considered a simulated orbit of a single planet and generated astrometry at 27 epochs (Table 3.1). For the sake of comparison, we performed the same fit using two other popular orbit fitting packages: `orvara` and `orbitize!`.

We followed the best practices laid out in the tutorials provided with each package. `orvara` and `orbitize!` use slightly different priors by default and cannot be made to match each other exactly without code modifications. `Orvara` uses an $\sqrt{e} \sin(\omega)$ and $\sqrt{e} \cos(\omega)$ parameterization while `orbitize!` uses separate uniform priors on both e and ω . Therefore, we ran our comparisons twice, first adopting priors similar to `orvara` and then priors similar to `orbitize!` In all cases, we used a log-uniform prior on semi-major axis between 0.1 and 1000 AU. We ran the `orvara` and `orbitize!` packages with settings recommended by their authors. These were 4 temperatures with 100 walkers for `orvara` and 20 temperatures with 1000 walkers for `orbitize!`. We ran Octofitter with 16 independent chains. We initialized the walkers used by `orvara` and `orbitize!` in a Gaussian ball around the true orbit. For `orbitize!`, we improved the convergence by randomly initializing half the walkers on the second mode of the ω marginal posterior. For Octofitter, on the other hand, we initialized it automatically

using our procedure of drawing from the priors and selecting the parameters with the highest posterior density.

We drew 20,000,000 samples from each walker using `orvara` and 100,000 from each walker using `orbitize!`. For Octofitter, we adapted the step size and mass matrix for 5,000 iterations and then drew a further 15,000 samples. The resulting posteriors are compared in Figure 3.4. To remove the burn-in phase, We discarded the first halves of the `orvara` and `orbitize!` chains and the first 5,000 samples of the Octofitter chains (the adaptation phase).

To measure how long each package takes to converge to a steady distribution, we followed a similar procedure to Ferrer-Chávez et al. (2021a). We divided each chain into 50 segments and assumed that in the final segment, the chain is fully converged. We then calculated the \hat{R} statistic between each segment and the final segment. We used the rank normalized and median folded version of the statistic as implemented in `MCMCDiagnosticTools.jl`¹⁶ to evaluate how well the samplers converged in the bulk and in the tails of the distribution. We considered it converged once \hat{R} became less than 1 ± 0.005 for all variables. We evaluated this on all walkers (`orvara` and `orbitize!`) and all chains (Octofitter) and took the median.

The most important result of this code comparison is that in the limit of large numbers of samples, all packages produce posterior distributions that largely agree with each other and that include the true orbit. The orbit paths in the plane of the sky are compared in Figure 3.4 and the parameters are compared in a corner plot in Figure 3.5. This should serve to further improve confidence in the results of all packages.

The sole exception is that the orbit posterior from Octofitter contains more samples from a high-eccentricity, high-semi-major axis branch of the posterior than the two other packages (Figure 3.5). Nonetheless, running SBC on this model with uncertainties and observation epochs from this dataset (Figure 3.3) reveals that Octofitter is acceptably calibrated and is not overestimating the spread of the posterior. It is not computationally feasible to perform the same SBC experiment with the other packages, but this likely indicates that the additional samples from the long tail of the posterior are representative of the true posterior and would eventually appear in the outputs of the other packages.

A second result is that for this problem, Octofitter converges to a steady distribution almost immediately using a single computer core, while other packages take many hours to converge (Figure 3.6). Since assessing the convergence of MCMC chains can be difficult, we suggest that with Octofitter, unsophisticated users are therefore less

¹⁶<http://turinglang.org/MCMCDiagnosticTools.jl/>




Figure 3.6 Comparison between packages of the average time until chains converged to a stationary distribution and the rate at which independent posterior samples are generated. The effective sample size (ESS) rate was measured separately using the bulk and tail methods of `MCMCDiagnostics.jl`. Note that the \hat{R} and ESS of the slowest variable for a given sampler are used as this is what ultimately limits sampling performance. These results are based on the astrometry presented in Table 3.1, and are expected to depend strongly on hardware, input data, and choice of priors.

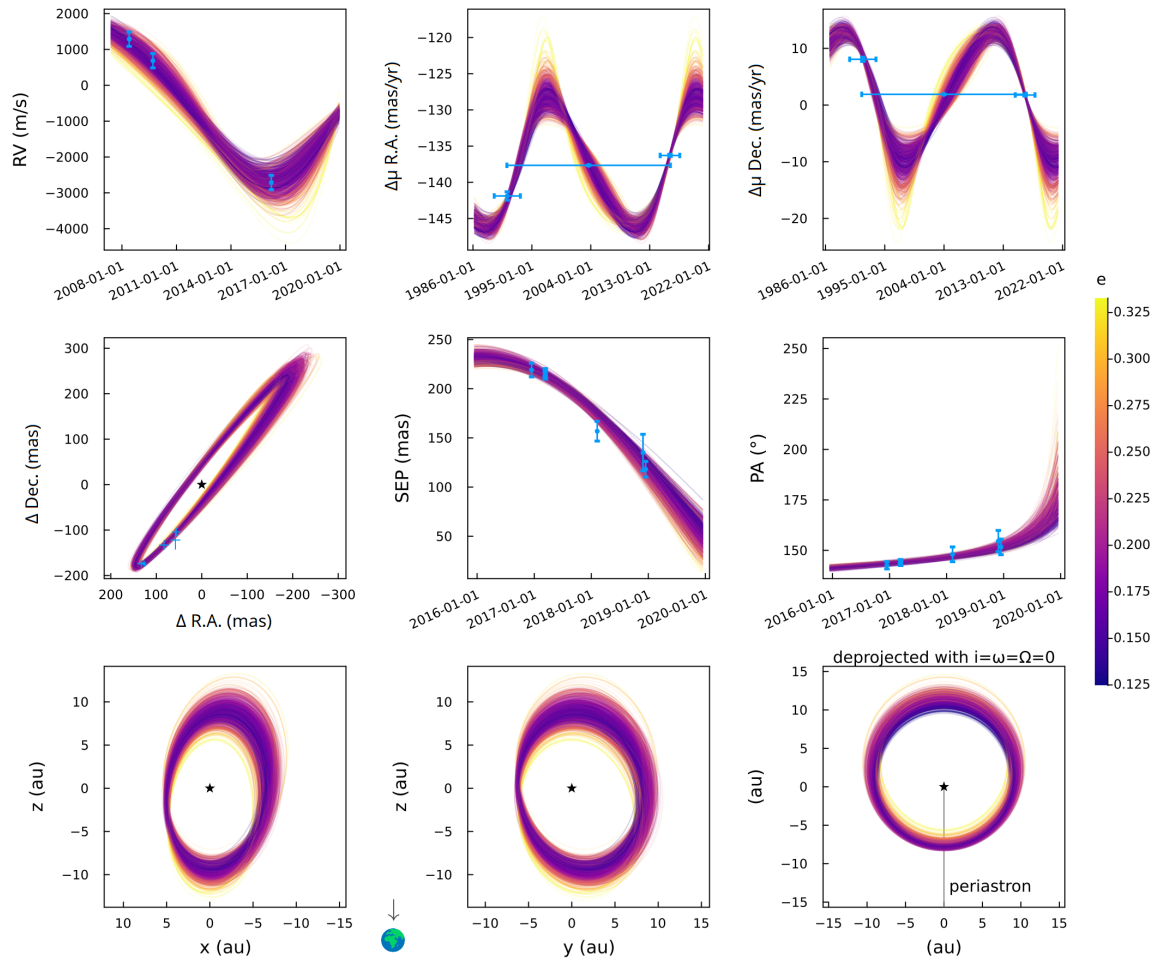


Figure 3.7 Sample plot output from Octofitter using data from the HD 91312 system. The top row visualizes the orbital posterior compared with velocity measurements of the host. The horizontal bars in the proper motion panels show the timespans over which the average velocity was measured. The middle row shows the posterior compared with astrometry measurements in the plane of the sky and in separation and position angle over time. The bottom row shows the orbital posterior in physical coordinates to compliment the astrometry plot. The rightmost panel shows a deprojected view of the system where orbits have been rotated face on and to place periastron at the bottom. The conventions used by Octofitter are described in Appendix A.

likely to use insufficiently converged results in their analyses. Of course, in addition to needing the chain to be converged, one must also generate a sufficient number of samples for the problem at hand. Interestingly, despite the much faster convergence, Octofitter is no faster than the other packages at producing independent samples. In most workflows, the bottleneck is ensuring convergence rather than a need to produce tens of thousands of samples, so Octofitter should still offer a decisive improvement in computation time.

The exact results are hardware specific, sensitive to the data, and depend on the choice of priors and parameterization. It is not feasible to fully account for all small differences between software packages, and it is likely that different approaches will perform better or worse depending on the problem at hand. An additional caveat is that the `ptemcee` sampler used by `orvara` and `orbitize!` scales across many cores which reduces the total sampling time. Octofitter supports running multiple chains in parallel but is not configured to use multiple cores to accelerate a single chain. Finally, given the convergence guarantees of Markov chain Monte Carlo, all packages would approach the same distribution given infinite time. This comparison, therefore, is meant only to illustrate the typical efficiency one might expect with Octofitter on similar problems and reasonable run times.

3.4.2 Demonstration with Relative Astrometry, RV, and PMA

We now demonstrate Octofitter on a real system with radial velocity, proper motion anomaly, and relative astrometry data.

The HD 91312 system consists of a $1.6M_{\odot}$ star orbited by a binary companion with a mass of approximately $0.3M_{\odot}$ and separation of roughly 10 AU (Chilcote et al., 2021). The companion was discovered by a direct imaging search targeting accelerating stars (Currie et al., 2021) from the HGCA. We now reproduce the orbital analysis of the discovery paper in order to demonstrate Octofitter’s modelling capabilities when applied to a combination of relative astrometry, radial velocity, and proper motion anomaly data.

For this analysis, we used the astrometry data from Chilcote et al. (2021), proper motion anomaly data from the HGCA (Brandt, 2021), and limited radial velocity data originally from Borgniet et al. (2019). As the original radial velocity data was not forthcoming, we measured it from the PDF plots submitted to the arXiv alongside the 2021 manuscript. We used similar priors as those described in the discovery paper,




Figure 3.8 GPI sequences used for simulations in this section. Each image was normalized to have the same average contrast at 200 mas separation from the star (just outside the edge of the mask). The images displayed above contain a simulated planet orbiting CCW at an average of just SNR 1 per epoch, spaced one month apart. Given this data, the model recovers the simulated planet at SNR 7.

namely log-normal priors on primary and companion mass and a uniform prior on the square root of eccentricity. These were chosen to match the analysis of the discovery paper for the purpose of comparing codes and demonstrating Octofitter rather than any physical motivation on our part.

The results of this orbit modelling are presented in Figure 3.7. Octofitter recovers the orbit of the companion with similar results to those presented by [Chilcote et al. \(2021\)](#). The radial velocity, proper motion anomaly, and relative astrometry are all consistent with the secondary companion having a mass of approximately of $300 M_{\text{Jup}}$, given the choice of priors described in the original discovery paper.

3.4.3 Demonstration with Images

We now present a series of simulations showing how this framework allows us to detect planets using multi-epoch direct images. We selected 50 sequences from the Gemini Planet Imager Exoplanet Survey ([Nielsen et al., 2020](#)), processed using a forward model matched filter (FMMF, [Ruffio et al., 2017](#)), that have a stellar I-band magnitude less than or equal to 6 and have no previously detected point sources. To be clear, we do not search these sequences for real companions but merely use them as realistic noise maps for our simulations. We normalized the contrasts of each sequence to the median




Figure 3.9 Octofitter’s ability to recover planets from simulations of circular, face-on orbits as a function of signal to ratio per epoch and number of epochs. **Left:** Recovered SNR. Cells above the yellow line would be detected with a 5σ threshold. The cell outlined in cyan corresponds to the data shown in Figure 3.2. **Right:** The same SNRs relative to an ideal \sqrt{N} improvement with the number of epochs. We find that the SNR grows as expected unless the final, combined SNR is below ≈ 5 . The recovered SNR falls off quickly below this value and levels off at approximately 1. This explains why the recovered SNR of 1 exceeds the injected value of 0.5 in the bottom left corner.

at 200 mas separation. This retains the true noise distribution but removes the effects of sequence-to-sequence variation on our results. We consider a hypothetical system at 20pc with a $1M_{\odot}$ star that is observed once per month for between 1 and 100 months (a little over eight years). Using these sequences and parameters, we generated simulated observations of a planet by injecting a synthetic PSF into the correct positions in each epoch. This input data is presented in Figure 3.8.

For all models, we adopted a $1 \pm 0.1M_{\odot}$ prior on M , a uniform $0 - 30M_{\text{jup}}$ prior on m , a Gaussian prior on parallax, a uniform prior on a , and a uniform circular prior on τ .

We test this model’s ability to detect planets in sequences of direct images for the most straightforward case: circular, face-on orbits. We injected the planet into 1, 5, 10, 25, 50, and 100 images with an average SNR ranging between 0 and 5 in each image. Finally, we applied our model to each of these simulated datasets to arrive at a grid of recovered SNR values as a function of number of epochs and SNR per epoch (Figure 3.9). We find that we are able to recover planet detections with arbitrarily low SNR per epoch despite orbital motion, provided we have a sufficient number of observations.

Figure 3.9 also shows how the recovered SNR compares with the SNR we would expect for combining images without orbital motion in the presence of uncorrelated

Gaussian noise. We find that the model detects the injected planets with near perfect \sqrt{N} scaling when the final, combined SNR is greater than ~ 5 . For instance, a planet injected into 100 epochs spaced one month apart, at an SNR per epoch of just 0.5 (below the noise), is still robustly detected at a final SNR of 5. We do note that since each image was taken of a different star, this ideal scaling is a best-case scenario. It is possible that repeated observations of the same target could lead to correlated residuals that reduce the final SNR, though the orbital motion of the planet should mitigate this in much the same way as angular differential imaging (Males et al., 2015; Marois et al., 2006a).

The left-hand columns of Figure 3.9 are of particular note. They show that the model fails to detect a point source injected into a single image at an SNR of 4. The failure to detect planets with expected significance below $\text{SNR} \approx 5$ can be understood by contrasting our detection criteria with the standard used in the field. In these models, we consider the overall SNR marginalized over all locations in the image (or orbits through a sequence), whereas typically, one looks at the maximum SNR at any given location. For example, in any SNR 5 detection of a planet, numerous other SNR 2 and 3 peaks exist. When looking at the SNR marginalized over all locations, these other peaks serve to reduce the final SNR. This makes the SNR calculated from the marginal photometry posterior a more stringent planet detection test.

3.4.4 Demonstration with Aperture Masking Interferometry and Radial Velocity

We now demonstrate Octofitter using a combination of simulated radial velocity data and aperture masking interferometry (AMI) visibilities. We considered a plausible scenario where a planet has been detected using radial velocity, had its orbit characterized, and is then followed up with a series of three observations with JWST NIRISS AMI (Doyon et al., 2012; Sivaramakrishnan et al., 2023). For this experiment, we connected the mass of the planet to its M_s band photometry using Sonora Bobcat models (Marley et al., 2021) for a fixed system age of 10 Myr. The simulated system had a true orbit with a semi-major axis of 2 AU, an eccentricity of 0.1, an inclination of 45° , a parallax distance of 100 mas, and a mass of $5 M_{\text{jup}}$. The star had a stellar absolute M_s -band magnitude of 2.6.

We simulated the radial velocity by adding Gaussian noise with a standard deviation of 25 m/s to 74 epochs of RV data generated from the above system. The data points were spaced by ~ 30 days between 2017 and 2023.

The AMI data was generated using Eq. 3.3 for a single companion following the given orbit. For this example, we used only the closure phase and did not include the squared visibilities. We again added Gaussian noise to the calculated closure phases, with a closing triangle dependent standard deviation taken from the closure phase uncertainties calculated with AMICAL (Soulain et al., 2020) from NIRISS AMI F480M data of (the presumed single stars) HD 101531 and HD 123991 calibrated against each other.

We applied Octofitter with a joint model of the RV data and three AMI epochs. We also compared this to χ^2 maps from the `Fouriever` framework (presented in Kammerer et al., 2023) of each individual epoch. The results of this experiment are shown in Figure 3.10.

Using the AMI data, the Octofitter model is able to constrain the inclination and, therefore, the mass of the planet. Due to the faint signal and presence of noise, the `Fouriever` results show spurious signals at each epoch. In comparison, Octofitter is able to connect the three epochs by a single higher-significance mode in orbit-parameter space.

This example illustrates how joint modelling across epochs can be used to increase the significance of AMI and other similar interferometric observations.

3.5 Conclusion

We have presented a new code, `Octofitter`, for modelling exoplanet relative astrometry, radial velocity, and proper motion anomaly, as well as performing non-traditional tasks like detecting moving point sources across images and interferometric observables.

- We demonstrated the simulation-based calibration procedure on a hypothetical orbit fitting task and found that for orbits parameterized with Thiele-Innes constants, `Octofitter` is acceptably calibrated.
- We compared the results of `Octofitter` to the popular packages `orvara` and `orbitize!` and found that all three arrive at similar posterior distributions.
- We showed that `Octofitter` can converge 1 to 2 orders of magnitude faster than other popular packages, making it computationally feasible to perform a variety of statistical checks like simulation-based calibration for individual datasets.

- We demonstrated a combined fit to relative astrometry, radial velocity, and proper motion anomaly of the HD 91312 system and found a companion mass that agreed with previous results.
- We demonstrated the ability to detect arbitrarily faint companions despite orbital motion using simulations and data from the Gemini Planet Imager Exoplanet Survey.
- We demonstrated a combined model of simulated radial velocity and multi-epoch JWST/NIRISS aperture masking interferometry.

Octofitter is a powerful new tool that will enable the community to broadly apply joint multi-epoch, multi-instrument, and multi-band direct imaging, interferometry, Doppler radial velocity, and astrometric motion. This will allow for the detection and characterization of fainter and lower-mass planets more similar to those in our solar system.

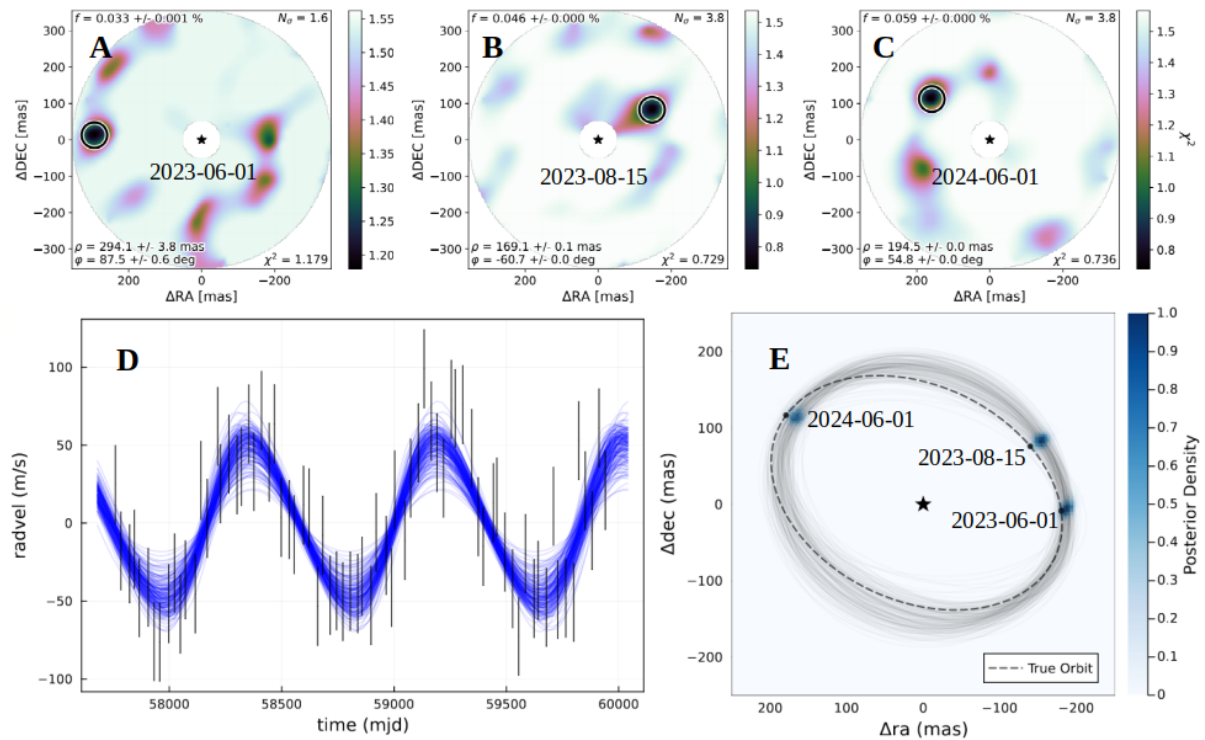


Figure 3.10 Joint modelling of AMI and RV data. **A-C:** Recovered χ^2 maps at each independent AMI epoch created using Fouriever (presented in Kammerer et al., 2023). **D and E:** joint radial velocity and AMI posteriors visualized as an RV time series (D) and orbit posterior in the plane of the sky (E).

Chapter 4

Deep Orbital Search for Additional Planets in the HR 8799 System

This chapter is based on material published in the *Astronomical Journal*, Volume 165, Issue 1, pp. 29 (2023). Authors: William Thompson, Christian Marois, Clarissa R. Do Ó, Quinn Konopacky, Jean-Baptiste Ruffio, Jason Wang, Andy J. Skemer, Robert J. De Rosa, and Bruce Macintosh. *My contribution to this work was as the lead author. Some of the observations were carried out by Quinn Konopacky and Christian Marois prior to my joining the project in 2018. I led the analysis with input from Jean-Baptiste Ruffio, Christian Marois, and Jason Wang. The stability analysis was completed by Clarissa R. Do Ó based on the general orbital modelling analysis. The remaining authors provided access to additional observations or comments on the manuscript.*

4.1 Introduction

The HR 8799 planetary system hosts four giant planets still glowing hot from their recent formation. It was the first multi-planetary system to be directly imaged (Marois et al., 2008a) in 2008. Since the subsequent detection of a fourth inner planet in 2010 (Marois et al., 2010b), it became the benchmark system in direct imaging. Extensive follow up observations of the four planets b, c, d, and e have characterized their orbits, compositions (Currie et al., 2011; Skemer et al., 2012; Konopacky et al., 2013; Skemer et al., 2014; Currie et al., 2014; Barman et al., 2015; Wertz et al., 2017; Greenbaum et al., 2018; GRAVITY Collaboration et al., 2019; Wang et al., 2020; Wahhaj et al., 2021; Wang et al., 2021; Sepulveda, Bowler, 2022), and now masses Brandt et al. (2021a).

Precision astrometric monitoring, orbit fitting, and dynamical modeling have found that the planets follow nearly coplanar orbits with relatively low eccentricity (0 to around 0.25) (Soummer et al., 2011; Bergfors et al., 2011; Currie et al., 2012; Sudol, Haghhighipour, 2012; Pueyo et al., 2015; Maire et al., 2015; Konopacky et al., 2016; Zurlo et al., 2016; Wertz et al., 2017; Wang et al., 2018; GRAVITY Collaboration et al., 2019). The orbits appear to form a near-resonant chain, with factor of two period multiples 1b:2c:4d:8e (e.g. Maire et al., 2015; Wang et al., 2018; Goździewski, Migaszewski, 2020). This is fortunate, since dynamical modeling shows that few configurations exist that are stable over millions of years besides these resonant chains.

These same orbital models show that the system’s stability is increasingly tenuous if the inner three planets have masses much above $8 M_{\text{jup}}$, which matches the masses of approximately $7 M_{\text{jup}}$ derived through bolometric luminosity and evolution models. Recently, Brandt et al. (2021a) combined previous sets of stable orbits from Wang et al. (2018) with careful modeling of the Hipparcos-Gaia proper motion anomaly to estimate the masses of c, d, and e as $9.6^{+1.9}_{-1.8} M_{\text{jup}}$. This dynamical mass measurement is in slight tension with the results from orbital stability and atmosphere modeling; however, any additional planets in the system would impact this dynamical mass measurement. In their orbit modeling Goździewski, Migaszewski (2014) found some stable configurations that extend the resonant chain down to a fifth inner planet of up to roughly $6 M_{\text{jup}}$ near 7.5 AU (1e:3f) or 9.5 AU (1e:2f), and a larger “generally stable” region for test particles below $\sim 6 - 7$ AU. Brandt et al do consider such a fifth planet, and place 3σ detectable mass limits at roughly $5.5 M_{\text{jup}}$ between 3 and 5 AU, $6 M_{\text{jup}}$ between 5 and 7 AU, and $7.5 M_{\text{jup}}$ near 9 AU, but these limits do not consider how the space of stable orbits used to fit the mass may change by adding a fifth planet.

Many groups have undertaken extensive direct imaging searches for additional planets in the system(e.g. Currie et al., 2014; Maire et al., 2015; Wahhaj et al., 2021). The most sensitive constraints on the mass of an additional inner planet come from Wahhaj et al. (2021) in Y, J, H (IFS) and K bands (IRDIS). Using BT-Setl models (Spiegel, Burrows, 2012) and assuming an age of 30 Myr, they place 5σ upper limits of $3.6 M_{\text{jup}}$ at 7.5 AU and $2.8 M_{\text{jup}}$ at 9.7 AU.

These limits still leave room in the semi-major axis–mass parameter space where a fifth inner planet could hide; however, there are significant challenges with further improving our sensitivities. Current observations at such separations are limited by quasi-static speckles (Marois et al., 2003). These speckles produce a non-Gaussian noise distribution that is highly correlated over time and sensitivity improves poorly

with increasing integration time. Observations thus use angular differential imaging (ADI; [Marois et al., 2006a](#)), at times in combination with spectral differential imaging (SDI; [Walker et al., 1999](#); [Racine et al., 1999](#); [Marois et al., 2000](#)), and reference star differential imaging (RDI; [Wahhaj et al., 2021](#)). These greatly improve sensitivity, but require observations to be scheduled near when the system transits the meridian or have a suitable reference star nearby. Compounding this issue, is orbital motion. Planets at smaller semi-major axes have much shorter orbital periods according to Kepler’s third law. At a separation of ~ 5 AU, a planet in this system would move fast enough that observations taken more than a few months apart would start to blur the planet. All told, this means that considerable integration time is required and that time is challenging to schedule within a few-month window necessary to freeze orbital motion.

Combining images in the presence of orbital motion was previously considered in [Males et al. \(2013\)](#) and [Nowak et al. \(2018\)](#). These approaches consider “de-orbiting” in that images are transformed and stacked to counteract orbital motion. In doing so, they find that they are able to increase the SNR of a faint candidate despite orbital motion. One challenge with these approaches is that it becomes difficult to quantify the significance of such a detection in a way that includes uncertainty in the candidate’s orbit. A second challenge is that the flux of a candidate can vary freely between epochs even if it is not consistent with later data. This later point may be why [Males et al. \(2013\)](#) find an increasing false positive rate with increasing orbital motion.

Separately, a large debris disk first noted in the star’s spectral energy distribution (SED; [Sadakane, Nishida, 1986](#); [Zuckerman, Song, 2004](#); [Rhee et al., 2007](#)) and then described by [Su et al. \(2009\)](#) lies beyond the known planets and extends outwards to perhaps as far as 1000 AU ([Matthews et al., 2014](#)). The disk is only marginally resolved and models do not yet constrain the inner edge, with estimates varying from 104_{-12}^{+8} AU ([Wilner et al., 2018](#)) to 145 AU ([Booth et al., 2016](#)) or 170 ± 40 AU ([Faramaz et al., 2021](#)). These works and the additional dynamical studies of [Goździewski, Migaszewski \(2018\)](#) and [Geiler et al. \(2019\)](#) consider several scenarios for what mechanism may have sculpted the inner edge of the disk, one of which is an additional outer planet between $0.1 M_{\text{jup}}$ and $3 M_{\text{jup}}$. According to [Faramaz et al. \(2021\)](#) the best limit on an additional outer planet in this regime is $1.25 M_{\text{jup}}$ by [Maire et al. \(2015\)](#); however, the contrast curves presented in that work end at 70 AU of projected separation. Further out, [Close, Males \(2009\)](#) set a lower limit of $\sim 3 M_{\text{jup}}$ between $\sim 200 - 600$ AU, yet it appears no lower limits have been published on additional outer planets between b and the start of the outer debris disk, or ~ 100 to 150 AU. Though it will not access the inner region

Table 4.1. Observations

Date (UT)	Mask	Integration (min)	FoV Rotation (°)	Seeing: DIMM, MASS, WRF (")
2009-08-01	none	53	167	NA
2009-10-31	none	34	158	0.6, 0.2, 0.8
2009-11-01	none	72	162	0.8, 0.2, 1.4
2010-07-21	corona400	26	156	0.8 - 0.5 - NA
2017-07-07	corona400	37	172	0.7, NA, 0.4
2017-07-11	corona400	57	179	NA, NA, 0.4
2017-07-12	corona400	90	179	NA, NA, 0.4
2017-07-13	corona400	73	178	0.5, 0.4, NA
2017-07-14	corona400	68	178	0.6, 0.4, NA
2020-08-23	none	62	169	0.6, 0.5, NA
2020-08-25	none	54	165	1.8, NA, NA
2020-08-27	none	43	170	0.8, 0.2, NA
2020-10-07	none	60	170	0.4, 0.2, NA
2020-11-17	none	53	178	0.8, 0.5, NA
2021-07-08	none	52	170	0.5, 0.2, NA
2021-07-09	none	65	173	0.5, 0.15, NA
2021-07-10	none	49	173	0.4, 0.2, NA
2021-07-11	none	57	172	0.7, 0.7, NA
Total				
12 yr baseline		1005	3069	

Note. — Observations grouped by year. The integration column gives the total science exposure time not including calibrations and overheads. Seeing information is summarized from the Mauna Kea Weather Center Archive where available. All sequences were captured with the L' filter using NIRC2 in narrow mode.

of HR 8799 with standard coronagraphic imaging, JWST is poised to place exquisite constraints on outer planets in this regime.

To search for these proposed additional inner and outer planets, we performed an extensive L' imaging campaign at Keck with NIRC2, re-processed archival NIRC2 data using direct S/N optimization, and used a joint Bayesian model of planet orbits and photometry to search for planets despite significant orbital motion. We present limits on the existence of any additional planets as well as a modest SNR candidate at approximately 4 – 5 AU worthy of further study.

4.2 Observations and Processing

4.2.1 Observations

To search for additional planets in the system, we conducted a campaign of deep L' imaging at Keck using the NIRC2 instrument (PI: K. Matthews) in 2017, 2020, and 2021.

Table 4.2. SNR in combined images

Year	SNR _b	SNR _c	SNR _d	SNR _e
2009	50	64	33	21
2010	28	48	27	13
2017	92	80	62	22
2020	56	90	75	27
2021	36	62	72	28

Our observations for this campaign totalled 14 quarter-nights. We took observations in pupil tracking mode so that the field of view rotated during each sequence, but the speckle pattern remained fixed. This allowed us to process the data with angular differential imaging (ADI; [Marois et al., 2006a](#)) to suppress the halo of quasi-static speckles ([Marois et al., 2003](#)). Each observation was scheduled such that HR 8799 would transit the meridian roughly half-way through the sequence.

We chose L' imaging since it balances the favourable contrast of young planets at longer wavelengths with the need to access tight inner working angles, as well as limiting noise from the thermal background.

The observing strategy evolved over the course of the campaign. The 2017 epoch was captured with the 400 mas diameter Lyot coronagraph, Differential Atmospheric Refraction (DAR) correction set to acquisition and track, and included dithering away from transit to improve background subtraction. For the 2020 epoch, we observed without a coronagraph and with limited dithering. This is because initial reductions showed that the quasi static speckles near the edge of the mask were not as stable as those in non-coronagraphic datasets and because we found that the speckle pattern was not stable between dither positions. For example, see the $\sim 3.5\times$ improvement in contrast between the otherwise very similar 2017 coronagraphic and 2020/2021 non-coronagraphic epochs near 200-300 mas separation (Figure 4.2 and Table 4.2.1). We did not switch to using the new optical vortex (PI: K. Matthews; [Serabyn et al., 2017](#)) installed in 2015 as it was available only in shared risk mode. Finally, for the 2021 dataset, we also chose to set DAR (Differential Atmospheric Refraction) to acquisition only and not acquisition and track to reduce the number of optics moving during our observations. Our 2021 dataset shown in Figure 4.1 achieved the deepest contrast close to the star. We captured background images at the start and/or ends of each sequence to reduce the thermal background. The exposure time of each sequence was adjusted to avoid saturating the first Airy ring of the stellar PSF. Individual exposures were




Figure 4.1 Combined images from each epoch, processed with SNAP and cropped to show the inner planets. The images are shown on a linear scale with correct throughput, normalized using the flux of planet d. Values outside the linear colourscale are shown in grey. Note that when using SNAP, the forward processed images should not be used for evaluating the noise. Close to the star, the PSFs of any potential planet can produce secondary positive and negative ghosts at other position angles. Instead, we calculate contrast curves using matching backwards-rotated reductions. The top panels are comprised of archival data, while the bottom panels are the new campaign. Significant orbital motion is visible between epochs. The full images including planet b are listed in Figure 4.18.

co-added by the detector to create 30-80s exposure images depending on conditions. Higher cadence observations better capture the moment-to-moment variation in the stellar PSF; however, they come at the expense of considerable dead time after each

image. For each sequence, we captured unsaturated non-coronagraphic images at the beginning and ends of each NIRC2 sequence to use as planet PSF templates and contrast calibrations.

4.2.2 Archival data selection

To this campaign, we added additional data from the Keck archive dating back to some of the first sequences taken of the system. We considered L' sequences captured by members of our collaboration. We hoped this data would increase our sensitivity, particularly to a planet whose orbit might have appeared closer to the star during our main campaign. Combined with our dedicated observing campaign, this brought the total integration time, not including overheads, bad frames, sky backgrounds, and other calibrations, to 16.75 hours of on-source data at L' . The sequences taken in 2009 include artifacts at wide separations due to nodding that we exclude from our models of the outer system.

Besides this data, the Keck archive contains on the order of a further 8 hours of L band observations from other researchers. We did not include these sequences due to the large manual effort required to reduce one-off observations captured with varying observing strategies and in some cases, unsuitable choices of focal plane mask that obscure the inner planets. While not complete, our sample contains on the order of 65% of all L' data that has been recorded of HR 8799 by NIRC2. If we assume that the SNR of the combined observations grows with the square root of the total exposure time, then reducing all remaining observations could in theory increase the SNR by up to 25%.

Future work could additionally combine data from other wavelengths and observatories. For instance, the Large Binocular Telescope's (LBT) LMIRCam has similar capabilities to NIRC2 at L band. Observations at other wavelengths e.g. by GPI and SPHERE could also be added, however, using these observations would require us to assume the color of the planets we search for *a-priori*.

For a full listing of the sequences used in this work, see Table 1.

4.2.3 ADI Reduction

Angular Differential Imaging (ADI) is a powerful technique for suppressing quasi-static speckles, but least-squares based algorithms (Lafrenière et al., 2007b; Marois et al.,




Figure 4.2 5σ contrast limits at each of the five L' epochs on a log-log scale. Blue circles mark the average separations and photometry from our results. The contrast achieved close to the star has improved over time, so later epochs have more weight in our modelling.

2014; Soummer et al., 2012) suffer from worsening contrast and planet self-subtraction at very small separations. We therefore developed a new technique for reducing differential imaging sequences called direct S/N optimization (Thompson, Marois, 2021). This technique offers improved contrast close to stars by solving a system of quadratic equations maximizing S/N , rather than linear equations minimizing noise. Additionally, it optimizes stacks of multiple images in a sequence simultaneously to reduce correlated residual noise. See the above reference for more information and a comparison of direct S/N optimization to LOCI on one of the 2020 sequences presented in this paper.

We used the Signal to Noise Analysis Pipeline (SNAP) to calibrate, align, and

reduce each sequence. We pre-processed all images by subtracting darks, flat fielding with Ks band dome flats, and subtracting sky backgrounds where available. We applied distortion corrections using the solutions of [Yelda et al. \(2010\)](#) and [Service et al. \(2016\)](#) for data taken after the 2015 NIRC2 servicing. We then high-pass filtered the images to further suppress thermal background noise using a 25 pixel standard deviation. This step has the side effect of suppressing any diffuse emission in the system, though we don't expect to detect any emission from the inner or outer debris disks in these observations. We aligned the NIRC2 images using an iterative cross-correlation procedure since NIRC2 does not possess "satellite spots", off axis faint copies of the stellar PSF used for astrometric and photometric calibration ([Sivaramakrishnan, Oppenheimer, 2006](#); [Marois et al., 2006b](#)). We aligned each image against the unsaturated PSF template and then stacked to create a master. We then cross-correlated each image against this master to improve their alignment, stacked them to create a new master, and then repeated the procedure a further two times. Finally, we rejected bad frames using a correlation threshold (usually 1-5 images per sequence).

We reduced the data using SNAP multi-target S/N optimization with batches of 10 images, and an optimized number of included reference images for each subtraction region. No parameters of the reduction were changed between sequences to prevent human bias. The S/N optimization procedure does not include a rejection distance / exclusion zone or other adjustable aggressiveness parameter; all images are used in the optimization including those in which the planet PSF overlaps.

As will be described in the following section, our models assume that the input images have well-calibrated planet throughput. The SNAP pipeline is throughput preserving for point sources as long as the instrument's PSF does not deviate significantly from the unsaturated templates captured before or after the sequence. The S/N optimization algorithm does not produce self-subtraction for the peaks of point sources, and over-subtraction is prevented by the use of optimization, buffer, and subtraction regions.

For each sequence, we also produced a matching, "backwards rotated" image in which the rotation direction of the ADI sequence was reversed. These backwards rotated images have the same noise distribution as the normal images, but do not contain any significant signal from the planets. This allows us to calculate a contrast curve for each sequence that is unbiased by planets. We verified that contrast curves of the backwards rotated sequences matched the regular sequences between planets b, c, d, and e where any additional planets are very unlikely to orbit.

Finally, we combined the processed results of each sequence using a contrast-weighted median stack of images taken less than 3 months apart¹⁷. During this period, any planet on a circular orbit with semi-major axis greater than or equal to 5 AU would move less than $\frac{1}{3} \lambda/D$. Conveniently, this resulted in one combined image for each year HR 8799 was observed. The SNRs of planets b, c, d, and e in each of these combined images are presented in Table 4.2.1.

4.2.4 Photometric calibration

Since NIRC2 does not have satellite spots for astrometric and photometric calibration, we reduced the data in units of contrast relative to the star, as measured by unsaturated images taken before and/or after the sequence. However, variable weather conditions and AO performance during a sequence between the saturated and unsaturated images led to an additional photometric error on the order of 10% between sequences which is greater than what we would expect from random variation at the recovered SNR. To compensate, we measured the flux of planets b, c, d, and e in each epoch and re-scaled the images to the SNR weighted average flux. This correction improved the photometric variation between epochs to the expected level, e.g. a 2% variation for planet c which is close to what we would expect for a planet with SNR of 60-90.

To calculate mass we convert from contrast to absolute magnitude using an apparent magnitude of 5.3 for HR 8799 in the L' filter (Marois et al., 2010b).

4.3 Modelling

During the 12 year baseline of our dataset, any planet with a semi-major axis less than ≈ 3500 AU would move more than $1\lambda/D$. This means that the data cannot be naïvely stacked to improve our sensitivity as the signal of a planet would not be aligned between epochs. Since we wish to consider a very large parameter space of inclined and eccentric orbits, simple approaches like rotating and scaling the images would not be effective.

We therefore create a probabilistic model of the system and jointly model the orbits and photometry of the planets. This combines the process of detecting candidates with orbit fitting (e.g. Blunt et al., 2020) without the intervening step of extracting candidate

¹⁷SNAP processes groups of 10 frames into individual reduced images to reduce correlated noise. We stacked these individual processed images across nights within the three month window

astrometry at each epoch. Since astrometry is extracted from images, it follows that if we can model images directly, then extracting astrometry and photometry as intermediate products is not necessary.

Crucially, this allows us to detect the signal of a planet that is too faint to see in a single epoch even if it has moved considerably between images. In fact, planets with arbitrarily low SNRs per image can grow to detectable levels given a sufficient number of epochs. The sensitivity and limitations of this method will be expanded upon in a dedicated publication.

We consider nine parameters:

- M , the total mass of the system;
- Π , the parallax of the system;
- a , the semi-major axis;
- e , the eccentricity;
- i , the orbital inclination;
- Ω , the longitude of ascending node;
- ω , the argument of periapsis;
- τ , the time of periastron passage following the convention of Orbitize! (Blunt et al., 2020);
- and L , the L' flux-ratio.

The physical, geometric, and orbital parameters define a unique Keplerian orbit through our images and a position at each epoch. We assume that the planet has the same photometry over all observations. Following Ruffio et al. (2018) we consider the log-likelihood of a planet having those parameters as

$$\log \mathcal{L} \propto \sum_i \frac{1}{2\sigma_{x,i}^2} (L^2 - 2LL_{x,i}) \quad (4.1)$$

where i is the epoch, x is the computed position at epoch i , $L_{x,i}$ is the measured photometry extracted at position x from the image i , and $\sigma_{x,i}^2$ is the variance in that photometry. The schematic in Figure 4.3 illustrates this procedure.




Figure 4.3 Schematic showing how we model planets across epochs with orbital motion. The solid black line shows the path of a hypothetical planet around a star, with a given set of parameters $(a, i, e, \Omega, \omega, \tau, M, \Pi)$. We calculate the position x in each image i using Kepler’s laws, and measure the flux at that location as $L_{x,i}$. We then calculate $\sigma_{x,i}$, the uncertainty in $L_{x,i}$, by measuring the contrast in an annulus at that separation from the matching backwards rotated image. $L_{x,i}$ is then compared to the model parameter L which is the same for all epochs using the likelihood function given by equation N.

To extract the photometry $L_{x,i}$ efficiently from the images, we perform a noise-weighted convolution by an Airy disk of $1\lambda/D$. Weighing the convolution locally at each pixel by the surrounding contrast is essential so that the peak SNR occurs at the location of the planet, rather than offset in the direction of lower noise. In regions with strongly sloped contrast curves (e.g. planet e) this correction prevents a position bias of up to 30 mas. We then look up the photometry at each coordinate using a bi-linear interpolation. We estimate the variance $\sigma_{x,i}^2$ using the contrast at that separation in each image. We measured the contrast curves themselves using matching backwards rotated ADI reductions so that the signals of any planets do not bias the contrast.

This likelihood function assumes that our convolved images are maximum likelihood estimates of planet photometry in the presence of Gaussian noise, and that contrast curves extracted from the backwards rotated noise maps provide good estimates of the variance in that estimate at each pixel ¹⁸. [Ruffio et al. \(2018\)](#) provides a derivation of this likelihood function and shows how non gaussianity does not significantly effect Bayesian upper limits by that definition. For detection thresholds on the other hand,

¹⁸One could compute the likelihood by injecting negative fake planets into the raw data; re-performing the post-processing for each epoch, position, and photometry; and examining the residuals; however, this would increase the compute time by a very large constant factor and is not computationally feasible at this time.

section 4.4.5 of this paper discusses how we correct for mildly non-Gaussian noise.

This approach is similar to [Mawet et al. \(2019\)](#) in which direct images of ϵ Eridani are combined with radial velocity data; however, we do not include radial velocities (the planets of HR 8799 orbit almost face on from our perspective) but instead combine images from multiple epochs as in [Skemer, Close \(2011\)](#). It also shares some similarities in concept with those of K-stacker ([Coroller et al., 2015](#); [Nowak et al., 2018](#)) and for example the recent results of [Le Coroller et al. \(2022\)](#).

Note also that we do not use a PCA ([Soummer et al., 2012](#)) or matched filter ([Ruffio et al., 2017](#)) based data processing as in [Ruffio et al. \(2018\)](#) and [Mawet et al. \(2019\)](#), neglect the effects of distorted planet PSFs, and simply perform a noise-weighted convolution of our data by an Airy disk of $1\lambda/D$.

Since the purpose of this work is to detect or place limits on the mass of any additional planets rather than precise orbital characterization, we do not consider any systematic errors from the instrument pointing or North angle. Error in registration or North angle could bias the orbital parameters and reduce our ability to recover planets. Thankfully, the results of [Yelda et al. \(2010\)](#) and [Service et al. \(2016\)](#) indicate that the North angle and platescale of NIRC2 are very stable over time, varying less than 0.6° and 0.1 mas/px respectively between 2010 and the service in 2015.

We consider separate models for the four known planets, an additional outer planet, and an additional inner planet. Results pertaining to each planet are coloured consistently across figures.

The purpose of four known planet models is to confirm we can recover their photometry. These models will additionally confirm that there are no significant North angle offsets between epochs that could impact our ability to detect additional planets. We choose uniform priors for the angular parameters ω and τ for each planet and broad but informative priors based on previous work for the orbital plane of the planets, eccentricities, and the planets' L band flux-ratios. We adopt uniform priors on the planet semi-major axes but truncate them in order to separate the planet models and prevent them from each sampling all four modes of a single posterior. For the orbital planes, we chose wide Gaussian priors based on previous modelling of the outer planets' orbits and planetary radial velocities by [Ruffio et al. \(2019\)](#). This constrains the direction of motion along our line of sight which roughly halves the size of the parameter space to explore.

We select priors on the flux-ratio that require it to be greater than or equal to zero, but not less than zero. This is because ADI processing introduces dark wings around

Table 4.3. Model priors

Parameter	Prior Distribution	Notes
M	$1.52 \pm 0.15 M_{\odot}$	Baines et al. (2012); Wang et al. (2018); Konopacky et al. (2016)
Π	24.46 ± 0.05 mas	GAIA Collaboration et al. (2021)
i	$20.8 \pm 4.5^{\circ}$	Ruffio et al. (2019)
Ω	$89 \pm 27^{\circ}$	Ruffio et al. (2019)
ω	Uniform circular*	
τ	Uniform circular*	
b		
L_b	$1 \times 10^{-4} \pm 1 \times 10^{-4}, L_b \in (0, 1)$	Marois et al. (2008a)
a_b	Uniform(50,180) AU	Images masked outside of 130-200px separation
e_b	Uniform	
c		
L_c	$2 \times 10^{-4} \pm 1 \times 10^{-4}$	Marois et al. (2008a)
a_c	Uniform(30,55) AU	Images masked outside of 80-120px separation
e_c	Uniform(0,1)	
d		
L_d	$1 \times 10^{-4} \pm 1 \times 10^{-4}, L_d \in (0, 1)$	Marois et al. (2008a)
a_d	Uniform(20,80) AU	Images masked outside of 56-80px separation
e_d	Uniform(0,1)	
e		
L_e	$1 \times 10^{-4} \pm 1 \times 10^{-4}, L_e \in (0, 1)$	Marois et al. (2010b)
a_e	Uniform(8,20) AU	Images masked outside of 80-120px separation
e_e	Uniform(0,1)	
Outer		
L_f	Uniform(0, 10^{-5})	
a_f	Uniform(100, 160) AU	Images masked outside of 180-500px separation
e_f	Uniform(0,1)	
Inner		
L_f	Uniform(0, 10^{-2})	
a_f	Uniform(1,14) AU	Images masked outside of 9-30px separation
e_f	Beta(1.1, 5)	Low-moderate eccentricity

Note. — * Parameterized using the arctangent of two standard normal distributions.

point sources. If a point source follows a Keplerian orbit through our images, the dark wings will nearly follow this same orbit leading to spurious detections of negative planets and/or reducing the significance of a detection by introducing false uncertainty in its flux-ratio. Besides enforcing positivity, we expect the exact shape of priors on the flux-ratio to have little effect on the posterior since we have more than enough data for the likelihood to overwhelm the prior.

For the inner and outer planet models, we again adopt broad but informative priors on the orbital plane of the system and a Beta distribution to prefer low eccentricities. We adopt a uniform distribution for semi-major axis a between 1 and 14 AU which constrains our search to orbits closer in than planet e. Finally, for the flux-ratio between the inner planet and star we adopt both a Uniform prior between 0 and 10×10^{-2} and a log normal prior centered on the expected pixel values. Given the amount of data, we expect the posterior to be relatively insensitive to this choice of prior (an assumption we will verify in Section 4.4.3). For the angular parameter ω which has a uniform prior, we in fact sample from a pair of Gaussian distributions, ω_x and ω_y centered at zero, and calculate $\omega = \tan^{-1}(\omega_y, \omega_x)$. This has the same distribution as a uniform prior on ω , but allows the sampler to easily wrap around past 0 and 2π . We do the same for τ , but restrict it to a domain of $[0, 1)$ by dividing by 2π .

Finally, for all models we adopted a Gaussian prior on host mass following [Konopacky et al. \(2016\)](#) and [Wang et al. \(2018\)](#) based on interferometric measurements of the stellar radius by [Baines et al. \(2012\)](#). For parallax Π of the system, we use a tight Gaussian prior from Gaia's EDR3 data release. We describe the priors further and sources for all parameters in Table 4.3.

4.3.1 Detection and Limits

To evaluate detections, we marginalize over all of the orbital parameters and inspect the flux-ratio (L) posterior. This histogram represents the posterior distribution of the planet's photometry regardless of the orbital parameters. Its central value is the best estimate of the planet's photometry, and the width of the distribution is the uncertainty in that estimate. A sharp peak that is well-separated from zero indicates a detection.

For ease of comparison with other methods, we summarize this posterior by calculating the SNR as the median divided by half the 84th – 17th percentile distance. This is analogous to the standard SNR calculation used to evaluate point source detections in single images, however it is marginalized over all plausible orbits making

it a stricter measure. In a traditional SNR map each point source is considered separately, even if there were, for example, many significant point sources with varying brightness. Here, this SNR is testing the hypothesis that there is a *single* planet with consistent flux.

We can also use a fully Bayesian approach to assessing detections. We can proceed by evaluating the relative probabilities of two models: a model of a planet with a finite brightness (M_1) and a model where there is no light from the planet (M_0). The “Bayes factor” is then the ratio of the marginal likelihoods of the data given the models times a prior on which model is more likely:

$$B_{M_1:M_0} = \frac{P(M_1|D) P(M_0)}{P(M_0|D) P(M_1)}$$

We adopt the standard prior that both models are *a-priori* equally likely, that is $P(M_1) = P(M_0)$. The Bayes factor between two models gives the relative probability of M_1 compared to M_0 . For example, if the Bayes factor $B_{M_1:M_0} = 10$ then given this data, it is ten times more likely that there is a planet than not.

Often Bayes factors are challenging to calculate numerically since MCMC based methods only produce samples proportional to the posterior density. However, in our case our two models are said to be “nested” since M_1 reduces to M_0 for $L = 0$. Since our prior on L is uncorrelated with the priors on the orbital parameters, we can calculate the Bayes factor between these nested models using the Savage-Dickey density ratio (Dickey, 1971; Koop, 2003). This allows us to calculate the Bayes factor $B_{M_0:M_1}$ as

$$B_{M_1:M_0} = \frac{P(L = 0)}{P(L = 0|D)}$$

That is, the prior on L evaluated at (or near, for numerical purposes) $L = 0$ divided by the marginal posterior of L evaluated at that same location.

A benefit of this approach is that we can assess detections without assuming the marginal flux-ratio posterior is Gaussian. This could occur even with perfectly Gaussian noise in the images if there is a strong dependence of the flux-ratio on one or more orbital parameters like a or e . Note however that the model M_1 itself assumes that the residual noise after post-processing is approximately Gaussian, an assumption we will verify in Section 4.4.5. Numerically, this approach requires posterior samples where the flux-ratio is near zero, which for significant detections may be far in the tails of the distribution. This calculation would require an impractically large number of samples for very significant detections that are well separated from zero. Of course, estimating

the Bayes factor accurately in order to assess a highly robust detection is somewhat moot. In any case, the use of nested sampling may allow one to reliably calculate the Bayes factor between these two models for such significant detections. For an example of using nested sampling and a Bayes factor to evaluate the presence of an exoplanet from aperture masking interferometry data, see [Blakely et al. \(2022\)](#).

A limitation of both the SNR and Bayes factor approaches to evaluating detections is that they consider only a single planet. If there were, for example, two planets in the data with different orbits and/or flux-ratios to the star, the overall SNR would suffer.

In these instances, the posterior must be examined more closely to disentangle the planets. For this situation, we attempt to reproduce a classical direct imaging SNR map which is not a direct output of this analysis method. To do so, we first select a given date. Natural choices could be the average date or date of the best input dataset. At this date, each planet drawn from the posterior has a well-defined spatial position which we calculate. We then examine the marginal flux-ratio distribution of the posterior draws that fall in each given pixel on that date. Finally, we calculate the SNR of these samples in the same way as above. These maps are built using the posterior so only include the most *a-posteriori* likely orbit and photometry parameters (no samples are available to perform this analysis along very unlikely orbits). We refer to these as “photometric accuracy” maps to distinguish them from traditional SNR maps.

Finally, we present Bayesian upper limits following [Ruffio et al. \(2018\)](#) by calculating the 84th percentile of this same marginal flux-ratio posterior. In our results, we present the Bayes factor in addition to our analog of the classic SNR whenever this calculation is feasible. The Bayes factor is arguably a more robust quantity, but it lacks a history and established threshold conventions of the SNR in the context of direct imaging.

4.3.2 Sampling

Sampling from images is a difficult problem since the direct imaging likelihood function has strong modes at the locations of planets and speckles surrounded by large flat regions where the likelihood is negligible. Compared to fitting orbits to astrometry points, sampled orbits that fall far away from any modes do not have gradients that pull subsequent samples towards a mode. Additionally, fitting near face-on orbits is challenging due to degeneracies between Ω , ω , and τ in our chosen parameterization. These effects combine to require small step sizes to explore near the mode and many steps to adequately explore the tails of the posterior. Our numerical tests showed that

simulated systems with companions injected with an overall SNR of approximately 3-6 are the most computationally demanding to sample, since the sampler must explore multiple peaks (SNRs closer to 1 are associated with posteriors that are broad and relatively smooth and above 6 the peak dominates).

Knowing this would be computationally demanding, we programmed our model in Julia (Bezanson et al., 2012). We used forward-mode automatic differentiation (Revels et al., 2016) to calculate the gradients of the log-posterior with respect to each parameter. This allowed us to use a higher order sampler without manually deriving gradients for each model. We sampled from the posterior using the No U-Turn Sampler (Hoffman, Gelman, 2014) variant of Hamiltonian Monte Carlo, as implemented in AdvancedHMC.jl (Xu et al., 2020). Internally, we used Bijectors.jl (Fjelde et al., 2020) to transform all constrained variables and priors to unconstrained distributions.

For each model, we used multiple independent chains with a maximum tree-depth of 13 steps. We initialized each independent chain by drawing 50,000 samples from the priors and picking the sample with the highest posterior density. After adapting the step size and mass matrix for 30,000 iterations and discarding the first 150,000 iterations, we ran each chain in increments of 150,000 iterations until converged. We thinned each chain by its auto-correlation time and further discarded occasional chains that failed to adapt and remained at their initial parameters.

We checked for convergence by inspecting trace plots, ensuring all parameters had Effective Sample Size (ESS) greater than 100, and that the Gelman, Rubin and Brooks diagnostic was less than 1.3 (Gelman, Rubin, 1992; Brooks, Gelman, 1998).

The code used in this paper is available in the Julia packages [PlanetOrbits.jl](https://seffal.github.io/PlanetOrbits.jl)¹⁹ and [Octofitter.jl](https://seffal.github.io/Octofitter.jl)²⁰.

4.3.3 Stability

We further evaluate the results of our inner planet model by testing them for orbital stability using the Python REBOUND WHFast integrator (Rein, Tamayo, 2015) to integrate sets of 5-planet orbits for 100,000 years. We determine whether the orbits are stable using the Mean Exponential Growth factor of Nearby Orbits (MEGNO; Cincotta et al., 2003) factor. The orbits that present a MEGNO of ≤ 2 for 100,000 years, which would indicate stability up until that time, are then integrated on a range of semi-major

¹⁹<https://seffal.github.io/PlanetOrbits.jl/dev/>

²⁰<https://seffal.github.io/Octofitter.jl/dev/>

axis and eccentricity for 1 Myr to find possibly stable neighbouring orbits.

We evaluate 5-planet solutions in two ways. First, we start by sampling 5-planet configurations from the posteriors for b, c, d, e and the candidate planet f. Then, we also analyzed the stability of the candidate planet f from our posteriors with the b,c,d and e planet parameters from the V_d model presented by (Goździewski, Migaszewski, 2014).

4.4 Results

In this section, we begin by describing the results of our models of the four known planets. Then, we describe our results of applying the same approach to search for any additional outer and inner planets.

4.4.1 Recovery of Known Planets

The four known planets are easily recovered by our models at very high significance. Figure 4.4 shows orbital paths drawn from the posterior of the four planet near-resonant model. Figure 4.5 shows the marginal posteriors of the photometry and selected orbital elements for each model. Despite not extracting astrometry points as an intermediate step, the orbital posteriors are consistent with previous studies (Wang et al., 2018).




Figure 4.4 Visualization of 1500 orbit draws from the posteriors of the b, c, d, e, and inner planet models over-plotted on the combined 2021 epoch. Directly modelling the photometry in our images allows us to simultaneously detect the known planets and produce orbital posteriors that agree with previous fits to extracted astrometry.

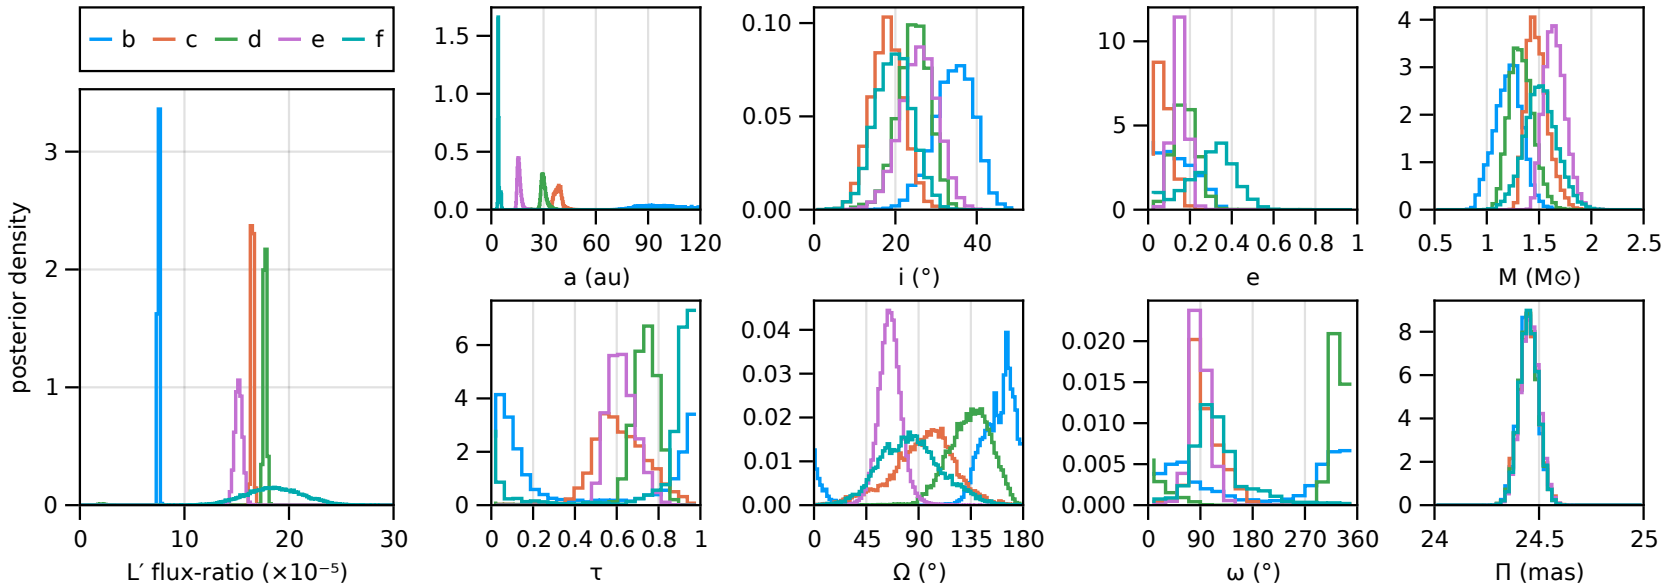


Figure 4.5 Marginal posteriors of photometry and orbital elements compared between the four known planets and the inner planet model.

The marginal L' histograms in Figure 4.5 show that the photometry posteriors are approximately Gaussian distributed and are consistent with previously reported values (Marois et al., 2008a, 2010b). The distributions are well separated from zero which indicate robust detections. We find SNRs of the planets b, c, d, and e from the combined observations of 122, 154, 125, and 32 respectively. Compared to the SNR measured at each individual epoch, these SNRs are greater by roughly a factor of $\sqrt{5}$ which is the ideal improvement in SNR we would expect by combining five observations limited by Gaussian noise.

Compared to previous studies of the systems orbital configuration using for example GPI (Wang et al., 2018), SPHERE (Wertz et al., 2017), and GRAVITY (GRAVITY Collaboration et al., 2019), these longer wavelength observations have reduced astrometric precision. We therefore present our orbital solutions here to show that the orbits derived by directly modelling the photometry in the images are consistent and to build confidence in this approach before applying it to search for additional unseen companions.

4.4.2 Limits on additional outer planets

Now that we have established that our technique recovers the four known planets, we turn our search outwards to search for any additional outer planets between 100 and 150 AU. For this search, we ignored the 2009 epoch due to nodding artifacts beyond the orbit of planet b. The posterior of this model contained one peak that we identified as a bright artifact on the far top edge of the 2017 epoch. We dropped samples with orbits that intersected that artifact before proceeding with our analysis.

We find no evidence for a fifth outer planet above an 85th percentile L' contrast of 4.6×10^{-6} . The overall SNR from this posterior is 0.7 and the log Bayes factor for an additional outer planet given this data is -1.6. Figure 4.7 presents our sensitivity as a function of semi-major axis as well as histogram of the full marginal photometry posterior. Using the system ages of Sepulveda, Bowler (2022) and COND models of Baraffe et al. (2003), this corresponds to a 1σ mass limit of roughly $0.9 M_{\text{jup}}$.

4.4.3 Evidence for a fifth inner planet

We now consider an additional planet interior to planet e. Against the full dataset, the model finds a mode close to the star. Figure 4.8 shows the marginal photometry

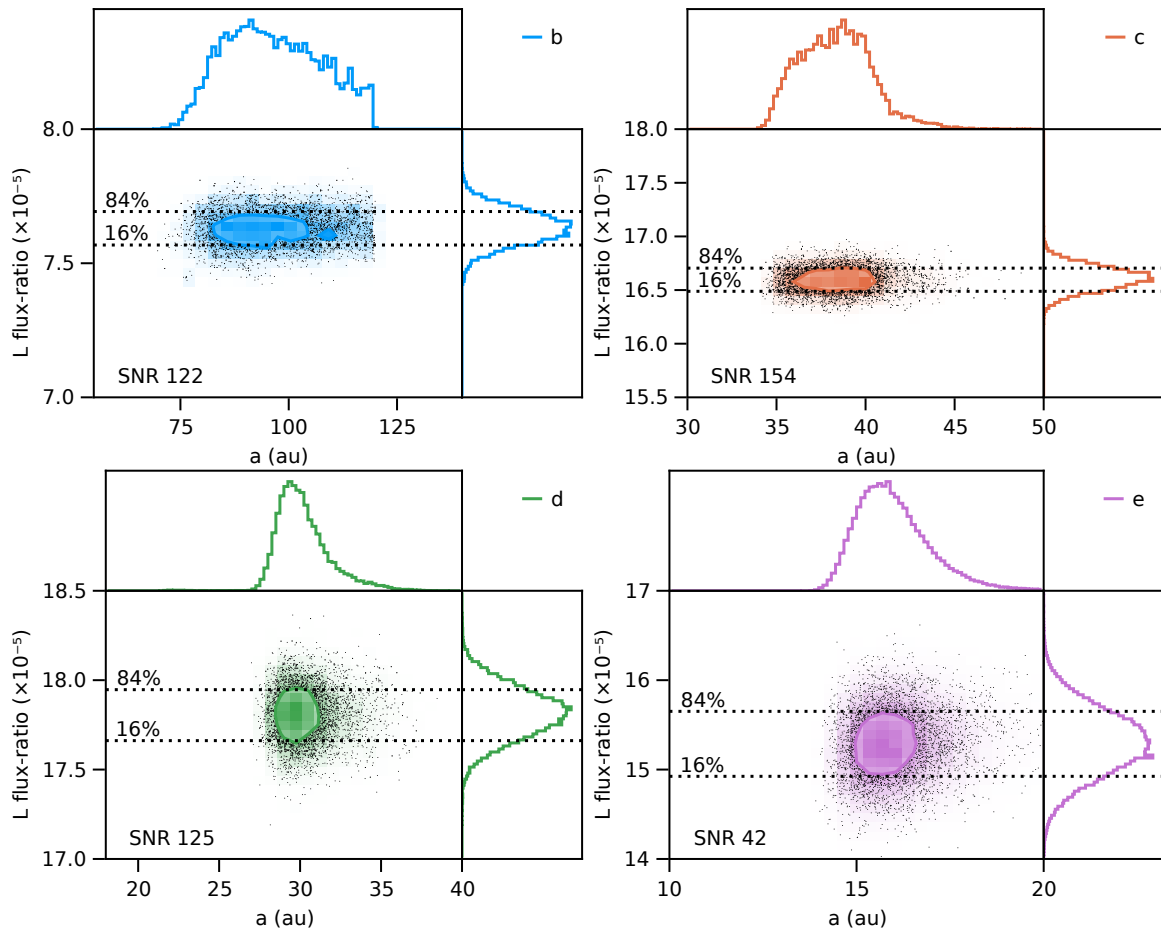


Figure 4.6 Marginal photometry vs. semi-major axis for planets b, c, d, and e. All four are detected at very high significance and SNRs greater than in any individual epoch. The scales are different between each panel.

vs. semi-major axis posteriors of our single inner planet model. The joint posterior of the inner planet model is multi-modal, with 2-3 families of plausible orbits that all pass through the same locations in 2020 and 2021 (Figures 4.9 and 4.16). This multi-modality is a result of weak photometric constraints in some epochs, leading the model to consider several plausible locations with consistent flux as the peaks found in higher quality epochs. Nonetheless, the marginal photometry posterior is roughly Gaussian and well constrained.

The marginal photometry posterior is well separated from zero, with a mean that is very similar to the photometry of planets c, d, and e. We find an SNR of 6.9 and a log Bayes factor of 18.

The marginal semi-major axis posterior is centred at 4.5 AU but is cut off below ~ 4 AU




Figure 4.7 Flux-ratio and semi-major axis marginal posteriors for an additional outer planet between planet b and the start of the outer debris disk. The black line gives the 84th percentile flux as a function of semi-major axis. No planet is detected and our sensitivity is not strongly dependent on any of the orbital parameters within this range of semi-major axis.




Figure 4.8 Flux-ratio and semi-major axis posteriors for an additional inner planet with a 1σ contour over-plotted. The dashed lines show 16% and 84% percentile limits. The recovered photometry is consistent with that of planets c, d, and e.




Figure 4.9 **Top:** Images from each of the five epochs. **Bottom:** 2D marginal position posterior at each epoch. This shows the *a-posteriori* most likely positions for an inner planet at each epoch. Note that the posterior is calculated jointly from all images; this figure merely visualizes it at individual epochs. Even epochs where the location is poorly constrained can still contribute to the SNR of the model as a whole by reducing the space of plausible orbits and planet photometry.

due to detector saturation so we cannot place a firm lower limit. The 85th percentile upper limit is 4.6 AU. The chains for this posterior are available at [10.5281/zenodo.6823071](https://zenodo.org/record/6823071).

Figure 4.9 shows images from our five epochs with positions calculated from orbits drawn from the posterior. The model places it NW of the star in 2021 and WNW in 2020. There are three plausible modes in 2017, and the location in earlier epochs is not well constrained.

A full corner plot showing the values, uncertainties, and covariance between all nine parameters is available in Figure 4.17.

In Figure 4.11, we draw orbits from the posterior and calculate their positions in 2021. We see that the model prefers a single location for the planet in 2021 North-North-West of the star. When we look at the median photometry of samples from the posterior that fall in this pixel, we find they are all roughly 2×10^{-4} in units relative contrast. Finally, when we look at the spread of the sample photometry, the “photometric accuracy”, we find again a cluster of SNR 5-9. These maps support the posterior being consistent with a single object rather than two or more.

Returning to the literature, various candidate point sources have previously been reported. The candidate reported by [Maire et al. \(2015\)](#) $3 - 4\sigma$ $0.2''$ due South of the

star in 2013, is not consistent with our results. Likewise, no compatible point sources are visible in the shorter wavelength (YJH) IFS data of [Wahhaj et al. \(2021\)](#). Their IRDIS data in K band (closer to L') does show a low SNR point source North-West of the star in 2019; however, again, its location may not be consistent with our orbital posterior to $> 1\sigma$. On the other hand, the point source reported by [Currie et al. \(2014\)](#), 4σ North-North-West of the star in 2012, may have a roughly correct position angle if the candidate's semi-major axis is ≈ 4.7 , though with a slightly greater separation. We did not include these sequences in our initial data selection (Section 4.2.2), so to add them afterwards knowing it may or may not have a compatible point source could introduce confirmation bias.

This is an intriguing result, but given the novelty of this analysis technique, lack of confirmation at other wavelengths and instruments, and points discussed later in this analysis, we do not yet consider this a robust detection.

4.4.4 Contribution of the 2021 Epoch

Figure 4.9 shows that the preferred location of model is the most localized in 2021. This was our best epoch and consists of observations taken over four nights. We now examine the impact of this epoch on the model in greater detail.

To begin, we injected a negative model planet into the 2021 sequences raw data prior to SNAP reduction. We placed the planet at the posterior expected position and photometry calculated from the full model of all five epochs. We then re-reduced the data with SNAP to produce Figure 4.10.

The panels in that figure show the image before and after injecting the negative planet model. The centre panel giving the difference between these reductions shows the flux removed by the planet model. The model reproduces much of the irregular structure directly around the star including a bright lobe opposite to the expected position. If the candidate is real, those effects can be understood as artifacts of the SNAP algorithm, tight separation, and rapid, near- 180° field of view rotation as the system transits the meridian from Maunakea. In fact, the angle between the two bright spots is just under the average field of view rotation in the 2021 sequences. Encouragingly, the opposite bright lobe is not picked up by the inner planet model meaning that a planet at that location and brightness in 2021 is not consistent with the other epochs. That said, the injected planet model is not a perfect match for the candidate. The negative side-lobes (artifacts of any ADI reduction) are somewhat mismatched with the model




Figure 4.10 The 2021 epoch before and after subtracting a model planet injected at the expected location and photometry of the single inner planet model fit to all 5 epochs. The area interior to 100 mas is excluded from the processing due to detector saturation. The bottom right image shows the flux removed by the negative planet injection. The structure is more complex than the typical dark wings from ADI processing due to the high FoV rotation and tight separation.

having a darker sidelobe to the East than and lighter to the West than the candidate. As a consequence, some flux remains to the North-East of the star after subtraction.

Next, we also run our model on the original unmodified data but fully exclude the 2021 epoch. In this case, we again find no detection though interestingly the earlier four epochs still predict a spot of high posterior density within $\sim 1\lambda/D$ of the location found in the full model. These results do not mean that the earlier four epochs do not contribute to the SNR of the planet candidate. They still contribute by ruling out large swaths of the orbital parameter space and the bright area South of the star in 2021.

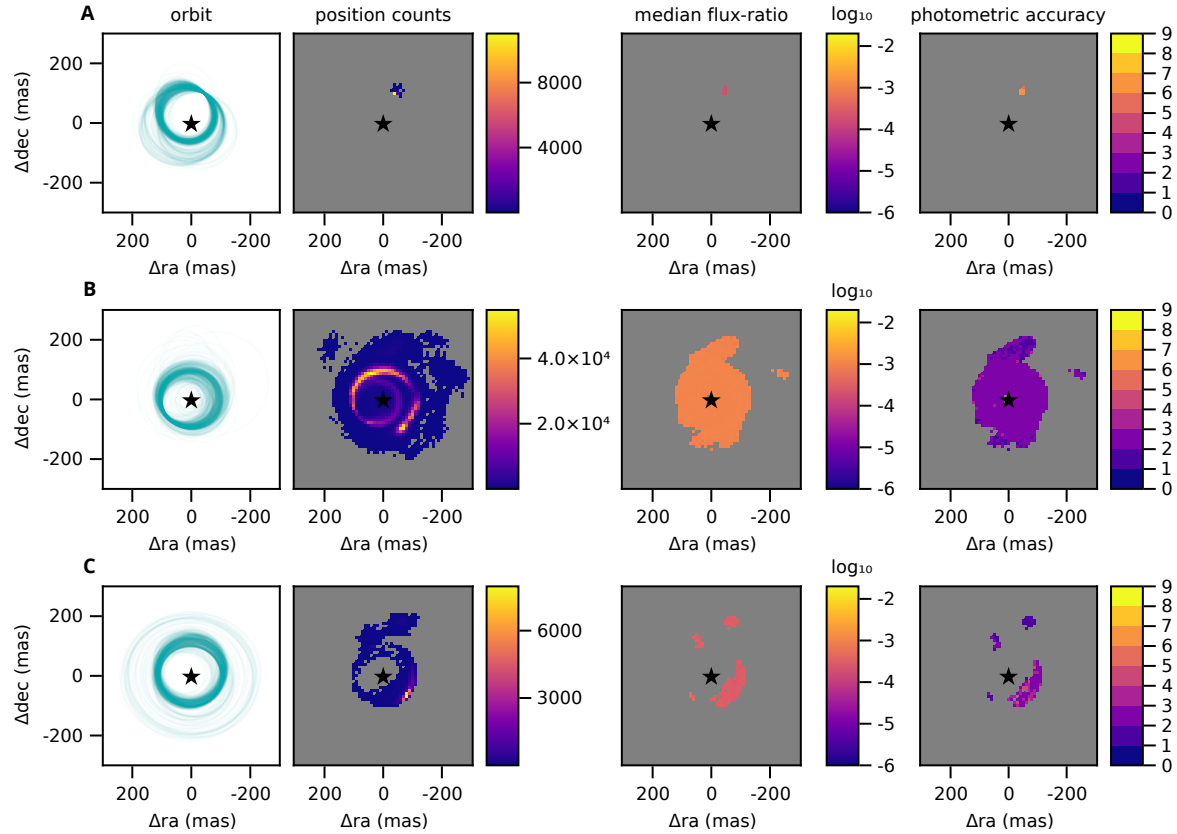


Figure 4.11 Comparison of the inner planet models against three different datasets. **A.** An inner planet model applied to the full dataset. **B.** Same as A, but ignoring the best epoch (2021). **C.** Same as A, but after subtracting the best fitting model found by A from the 2021 epoch. The four panels show orbits drawn from the posterior, posterior density of planet position calculated at the 2021 epoch, median photometry if the planet were at that position in 2021, and photometric accuracy at those positions i.e. the SNR of a planet if it were on an orbit that would pass through that pixel in 2021. Pixels are left blank where there are only a negligible number of samples.

4.4.5 Noise Distribution and Sample Size

When evaluating a candidate in direct imaging, we should consider non-Gaussian noise and small sample statistics. Both act to increase the False Positive Fraction (FPF) and reduce our confidence in a detection. We calculate penalty factors for both of these effects on 3σ and 5σ Gaussian equivalent FPFs.

First, we consider non-Gaussian noise. Following [Marois et al. \(2008b\)](#), we compare the distribution of our data at the separation of the inner planet candidate to a Gaussian and estimate its effect on both detection thresholds. In the top panel of [Figure 4.12](#), we plot the confidence levels of a Gaussian distribution and of standardized pixel data from the inner $1\lambda/D$ annulus of all five backwards-rotated photometry maps. We extrapolate our data by fitting an exponential, and find that to reach a confidence level equivalent to a Gaussian at 3σ and 5σ , we should penalize our SNR by factors of 1.07 and 1.35 respectively. The bottom panel of [Figure 4.12](#) shows the log Probability Density Function (PDF) of a Gaussian distribution and an Empirical Probability Density Function (EPDF) of standardized pixel data from the inner $1\lambda/D$ annulus of the backwards-rotated photometry maps. By expanding the Gaussian by these factors, it fully encompasses our data at a z-score of 3. The backwards rotated noise maps do not contain any pixels at z-scores beyond 3.5 (the forwards reduction does of course, since it contains the signal of the candidate) but the extrapolation appears valid and conservative. These factors are relatively small indicating that the residual noise in each epoch is close to Gaussian distributed. This is not surprising due to the central limit theorem since, besides 2010, each epoch is a stack of 3-5 sequences with uncorrelated noise.

Next, we consider the effect of small sample statistics near the star. The contrast curves underlying our model at each epoch are calculated from as few as five independent samples at a separation of roughly $2\lambda/D$. If each epoch contributed to the SNR of our model in equal measure, then we might consider the noise sample to be five times larger, reducing the effects of small sample statistics. However, the 2021 epoch contributes significantly to the overall figure. It's not yet clear how to correct this model for small sample statistics, but we can take a conservative approach by considering the final SNR to come only from a single image (instead of 18 sequences) and applying the correction factors of [Mawet et al. \(2014\)](#), Table 1, at $2\lambda/D$. This gives penalty factors of 1.35 and 2.2 for 3σ and 5σ FPFs respectively.

Combining these two sets of penalty factors, we should in fact apply 4.3σ and 14.9σ thresholds to reach large sample size, Gaussian equivalent FPFs for 3σ and 5σ

respectively. The candidate does easily meet the 3σ threshold but clearly falls well below a 5σ threshold for detection. This correction is perhaps overly pessimistic but reflects our goal to communicate this signal is a candidate worthy of additional study rather than an unambiguous detection.

4.4.6 Mass and Proper Motion Anomaly

The photometry of the inner planet candidate is consistent with planets c, d, and e, and brighter than b. Without photometry at other wavelengths, we therefore assume that it would have a similar mass to the inner planets c, d, and e.

The masses of the HR8799 planets have been estimated using several approaches. From the beginning, luminosity modelling of the planets has suggested masses of approximately $7 M_{\text{jup}}$ for the inner planets (Marois et al., 2008a). Orbital stability analyses by Wang et al. (2018) support the inner planets having masses up to approximately $7 M_{\text{jup}}$, though small islands of stability may exist for high masses. Sepulveda, Bowler (2022) on the other hand, model the orbits of the planets in order to constrain the dynamical mass of the star. By combining this stellar mass with stellar and planet evolution models, they find that the inner planets c, d, and e likely have masses in the range of $4.1 - 7.0 M_{\text{jup}}$ where age is the dominant contributor to the uncertainty. In contrast to these estimates, Brandt et al. (2021a) find somewhat higher masses of $9.6^{+1.9}_{-1.8} M_{\text{jup}}$ by combining the stable orbits found by Wang et al. (2018) with the Hipparcos-Gaia Catalog of Accelerations (HGCA; Brandt, 2021) which calibrates the Hipparcos catalogue against Gaia's (GAIA Collaboration et al., 2021). The addition of a fifth massive planet would alter the solution space for all three methods incorporating orbital dynamics; however, the method based on proper motion anomaly would be the most impacted. We now examine if this potential addition could account for the slight tension between these mass estimates.

Taking a more basic approach than that of Brandt et al. (2021a), we model the proper motion anomaly of the star by assuming that the Gaia and Hipparcos missions each independently measure the position and instantaneous proper motion 25 times spaced equally throughout their respective missions. We then compare these quantities against observations for a four planet and five planet model. For the four planet model, we use the stable orbital parameters of Goździewski, Migaszewski (2014) and fixed masses of $5 M_{\text{jup}}$ and $7 M_{\text{jup}}$ for b, and c, d, e respectively. For the five planet model, we use the same parameters but add a fifth inner planet from our orbital posterior with a mass of




Figure 4.12 **Top:** Confidence levels of a Gaussian distribution and empirical confidence level from the inner $1\lambda/D$ annulus of our images extrapolated to $1 - 10^{-6}$ using an exponential function. We increase our SNR thresholds to account for this slight deviation from Gaussian noise. **Bottom:** Empirical log PDF of standardized pixel intensities compared to a Gaussian model, and Gaussians scaled to cover the tails of the distribution.




Figure 4.13 Stellar astrometric motion predicted at the Hipparcos and Gaia epochs for the stable 4 planet V_d solution of [Goździewski, Migaszewski \(2014\)](#) (blue square) and with the addition of the fifth inner candidate (coloured by the semi-major axis of the candidate). The red error bars show the stellar astrometric motion at the Hipparcos and Gaia epochs as calculated by the HGCA ([Brandt, 2021](#)). The black marker shows the long term proper motion from the HGCA calculated from the difference in position between both epochs. The planets c, d, e, and the candidate f are assumed to have masses of $7 M_{\text{jup}}$, while b is assumed to have a mass of $5 M_{\text{jup}}$.

$7 M_{\text{jup}}$.

Figure 4.13 shows the result of this comparison. Proper motion anomaly at the Hipparcos epoch is consistent with either the four or five planet model, but agreement at the Gaia epoch is significantly improved by the addition of the candidate at the most likely semi-major axis.

Though far from conclusive, this analysis demonstrates that the inclusion of the inner candidate could fully account for the observed proper motion anomaly if the candidate and inner planets c, d, and e have masses close to $7 M_{\text{jup}}$. This is slightly at odds with the [Brandt et al. \(2021a\)](#) result since they use the HGCA data to conclude that any additional inner planet more massive than $\approx 6 M_{\text{jup}}$ between 3 and 8 AU are unlikely. The discrepancy could be due to their more sophisticated modelling or the lower masses used in our experiment for the known planets.

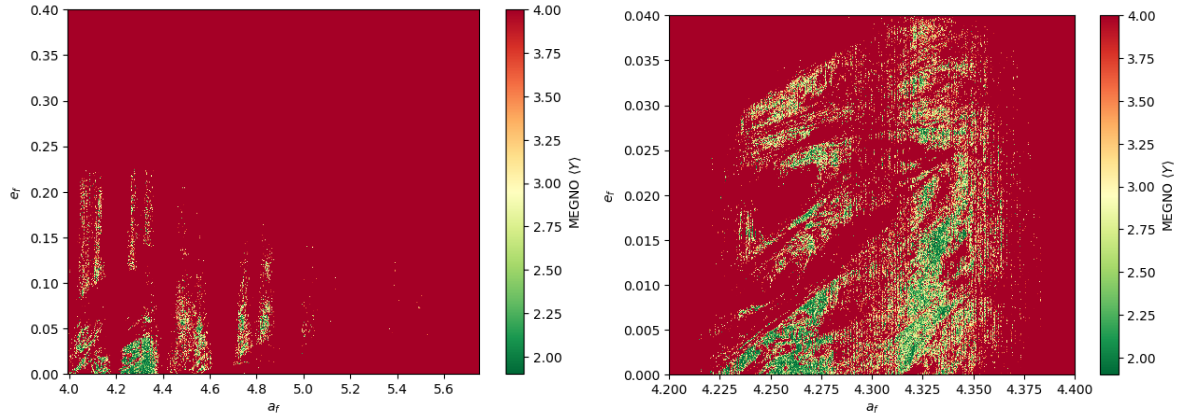


Figure 4.14 The MEGNO stability map for the coplanar ‘f’ candidate using the parameters from Table 4.4.7, for an integration time of 0.5 Myr (left) and 1.0 Myr (right). Most of the configurations become unstable (MEGNO > 2) within this integration period.

4.4.7 Stability

We now consider how the addition of an inner planet would impact the stability of the system. For these tests we adopt a smaller than realistic mass for the candidate inner planet of $3.75M_{\text{Jup}}$ to ease the search for stable orbits. For our 5-planet models obtained entirely from the posteriors presented on Figure 4.5, we found that no configuration is stable for the age of the system of 10–23 Myr (Sepulveda, Bowler, 2022). When using the V_d Model parameters presented on Goździewski, Migaszewski (2014) for planets b, c, d and e we find one unconstrained orbit for ‘f’ that remains stable for 0.4 Myr. A grid search around this sample’s semi-major axis and eccentricity revealed similar orbits that are stable for up to 0.75 Myr. A grid search over co-planar orbits found regions that remain stable for up to 1.5 Myr. Most of the other configurations found from the posteriors become unstable within the first 0.5 Myr in our N-body simulations. These configurations of stable orbits are presented in Table 4.4.7 and MEGNO stability maps of the co-planar orbits are presented in Figure 4.14. These simulations do yet not explore changing the masses or orbital parameters of planets b, c, d, and e. It is likely that regions of greater stability could be found if these parameters were also explored in a future analysis.

Table 4.4. Stable orbits

Planet	Mass (M_J)	a (AU)	e	i ($^\circ$)	Ω ($^\circ$)	ω ($^\circ$)	τ ($^\circ$)
b	6.527020	68.597249	0.017425	27.502	63.953507	37.631852	0.872678
c	11.868008	39.486207	0.054102	27.502	63.953507	89.946387	0.403048
d	7.178005	25.705213	0.137796	27.502	63.953507	33.186128	0.139618
e	6.298260	15.660910	0.168239	27.502	63.953507	110.074917	0.902362
coplanar 'f'	3.75000	4.325000	0.068600	27.502	63.953507	145.767982	0.833691
non-coplanar 'f'	3.75000	4.510300	0.036997	14.986	82.304917	119.041162	0.837036

Note. — Stable five planet solutions. The non-coplanar solution is drawn from our orbital posterior and is stable for 0.75Myr. The coplanar solution was found using a search grid near the orbital parameters of the posterior of the inner planet model. The stellar mass is $1.716162 M_\odot$. The parameters of planets b, c, d, and e are from the [Goździewski, Migaszewski \(2014\)](#) V_d model.




Figure 4.15 85th percentile upper limit of the marginal L band flux-ratio posterior for an additional fifth inner planet (black). The turquoise lines show the same, but after subtracting a planet model from the raw data of 2021 epoch (C in Figure 4.11). The grey lines show the traditional 1σ projected contrast of individual epochs from backwards rotated reductions.

4.4.8 Sensitivity to planets besides the 4-5 AU candidate

Figure 4.15 presents our sensitivity to additional inner planets besides the candidate near 4-5 AU. We show the 85th percentile of the marginal L flux-ratio posterior, conditioned on different ranges of semi-major axis. We show the sensitivity for both the regular dataset, and the dataset in which we subtracted the candidate from the 2021 epoch. Ignoring the candidate presented above, assuming the system ages of [Sepulveda, Bowler \(2022\)](#) (10-23 Myr) and extrapolating the COND models of [Baraffe et al. \(2003\)](#), this translates to 1σ upper mass limits of $4.3 M_{\text{jup}}$ and $3.0 M_{\text{jup}}$ for planets with orbits that have time-averaged projected separations of 150 mas and 250 mas respectively. Note that this sensitivity versus average projected separation is not directly equivalent to contrast versus separation at a given epoch as is usually quoted in the literature. In general, the sensitivity is a function of all orbital parameters. For instance, it improves with higher eccentricity at small semi-major axes since such a planet would spend more time away from the star.

4.5 Conclusion

In this paper we presented a deep targeted search in the HR 8799 system for additional planets using data at 3.8 microns.

- We observed HR 8799 for 14 quarter nights with NIRC2 between 2017 and 2021
- We further gathered a further 4 quarter nights from the Keck archive for a 12 year baseline.
- We processed the data using direct S/N optimization ([Thompson, Marois, 2021](#)) to improve contrasts at very small separations
- We used the Hamiltonian Monte Carlo method to explore both the space of possible orbits and flux from planets, allowing us to search for planets in the presence of orbital motion. This approach differs from the K-Stacker ([Nowak et al., 2018](#)) technique since it requires the flux to be consistent across epochs, gives the covariance between the planet's flux and orbital parameters, and ultimately gives a detection that can be evaluated against uncertainties in both flux and orbital parameters.

Using these methods, we presented the results of our search:

- We demonstrated this technique by re-detecting the four known planets b, c, d, and e at very high SNR
- We presented limits on the flux of any additional outer planets between b and the start of the debris disk and did not find any additional outer planets above a 5σ significance threshold, or a mass limit of roughly $0.9 M_{\text{jup}}$.
- We found a modest SNR candidate interior to the orbit of planet e. This planet would have a semi-major axis of approximately 4-5 AU, and L-band photometry similar to the inner planets c, d, and e. According to the analysis [Sepulveda, Bowler \(2022\)](#) this would imply a mass of 4.1 to $7.0 M_{\text{jup}}$.

For this candidate planet, we considered several factors.

- We verified that subtracting a planet model from the raw data at the appropriate position and intensity in the best epoch and re-running the model removes the detection.
- We adjusted our detection thresholds to account for slightly non-Gaussian noise and a conservative treatment of the impacts of small sample statistics.
- We showed that the addition of a fifth inner planet could improve the agreement between the dynamical mass measurements of the planets and mass estimates/limits from atmosphere models and orbital stability modelling.
- We performed rejection sampling with REBOUND using the inner candidate orbital posterior combined with a four planet solution known to be stable. We were successful in finding five planet orbital solutions that were stable for 0.75Myr using a lowered planet mass. A grid search over nearby co-planar orbital parameters found small families of orbits that are stable for up to 1.5Myr.

Overall, we found that the inner candidate at SNR 6.9 easily met a 3σ equivalent FPF threshold (4.3σ) but does not meet a 5σ equivalent FPF threshold (14.9σ). This is primarily because of uncertainty in the contrast at each epoch due to the limited sample size at small separations from the star. We consider this evidence intriguing, but caution that these results fall short of a conclusive detection.

[Wahhaj et al. \(2015\)](#) presented the most sensitive limits on the K-band flux of an additional inner planet. They found a 5σ upper limit of 3×10^{-5} relative to the star at 100 mas separation. Adjusting this figure for small sample statistics in the same

was as our data gives an upper limit of 7×10^{-5} . As such, their non-detection of this candidate at K-band is a point against the candidate. If on the other hand the candidate is confirmed, this would give the candidate a very red colour compared to c, d, and e ($K - L > 2$). If real, such a red colour might be caused by the different environment much closer to the star. One possible explanation might lie in photochemical hazes as proposed for 51 Eri b (Zahnle et al., 2016; Macintosh et al., 2015) due to such a planet receiving more than an order of magnitude more light from the star than planet e.

Additional follow up observations would be necessary to confirm this candidate; however, assuming an ideal \sqrt{N} growth in SNR and a decreasing penalty for small sample statistics, this would require a further 12 quarter nights with NIRC2 of similar quality to 2021. Instead, the best chance at confirming or rejecting this candidate may come from upcoming instruments with improved contrast at 100-250mas. Followup observations from GRAVITY (GRAVITY Collaboration et al., 2017) might be possible, but would be challenging given the remaining uncertainty in the candidate's orbit. Another avenue that may be worth exploring is searching for the candidate with a fibre-fed spectrograph like the Keck Planet Imager and Characterizer (Delorme et al., 2021), though this again requires a well-determined orbit.

Regardless of if this candidate is confirmed, we demonstrated the utility of searching for planets in direct images by combining orbit modelling and planet detection. This approach could considerably loosen scheduling requirements when searching for rapidly moving targets like planets around Alpha Cent, in addition to making the best use of direct imaging archives.




Figure 4.16 Position posterior density for the inner planet model at different epochs. The position is shown on August 15th of each year between 2009 and 2024.

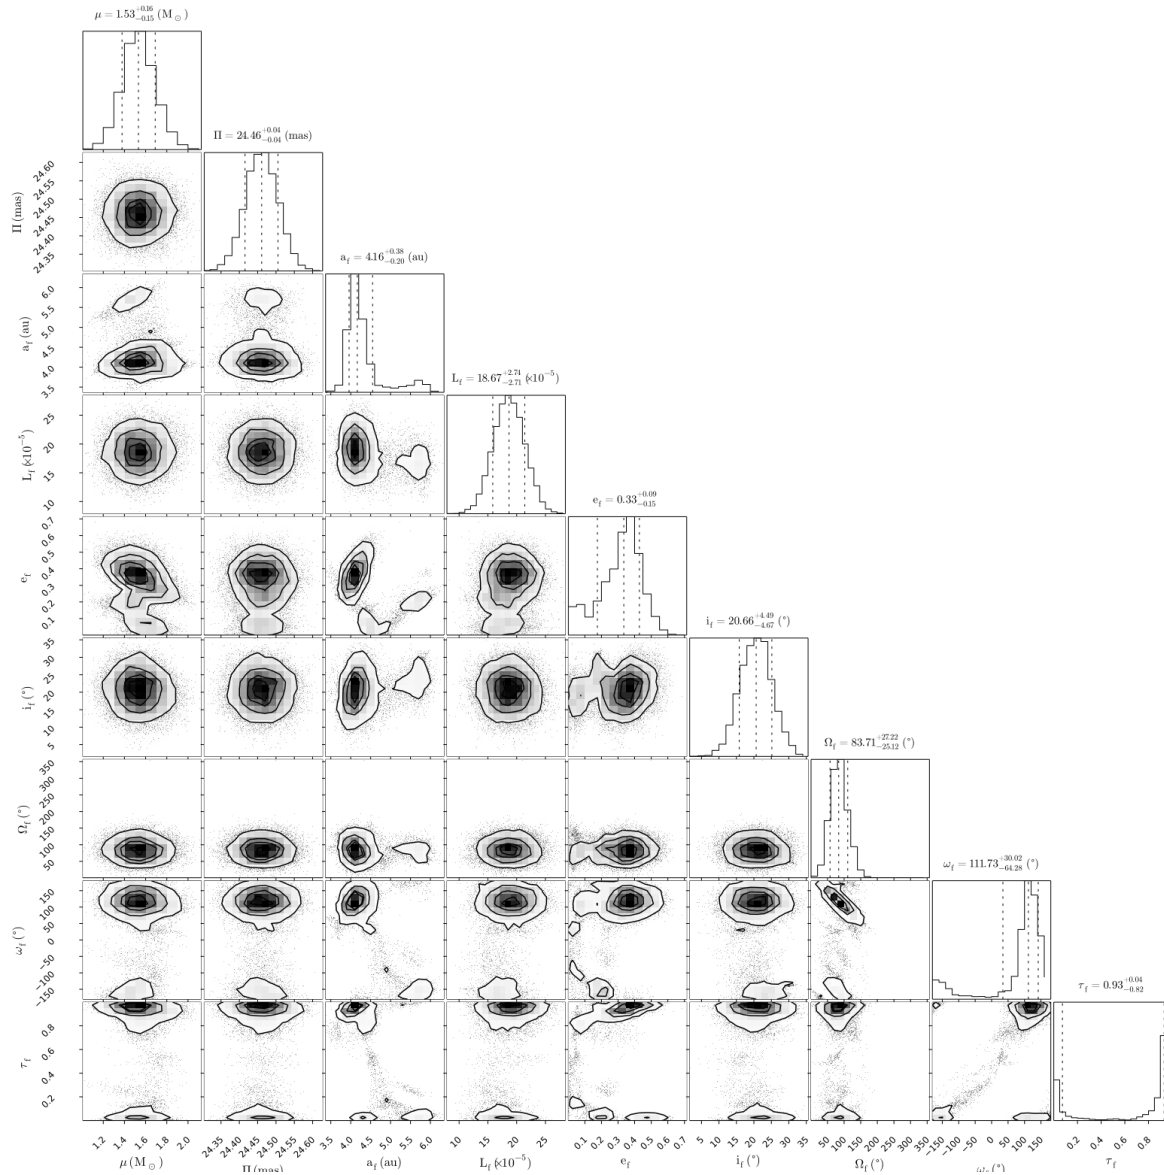


Figure 4.17 Corner plot showing the posterior of the inner planet model applied to all epochs. Note that the angular parameters i , Ω , ω , and τ are periodic. Chains are available at 10.5281/zenodo.6823071.




Figure 4.18 Full combined image from each epoch. The top two panels consist of archival data while the bottom three are from the new campaign. The artifacts south of the star in the 2009 image are from a chopping strategy used to subtract the thermal background. The approximate location of the inner edge of the outer debris disk is outlined in blue. These processed images are available at [10.5281/zenodo.6823071](https://zenodo.org/record/6823071).

Chapter 5

Performance of the Fast Atmospheric Self Coherent Camera and an Improved Measurement Algorithm

This chapter is based on material published in the Proceedings of the SPIE Volume 12185, Adaptive Optics Systems VIII; 121852C (2022). Authors: William Thompson, Christian Marois, Garima Singh, Olivier Lardière, Benjamin Gerard, Qiang Fu, and Wolfgang Heidrich. *My contribution to this work was as the lead author. I completed the software development, developed the control loops, and performed the data analysis. C. Marois is the PI for the NEW-EARTH Lab and supported the laboratory demonstration. G. Singh contributed to the text and the low order wavefront sensor loops necessary for this demonstration. O. Lardière designed and built the optics for the NEW-EARTH Lab. C. Marois, G. Singh, O. Lardière, and B. Gerard all contributed to discussions B. Gerard, Q. Fu, and W. Heidrich contributed the focal plane mask.*

5.1 Introduction

In order to directly image faint planets as they orbit their host stars, telescope instruments must correct the incoming wavefront with high accuracy. The two primary sources of wavefront error are atmospheric speckles, caused by our turbulent atmosphere's changing index of refraction, and instrumental, also known as quasi-static speckles

([Marois et al., 2003](#)). Once the wavefront is corrected, a Focal Plane Mask (FPM) and Lyot stop are typically used to block the light from the star while letting through as much light from the planet as possible.

The Self Coherent Camera (SCC) is a technique for focal plane wavefront sensing in the context of coronagraphy pioneered by [Baudoz et al. \(2005\)](#). Rather than merely discarding diffracted light at the Lyot stop, a small amount of the light is allowed to pass through a second much smaller aperture placed off-axis, referred to as the “pinhole”. This light creates a reference beam that propagates to the focal plane where it interferes much like a traditional double-slit experiment with any starlight that makes it past the FPM. Given that the pinhole is quite small, it acts as a filter that produces a near-perfect reference beam that only passes a small amount of low-order aberration.

The resulting fringe patterns between the main beam and the reference beam encode the relative phase offset between the two arms of the common-path interferometer. This information can be processed to extract the full complex electric field at the focal plane to enable active control and post-processing with Coherent Differential Imaging (CDI; [Guyon, 2004](#); [Baudoz et al., 2006](#); [Bordé, Traub, 2006](#); [Give'On et al., 2007](#); [Serabyn et al., 2011](#); [Sauvage et al., 2012](#); [Gerard et al., 2018a](#); [Potier, al. et, 2022](#)).

Since the light in the reference beam picked off by the FPM comes only from the star, the resulting interference pattern is only present between the reference beam and residual starlight. Light from a planet or disk travels through only one arm of the interferometer (the main pupil) so is not fringed.

A key improvement to this technique was developed by [Gerard et al. \(2018a,b\)](#): the Fast Atmospheric Self-Coherent Camera Technique (FAST). In FAST, the FPM is modified to direct the stellar core further off-axis and concentrate the light directly on the pinhole. This increases the light through the pinhole significantly to allow a better match between the intensities of residual speckles and the reference beam. The result is that fringe pattern can be detected in the same amount of time as the speckle itself, dramatically increasing the possible correction rate.

In this work, we present results of our laboratory testing of FAST: both active Deformable Mirror (DM) correction and CDI post-processing. We also present an algorithm for extracting SCC measurements that is numerically faster, tolerant of saturated pixels, and simpler to implement. This work is an important step towards the Subaru Pathfinder Instrument for Detection of Exoplanets and Removal of Speckles (SPIDERS; [Lardière, al. et, 2022](#); [Marois, al. et, 2022](#)) and the Gemini Planet Imager's CAL2 upgrade ([Marois et al., 2020](#)). Both of these instruments under development will

include an SCC with very similar configurations to the laboratory setup we present here. Achieving not only the best possible performance from the SCC, but also robust and reliable operation are critical to meeting the science goals of these projects.

5.2 Experimental Setup

We have been working to test FAST at the NRC Extreme Wavefront control for Exoplanet Adaptive optics Research Topics at Herzberg (NEW-EARTH) laboratory. The following describes our experimental setup in brief, and the full details are available in a previous conference proceeding by [Lardière et al. \(2020\)](#).

The Victoria Pathfinder for Exoplanet Research (VIPER) bench in the NEW-EARTH lab contains: an ALPAO 468 Deformable Mirror (DM), tilt-gaussian-vortex (TGV) FAST FPM, reflective Lyot stop with pinhole for FAST, Lyot-stop Low-Order Wavefront Sensor (LLOWFS; [Singh et al., 2014, 2015](#)), and C-RED2 science camera / focal-plane wavefront sensor. The TGV FPM is a fresnel design manufactured using additive lithographic fabrication by KAUST with four discrete layers. The bench also includes an imaging Fourier transform spectrograph to be described in future papers. The tests presented in this proceeding used an unpolarized white SuperK COMPACT light source from NKT Photonics with a 2% bandpass filter rather than a narrow laser. This reproduces the slight blurring of fringes in a realistic bandpass.

The LLOWFS arm is responsible for controlling low-order Zernike modes and keeping the star centred on the FPM. It sees all photons that are blocked by the coronagraph, so has access to significantly more light and accordingly can run at similar speeds to traditional Adaptive Optics loops. The C-RED2 is placed downstream of the Lyot stop where it functions as both the science detector and focal-plane wavefront sensor. The FPM is designed such that the reference pinhole is quite far off-axis. The resulting fringe pattern between main pupil and the reference beam is at high spatial (less than $1/4 \lambda/D$). This fully separates information in the fringes from any astrophysical signal.

The ALPAO-468 Deformable Mirror (DM) has a circular arrangement of actuators providing us with a $24\lambda/D$ diameter control region. In its current setup, the reflective Lyot stop blocks the two outer annuli of actuators. These actuators can therefore not be sensed by the SCC, but are still seen with limited sensitivity by the LLOWFS ([Singh, et al. 2022](#)). We operate the SCC in both full and half dark hole modes, allowing us

to correct phase errors across the entire control region or phase and amplitude errors across any one half of the control region.

We placed a Thorlabs optical chopper behind the reflective Lyot stop's reference pinhole. This allows us to record fringed and unfringed pairs of images. With the pinhole placed off-axis, it is possible to operate the SCC without a chopper; however, we found that operating on differential images improves the stability of system and allows us to safely increase the gain. The current optical chopper is limited to at least a 200Hz on/off chopping rate. The SPIDERS pathfinder described in [Lardière, al. et \(2022\)](#) and [Marois, al. et \(2022\)](#) will employ a custom-built optical chopper that can smoothly operate at lower speeds ([Johnson, al. et, 2022](#)).

5.3 VENOMS: A New Software Toolkit for Adaptive Optics Laboratories

In order to test the FAST SCC, we needed to first develop a software control system for the NEW-EARTH Lab. We needed the flexibility to test multiple different algorithms in a laboratory environment, while at the same time we needed to be able to run them at realistic speeds. This is because at very high contrasts, the instrument itself is a source of variation and turbulence. The loops needed to run quickly enough that this evolution would not negatively impact results and more closely mimic the environment and loop speed of a real adaptive optics system.

For this reason, I developed a new control system software called the Versatile and Efficient New-earth-lab Operating and Monitoring Software (VENOMS, [Figure 5.1](#)). I now describe the software architecture and performance of this system.

VENOMS was developed using the Julia programming language ([Bezanson et al., 2012](#)). I chose Julia for this project because it can approach C-like performance while presenting a convenient high level syntax for implementing control algorithms similar to Python and MATLAB.

A conceptual schematic of VENOMS is presented in [Figure 5.2](#). The system is driven off of a single text based TOML²¹ configuration file. Inside this configuration file, the user lists any components they wish to initialize as well as any global system parameters. VENOMS loads this configuration file on startup to determine the list of available components.

²¹<https://toml.io/en/>

VENOMS uses a trait-based component system. Components represent logical blocks in the system. Using traits, any object can declare itself to be a VENOMS component. All components must be assigned a unique name in the configuration file which serves as an identifier throughout the application.

At every screen refresh, the system loops over all active components and calls their `draw` method. This `draw` method is responsible for filling a graphical user interface panel representing that component.

Components can opt into additional functionality by applying the `Device` and/or `Loop` traits. The `Device` trait is a contract that indicates the component supplies a `connect` and `disconnect` method. Components implementing the `Device` trait are listed in a panel of available devices along with buttons for the user to connect and disconnect the device. The `Loop` trait indicates that a component has an associated event loop and supplies a `setup` and `runloop` method. If the user or system initiates a loop, it is added to a list of currently running loops listed in the user interface. The loop is wrapped with a `LoopController` which is responsible for launching and monitoring a thread dedicated to that function (technically, a Julia `Task`). A system of edge-triggered events are handled by the loop controllers. This allows individual loops to post results or wait for results from zero or more other loops.

In testing, I found that VENOMS was able to sustain loop speeds of well over 1000 updates per second. This far exceeds the requirements we had for testing the self coherent camera (SCC) but was useful for testing the low order wavefront sensor at roughly 600 Hz. Median system latency was measured to be less than 1 ms in the best cases. That said, the long tail (> 99.5%) latency was quite high, and potentially as high as 300 ms. This was caused by Julia's stop-the-world garbage collector which must run periodically to free unused memory. Future developments could move VENOMS from a multi-threaded to a multi-process application. Splitting the application into small processes would reduce the total memory (heap size) which must be swept by the garbage collector at any one time and not require all threads to pause at the same time.

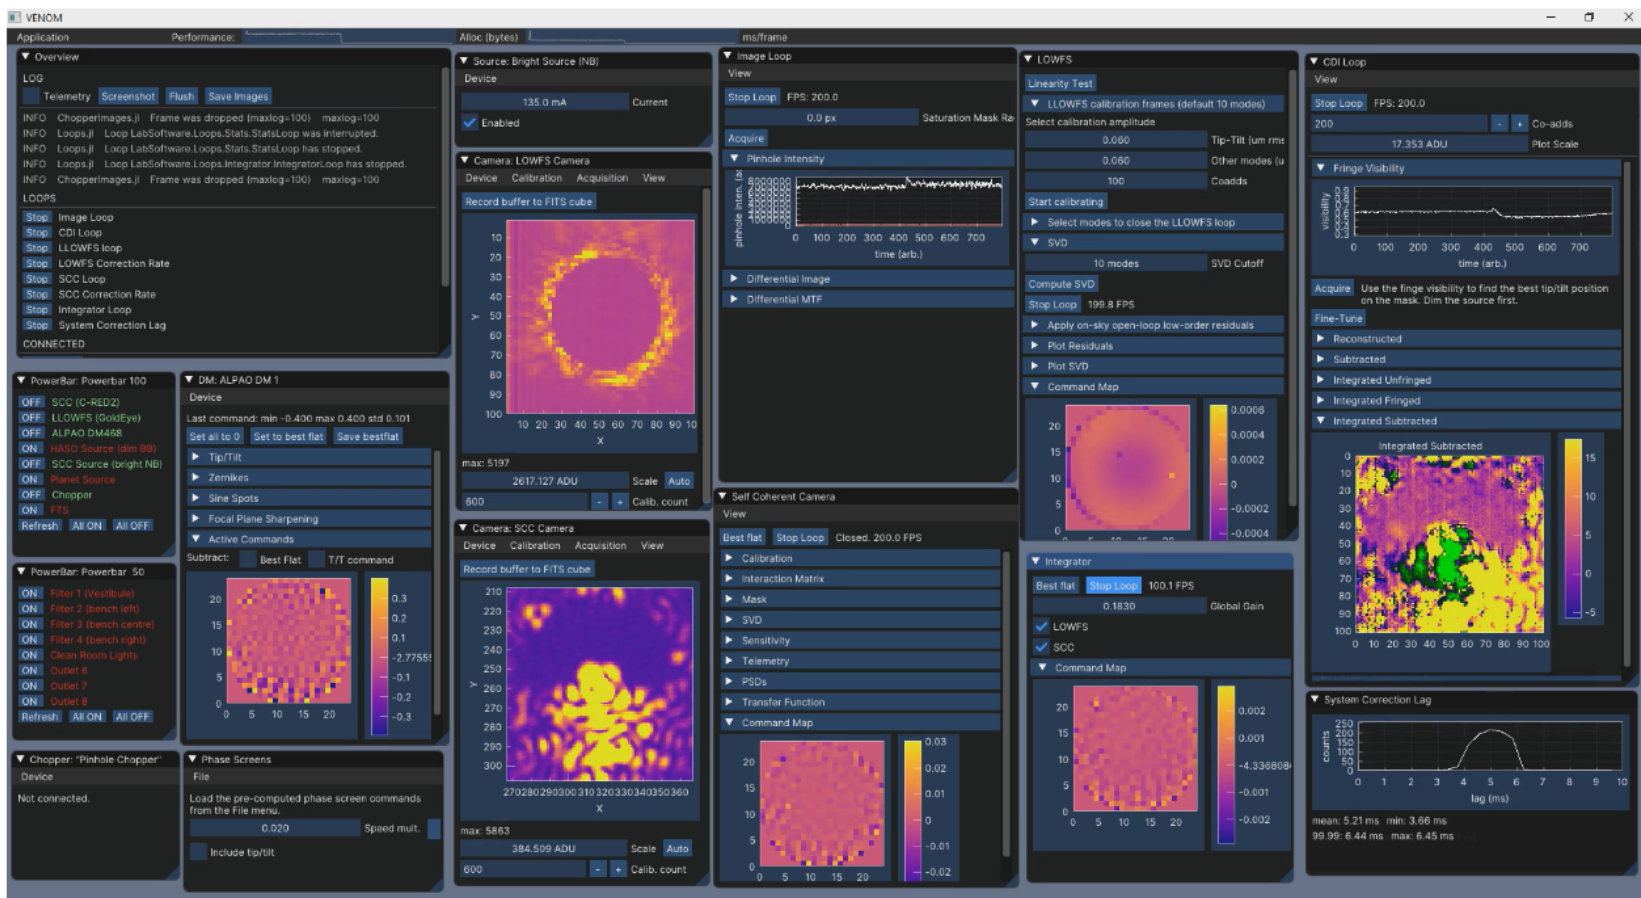


Figure 5.1 Screenshot of the VENOMS application showing closed loop operation.

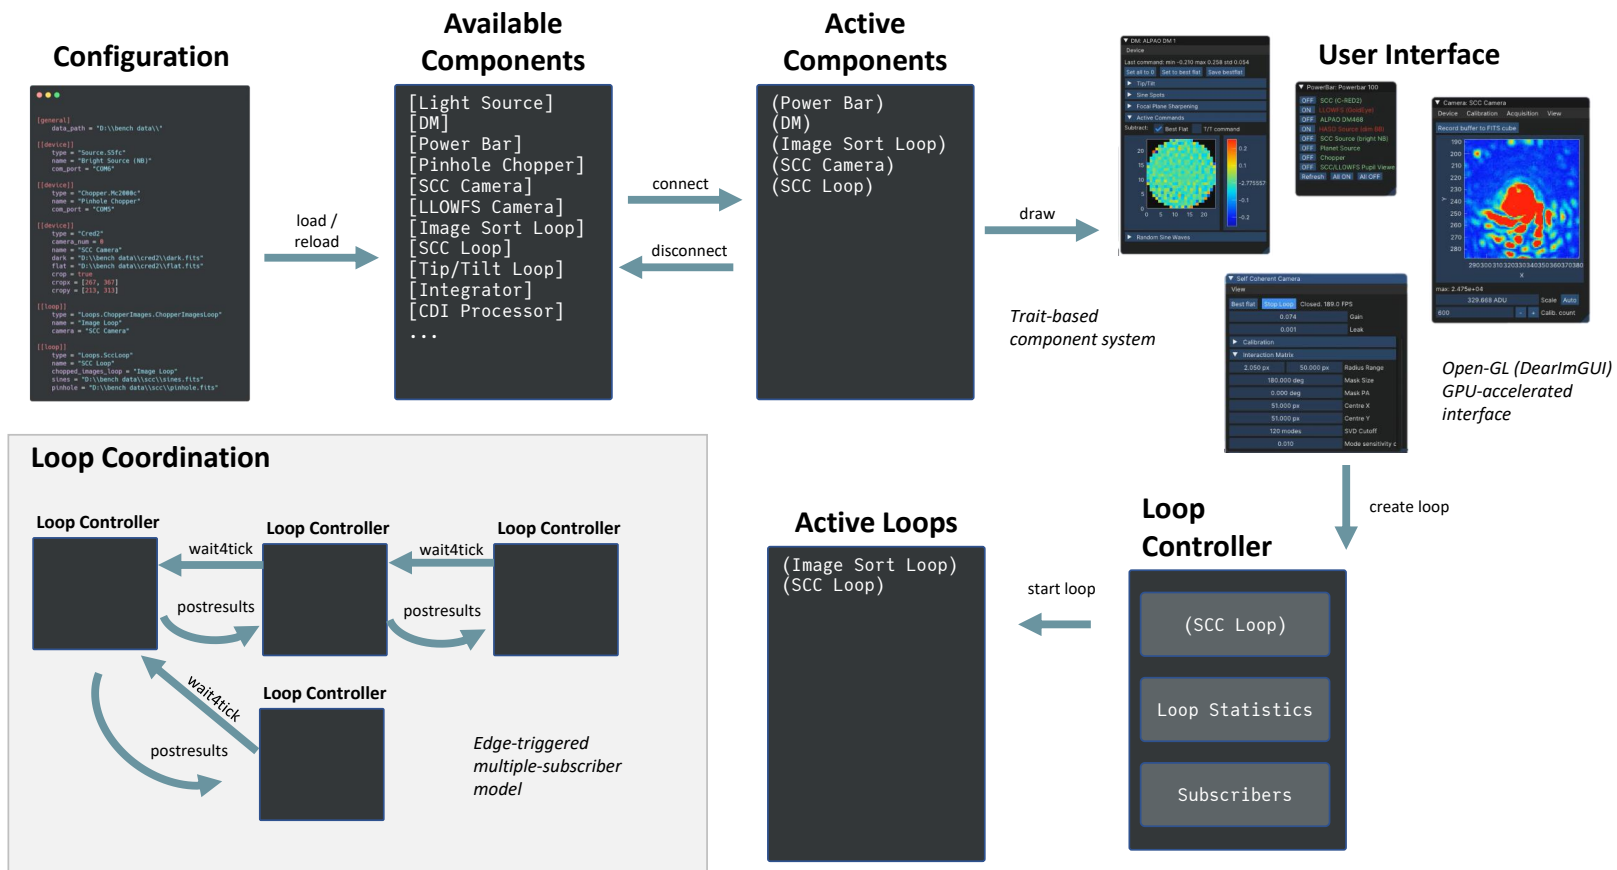


Figure 5.2 Conceptual block diagram describing the architecture of the VENOMS real time lab control software.

5.4 SCC: Measurement Algorithm and Results

We now describe how we implemented the SCC control loop on the VIPER bench and a simplified algorithm used for SCC measurement.

5.4.1 Calibration

Just as in previous SCC implementations, we calibrate the SCC using Fourier modes applied to the DM. It is possible to calibrate an SCC with any combination of DM commands as long as they form a complete and independent basis. In practice, Fourier modes are one of the best choices of calibration basis since they concentrate the impact of the command in just two $\sim \lambda/D$ spots. This increases the SNR of the poke and speeds up calibration compared to, for example, poking individual actuators. Since the outer two annuli of actuators are blocked by the Lyot stop, we do not control those actuators with the SCC. We also remove any tip, tilt, and piston components from the Fourier modes.

We choose the amplitude of the calibration Fourier modes to be as small as possible without making the calibration sequence impractically slow. Smaller amplitude calibrations should improve the linearity of the SCC (Gerard, 2020) but also decrease the calibration signal and thus increase the integration time needed during calibration. We apply each Fourier mode in a push-pull sequence. Figure 5.3 shows the fringed and unfringed images recorded for a push and pull of one mode as well as several combinations of these images. It is sufficient to use only fringed images in the calibration (i.e., push - pull, as outlined by the red box) however, we find that in practice the DM does not apply a perfectly symmetric change in the wavefront between push and pull. This results in some unfringed structure which is undesirable in our calibration matrix. Instead, we calibrate the SCC by calculating

$$(\text{push, fringed} - \text{push, unfringed}) - (\text{pull, fringed} - \text{pull, unfringed}) = \quad (5.1)$$

$$\text{push, fringed} - \text{pull, fringed} - \text{push, unfringed} + \text{pull, unfringed} = \quad (5.2)$$

$$(\text{push, fringed} - \text{pull, fringed}) - (\text{push, unfringed} - \text{pull, unfringed}). \quad (5.3)$$

This does a better job isolating just the change in fringes (Figure 5.3, outlined by blue box). The remaining non-fringe structure is only a result of turbulence or changing source intensity, and is typically only visible near the core. We finally suppress this

structure using a high-pass filter of λ/D (or equivalently, fourier filtering to mask out the center of the modulation transfer function). Note however that this filtering is only performed during calibration, and would also not be required in a perfectly stable environment.

In order to show the Fourier modes clearly in Fig. 5.3, we first applied a static full dark hole correction. This isn't necessary in general, though calibrating iteratively in this way may improve stability when bringing up a system with significant non-common path errors. A downside of this calibration and measurement algorithm is that it does not separate out and track the reference beam Point Spread Function (PSF). A changing reference beam PSF would have the effect of spatially modulating the gain of the SCC. For systems where the reference PSF is not stable, it could still be estimated using differential images with Fourier filtering or fit using a PSF model. Here, we assume that any reference PSF during on-sky operation is on average the same as the one used to capture the calibration sequence modulo changing source intensity. This assumption should hold on-sky thanks to the LLOWFS.

Figure 5.4 shows the results of this calibration procedure for three Fourier modes. The sub-figure **B** corresponds to the images shown in the previous figure. In previous works (Gerard, 2020) it was recommended to apply a circular mask around bright peaks to reduce crosstalk between modes. We do not currently apply those masks to our calibration matrix, but future testing with atmospheric residuals is warranted to fully evaluate their impacts.

In order to select the SCC's correction region, we create a mask with the same dimensions as the image with either a full dark hole or a chosen half-dark hole. With the TGV FPM, we limit this mask to begin at $2\lambda/D$ in separation from the star and end at approximately $11\lambda/D$. It is important to fully exclude pixels outside the control region because the DM, while not having the necessary actuators to properly correct areas outside the control region, can still apply small corrections using second-order effects (e.g. second order sine waves out to $24\lambda/D$ visible in some calibration images). Rather than finding a compromise between a very poor correction outside the control region and a slightly worse correction inside the control region, we prefer that the SCC focus entirely on the control region where a full correction is possible.

To generate the SCC interaction matrix, we flatten each image and select just the masked pixels in our desired dark hole region. We then concatenate them into a $N_{\text{modes}} \times N_{\text{selected pixels}}$ matrix and take the pseudo-inverse. In order to regularize the matrix inversion and reduce noise propagation, we employ a Tikhonov regularization.

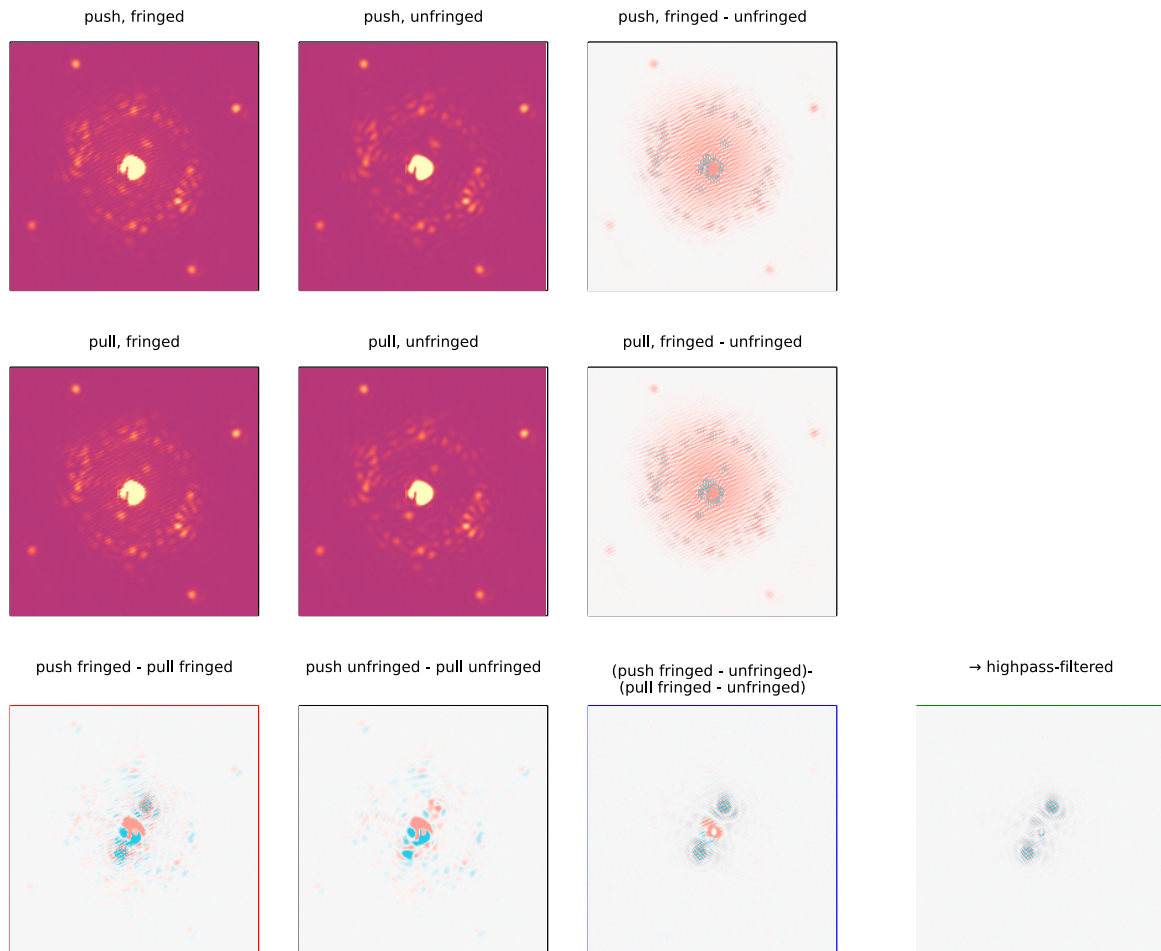


Figure 5.3 Images illustrating the calibration of FAST at the NEW-EARTH lab's VIPER bench. In these images, a slice of the interaction matrix for one Fourier mode is recorded (outlined with a green box). The push and pull differential images (third column, top and middle) also show the envelope of the reference PSF (outlined with a red box). The four spots visible outside the control region are caused by print-through on the DM membrane.




Figure 5.4 Prepared slices from the interaction matrix for three different sinusoidal modes. The DM is not capable of perfectly reproducing sinewaves so there is some noticeable cross talk (central region in C), particularly with sinewaves much greater than 6 cycles/pupil. We do not apodize the beam in VIPER so the Fourier modes show Airy rings.

Finally, this inverted matrix is multiplied with a matrix consisting of the commands used to calibrate the SCC. This results in a $N_{\text{pixel}} \times N_{\text{actuator}}$ matrix.

As an additional step to prevent noise propagation, we calculate the sensitivity of SCC towards each mode. We then apply a cutoff to completely remove modes that are poorly sensed by the SCC— typically those that fall at the edges of the control region or SCC mask (to be described in the following section).

5.4.2 Measurement

Once we have the SCC’s interaction matrix, we can take a measurement from the SCC using a single matrix-vector multiply. We work with differential images (fringed - unfringed). Operating on differential images removes the effects of changing detector cosmetics and dark current. As an additional benefit, the unfringed frames can serve as science images since they are devoid of photon noise from the reference beam. In our experience the reduction in correction rate is a worthwhile trade-off to increase loop stability. As an additional benefit, we can choose after the fact to use both the fringed and unfringed frames as our science images if we are speckle limited or only the unfringed frames if we are limited by photon noise of the reference beam.

First, we collect into a vector the pixels in the masked dark hole region (as was chosen when generating the interaction matrix). Next, we multiply this vector of pixels

by the interaction matrix described above to get a vector of DM actuator commands. Finally, we employ a basic integrator in closed loop. The DM command is updated by:

$$\text{cmd} = \text{cmd} - g \times \text{measurement}$$

We apply this algorithm to differential images but it is perfectly valid to use it on fringed images only. This is because the fringes isolated by the above calibration procedure are symmetric in intensity around zero: any smooth structure is removed by the push/pull fringed/unfringed double difference and by the high-pass filter. As such, they can only fit the fringes in an image and not any smooth incoherent structure.

This algorithm is computationally efficient, consisting only of a pixel gathering or masking operation followed by a single matrix-vector multiply. Additionally, by avoiding Fourier transforms, it is robust to saturated pixels and some detector artifacts. In differential images, saturated regions become zero and simply contribute no information to the SCC. If applied to fringed images instead of differential images, the same holds because saturated regions cannot be fit using a fringe pattern with symmetric intensity.

Robustness to saturated pixels is important for a few reasons. First, all real detectors have some level of read noise so it is necessary to increase the detector's exposure time until the fringe signal of a typical speckle is above this noise level. At the same time, physical detectors, especially those with very low read noise, have limited dynamic range. For example, a 14-bit analogue to digital converter (ADC) has roughly 4 orders of magnitude dynamic range. If for algorithmic reasons we were prevented from ever saturating the detector, then considering an FPM that suppresses core light by least $100\times$, the faintest easily detected speckle would be just 10^{-6} . In our particular implementation, we limited spatial frequency of modes controlled by the SCC. Approaching the edge of the control region, the sampling of the DM falls to just under one sinewave every two actuators. This makes it challenging for the DM to accurately construct the sine waves at higher spatial frequencies. We found that by limiting the DM's correction to an outer separation of $10\lambda/D$ (or even 6) results in improved contrast closer in.

Since the slowest changing quasi-static speckles have the largest impact on direct imaging performance in most conditions, we aim for the SCC to provide a strong DC correction. As such, we eschewed using a leaky integrator. Without a leak, this implementation was more susceptible to poorly sensed modes building up in the integrator. To prevent this kind of instability, we were careful to remove modes that were poorly sensed by the SCC.

One unexpected noise source was evolving dark current patterns on the CRED-2 detector. The detector exhibits a pattern of columns that slowly change in intensity over a period of seconds to minutes. This makes them difficult to remove with a dark calibration. Working with differential images does remove most of this evolving dark current but still leads to a small net linear drift recorded by the SCC. If this signal propagates into a mode that is poorly sensed by the SCC, a command can slowly build up in the integrator.

To limit their effect, we rotated the detector by 30 degrees. This ensures the column pattern (and a much fainter row-oriented pattern) do not match the orientation of the fringes. As a further mitigation to evolving detector bias/dark current, we intend in future setups to alternate the fringed - unfringed subtraction sequence in time. This will use each frame twice at twice the speed and the double-difference will remove any net linear drift in the detector bias.

5.4.3 Laboratory Results

We achieved promising results with the SCC on the VIPER bench. Figure 5.5 shows the results of closing the SCC loop on static errors at a gain of 0.3. The linked video was slowed down by a factor of 10 so that individual iterations are visible. The loop corrects the half dark hole down to nearly the level of read noise in less than 10 iterations (with the bulk of the correction taking just 2 iterations). In this video, the contrast is measured per frame at 400 frames per second. Accordingly, the contrast inside the half dark hole plateaus near the limit of the detector read noise.

Figure 5.6 shows the contrast on the bright and dark halves of the image for a stack of 300 unfringed frames for a total of 300 ms integration. In this time, we achieve a contrast of $\sim 1 \times 10^{-6}$ inside the dark hole before CDI, or roughly a 500 \times improvement compared to the light side.

Finally, we show in 5.7 how the SCC loop generates a full dark hole despite starting with most of the control region saturated thanks to the improved measurement algorithm. The SCC does not have access to any information from saturated pixels so the correction must proceed inwards from unsaturated areas. This process is slow (it takes on the order of 100 iterations) but robust and, unlike algorithms that require Fourier transforms, the loop remains stable.

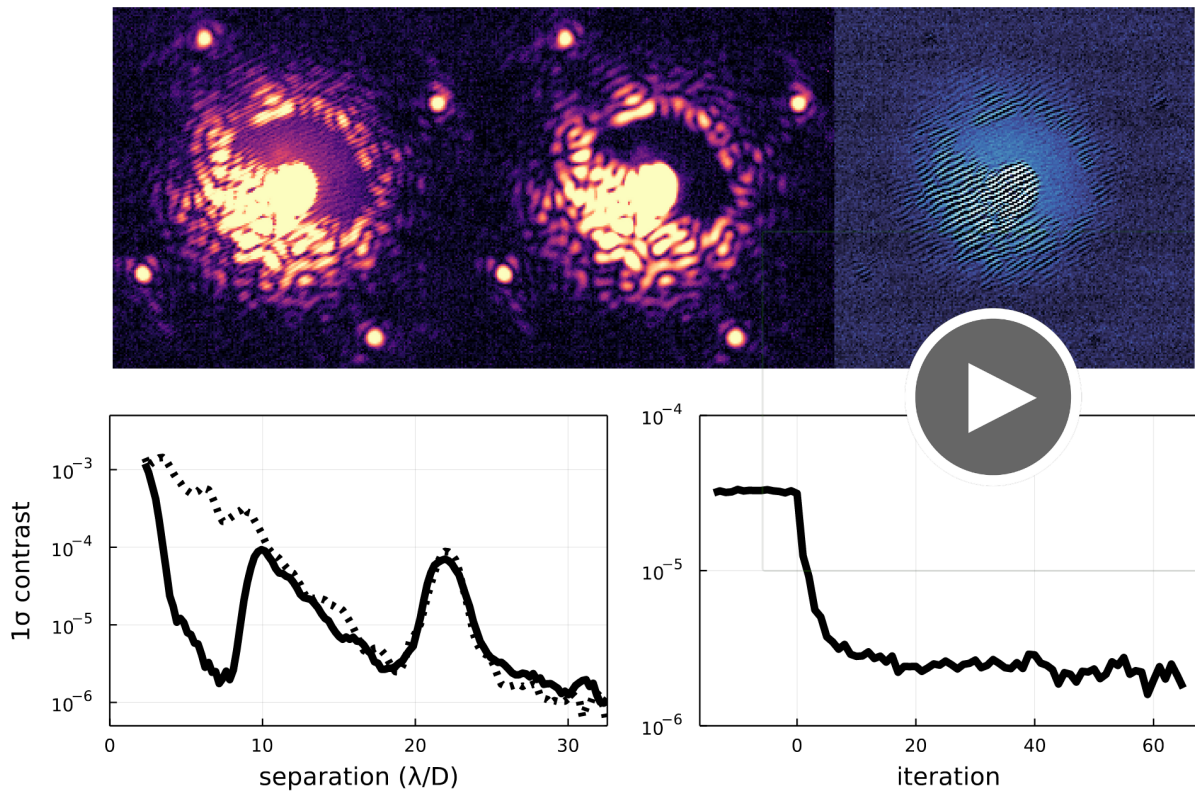


Figure 5.5 The SCC loop closing on the bench's static aberrations and local bench turbulence. **Left:** fringed SCC image. **Centre:** unfringed SCC image. **Right:** differential SCC image. **Bottom left:** contrast in the half dark-hole side of the unfringed image as a function of separation at the current iteration. **Bottom right** best contrast in the half dark hole of the unfringed image over time. Some faint incoherent structure is visible above the star in the fringed and unfringed image. These individual short exposure images are limited by read noise so the contrast appears to plateau after around 15 iterations. Video available at <http://dx.doi.org/10.1117/12.2630415.1>




Figure 5.6 Stack of unfringed SCC images with a higher source intensity than Figure 5.5 shown with an arcsin stretch.

5.5 Post Processing with Coherent Differential Imaging

In addition to real time correction, the SCC allows us to reconstruct the stellar PSF at each iteration. This enables a type of differential imaging in the spirit of Reference star Differential Imaging (RDI; [Wahhaj et al., 2021](#)), Angular Differential Imaging (ADI; [Marois et al., 2006a](#)), and Spectral Differential Imaging (SDI; [Walker et al., 1999](#); [Racine et al., 1999](#); [Marois et al., 2000](#)) called Coherent Differential Imaging (CDI). CDI is possible because the SCC fringes allow us to differentiate between light that travelled through only the main pupil (unfringed in the science images) and light that travelled through both main pupil and the reference pinhole (fringed in the science images).

CDI as a current concept was developed over a series of papers ([Guyon, 2004](#); [Baudoz et al., 2006](#); [Bordé, Traub, 2006](#); [Give'On et al., 2007](#); [Serabyn et al., 2011](#); [Sauvage et al., 2012](#); [Gerard et al., 2018a](#); [Potier, al. et, 2022](#)) that are well-summarized in ([Gerard, 2020](#)). Simulations from some of those works indicated that CDI using FAST could approach roughly 100× improvement on a 0th magnitude star. In CDI, one reconstructs the stellar PSF using wavefront measurements from the SCC. The SCC only measures coherent light from the star so this reconstruction is free of contamination from other astrophysical sources. Since these wavefront measurements are done on the science detector it is in principal easy to match the reconstructed PSF with the target image.

One unexpected challenge we have encountered in achieving these simulated CDI results with the SCC is the interpretation of the sidelobe as a shifted copy of the stellar




Figure 5.7 Sequence in time (iteration increasing downwards) of the SCC converging on a full dark hole despite having more than half the control region initially saturated. From left to right: fringed, unfringed, and difference. Convergence is slow in saturated regions since the SCC must work its way inwards from surrounding unsaturated pixels, but the loop nonetheless makes progress and remains stable.




Figure 5.8 Best fitting pixel MTF and corresponding pixel PSF of our C-RED2 on a log scale. The outline of a fringed image MTF is plotted overtop the pixel MTF in white. This shows how the sidelobes are convolved by a different, non-symmetric part of the pixel MTF than the main pupil.

electric field from the main pupil. In fact, the sidelobe and main pupil differ because they are measured by the detector on different spatial scales, and thus are convolved by different parts of the pixel Modulation Transfer Function (MTF). The pixel MTF is a broad sinc function that results from the finite size of pixels in the focal plane. Effectively, detectors are less sensitive to very high spatial frequency structures than low spatial frequency structures. Effectively, we must reconstruct not just the wavefront but the wavefront as sampled by the science detector. Since we are aiming to subtract speckles with a high degree of accuracy, this complication from the pixel MTF quickly becomes significant.

5.5.1 Fourier Transform and Model Based Algorithm

The first solution we present is to begin by deconvolving the science images by the pixel PSF. This leads to a slight sharpening of the science image and a frequency dependent amplitude correction of the sidelobe used to construct the reference image. For the C-RED2, we modelled the pixel MTF as a sinc function (Figure 5.8). We captured a sequence of fringed and unfringed images and optimized the radial extent of the pixel MTF to get the best subtraction. As expected, the best subtraction occurred when the sinc function was scaled to have its first null at the far corners of the MTF. Instead of adopting this model based approach, future work could directly measure the pixel MTF using a test pattern or knife-edge test. The results of this modeling are shown in




Figure 5.9 Fourier transform based CDI subtraction. Left: unfringed image. Right: post-processed image.




Figure 5.10 Contrast improvement shown in Figure 5.9.

Figures 5.9 and 5.10. The downside of this approach is that it requires a deconvolution step. Deconvolutions are often avoided in image processing since they can propagate noise across spatial frequencies.

The reference pinhole PSF is required in order to reconstruct the stellar PSF from the differential images. This pinhole is not infinitely small, so the resulting PSF is affected by some low-order aberrations that may evolve throughout a sequence. Since we use an optical chopper, it's possible to extract the pinhole reference PSF from the low frequency components of the differential images. However, we find that this signal

is quite noisy in individual pairs of images. We instead use a synthetic model of the reference PSF as an Airy disk (with a possible central obscuration) and optimize it's location, scale, and position for sub-sequences of fringed and unfringed images. In our Julia-based software, this is done easily using the package `PSFModels.jl`²².

5.5.2 Self-Consistent Image Plane Algorithm

The CDI solution described above is effective, but like the original SCC measurement algorithm, it requires a series of Fourier transforms that may propagate noise from saturated pixels. It furthermore requires a deconvolution step which is not ideal for noisy data.

A future approach we aim to test in the NEW-EARTH lab is an image domain algorithm in the same vein as the new SCC measurement algorithm presented in this paper. Our proposed CDI approach is to leverage the SCC calibration matrix inside the control region of the DM. We would first apply the image-based SCC measurement algorithm to find a new modal basis. Using this modal measurement, we would then compute a linear combination of unfringed sine wave spots recorded simultaneously with the SCC calibration matrix.

That is,

$$(\text{push, unfringed} - \text{neutral, unfringed}) - (\text{pull, unfringed} - \text{neutral, unfringed}) \quad (5.4)$$

Effectively, we would directly match the fringes in an image with unfringed speckles. Since we would perform this correction in software rather than with a hardware DM, we could add additional flexibility by splitting the two spots from the DM Fourier spots into independent modes. This PSF reconstruction would self-consistently incorporate the reference beam PSF and pixel MTF. Two limitations of this approach would be that it only works within the control region and that it does not respond to a changing reference beam PSF. Thankfully this PSF is spatially filtered and in our case, stabilized by the LLOWFS. In order to expand the region of subtraction outside the control region of the DM, it may be possible to use on-sky sequences where turbulence naturally produces speckle variations at these separations. To build such a CDI speckle library, a second form of diversity would be required like ADI (Marois et al., 2006a) or RDI (Wahhaj et al., 2021) to ensure that any planets or other astrophysical sources do

²²<https://github.com/JuliaAstro/PSFModels.jl>




Figure 5.11 Initial results of an image-domain based CDI subtraction algorithm showing tolerance to saturated pixels and lack of ringing at nulls of the reference PSF.

not contaminate the calibration. This could allow the image-based CDI subtraction algorithm to function over the full field of view illuminated by the reference beam.

Though not yet fully developed, an example of this approach is presented in Figure 5.11. We aim to further test this approach at the NEW-EARTH lab in the coming months.

5.6 Conclusion

We have presented the current performance of FAST in the NEW-EARTH lab, and in addition, improved algorithms for extracting SCC measurements and CDI reconstructions. On the VIPER bench, we have routinely achieved raw 1σ contrasts of 5×10^{-7} in relatively uncontrolled conditions by combining the SCC and CDI. These improved algorithms will be important for robustly implementing FAST on-sky in the SPIDERS pathfinder and Gemini Planet Imager CAL2 upgrade.

Chapter 6

A Direct Imaging Fourier Transform Spectrograph

This chapter is based on work completed for SPIDERS, the Subaru Pathfinder Instrument for Detecting Exoplanets and Extracting Spectra. Authors: William Thompson, Adam Johnson, Olivier Lardière, Christian Marois, and Frédéric Grandmont. *My contribution to this work was writing the text, developing the original concept, completing the simulations, performing the experiments, developing the data reduction software, and performing the analysis. A. Johnson contributed the mechanical design and drive electronics. O. Lardière performed the optical design and alignment. C. Marois, F. Grandmont, A. Johnson, and O. Lardière contributed to discussions.*

6.1 Motivation

One of the main strengths of direct imaging and related techniques is the ability to study planets spectroscopically—that is, to measure the brightness of a planet at many wavelengths. Spectrographs can be described along four axes:

- their wavelength range, or bandwidth;
- their spectral resolution, R ;
- their field of view;
- and their noise properties.

To date, spectrographs in dedicated direct imaging instruments have fallen into two families.

The first family consists of high resolution, fibre-fed or long-slit spectrographs. These cross-dispersing spectrographs provide very high spectral resolution and bandwidth from a single location in the sky. They are effective instruments for characterizing known planets, provided that their positions can be accurately predicted in advance to within a small fraction of a λ/D .

The second family consists of low spectral-resolution imaging spectrographs, also known as Integral Field Spectrographs (IFSs). These instruments capture a spectrum from each point in a grid of locations in the sky. An IFS is integrated into most major direct imaging instruments including the Gemini Planet Imager ([Macintosh et al., 2014](#)), VLT-SPHERE ([Beuzit et al., 2019](#)), and SCExAO/CHARIS ([McElwain et al., 2012](#)). Because these instruments measure spectra at many locations, the planet does not have to be located or even known in advance. This makes IFSs useful for detection in addition to planet characterization. An IFS can be used to detect planets more effectively than a simple imager by searching for the spectral signal of a planet and/or by using Spectral Differential Imaging (SDI, see Section 1.5.5). Unfortunately, the potential of an IFS to search for planets using their spectral signals is limited by their low resolution. With a spectral resolution as low as R 50, only large molecular absorption bands stand out.

As it turns out, the limitations of both families (field of view or spectral resolution) have a common cause. Both designs are, at their heart, dispersive spectrographs and therefore require at least one detector pixel per spatial- and spectral-resolution element. With these designs, we are therefore limited by the size of astronomical detectors available for purchase. At the time of writing, appropriate detectors are available with at most 10s of millions of pixels. Very large format imagers like those for the Vera C. Rubin Observatory use rafts of multiple detectors ([Gilmore et al., 2008](#)) which come with complications and at considerable cost.

We now return to the Self Coherent Camera (SCC, Section 1.16). In addition to acting as an active wavefront sensor, an SCC allows one to perform Coherent Differential Imaging (CDI) as a post-processing step (see Chapter 5 for details). CDI is very powerful compared to ADI and SDI as it works using data from a single instant and wavelength. ADI is limited by speckle evolution while the pupil rotates and SDI is limited by chromatic evolution of the wavefront.

The SCC and therefore CDI were originally envisaged as operating only in a narrow bandpass. The multi-reference SCC concept extends this to a slightly wider bandpass

in order to allow more light to reach the detector (Delorme et al., 2016); however, it still only records light from the planet with using a single bandpass and makes tradeoffs with respect to the reference beam intensity (Gerard, 2020).

It is clearly desirable to develop a system that combines CDI with spectroscopy.²³ This would allow both access to more light like in the multi-reference SCC concept and, uniquely, allow for improved spectroscopic sensitivity by removing residual starlight from each wavelength slice in post-processing. It is also worth noting the CDI is complementary to both ADI and SDI. All three could in theory be freely combined on the same dataset to achieve even deeper and more robust contrasts.

Unfortunately, a traditional dispersive IFS cannot effectively support the SCC or CDI. The SCC embeds the phase of the speckle electric field in high spatial-frequency fringes. To resolve these fringes, the detector must therefore super-Nyquist sample the scene. This would place at least four times greater pressure on already pixel-starved dispersive integral field spectrographs.

Besides the challenge of combining CDI with spectroscopy, an orthogonal challenge with dispersive integral field spectrographs may be chromaticity. Current instruments that support SDI are limited by chromatic evolution of the wavefront between spectral channels (Gerard et al., 2019a). Simulations of known aberrations GPI's optics (Marois et al., 2012) do not fully account for this chromaticity. By contrast, laboratory experiments testing the use of tunable filters for SDI show no such loss in performance (Ingraham, 2013). Accordingly, a plausible culprit may be lenslet integral field spectrographs themselves.

Lenslet based integral field spectrographs (sometimes known as IFUs) produce a grid of micro-spectra which are spread diagonally across the detector. To extract a spectral cube (2D intensity vs a third spectral axis), one must use software to interpolate along these micro-spectra. At high contrasts, it is plausible that pixel response, inter-pixel cross talk, and imperfect interpolation along the micro-spectra could lead to a loss of correlation between spectral channels. It is possible that this uncorrelated noise in turn limits the ability of SDI to suppress speckles. It is plausible that an integral field spectrograph built around a simpler extraction procedure could result in improved SDI processing.

²³Deformable mirrors can only perfectly correct a single wavelength so there is less reason to perform active wavefront control with multiple wavelengths.




Figure 6.1 Schematic of an Imaging Fourier Transform Spectrograph (IFTS) using a flexural bearing and corner-cube design (Grandmont, 2006).

6.2 Proposed Solution

I propose that an Imaging Fourier Transform Spectrograph (IFTS) would solve the challenges of accessing higher spectral resolution, combining CDI with spectroscopy, and of the supposed IFS chromaticity. An IFTS is an imaging spectrograph built using a Michelson Interferometer. Unlike a dispersive IFS, an IFTS records an entire spectrum on each pixel of the detector. This fundamentally removes the field of view – spectral resolution trade off.

In an IFTS, one directs a collimated beam of broad band light into an interferometer consisting of a 50/50 beam splitter and two mirrors with a varying optical path length difference. When one varies the length of one arm, light travels further down one arm than the other. When it then recombines through the beam splitter, it interferes with a delayed copy of itself. For a given optical path difference in nanometers, the phase difference in radians will depend on the wavelength of light. As such, scanning in optical path difference creates an intensity modulation that depends on the spectral content of the scene. This intensity vs. optical path difference (OPD) measurement is called an

interferogram. One can recover the original spectrum simply by taking the discrete Fourier transform of the interferogram. Each pixel records its own interferogram and therefore its own spectrum.

An FTS has some similarities with a tunable filter (e.g. [Ingraham \(2013\)](#)), in that each detector pixel records an independent spectrum. It differs in that all wavelengths are observed simultaneously. This prevents any speckle evolution between wavelength slices, and results in a much higher observing efficiency. This difference also results in different noise behaviour. On the one hand, an FTS is much less sensitive to dark current and read noise than a tunable filter or dispersive spectrograph. Yet for the same reasons, it is more impacted by photon noise and photon-noise propagates between wavelength slices.

A schematic of an IFTS is presented in [Figure 6.1](#). This design is based on a pair of corner-cube retroreflectors. Unlike a traditional Michelson interferometer, both arms are fixed together, but pivot about a fixed beam-splitter. For a detailed review on corner-cube IFTS, see [Grandmont \(2006\)](#).

The following subsection presents a simulated high resolution spectrum of an exoplanet (from Sonora models, [Marley et al. \(2021\)](#)), the resulting interferogram, and the final recovered spectrum.

The benefits of an IFTS are significant. First, it is easy to accommodate the sampling requirements of CDI. Any detector capable of operating an SCC in a single-band should suffice. This would make it possible to measure the electric field of residual starlight speckles across multiple wavelengths. Second, it would allow for an IFS with 2-3 orders of magnitude higher spectral resolution. This would allow one to use cross-correlations with exoplanet atmosphere models to search for new planets and to detect particular compounds. Planetary RV should also be possible ([Ruffio et al., 2023](#)). These are well beyond the capabilities of low-resolution IFSs used in dedicated direct imaging instruments.

IFTs have previously been deployed as astronomical instruments, though never behind an adaptive optics system or for high contrast applications. One example is the SITELE spectrographs at the Canada France Hawaii Telescope ([Drissen et al., 2019](#)). SITELE offers seeing limited, wide-field imaging-spectroscopy in the visible.




Figure 6.2 Template spectrum used for simulations. The overall intensity is normalized such that the planet has a J band relative contrast of 2×10^{-4} around a fifth magnitude star.

6.3 Simulations

In order to examine how an IFTS would perform in a direct imaging instrument, I carried out a series of simulations. I used cloud-free Sonora Bobcat models (Marley et al., 2021) as a theoretical template spectrum. I selected a spectrum for a planet with solar metallicity and an effective temperature of 1500 K. I normalized the spectrum to have a J band relative contrast of 2×10^{-4} around a fifth magnitude star. This is somewhat optimistic as it would correspond to a very bright planet or a brown dwarf. I considered that the star had a flat spectrum that had been suppressed to a 1σ contrast of 1×10^{-7} . This is the raw contrast that is anticipated by a GPI CAL2-like system Marois et al. (2022); Gerard (2020), but would be optimistic for SPIDERS. I then considered that the planet was observed through an ideal optical filter allowing only light between 1.25 and 1.4 micron to pass. This template is shown in Figure 6.2.

I further considered realistic illumination levels on an 8 m class telescope, 70%

instrument throughput, and 65% modulation efficiency. I simulated a detector with very low read noise ($0.5 e^-$ /frame), 600 frames per second read out, and a scan step of 316 nm (half the wavelength of a Helium-Neon laser). This path length and step size results an average spectral resolution of $R \approx 30,000$.

This scan speed was selected so that the intensity modulation we induce occurs at a higher temporal frequency than the majority of atmospheric turbulence. An alternative design would be to scan much more slowly so that variations in atmospheric speckles average during each exposure. Note that for space applications, a much wider range of scan speeds would be acceptable. The results of these simulations are presented in Figures 6.3, 6.4, 6.5, and 6.6.

Beginning with the interferograms, Figure 6.3 shows an ideal case without any form of noise. The spike at zero-path length difference is the location where all wavelengths interfere constructively (called the “centre-burst”). The narrow trough on either side of the centre-burst is the location where most of the light interferes destructively. If the template were a single narrow line, the interferogram would appear as a sine wave; however, since the template is relatively flat, most of the power is concentrated in a single peak in the interferogram. The second panel shows the same case including realistic (if slightly optimistic) noise levels. The simulations indicate that read noise has a negligible impact on the spectrum. Instead, quantization noise from the small number of Poisson distributed photons dominates.

Figure 6.4 shows spectra extracted from the simulation including realistic noise. These spectra have an average resolution of $R \approx 30,000$. Due to the sheer faintness of planets, individual wavelength slices are noisy at this resolution. That said, the bottom panel illustrates a long integration of 4 hours. We see that the high resolution scan approaches the input spectrum. Despite the noise, Figure 6.5 shows that applying a cross-correlation between a 1000 s integration and the input template still results in a robust detection of the planet.

Finally, Figure 6.6 demonstrates a 3.5 minute integration of low-resolution scans ($R \approx 300$). We see that the FTS recovers the broad features of the input template as expected. For the sake of comparison, a simulation of a low resolution dispersive spectrograph like that of the Gemini Planet Imager (Macintosh et al., 2014) is over-plotted.

The overall results of these simulations indicate that an IFTS would be a compelling alternative to a dispersive IFS for high contrast imaging.

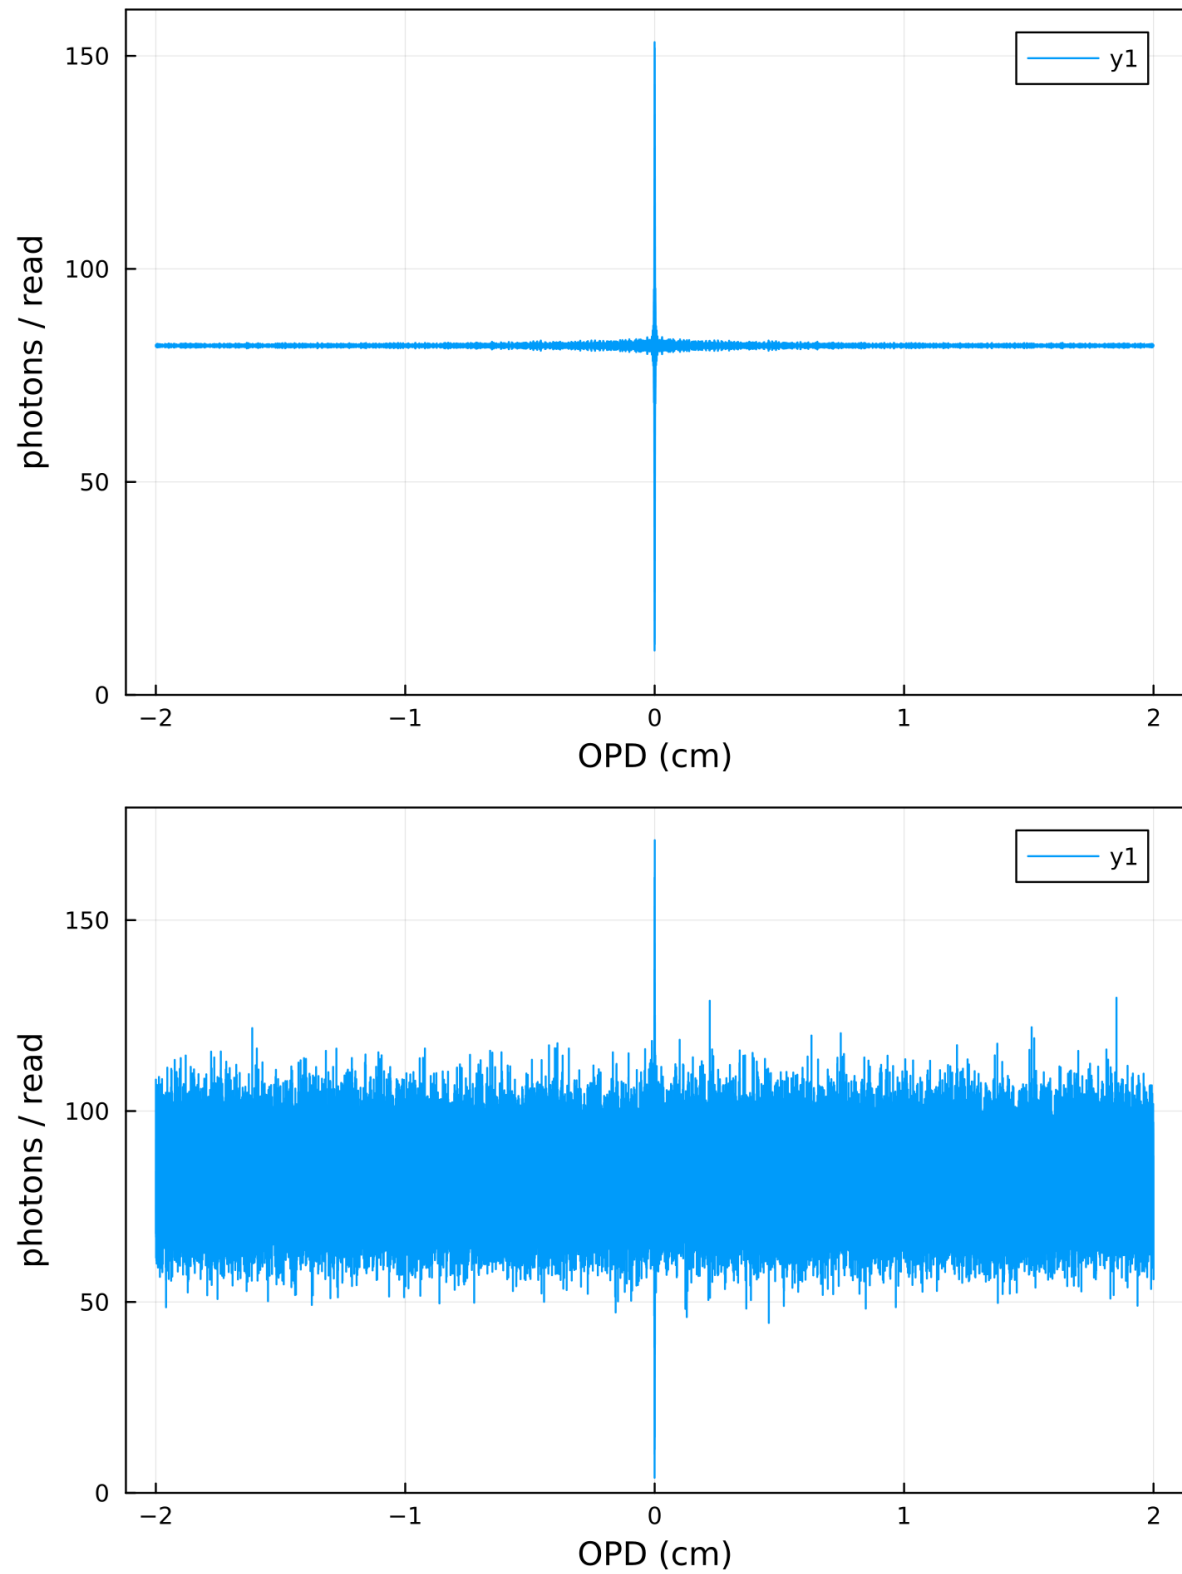


Figure 6.3 Simulated FTS interferogram for the parameters described in the text. The top panel shows an ideal, noiseless interferogram while the bottom panel shows an interferogram with photon and read noise.

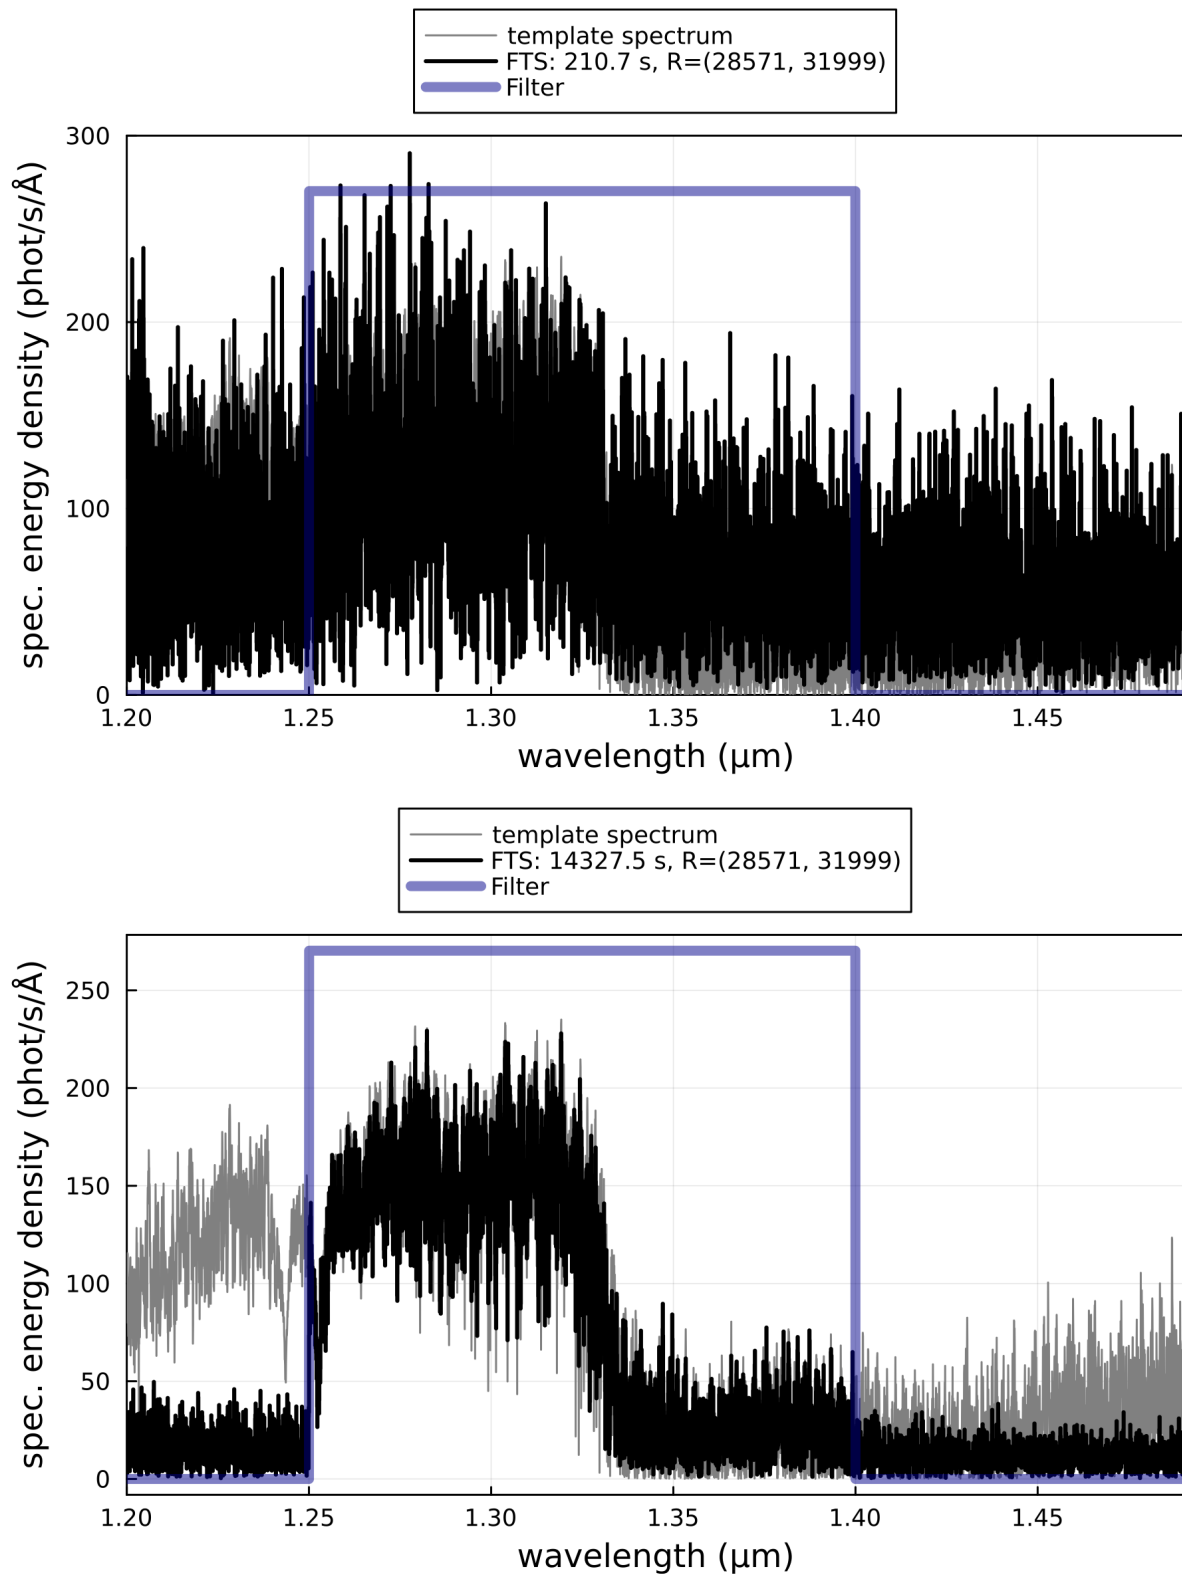


Figure 6.4 Simulated FTS spectrum for the parameters described in the text. The top panel shows a single $1/2\text{cm}^{-1}$ scan ($R \approx 30,000$). For that resolution, a single scan takes 121 s. The bottom panel shows the same results stacked over four hours.




Figure 6.5 Cross-correlation of the template spectrum against a 1000 s integration of R 30,000 scans from the top panel of 6.4. This demonstrates that the planet is still recovered robustly despite the noise. The SNR after performing a cross correlation depends on how accurately the model spectrum matches the planet's.

6.4 Implementation

To build an imaging Fourier transform spectrograph, we repurposed a FTPA2000-300 series FT-NIR spectrometer. This line of spectrometers are used for recording the transmission spectra of gas samples. They contain a light source, Fourier Transform Spectrograph, and associated control electronics.

We disassembled this device, keeping the FTS optics, voice-coil drive, and metrology laser system and discarding the light source, control electronics, and detector. Figure 6.7 shows an image of the disassembled FTS corner cubes and beam splitter. We built a new mechanical frame so that it could be fixed to an optical bench and optically aligned. Finally, we developed a new control electronics board that fully replaced the existing control system. This board interprets the signals of the remaining laser metrology system and drives a current through the devices voice coil actuator using a control




Figure 6.6 Simulations of low resolution FTS scans. As an alternative to the R 30,000 scan, this shows a stack of 100 scans captured $R \approx 300$ for a total integration time of 3.5 min. This illustrates the flexible resolution of FTSs. The scan length should be chosen to maximize the amount of information retrieved.

loop. This custom control solution allowed us to scan the FTS at our own choice of scan speed, scan length, and trigger interval. These changes were essential since the original drive electronics scanned too quickly and would have required our camera detector to run at several thousand FPS.

For the tests presented in this chapter, I used a First Light Imaging C-RED2 NIR detector. We configured our FTS controller to trigger the detector every 400 nm of optical path difference.

We used the resulting IFTS as a second science path in the SPIDERS pathfinder instrument (Marois, *al. et*, 2022). This places the IFTS downstream of a coronagraph shared with an Self Coherent Camera (SCC) focal plane wavefront sensor. A diagram of SPIDERS illustrating the SCC and IFTS paths is presented in Figure 6.8.




Figure 6.7 Image of the FTPA2000-300 series corner cube mirrors and beam splitter. In this perspective, the input beams enters into the page towards corner cube 2, and the output beam exits to the right. Photo courtesy of O. Lardière.

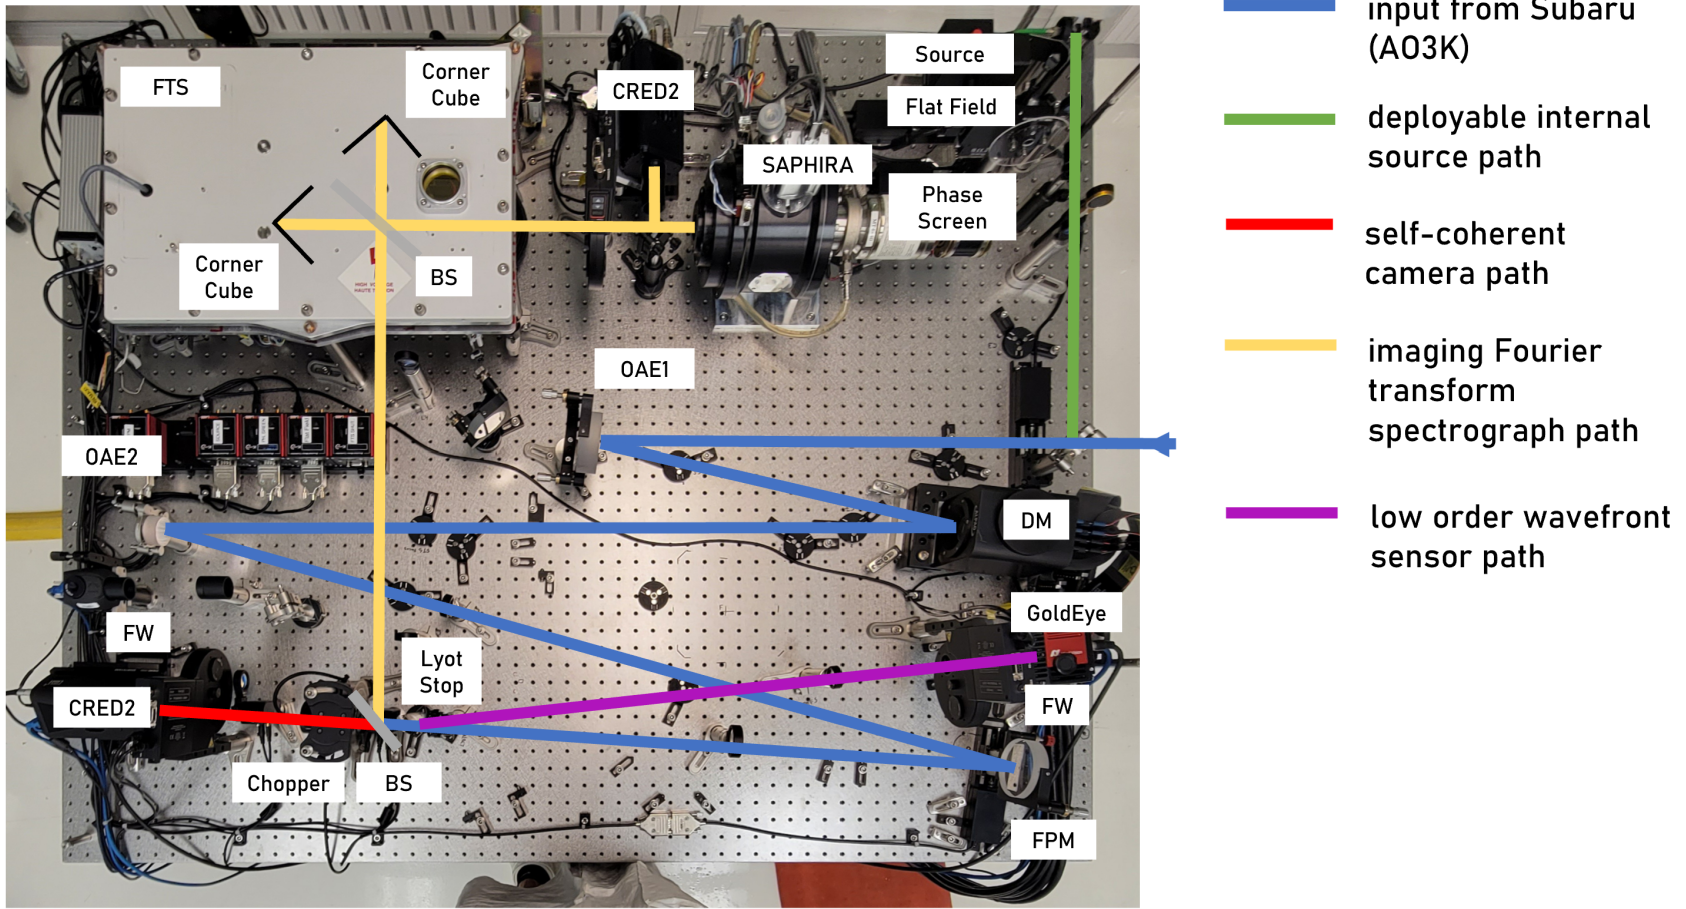


Figure 6.8 Diagram of the SPIDERS instrument showing the IFTS and SCC optical paths. The second CRED2 at the top of this image was used for these tests . Note that the FTS assembly is flipped and rotated vs. the image in Figure 6.7. Photo courtesy of O. Lardière.




Figure 6.9 Interferogram recorded from the IFTS tests. For this plot, the data is summed over all camera pixels.

6.5 Results

I now present preliminary results from the IFTS developed for SPIDERS. These experiments used an internal, pulsed white light source. The light source and camera (a First Light Imaging CRED2) were both triggered from the IFTS control electronics. A short section of the interferogram was used resulting in a spectral resolution of approximately R 350.

Before capturing the IFTS data, I created a half dark hole using the SCC arm of the instrument with an H band apodized Lyot coronagraph and a narrow band filter. I paused the SCC control loop after the dark hole was created.

Figure 6.10 shows image slices and a spectrum taken from a spectral data cube. The bottom panel shows the spectrum measured by summing over all pixels of each wavelength slice. A significant ripple in intensity is visible that is not yet explained at the time of writing, but may be a result of an optical ghost that travelled multiple times through the interferometer. Despite the unexpected intensity variations, the images are




Figure 6.10 IFTS images from three different wavelength slices marked in the spectrum below. The cube is captured simultaneously across I, Z, Y, J, and H band.

very clean and show the expected magnification with wavelength.

The first notable result is that the dark hole is relatively achromatic across this bandpass. This is highly significant as it means that narrowband correction by the SCC improves the surrounding wavelengths as well.

A second notable result is that the SCC fringes are sharply resolved in each wavelength slice as hoped.

Using these results, I attempted to perform SDI and CDI across all wavelength slices. For all cases, I optimized the subtraction in a rectangular region above the star that included the half dark hole and the surrounding speckles.

Figure 6.11 presents preliminary SDI results. For this experiment, I applied a simple global LOCI (Marois et al., 2010a) algorithm that rejected slices taken 10% or closer in wavelength. I find that the extracted images are relatively achromatic over the wide 800 nm to 1650 nm bandpass. The contrast improvement, plotted in Figure 6.12, indicates that residual speckles outside the half dark hole are suppressed by up to nearly 40 \times . This SDI result is highly promising as it exceeds the SDI improvement seen in GPI.




Figure 6.11 Results of applying SDI post processing to the IFTS datacube. The left panels show a broad band image stacked across all wavelengths shown in Figure 6.10 (IZYJH bands), while the right panels show the residuals after applying SDI. The bottom row presents the same data with intensity stretched by a factor of ten.




Figure 6.12 Contrast improvement from SDI using an IFTS data cube.

Thanks to the super-Nyquist sampling available in this design, the SCC fringes are resolved in each wavelength slice. This allows for the use of both SDI and CDI. Figure 6.13 then presents the application of both SDI and CDI on a given wavelength slice. For this case, I used the same least squares algorithm to the global SDI described above but added the CDI reconstructed image as an additional reference.

Both SDI and CDI perform very well on their own, and see a modest further improvement by combining the two. The level of this improvement is limited by the read noise that remains in the images—that is, there were effectively no further speckles to remove. In a real world situation, it is likely the CDI will have a greater impact than SDI at close separations. This is because the performance of SDI drops with decreasing separation, while the performance of CDI does not.

More experiments with realistic illumination levels, integration times, and atmospheric residuals will be needed to confirm that these techniques will work in a real world situation.




Figure 6.13 Comparison of CDI alone, SDI alone, and CDI and SDI working together. Both CDI and SDI offer a significant improvement in contrast and a modest further improvement is seen when combining them.

6.6 Outlook

The simulation and lab results in this chapter show that an IFTS is a very promising spectrograph design for direct imaging. It can easily accommodate the sampling requirements of the SCC and CDI, while delivering much higher spectral resolution than a traditional dispersive IFU. That said, the very high spectral resolution provided by fibre-fed spectrographs continues to make them the best choice for detailed characterization of known planets. This work demonstrates that a dark hole, constructed with focal plane wavefront sensing (the FAST SCC in this case), can be relatively achromatic, at least within the instrument itself. It also demonstrated that an IFTS can be used to improve contrasts with SDI beyond previous instruments, and with multi-spectral CDI. This later technique is particularly promising since its performance does not degrade at close separations from the star, nor does it negatively impact the signals of extended sources like debris disks.

Future instruments could incorporate an IFTS in the same path as the SCC to avoid splitting the light. This may be possible using a Wynne-corrector, an optical element that counteracts the magnification with wavelength effect. In such an instrument, each raw camera frame could be used to drive the SCC in broad band and real time (since all fringes would be aligned) all while the FTS is scanning. The resulting images could then be post processed to recover high resolution spectra.

Chapter 7

Extremely Bright Orbital Guide Beacons for Extremely Large Telescopes

This chapter is based on material published in the proceedings of the Adaptive Optics for Extremely Large Telescopes 6 conference in 2019. Authors: William Thompson and *My contribution to this work was as the lead author. I performed the calculations and contributed the retroreflector concept. Christian Marois contributed to the general concept and provided comments on the text.*

7.1 Introduction

Ground-based direct imaging of rocky exoplanets in reflected light will require unprecedented Adaptive Optics (AO) performance and near sky background limited sensitivity. An ideal AO solution would mitigate lag, very high order aberrations, and chromatic effects between guiding and science wavelengths. In the most extreme cases, we would like to correct wavefront errors before even a single aberrated photon arrives at the science camera so we are no longer limited by light from the star. For this scenario, natural guide stars are too dim to measure wavefront errors with sufficient speed and accuracy. Laser guide “stars” that excite the sodium layer of our atmosphere, on the other hand, only sense a single out-of-band wavelength and suffer from spot elongation, the cone effect, and are still relatively dim. We propose that for high value targets like α -Cent, a constellation of artificial guide beacons could be launched into orbit to act as

reference sources for an extreme adaptive optics system. Satellite based light sources have been used for photometric calibration (Albert, 2011) and have previously been proposed for guiding adaptive optics systems (Greenaway, 1991; Marlow et al., 2016). More recently, satellite laser guide stars have been proposed as a way of reducing the stability requirements of segmented space telescopes (Douglas et al., 2019). Here we discuss the possibilities and considerations of orbiting guide beacons for high contrast imaging of exoplanets from the ground. Such guide beacons could be very bright, allowing for extremely fast, high order corrections. They could be operated at any convenient wavelength, including in the science band. By using multiple wavelengths simultaneously it would in be possible in theory to perform tomographic reconstruction of the atmosphere, telescope, and optics; measure chromatic effects; and perform predictive control. Finally, such guide beacons could in addition serve as photometric calibration sources for detecting variability in the light from companions unaffected by stellar variability. We outline two basic concepts for the orbital guide beacons themselves, describe a range of eccentric orbits that could periodically place the beacons near a target for a sufficient period to undertake observations, and give basic calculations on their brightness seen from the Earth.

7.2 Guide Beacons

The primary purpose of a guide beacon is to act as a bright point source at a high altitude. This could be accomplished in two ways: either the beacon could be actively powered and shine light towards an observatory on the ground, or it could be a passive optical device that reflects light emitted from a ground source back down to the observatory. Beginning with an active guide beacon, a small-sat could be equipped with a laser, 10 cm launch telescope, and fine pointing capabilities. The small-sat would actively orient the launch telescope back towards the observatory using commercially available fine pointing solutions, perhaps with a low power laser uplink from the ground. The satellite could additionally be equipped with basic maneuvering capability for station-keeping, but we emphasize that it would not need to expend propellant during observations as in some occulting star-shade concepts. It would most likely not be feasible for photovoltaic panels to power the laser directly (see the brightness section for a discussion on laser power), so battery storage sufficient for 1-2 hours of observing time every few days would be necessary. Passive guide beacons, on the other hand, could be constructed

to act as retroreflectors. A retroreflector is a simple optical device that returns light in the incident direction. Small retroreflectors are present on many satellites in low Earth orbit where they are used for laser-ranging. Using a retroreflector, the passive guide beacon would return light emitted from a co-located launch telescope on the ground back towards the observatory. By placing the retroreflectors on multiple sides of the beacon, no active stabilization or electronics would be necessary reducing mass and cost. In this case, a larger constellation could be launched to overcome the lack of station keeping capabilities.

7.3 Orbital Parameters

For this concept to be practical, an orbit for the guide beacon must be chosen to meet three requirements. First, the guide beacon must remain near the line of sight between a target and an observatory for long enough for observations to be carried out. Second, the beacon must return to that same line of sight at some regular period so there are multiple chances to observe the target. Finally, for the sake of reducing cost and complexity, the guide beacon should not have to actively expend propellant to maintain alignment during observations. Beginning with the conjunction period and desire not to expend propellants to extend observations, these requirements largely rule out low Earth orbits due to their very low conjunction durations on the order of seconds. Instead, the beacons could be placed into highly eccentric orbits with an apogee directly on the line of sight between the observatory and target. A high eccentricity orbit would minimize the orbital velocity at the apogee. To carry out multiple nights of observations of the same target, the beacon should be placed into an orbit with a period that is close to an integer multiple of one sidereal day. This will bring the guide beacon back in front of the target periodically. Observing times would be adjusted by approximately four minutes per night to account for the Earth orbiting the Sun.

We ran basic orbital simulations to find the orbital parameters that maximize the duration a guide beacon would stay in view. We consider both an observatory located at the equator with a target declination of zero degrees, and Gemini South with a declination of α -Cen, a high value target for high-contrast imaging. We calculate the duration in view assuming a 35 arcsecond diameter field of view consistent with the NFIRAOS laser guide stars. The actual tolerance requirements will be studied in a later analysis. The gravitational influence of the moon will in the very least cause the orbit to

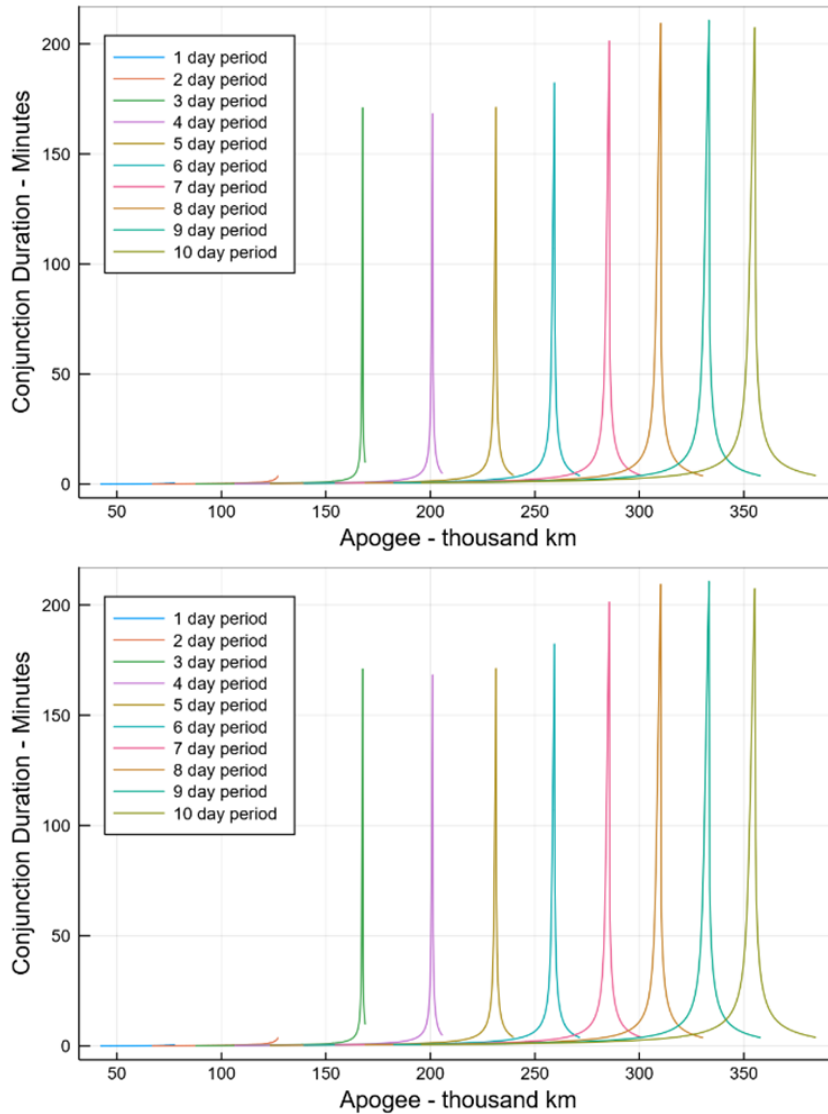


Figure 7.1 These curves show the duration a guide beacon could remain within 35 arcseconds of the target. The different lines are given for different possible orbital periods which are multiples of one sidereal day and will therefore allow for repeat observations of the target. Top: a zero declination target seen from an observatory at the equator. Bottom: α -Cent seen from Gemini South.

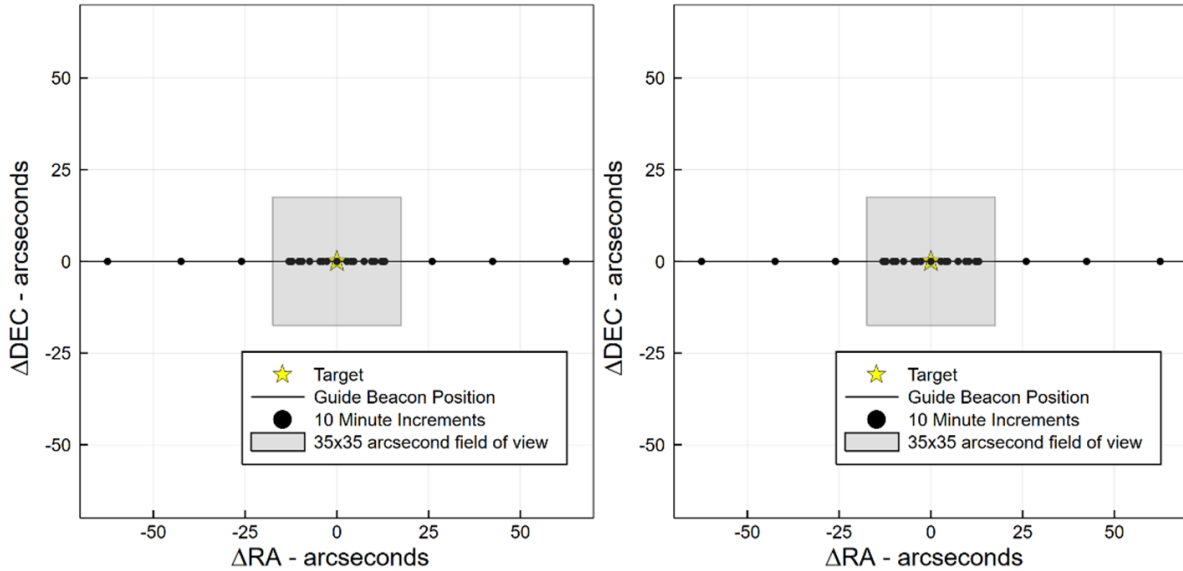


Figure 7.2 These plots show the projection across the sky that a guide beacon could take in front of the target for two example orbits with ten day periods. Left: a zero declination target seen from an observatory at the equator. Right: α -Cent seen from Gemini South.

precess, but for this analysis we consider only the gravity of the Earth. Finally, we only consider orbits with periods that are integer multiples of one sidereal day to allow for multiple nights of observations during a semester. The results presented in Figure 7.1 show that the duration in which the beacon would remain close to the target depends on the observatory location, target declination, and of course, the orbital parameters of the guide beacon. For any given period (and corresponding semi-major axis) there is an optimal eccentricity that maximizes the conjunction duration which can be seen as peaks in Figure 7.1. The main challenge is to find an orbit that has sufficiently low velocity compared to the observatory's motion around the Earth. This improved with either a wider orbit, or with an observatory closer to the equator. The conjunction duration also depends on the declination of the target, with targets closer to the celestial equator allowing for longer conjunctions.

Figure 7.2 shows the paths across the sky a guide beacon could take for two scenarios. For low declination targets, the limiting effect is the relative horizontal motion between the observatory and the guide beacon, whereas for high declination targets, the limiting effect is the vertical motion from the inclined orbit necessary to place the beacon in front of the target. Note that for some targets further away from the celestial equator, there are orbits that exhibit retrograde motion. There are many options for constellation




Figure 7.3 These plots show the motion of a constellation of five beacons moving past a target over 30 minutes calculated for a target with a declination of 10 degrees seen from Maunakea. Left: -15 mins Centre: 0 mins Right: +15 mins.

formations, one example of which is presented in Figure 7.3. These best paths of the formations will depend on the observatory, target, and adaptive optics system, so we leave a full study for a later analysis.

7.4 Brightness

Figure 7.4 shows the brightness of guide beacons calculated for both the active and passive reflector cases. In the first case, we consider a 10 cm diameter launch telescope on the beacon and a laser with a 20W output. In the second case, we consider a laser with a 2 kW output launched from a 10 m telescope with adaptive optics on the ground and reflected off of a 10 cm diameter retroreflector. These choices are all on the higher end of plausible to demonstrate the idea, but could be scaled down and still provide a bright reference source. In both cases, diffraction dictates that a guide beacon with a 10 cm diameter and an apogee near the orbit of the moon would have a spot size of approximately 2 km by the time its light reaches the Earth (for a wavelength of 800 nm). We find that the guide beacons could possess apparent magnitudes as bright as -5 to -10. With such a bright, in-band reference source, a wavefront sensor could be sensitive to high order aberrations even with very short exposure times. Though these sources would be very bright by astronomical standards, there is no risk to human eye safety since they would still appear dimmer than the full moon, and only active for short periods of time.




Figure 7.4 The I-band apparent magnitude of a beacon with the properties described above as a function of distance, for two cases: a 20 W laser on an active guide beacon, and a 2 kW laser launched from the ground at a passive retroreflector.

7.5 Outlook

Though building and launching orbiting guide beacons would clearly require a significant development effort, they could consist mostly of commercially available components. In comparison to orbiting star-shade concepts, this effectively shifts the research effort from space to the ground. Taking full advantage of this concept from the ground will necessitate the development of very high frame rate detectors, high density deformable mirrors, high performance coronagraphs, and fast real-time controllers. Techniques will need to be developed for closing an AO loop and performing tomographic reconstruction on moving guide beacons. An opportunity for early testing could be to use existing satellites that carry retroreflectors in low Earth orbit. While this paper has focused on presenting the orbital guide beacon concept for high-contrast imaging, end-to-end simulations will be required to understand the limits of high-contrast coronagraphic adaptive optics systems that are not starved for light and evaluate the

full potential of this concept.

Chapter 8

Conclusions

To understand if our solar system is typical, astronomers must study a sample of planets around Sun-like stars, both indirectly with light from their stars, and directly with light emitted, absorbed, or reflected from the planet. The key missing ingredients needed to make this possible are improved tools for isolating planets' light from that of their stars and better linking that information with indirect signals.

Instrumentation

In this dissertation, I presented several ways in which we can better isolate light from planets by suppressing quasi-static speckles.

The fast atmospheric self-coherent camera technique (FAST SCC) promises to suppress glare from stars in real time using interferometry and techniques from adaptive optics. I developed improved techniques for applying the FAST SCC in real world conditions and demonstrated this technique in a laboratory setting. The technique suppressed quasi-static speckles by a factor of approximately 400× which reached the stroke accuracy limit of our deformable mirror.

Coherent differential imaging (CDI) is a way to use electric field measurements from an SCC to digitally subtract quasi-static speckles while preserving the signals of planets and disks. I presented an improvement to this technique based on correcting the detector pixel response, leading to a roughly 20× suppression of quasi-static speckles.

Access to high resolution spectra from exoplanets is necessary in order to study them in depth. I presented a concept for an imaging Fourier Transform Spectrograph (IFTS) that promises to deliver medium to high spectral resolution ($R \approx 20,000$) across entire images in a way that can be combined with coherent differential imaging. I

demonstrated this concept using simulations and laboratory data, including broadband spectroscopy of an SCC dark hole. I further demonstrated the combination of coherent and spectral differential imaging for a nearly 40× improvement in quasi-static speckles.

Finally, I proposed a speculative concept for retro-reflecting orbital guide beacons, which could provide a path for imaging Earth-like planets from ground-based observatories.

Analysis Techniques

In addition to improved instrumentation for the future, I developed improved analysis techniques that are useful on current and archival data.

I re-examined the mathematical basis for differential imaging and developed a formalism for planet signal-to-noise ratio in the presence of self-subtraction. Using this formalism, I presented a new algorithm for reducing angular and spectral differential imaging data and demonstrated a 3 – 5× reduction in speckle noise close to stars.

All direct imaging techniques eventually hit a limit set by the orbital motion of planets. Integrating longer than this limit results in smearing of the planet signal and will prevent the detection of fainter planets. Though approaches for addressing this problem already existed in the literature, I developed it further into a general solution for detecting planets and modelling their orbits.

I further developed this approach into a general way to combine direct imaging with aperture masking interferometry, radial velocity, and proper motion anomaly data, merging many different exoplanet techniques into a single more powerful tool.

Additionally, I demonstrated the application of simulation based calibration to orbit modelling. This provides additional checks that a statistical orbit model is accurately sampled.

HR 8799

I completed a large observation campaign of the HR 8799 system. These observations included nine new quarter-nights, five previously unpublished quarter-nights, and four archival quarter-nights for a total of 16.75 hours of science integration time. This amounts to the second-largest targeted direct imaging campaign yet assembled, after NEAR ([Wagner et al., 2021](#)).

I applied a some of the new analysis techniques described above including image processing and multi-epoch orbital detection. Using these, I placed new limits on the

existence of any additional outer planets in the system and detected a candidate fifth planet far interior to the known four. This candidate planet would resolve tension between the dynamical and model-estimated masses of the four other planets. It furthermore fits well with the [Su et al. \(2009\)](#) prediction for the location of the system's inner debris disks.

Directly detecting massive planets inwards of 5 AU is a key step towards detecting their lower mass, Jupiter and Saturn-like counterparts.

Future Work

The techniques developed in this dissertation have broad applicability to the direct imaging community. I now briefly list a few specific targets and efforts that follow directly from this work.

Confirmation of HR 8799 f The candidate fifth planet will be challenging to confirm. The best current facility would be JWST's NIRISS aperture masking interferometry. This would require approximately three hours of telescope time. In the near future, the SPIDERS instrument I contributed to aims achieve this level of sensitivity from the ground.

51 Eri Planet Search The 51 Eri system hosts one known young Jovian planet [Macintosh et al. \(2015\)](#). The Gemini Planet Imager Exoplanet Survey (GPIES) imaged this target more than five times across multiple bands. The GRAVITY interferometer ([GRAVITY Collaboration et al., 2017](#)) has since tightly constrained the orbit of planet b (Jason Wang, private communication; whereistheplanet.com). Combining these constraints with GPIES images, Gaia-Hipparcos proper motion anomaly, and ground based radial velocity curves could find or place limits on the existence of any additional planets.

On-Sky Testing The SPIDERS pathfinder ([Marois et al., 2022](#)) has been built to test the SCC, CDI, IFTS, and spectro-coherent differential imaging at the Subaru telescope. Validating these techniques on-sky will be a crucial step before they can be broadly applied as part of a next generation direct imaging survey. If proven, a SPIDERS could be used to confirm or rule out the HR 8799 candidate.

Merging the IFTS and SCC

This dissertation explored using an IFTS in conjunction with an SCC by splitting the light between two different science paths. This was required because the self-coherent camera must operate in a single narrow band. Operating in a broader band would allow in much more light, enabling faster response to changing speckles from the instrument and atmosphere. The problem with using a wider filter is that the PSF magnification effect blurs the SCC fringes. One proposed solution to this problem is to use multiple reference pinholes that create fringes with different angles (Delorme et al., 2016). This limits the blurring over moderate band passes, at the expense of reduced fringe visibility. Instead, a better solution could be to use a Wynne-corrector to de-magnify the PSF as a function of wavelength. This might be able to directly align the fringes across wavelengths and allow a cleaner broad band correction.

Furthermore, it may be possible to combine a Wynne-corrected SCC and an IFTS into a single instrument arm. Thanks to the SCC fringes being aligned across all wavelengths, the image feed could be used to directly close the loop while the IFTS is scanning. This would avoid splitting the light between two optical paths.

Final Remarks

This dissertation has demonstrated improvements in speckle suppression that, taken at face value, could result in nearly five orders of magnitude increased sensitivity to directly imaged exoplanets. It has further developed techniques for compensating for orbital motion, removing one of the sensitivity limits of direct imaging. Of course, such an improvement will not be realized in practice—while the methods presented all complement each other, it is likely that they will eventually reach common limits caused by instrument stability, lack of photons, incoherent light, and astrophysical noise sources. Applying combinations of these techniques together in real world conditions will reveal where these joint limits lie. Nonetheless, these developments bring a substantial improvement in our ability to suppress the quasi-static speckle glare.

Thankfully, the future isn't bright for directly imaged planets!

Bibliography

Albert J. SATELLITE-MOUNTED LIGHT SOURCES AS PHOTOMETRIC CALIBRATION STANDARDS FOR GROUND-BASED TELESCOPES // *AJ.* XII 2011. 143, 1. 8.

Armitage Philip J. Astrophysics of Planet Formation. I 2020. 2.

Auvergne M., Bodin P., Boisnard L., Buey J. T., Chaintreuil S., Epstein G., Jouret M., Lam-Trong T., Levacher P., Magnan A., Perez R., Plasson P., Plessier J., Peter G., Steller M., Tiphène D., Baglin A., Agogué P., Appourchaux T., Barbet D., Beaufort T., Bellenger R., Berlin R., Bernardi P., Blouin D., Boumier P., Bonneau F., Briet R., Butler B., Cautain R., Chiavassa F., Costes V., Cuvillo J., Cunha-Parro V., de Oliveira Fialho F., Decaudin M., Defise J. M., Djalal S., Docclo A., Drummond R., Dupuis O., Exil G., Fauré C., Gaboriaud A., Gamet P., Gavalda P., Grolleau E., Gueguen L., Guivarc'h V., Guterman P., Hasiba J., Huntzinger G., Hustaix H., Imbert C., Jeanville G., Johlander B., Jorda L., Journoud P., Karioty F., Kerjean L., Lafond L., Lapeyrere V., Landiech P., Larqué T., Laudet P., Le Merrer J., Leporati L., Leruyet B., Levieuge B., Llebaria A., Martin L., Mazy E., Mesnager J. M., Michel J. P., Moalic J. P., Monjoin W., Naudet D., Neukirchner S., Nguyen-Kim K., Ollivier M., Orcesi J. L., Ottacher H., Oulali A., Parisot J., Perruchot S., Piacentino A., Pinheiro da Silva L., Platzer J., Pontet B., Pradines A., Quentin C., Rohbeck U., Rolland G., Rollenhagen F., Romagnan R., Russ N., Samadi R., Schmidt R., Schwartz N., Sebbag I., Smit H., Sunter W., Tello M., Toulouse P., Ulmer B., Vandermarcq O., Vergnault E., Wallner R., Wautier G., Zanatta P. The CoRoT Satellite in Flight: Description and Performance // *Astronomy and Astrophysics.* X 2009. 506. 411–424.

Baba Naoshi, Murakami Naoshi. A Method to Image Extrasolar Planets with Polarized Light // *PASP.* XII 2003. 115, 814. 1363.

Baines Ellyn K., White Russel J., Huber Daniel, Jones Jeremy, Boyajian Tabetha, McAlister Harold A., ten Brummelaar Theo A., Turner Nils H., Sturmann Judit, Sturmann Laszlo, Goldfinger P. J., Farrington Christopher D., Riedel Adric R., Ireland Michael, von Braun Kaspar, Ridgway Stephen T. THE CHARA ARRAY ANGULAR DIAMETER OF HR 8799 FAVORS PLANETARY MASSES FOR ITS IMAGED COMPANIONS // *ApJ.* XII 2012. 761, 1. 57.

Baraffe I., Chabrier G., Barman T. S., Allard F., Hauschildt P. H. Evolutionary Models for Cool Brown Dwarfs and Extrasolar Giant Planets. The Case of HD 209458 // *Astronomy and Astrophysics,* v.402, p.701-712 (2003). V 2003. 402. 701.

- Barman Travis S., Konopacky Quinn M., Macintosh Bruce, Marois Christian.* Simultaneous Detection of Water, Methane, and Carbon Monoxide in the Atmosphere of Exoplanet HR8799b // *The Astrophysical Journal*. V 2015. 804. 61.
- Baudoz P., Boccaletti A., Baudrand J., Rouan D.* The Self-Coherent Camera: A New Tool for Planet Detection // *Proceedings of the International Astronomical Union*. X 2005. 1, C200. 553–558.
- Baudoz P., Boccaletti A., Baudrand J., Rouan D.* The Self-Coherent Camera: a new tool for planet detection // *IAU Colloq. 200: Direct Imaging of Exoplanets: Science & Techniques*. I 2006. 553–558.
- Berger J. P.* An introduction to visibility modeling // *EAS Publications Series*. 6. I 2003. 23. (EAS Publications Series).
- Bergfors C., Brandner W., Janson M., Köhler R., Henning T.* VLT/NACO Astrometry of the HR 8799 Planetary System. L'-Band Observations of the Three Outer Planets // *Astronomy and Astrophysics*. IV 2011. 528. A134.
- Berry Matthew, Healy Liam.* The Generalized Sundman Transformation for Propagation of High-Eccentricity Elliptical Orbits. 2002.
- Best William M. J., Dupuy Trent J., Liu Michael C., Siverd Robert J., Zhang Zhoujian.* The UltracoolSheet: Photometry, Astrometry, Spectroscopy, and Multiplicity for 3000+ Ultracool Dwarfs and Imaged Exoplanets. XI 2020.
- Best William M. J., Liu Michael C., Magnier Eugene A., Dupuy Trent J.* A Volume-Limited Sample of L and T Dwarfs Defined by Parallaxes. VII 2018. 62.
- Best William M. J., Liu Michael C., Magnier Eugene A., Dupuy Trent J.* A Volume-limited Sample of Ultracool Dwarfs. I. Construction, Space Density, and a Gap in the L/T Transition // *The Astronomical Journal*. I 2021. 161. 42.
- Beuzit J.-L., Vigan A., Mouillet D., Dohlen K., Gratton R., Boccaletti A., Sauvage J.-F., Schmid H. M., Langlois M., Petit C., Baruffolo A., Feldt M., Milli J., Wahhaj Z., Abe L., Anselmi U., Antichi J., Barette R., Baudrand J., Baudoz P., Bazzon A., Bernardi P., Blanchard P., Brast R., Bruno P., Buey T., Carbillet M., Carle M., Cascone E., Chapron F., Charton J., Chauvin G., Claudi R., Costille A., De Caprio V., de Boer J., Delboulbé A., Desidera S., Dominik C., Downing M., Dupuis O., Fabron C., Fantinel D., Farisato G., Feautrier P., Fedrigo E., Fusco T., Gigan P., Ginski C., Girard J., Giro E., Gisler D., Gluck L., Gry C., Henning T., Hubin N., Hugot E., Incorvaia S., Jaquet M., Kasper M., Lagadec E., Lagrange A.-M., Le Coroller H., Le Mignant D., Le Ruyet B., Lessio G., Lizon J.-L., Llored M., Lundin L., Madec F., Magnard Y., Marteaud M., Martinez P., Maurel D., Ménard F., Mesa D., Möller-Nilsson O., Moulin T., Moutou C., Origné A., Parisot J., Pavlov A., Perret D., Pragt J., Puget P., Rabou P., Ramos J., Reess J.-M., Rigal F., Rochat S., Roelfsema R., Rousset G., Roux A., Saisse M., Salasnich B., Santambrogio E., Scuderi S., Segransan D., Sevin A., Siebenmorgen R., Soenke C., Stadler E., Suarez M., Tiphène D., Turatto M., Udry S., Vakili F., Waters L. B. F. M., Weber L., Wildi F., Zins G., Zurlo A.* SPHERE: The Exoplanet Imager for the Very Large Telescope // *Astronomy and Astrophysics*. XI 2019. 631. A155.

- Bezanson Jeff, Karpinski Stefan, Shah Viral B., Edelman Alan. Julia: A Fast Dynamic Language for Technical Computing // arXiv e-prints. IX 2012.
- Bisnovatyi-Kogan G. S. Planetary System around the Pulsar PSR 1257+12. // Astronomy and Astrophysics. VIII 1993. 275. 161–162.
- Blakely Dori, Francis Logan, Johnstone Doug, Soullain Anthony, Tuthill Peter, Cheetham Anthony, Sanchez-Bermudez Joel, Sivaramakrishnan Anand, Dong Ruobing, Marel Nienke van der, Cooper Rachel, Vigan Arthur, Cantalloube Faustine. Two Rings and a Marginally Resolved, 5 Au Disk around LkCa 15 Identified via Near-infrared Sparse Aperture Masking Interferometry // ApJ. V 2022. 931, 1. 3.
- Blunt Sarah, Wang Jason J., Angelo Isabel, Ngo Henry, Cody Devin, De Rosa Robert J., Graham James R., Hirsch Lea, Nagpal Vighnesh, Nielsen Eric L., Pearce Logan, Rice Malena, Tejada Roberto. Orbitize!: A Comprehensive Orbit-fitting Software Package for the High-contrast Imaging Community // The Astronomical Journal. III 2020. 159. 89.
- Bonsor Amy, Lichtenberg Tim, Drżkowska Joanna, Buchan Andrew M. Rapid Formation of Exoplanetesimals Revealed by White Dwarfs // Nat Astron. I 2023. 7, 1. 39–48.
- Booth Mark, Jordán Andrés, Casassus Simon, Hales Antonio S., Dent William R. F., Faramaz Virginie, Matrà Luca, Barkats Denis, Brahm Rafael, Cuadra Jorge. Resolving the Planetesimal Belt of HR 8799 with ALMA // Mon. Not. R. Astron. Soc: Lett. VII 2016. 460, 1. L10–L14.
- Bordé Pascal J., Traub Wesley A. High-Contrast Imaging from Space: Speckle Nulling in a Low-Aberration Regime // ApJ. II 2006. 638, 1. 488–498.
- Borgniet S., Lagrange A.-M., Meunier N., Galland F., Arnold L., Astudillo-Defru N., Beuzit J.-L., Boisse I., Bonfils X., Bouchy F., Debondt K., Deleuil M., Delfosse X., Desort M., Díaz R. F., Eggenberger A., Ehrenreich D., Forveille T., Hébrard G., Loeillet B., Lovis C., Montagnier G., Moutou C., Pepe F., Perrier C., Pont F., Queloz D., Santerne A., Santos N. C., Ségransan D., Silva R. da, Sivan J. P., Udry S., Vidal-Madjar A. Extrasolar Planets and Brown Dwarfs around AF-type Stars - X. The SOPHIE Sample: Combining the SOPHIE and HARPS Surveys to Compute the Close Giant Planet Mass-Period Distribution around AF-type Stars // A&A. I 2019. 621. A87.
- Borucki William J., Koch David, Basri Gibor, Batalha Natalie, Brown Timothy, Caldwell Douglas, Caldwell John, Christensen-Dalsgaard Jørgen, Cochran William D., DeVore Edna, Dunham Edward W., Dupree Andrea K., Gautier Thomas N., Geary John C., Gilliland Ronald, Gould Alan, Howell Steve B., Jenkins Jon M., Kondo Yoji, Latham David W., Marcy Geoffrey W., Meibom Søren, Kjeldsen Hans, Lissauer Jack J., Monet David G., Morrison David, Sasselov Dimitar, Tarter Jill, Boss Alan, Brownlee Don, Owen Toby, Buzasi Derek, Charbonneau David, Doyle Laurance, Fortney Jonathan, Ford Eric B., Holman Matthew J., Seager Sara, Steffen Jason H., Welsh William F., Rowe Jason, Anderson Howard, Buchhave Lars, Ciardi David, Walkowicz Lucianne, Sherry William, Horch Elliott, Isaacson Howard, Everett Mark E., Fischer Debra, Torres Guillermo, Johnson John Asher, Endl Michael, MacQueen Phillip, Bryson Stephen T., Dotson Jessie, Haas Michael, Kolodziejczak Jeffrey, Van Cleve Jeffrey, Chandrasekaran Hema, Twicken Joseph D., Quintana Elisa V., Clarke Bruce D., Allen Christopher,

- Li Jie, Wu Haley, Tenenbaum Peter, Verner Ekaterina, Bruhweiler Frederick, Barnes Jason, Prsa Andrej.* Kepler Planet-Detection Mission: Introduction and First Results // *Science*. II 2010. 327, 5968. 977–980.
- Bottom Michael, Femenia Bruno, Huby Elsa, Mawet Dimitri, Dekany Richard, Milburn Jennifer, Serabyn Eugene.* Speckle Nulling Wavefront Control for Palomar and Keck // *Adaptive Optics Systems V*. 9909. 2016. 1507–1522.
- Bowler Brendan P., Blunt Sarah C., Nielsen Eric L.* Population-Level Eccentricity Distributions of Imaged Exoplanets and Brown Dwarf Companions: Dynamical Evidence for Distinct Formation Channels // *AJ*. I 2020. 159, 2. 63.
- Bowler Brendan P., Tran Quang H., Zhang Zhoujian, Morgan Marvin, Ashok Katelyn B., Blunt Sarah, Bryan Marta L., Evans Analis E., Franson Kyle, Huber Daniel, Nagpal Vighnesh, Wu Ya-Lin, Zhou Yifan.* Rotation Periods, Inclinations, and Obliquities of Cool Stars Hosting Directly Imaged Substellar Companions: Spin-Orbit Misalignments Are Common // *The Astronomical Journal*. IV 2023. 165. 164.
- Brandt, G. M. , Michalik Daniel, Brandt Timothy D., Li Yiting, Dupuy Trent J., Zeng Yunlin.* htof: A New Open-source Tool for Analyzing Hipparcos, Gaia, and Future Astrometric Missions // *The Astronomical Journal*. XII 2021. 162, 6. 230.
- Brandt G. Mirek, Brandt Timothy D., Dupuy Trent J., Michalik Daniel, Marleau Gabriel-Dominique.* The First Dynamical Mass Measurement in the HR 8799 System // *arXiv:2105.12820 [astro-ph]*. V 2021a.
- Brandt, T. D. .* The Hipparcos-Gaia Catalog of Accelerations: Gaia EDR3 Edition // *The Astrophysical Journal Supplement Series*. VI 2021. 254, 2. 42. *arXiv: 2105.11662*.
- Brandt, T. D. , Dupuy Trent J., Li Yiting, Brandt, G. M. , Zeng Yunlin, Michalik Daniel, Gagliuffi Daniella C. Bardalez, Raposo-Pulido Virginia.* orvara: An Efficient Code to Fit Orbits using Radial Velocity, Absolute, and/or Relative Astrometry // *arXiv:2105.11671 [astro-ph]*. V 2021. *arXiv: 2105.11671*.
- Brandt Timothy D.* The Hipparcos-Gaia Catalog of Accelerations: Gaia EDR3 Edition // *ApJS*. VI 2021. 254, 2. 42.
- Brandt Timothy D., Dupuy Trent J., Li Yiting, Brandt G. Mirek, Zeng Yunlin, Michalik Daniel, Gagliuffi Daniella C. Bardalez, Raposo-Pulido Virginia.* Orvara: An Efficient Code to Fit Orbits Using Radial Velocity, Absolute, and/or Relative Astrometry // *arXiv:2105.11671 [astro-ph]*. V 2021b.
- Brennan Pat.* The Planet That Wasn't There (Oh, Wait, It Was). III 2018.
- Brooks Stephen P., Gelman Andrew.* General Methods for Monitoring Convergence of Iterative Simulations // *Journal of Computational and Graphical Statistics*. XII 1998. 7, 4. 434–455.
- Chilcote Jeffrey, Tobin Taylor, Currie Thayne, Brandt Timothy D., Groff Tyler D., Kuzuhara Masayuki, Guyon Olivier, Lozi Julien, Jovanovic Nemanja, Sahoo Ananya, Deo Vincent, Akiyama Eiji, Janson Markus, Knapp Jill, Kwon Jungmi, McElwain Michael W., Nishikawa Jun, Wagner Kevin, Helminiak Krzysztof, Skaf Nour, Tamura Motohide.* SExAO/CHARIS Direct Imaging of A Low-mass Companion At A Saturn-like Separation from an Accelerating Young A7 Star // *The Astronomical Journal*. XII 2021. 162. 251.

- Cincotta P. M., Giordano C. M., Simó C.* Phase Space Structure of Multi-Dimensional Systems by Means of the Mean Exponential Growth Factor of Nearby Orbits // *Physica D Nonlinear Phenomena*. VIII 2003. 182. 151–178.
- Close Laird M., Males Jared R.* A SEARCH FOR WIDE COMPANIONS TO THE EXTRASOLAR PLANETARY SYSTEM HR 8799 // *ApJ*. XII 2009. 709, 1. 342–348.
- Cook Samantha R., Gelman Andrew, Rubin Donald B.* Validation of Software for Bayesian Models Using Posterior Quantiles // *Journal of Computational and Graphical Statistics*. IX 2006. 15, 3. 675–692.
- Coroller H. Le, Nowak M., Arnold L., Dohlen K., Fusco T., Sauvage J. F., Vigan A.* K-Stacker, a New Way of Detecting and Characterizing Exoplanets with High Contrast Imaging Instruments // *arXiv:1510.06331 [astro-ph]*. X 2015.
- Cowan Nicolas B., Fujii Yuka.* Mapping Exoplanets // *Handbook of Exoplanets*. Cham: Springer International Publishing, 2017. 1–16.
- Currie Thayne, Biller Beth, Lagrange Anne-Marie, Marois Christian, Guyon Olivier, Nielsen Eric, Bonnefoy Mickael, De Rosa Robert.* Direct Imaging and Spectroscopy of Extrasolar Planets. V 2022.
- Currie Thayne, Brandt G. Mirek, Brandt Timothy D., Lacy Brianna, Burrows Adam, Guyon Olivier, Tamura Motohide, Liu Ranger Y., Sagynbayeva Sabina, Tobin Taylor, Chilcote Jeffrey, Groff Tyler, Marois Christian, Thompson William, Murphy Simon J., Kuzuhara Masayuki, Lawson Kellen, Lozi Julien, Deo Vincent, Vievard Sebastien, Skaf Nour, Uyama Taichi, Jovanovic Nemanja, Martinache Frantz, Kasdin N. Jeremy, Kudo Tomoyuki, McElwain Michael, Janson Markus, Wisniewski John, Hodapp Klaus, Nishikawa Jun, Helminiak Krzysztof, Kwon Jungmi, Hayashi Masahiko.* Direct Imaging and Astrometric Detection of a Gas Giant Planet Orbiting an Accelerating Star // *Science*. IV 2023. 380, 6641. 198–203.
- Currie Thayne, Brandt Timothy, Kuzuhara Masayuki, Chilcote Jeffrey, Cashman Edward, Liu R. Y., Lawson Kellen, Tobin Taylor, Brandt G. Mirek, Guyon Olivier, Lozi Julien, Deo Vincent, Vievard Sebastien, Ahn Kyohoon, Skaf Nour.* A New Type of Exoplanet Direct Imaging Search: The SCEAO/CHARIS Survey of Accelerating Stars. IX 2021.
- Currie Thayne, Burrows Adam, Girard Julien H., Cloutier Ryan, Fukagawa Misato, Sorahana Satoko, Kuchner Marc, Kenyon Scott J., Madhusudhan Nikku, Itoh Yoichi, Jayawardhana Ray, Matsumura Soko, Pyo Tae-Soo.* DEEP THERMAL INFRARED IMAGING OF HR 8799 Bcde: NEW ATMOSPHERIC CONSTRAINTS AND LIMITS ON A FIFTH PLANET // *ApJ*. X 2014. 795, 2. 133.
- Currie Thayne, Burrows Adam, Itoh Yoichi, Matsumura Soko, Fukagawa Misato, Apai Daniel, Madhusudhan Nikku, Hinz Philip M., Rodigas T. J., Kasper Markus, Pyo T. S., Ogino Satoshi.* A Combined Subaru/VLT/MMT 1-5 Mm Study of Planets Orbiting HR 8799: Implications for Atmospheric Properties, Masses, and Formation // *The Astrophysical Journal*. III 2011. 729. 128.
- Currie Thayne, Fukagawa Misato, Thalmann Christian, Matsumura Soko, Plavchan Peter.* Direct Detection and Orbit Analysis of the Exoplanets HR 8799 Bcd from Archival 2005 Keck/NIRC2 Data // *ApJ*. VIII 2012. 755, 2. L34.

- De Rosa Robert J., Nielsen Eric L., Wahhaj Zahed, Ruffio Jean-Baptiste, Kalas Paul G., Peck Anne E., Hirsch Lea A., Roberson William.* Direct Imaging Discovery of a Super-Jovian around the Young Sun-like Star AF Leporis. II 2023.
- Delorme J.-R., Galicher R., Baudoz P., Rousset G., Mazoyer J., Dupuis O.* Focal Plane Wavefront Sensor Achromatization : The Multireference Self-Coherent Camera // *A&A.* IV 2016. 588. A136.
- Delorme Jacques-Robert, Jovanovic Nemanja, Echeverri Daniel, Mawet Dimitri P., Wallace James Kent, Bartos Randall D., Cetre Sylvain, Wizinowich Peter L., Ragland Sam, Lilley Scott J., Wetherell Edward, Doppmann Greg, Wang Jason J., Morris Evan C., Ruffio Jean-Baptiste, Martin Emily C., Fitzgerald Michael P., Ruane Garreth J., Schofield Tobias, Suominen Nick, Calvin Benjamin, Wang Eric, Magnone Kenneth G., Johnson Christopher A., Sohn Ji Man, Lopez Ronald A., Bond Charlotte Z., Pezzato Jacklyn, Llop-Sayson Jorge, Chun Mark R., Skemer Andrew J.* Keck Planet Imager and Characterizer: A Dedicated Single-Mode Fiber Injection Unit for High-Resolution Exoplanet Spectroscopy // *JATIS.* VIII 2021. 7, 3. 035006.
- Desch S. J.* Mass Distribution and Planet Formation in the Solar Nebula // *ApJ.* XII 2007. 671, 1. 878.
- Dickey James M.* The Weighted Likelihood Ratio, Linear Hypotheses on Normal Location Parameters // *The Annals of Mathematical Statistics.* II 1971. 42, 1. 204–223.
- Douglas E. S., Males J. R., Clark J., Guyon O., Lumbres J., Marlow W., Cahoy K. L.* Laser Guide Star for Large Segmented-aperture Space Telescopes. I. Implications for Terrestrial Exoplanet Detection and Observatory Stability // *AJ.* I 2019. 157, 1. 36.
- Doyon René, Hutchings John B., Beaulieu Mathilde, Albert Loic, Lafrenière David, Willott Chris, Touahri Driss, Rowlands Neil, Maszkiewicz Micheal, Fullerton Alex W.* The JWST Fine Guidance Sensor (FGS) and near-Infrared Imager and Slitless Spectrograph (NIRISS) // *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave.* 8442. 2012. 1005–1017.
- Drissen Laurent, Martin Thomas, Rousseau-Nepton Laurie, Robert Carmelle, Martin R. Pierre, Baril Marc, Prunet Simon, Joncas Gilles, Thibault Simon, Brousseau Denis, Mandar Julie, Grandmont Frédéric, Yee Howard, Simard Luc.* SITELE: An Imaging Fourier Transform Spectrometer for the Canada-France-Hawaii Telescope // *Monthly Notices of the Royal Astronomical Society.* V 2019. 485, 3. 3930–3946.
- Dupuy Trent J., Kraus Adam L.* Distances, Luminosities, and Temperatures of the Coldest Known Substellar Objects // *Science.* IX 2013. 341. 1492–1495.
- Dupuy Trent J., Liu Michael C.* The Hawaii Infrared Parallax Program. I. Ultracool Binaries and the L/T Transition // *The Astrophysical Journal Supplement Series.* VIII 2012. 201. 19.
- Eastman Jason D., Rodriguez Joseph E., Agol Eric, Stassun Keivan G., Beatty Thomas G., Vanderburg Andrew, Gaudi B. Scott, Collins Karen A., Luger Rodrigo.* EXOFASTv2: A Public, Generalized, Publication-Quality Exoplanet Modeling Code. VII 2019.
- FITS Working Group .* Definition of the Flexible Image Transport System (FITS). 2016.

- Faramaz Virginie, Marino Sebastian, Booth Mark, Matrà Luca, Mamajek Eric E., Bryden Geoffrey, Stapelfeldt Karl R., Casassus Simon, Cuadra Jorge, Hales Antonio S., Zurlo Alice.* A Detailed Characterization of HR 8799's Debris Disk with ALMA in Band 7 // *The Astronomical Journal*. VI 2021. 161. 271.
- Fernandes Rachel B., Mulders Gijs D., Pascucci Ilaria, Mordasini Christoph, Emsenhuber Alexandre.* Hints for a Turnover at the Snow Line in the Giant Planet Occurrence Rate // *ApJ*. III 2019. 874, 1. 81.
- Ferrer-Chávez Rodrigo, Blunt Sarah, Wang Jason J.* Algorithmic Speedups and Posterior Biases from Orbit Fitting of Directly Imaged Exoplanets in Cartesian Coordinates // *Res. Notes AAS*. VII 2021a. 5, 7. 162.
- Ferrer-Chávez Rodrigo, Wang Jason J., Blunt Sarah.* Biases in Orbital Fitting of Directly Imaged Exoplanets with Small Orbital Coverage // *The Astronomical Journal*. V 2021b. 161. 241.
- Fjelde Tor Erlend, Xu Kai, Tarek Mohamed, Yalburgi Sharan, Ge Hong.* Flexible Transformations for Probability Distributions // *Proceedings of The 2nd Symposium on Advances in Approximate Bayesian Inference*. 118. 2020. 1–17. (Proceedings of Machine Learning Research).
- Fletcher Roger.* Practical Methods of Optimization // *Perth Meridien Observations*. 1987.
- Foo Gregory, Palacios David M., Swartzlander Grover A.* Optical Vortex Coronagraph // *Optics letters*. 2005. 30, 24. 3308–3310.
- Foreman-Mackey Daniel, Hogg David W., Lang Dustin, Goodman Jonathan.* Emcee: The MCMC Hammer // *PASP*. II 2013. 125, 925. 306.
- Foreman-Mackey Daniel, Luger Rodrigo, Agol Eric, Barclay Thomas, Bouma Luke G., Brandt Timothy D., Czekala Ian, David Trevor J., Dong Jiayin, Gilbert Emily A., Gordon Tyler A., Hedges Christina, Hey Daniel R., Morris Brett M., Price-Whelan Adrian M., Savell Arjun B.* exoplanet: Gradient-based probabilistic inference for exoplanet data & other astronomical time series. 6 2021.
- Fulton Benjamin J., Petigura Erik A., Blunt Sarah, Sinukoff Evan.* RadVel: The Radial Velocity Modeling Toolkit // *Publications of the Astronomical Society of the Pacific*. IV 2018. 130. 044504.
- Fulton Benjamin J., Rosenthal Lee J., Hirsch Lea A., Isaacson Howard, Howard Andrew W., Dedrick Cayla M., Sherstyuk Ilya A., Blunt Sarah C., Petigura Erik A., Knutson Heather A., Behrman Aida, Chontos Ashley, Crepp Justin R., Crossfield Ian J. M., Dalba Paul A., Fischer Debra A., Henry Gregory W., Kane Stephen R., Kosiarek Molly, Marcy Geoffrey W.* California Legacy Survey II. Occurrence of Giant Planets Beyond the Ice Line // *arXiv:2105.11584 [astro-ph]*. V 2021.
- GAIA Collaboration, Brown A. G. A., Vallenari A., Prusti T., de Bruijne J. H. J., Babusiaux C., Biermann M., Creevey O. L., Evans D. W., Eyer L., Hutton A., Jansen F., Jordi C., Klioner S. A., Lammers U., Lindgren L., Luri X., Mignard F., Panem C., Pourbaix D., Randich S., Sartoretti P., Soubiran C., Walton N. A., Arenou F., Bailer-Jones C. a. L., Bastian U., Cropper M., Drimmel R., Katz D., Lattanzi M. G., van Leeuwen F., Bakker J., Cacciari C., Castañeda J., De Angeli F., Ducourant C., Fabricius C., Fouesneau M., Frémat Y., Guerra R., Guerrier A., Guiraud J., Jean-Antoine Piccolo A., Masana E., Messineo R., Mowlavi N., Nicolas C.,*

Nienartowicz K., Pailler F., Panuzzo P., Riclet F., Roux W., Seabroke G. M., Sordo R., Tanga P., Thévenin F., Gracia-Abril G., Portell J., Teyssier D., Altmann M., Andrae R., Bellas-Velidis I., Benson K., Berthier J., Blomme R., Brugaletta E., Burgess P. W., Busso G., Carry B., Cellino A., Cheek N., Clementini G., Damerджи Y., Davidson M., Delchambre L., Dell'Oro A., Fernández-Hernández J., Galluccio L., García-Lario P., Garcia-Reinaldos M., González-Núñez J., Gosset E., Haigron R., Halbwachs J.-L., Hambly N. C., Harrison D. L., Hatzidimitriou D., Heiter U., Hernández J., Hestroffer D., Hodgkin S. T., Holl B., Janßen K., Jevardat de Fombelle G., Jordan S., Krone-Martins A., Lanzafame A. C., Löffler W., Lorca A., Manteiga M., Marchal O., Marrese P. M., Moitinho A., Mora A., Muinonen K., Osborne P., Pancino E., Pauwels T., Petit J.-M., Recio-Blanco A., Richards P. J., Riello M., Rimoldini L., Robin A. C., Roegiers T., Rybizki J., Sarro L. M., Siopis C., Smith M., Sozzetti A., Ulla A., Utrilla E., van Leeuwen M., van Reeve W., Abbas U., Abreu Aramburu A., Accart S., Aerts C., Aguado J. J., Ajaj M., Altavilla G., Álvarez M. A., Álvarez Cid-Fuentes J., Alves J., Anderson R. I., Anglada Varela E., Antoja T., Audard M., Baines D., Baker S. G., Balaguer-Núñez L., Balbinot E., Balog Z., Barache C., Barbato D., Barros M., Barstow M. A., Bartolomé S., Bassilana J.-L., Bauchet N., Baudesson-Stella A., Becciani U., Bellazzini M., Bernet M., Bertone S., Bianchi L., Blanco-Cuaresma S., Boch T., Bombrun A., Bossini D., Bouquillon S., Bragaglia A., Bramante L., Breedt E., Bressan A., Brouillet N., Bucciarelli B., Burlacu A., Busonero D., Butkevich A. G., Buzzi R., Caffau E., Cancelliere R., Cánovas H., Cantat-Gaudin T., Carballo R., Carlucci T., Carnerero M. I., Carrasco J. M., Casamiquela L., Castellani M., Castro-Ginard A., Castro Sampol P., Chaoul L., Charlot P., Chemin L., Chiavassa A., Cioni M.-R. L., Comoretto G., Cooper W. J., Cornez T., Cowell S., Crifo F., Crosta M., Crowley C., Dafonte C., Dapergolas A., David M., David P., de Laverny P., De Luise F., De March R., De Ridder J., de Souza R., de Teodoro P., de Torres A., del Peloso E. F., del Pozo E., Delbo M., Delgado A., Delgado H. E., Delisle J.-B., Di Matteo P., Diakite S., Diener C., Distefano E., Dolding C., Eappachen D., Edvardsson B., Enke H., Esquej P., Fabre C., Fabrizio M., Faigler S., Fedorets G., Fernique P., Fienga A., Figueras F., Fouron C., Fragkoudi F., Fraile E., Franke F., Gai M., Garabato D., Garcia-Gutierrez A., García-Torres M., Garofalo A., Gavras P., Gerlach E., Geyer R., Giacobbe P., Gilmore G., Girona S., Giuffrida G., Gomel R., Gomez A., Gonzalez-Santamaria I., González-Vidal J. J., Granvik M., Gutiérrez-Sánchez R., Guy L. P., Hauser M., Haywood M., Helmi A., Hidalgo S. L., Hilger T., Hładczuk N., Hobbs D., Holland G., Huckle H. E., Jasniewicz G., Jonker P. G., Juaristi Campillo J., Julbe F., Karbevská L., Kervella P., Khanna S., Kochoska A., Kontizas M., Kordopatis G., Korn A. J., Kostrzewa-Rutkowska Z., Kruszyńska K., Lambert S., Lanza A. F., Lasne Y., Le Champion J.-F., Le Fustec Y., Lebreton Y., Lebzelter T., Leccia S., Leclerc N., Lecoœur-Taïbi I., Liao S., Licata E., Lindstrøm E. P., Lister T. A., Livanou E., Lobel A., Madrero Pardo P., Managau S., Mann R. G., Marchant J. M., Marconi M., Marcos Santos M. M. S., Marinoni S., Marocco F., Marshall D. J., Martin Polo L., Martín-Fleitas J. M., Masip A., Massari D., Mastrobuono-Battisti A., Maze T., McMillan P. J., Messina S., Michalik D., Millar N. R., Mints A., Molina D., Molinaro R., Molnár L., Montegriffo P., Mor R., Morbidelli R., Morel T., Morris D., Mulone A. F., Munoz D., Muraveva T., Murphy C. P., Musella I., Noval L., Ordénovic C., Orrù G., Osinde J., Pagani C., Pagano I., Palaversa L., Palicio P. A., Panahi A., Pawlak M., Peñalosa Esteller X., Penttilä A., Piersimoni A. M., Pineau F.-X., Plachy E., Plum G., Poggio E., Poretti E., Poujoulet E., Prša A., Pulone L., Racero E., Ragaini S., Rainer M., Raiteri C. M., Rambaux N., Ramos P., Ramos-Lerate M., Re Fiorentin P., Regibo S., Reylé C., Ripepi V., Riva A., Rixon G., Robichon N., Robin C., Roelens M., Rohrbasser L., Romero-Gómez M., Rowell N., Royer F., Rybicki K. A., Sadowski G., Sagristà Sellés A., Sahlmann J., Salgado J., Salguero E., Samaras N., Sanchez Gimenez V., Sanna N., Santoveña R., Sarasso M., Schultheis M., Sciacca E., Segol M., Segovia J. C., Ségransan D., Semeux D., Shahaf S., Siddiqui H. I., Siebert A., Siltala L., Slezak E., Smart R. L., Solano E., Solitro F., Souami D., Souchay J., Spagna A., Spoto F., Steele I. A., Steidelmüller H.,

- Stephenson C. A., Süveges M., Szabados L., Szegedi-Elek E., Taris F., Tauran G., Taylor M. B., Teixeira R., Thuillot W., Tonello N., Torra F., Torra J., Turon C., Unger N., Vaillant M., van Dillen E., Vanel O., Vecchiato A., Viala Y., Vicente D., Voutsinas S., Weiler M., Wevers T., Wyrzykowski Ł, Yoldas A., Yvard P., Zhao H., Zorec J., Zucker S., Zurbach C., Zwitter T. Gaia Early Data Release 3. Summary of the Contents and Survey Properties // *Astronomy and Astrophysics*. V 2021. 649. A1.
- GRAVITY Collaboration , Abuter R., Accardo M., Amorim A., Anugu N., Ávila G., Azouaoui N., Benisty M., Berger J. P., Blind N., Bonnet H., Bourget P., Brandner W., Brast R., Buron A., Burtscher L., Cassaing F., Chapron F., Choquet É, Clénet Y., Collin C., Coudé Du Foresto V., de Wit W., de Zeeuw P. T., Deen C., Delplancke-Ströbele F., Dembet R., Derie F., Dexter J., Duvert G., Ebert M., Eckart A., Eisenhauer F., Esselborn M., Fédou P., Finger G., Garcia P., Garcia Dabo C. E., Garcia Lopez R., Gendron E., Genzel R., Gillessen S., Gonte F., Gordo P., Grould M., Grözinger U., Guieu S., Haguenaue P., Hans O., Haubois X., Haug M., Haussmann F., Henning Th, Hippler S., Horrobin M., Huber A., Hubert Z., Hubin N., Hummel C. A., Jakob G., Janssen A., Jochum L., Jocou L., Kaufer A., Kellner S., Kendrew S., Kern L., Kervella P., Kiekebusch M., Klein R., Kok Y., Kolb J., Kulas M., Lacour S., Lapeyrère V., Lazareff B., Le Bouquin J.-B., Lèna P., Lenzen R., Lévêque S., Lippa M., Magnard Y., Mehrgan L., Mellein M., Mérand A., Moreno-Ventas J., Moulin T., Müller E., Müller F., Neumann U., Oberti S., Ott T., Pallanca L., Panduro J., Pasquini L., Paumard T., Percheron I., Perraut K., Perrin G., Pflüger A., Pfuhl O., Phan Duc T., Plewa P. M., Popovic D., Rabien S., Ramírez A., Ramos J., Rau C., Riquelme M., Rohloff R.-R., Rousset G., Sanchez-Bermudez J., Scheithauer S., Schöller M., Schuhler N., Spyromilio J., Straubmeier C., Sturm E., Suarez M., Tristram K. R. W., Ventura N., Vincent F., Waisberg I., Wank I., Weber J., Wieprecht E., Wiest M., Wiezorrek E., Wittkowski M., Woillez J., Wolff B., Yazici S., Ziegler D., Zins G. First Light for GRAVITY: Phase Referencing Optical Interferometry for the Very Large Telescope Interferometer // *Astronomy and Astrophysics*. VI 2017. 602. A94.
- GRAVITY Collaboration , Lacour S., Nowak M., Wang J., Pfuhl O., Eisenhauer F., Abuter R., Amorim A., Anugu N., Benisty M., Berger J. P., Beust H., Blind N., Bonnefoy M., Bonnet H., Bourget P., Brandner W., Buron A., Collin C., Charnay B., Chapron F., Clénet Y., Coudé du Foresto V., de Zeeuw P. T., Deen C., Dembet R., Dexter J., Duvert G., Eckart A., Förster Schreiber N. M., Fédou P., Garcia P., Garcia Lopez R., Gao F., Gendron E., Genzel R., Gillessen S., Gordo P., Greenbaum A., Habibi M., Haubois X., Haußmann F., Henning Th., Hippler S., Horrobin M., Hubert Z., Jimenez Rosales A., Jocou L., Kendrew S., Kervella P., Kolb J., Lagrange A.-M., Lapeyrère V., Le Bouquin J.-B., Léna P., Lippa M., Lenzen R., Maire A.-L., Mollière P., Ott T., Paumard T., Perraut K., Perrin G., Pueyo L., Rabien S., Ramírez A., Rau C., Rodríguez-Coira G., Rousset G., Sanchez-Bermudez J., Scheithauer S., Schuhler N., Straub O., Straubmeier C., Sturm E., Tacconi L. J., Vincent F., van Dishoeck E. F., von Fellenberg S., Wank I., Waisberg I., Widmann F., Wieprecht E., Wiest M., Wiezorrek E., Woillez J., Yazici S., Ziegler D., Zins G. First Direct Detection of an Exoplanet by Optical Interferometry: Astrometry and K -Band Spectroscopy of HR 8799 e // *A&A*. III 2019. 623. L11.
- Galicher R., Baudoz P., Delorme J.-R., Mawet D., Bottom M., Wallace J. K., Serabyn E., Shelton C. Minimization of Non-Common Path Aberrations at the Palomar Telescope Using a Self-Coherent Camera // *Astronomy and Astrophysics*. XI 2019. 631. A143.
- Galicher R., Marois C., Macintosh B., Zuckerman B., Barman T., Konopacky Q., Song I., Patience J., Lafrenière D., Doyon R., Nielsen E. L. The International Deep Planet Survey. II. The Frequency of Directly Imaged Giant Exoplanets with Stellar Mass // *Astronomy & Astrophysics*. X 2016. 594.

Astrometry and Photometry in High Contrast Imaging. // . IX 2011. P25.

Garstka Michael, Cannon Mark, Goulart Paul. COSMO: A Conic Operator Splitting Method for Convex Conic Problems // arXiv:1901.10887 [math]. I 2019.

Geiler Fabian, Krivov Alexander, Booth Mark, Löhne Torsten. The Scattered Disc of HR 8799 // Monthly Notices of the Royal Astronomical Society. II 2019. 483, 1. 332–341.

Gelman Andrew, Rubin Donald B. Inference from Iterative Simulation Using Multiple Sequences // Statistical Science. XI 1992. 7, 4. 457–472.

Gerard Benjamin L., Marois Christian, Currie Thayne, Brandt Timothy, Chilcote Jeffrey K., Draper Zachary H., Groff Tyler, Guyon Olivier, Hayashi Masahiko, Jovanovic Nemanja, Knapp Gillian R., Kudo Tomoyuki, Kwon Jungmi, Lozi Julien, Martinache Frantz, McElwain Michael, Tamura Motohide, Uyama Taichi. A Chromaticity Analysis and PSF Subtraction Techniques for SCEXAO/CHARIS Data // The Astronomical Journal. VII 2019a. 158. 36.

Gerard Benjamin L., Marois Christian, Galicher Raphaël. Fast Coherent Differential Imaging on Ground-based Telescopes Using the Self-coherent Camera // AJ. VIII 2018a. 156, 3. 106.

Gerard Benjamin L., Marois Christian, Galicher Raphaël, Baudoz Pierre, Patapis Polychronis, Kühn Jonas. The Fast Atmospheric Self-Coherent Camera Technique: Laboratory Results and Future Directions // arXiv:1910.04554 [astro-ph]. X 2019b.

Fast Focal Plane Wavefront Sensing on Ground-Based Telescopes. // . VII 2018b. 0703. 1070351.

Gerard Benjamin Lionel. Exoplanet Imaging Speckle Subtraction: Current Limitations and a Path Forward. 2020.

The LSST Camera Overview: Design and Performance. // . VII 2008. 7014. 70140C.

Give'On Amir, Belikov Ruslan, Shaklan Stuart, Kasdin Jeremy. Closed loop, DM diversity-based, wavefront correction algorithm for high contrast imaging systems // Optics Express. I 2007. 15, 19. 12338.

Give'on Amir, Kern Brian, Shaklan Stuart, Moody Dwight C., Pueyo Laurent. Broadband Wavefront Correction Algorithm for High-Contrast Imaging Systems // Astronomical Adaptive Optics Systems and Applications III. 6691. IX 2007. 63–73.

Goldstein Herbert, Poole Charles P., Safko John L. Classical Mechanics. San Francisco Munich: Addison Wesley, 2008. 3. ed., [Nachdr.].

Goździewski Krzysztof, Migaszewski Cezary. Multiple Mean Motion Resonances in the HR 8799 Planetary System // Monthly Notices of the Royal Astronomical Society. VI 2014. 440. 3140–3171.

Goździewski Krzysztof, Migaszewski Cezary. The Orbital Architecture and Debris Disks of the HR 8799 Planetary System // The Astrophysical Journal Supplement Series. IX 2018. 238. 6.

- Goździewski Krzysztof, Migaszewski Cezary. An Exact, Generalized Laplace Resonance in the HR 8799 Planetary System // *ApJL*. X 2020. 902, 2. L40.
- Grandmont Frédéric. Développement d'un Spectromètre Imageur à Transformée de Fourier Pour l'astronomie / Développement d'un Spectromètre Imageur à Transformée de Fourier Pour l'astronomie / Development of an Imaging Spectrometer with Fourier Transform for Astronomy; I 2006.
- Grant Michael, Boyd Stephen, Ye Yinyu. Disciplined Convex Programming // *Global Optimization*. 84. Boston: Kluwer Academic Publishers, 2006. 155–210.
- Greenaway Alain H. Satellite-Borne Laser for Adaptive Optics Reference // *Space Astronomical Telescopes and Instruments*. 1494. IX 1991. 386–393.
- Greenbaum Alexandra Z., Pueyo Laurent, Ruffio Jean-Baptiste, Wang Jason J., Rosa Robert J. De, Aguilar Jonathan, Rameau Julien, Barman Travis, Marois Christian, Marley Mark S., Konopacky Quinn, Rajan Abhijith, Macintosh Bruce, Ansdell Megan, Arriaga Pauline, Bailey Vanessa P., Bulger Joanna, Burrows Adam S., Chilcote Jeffrey, Cotten Tara, Doyon Rene, Duchêne Gaspard, Fitzgerald Michael P., Follette Katherine B., Gerard Benjamin, Goodsell Stephen J., Graham James R., Hibon Pascale, Hung Li-Wei, Ingraham Patrick, Kalas Paul, Larkin James E., Maire Jérôme, Marchis Franck, Metchev Stanimir, Millar-Blanchaer Maxwell A., Nielsen Eric L., Norton Andrew, Oppenheimer Rebecca, Palmer David, Patience Jennifer, Perrin Marshall D., Poyneer Lisa, Rantakyö Fredrik T., Savransky Dmitry, Schneider Adam C., Sivaramakrishnan Anand, Song Inseok, Soummer Rémi, Thomas Sandrine, Wallace J. Kent, Ward-Duong Kimberly, Wiktorowicz Sloane, Wolff Schuyler. GPI Spectra of HR 8799 c, d, and e from 1.5 to 2.4 μ m with KLIP Forward Modeling // *AJ*. V 2018. 155, 6. 226.
- Guyon Olivier. Imaging Faint Sources within a Speckle Halo with Synchronous Interferometric Speckle Subtraction // *ApJ*. XI 2004. 615, 1. 562–572.
- HI4PI Collaboration, Ben Bekhti N., Flöer L., Keller R., Kerp J., Lenz D., Winkel B., Bailin J., Calabretta M. R., Dedes L., Ford H. A., Gibson B. K., Haud U., Janowiecki S., Kalberla P. M. W., Lockman F. J., McClure-Griffiths N. M., Murphy T., Nakanishi H., Pisano D. J., Staveley-Smith L. HI4PI: A full-sky H I survey based on EBHIS and GASS // *A&A*. X 2016. 594. A116.
- Hardegree-Ullman Kevin K., Apai Dániel, Bergsten Galen J., Pascucci Ilaria, López-Morales Mercedes. Bioverse: A Comprehensive Assessment of the Capabilities of Extremely Large Telescopes to Probe Earth-like O₂ Levels in Nearby Transiting Habitable Zone Exoplanets. IV 2023.
- Hartkopf William I., McAlister Harold A., Franz Otto G. Binary Star Orbits from Speckle Interferometry. II. Combined Visual/Speckle Orbits of 28 Close Systems // *The Astronomical Journal*. IX 1989. 98. 1014.
- Hashimoto Jun, Aoyama Yuhiko, Konishi Mihoko, Uyama Taichi, Takasao Shinsuke, Ikoma Masahiro, Tanigawa Takayuki. Accretion Properties of PDS 70b with MUSE // *AJ*. IV 2020. 159, 5. 222.
- Hedman M. M., Nicholson P. D. Kronoseismology: Using Density Waves in Saturn's C Ring to Probe the Planet's Interior // *The Astronomical Journal*. VII 2013. 146. 12.

- Hoch Kielan K. W., Konopacky Quinn M., Theissen Christopher A., Ruffio Jean-Baptiste, Barman Travis S., Perrin Marshall D., Macintosh Bruce, Marois Christian. Assessing the C/O Ratio Formation Diagnostic: A Potential Trend with Companion Mass. XII 2022.
- Hoffman Matthew D., Gelman Andrew. The No-u-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo // *Journal of Machine Learning Research*. 2014. 15, 47. 1593–1623.
- Householder Aaron, Weiss Lauren. The Inconsistent use of ω in the RV Equation. 2023.
- Iglesias-Marzoa Ramón, López-Morales Mercedes, Jesús Arévalo Morales María. The Rvfit Code: A Detailed Adaptive Simulated Annealing Code for Fitting Binaries and Exoplanets Radial Velocities // *Publications of the Astronomical Society of the Pacific*. VI 2015. 127. 567.
- Détection et Caractérisation de Naines Brunes et Exoplanètes Avec Un Filtre Accordable Pour Applications Dans l'espace. // . III 2013.
- J.H. . Wobbly Evidence // *Scientific American*. 1988. 259, 5. 24–26.
- Johnson Adam B., *al. et.* Blinking the fringes, initial development and results of the Ultra-Low Speed Optical Chopper for the Self-Coherent Camera // *Adaptive Optics Systems VIII*. 12185-271. 2022. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series,).
- Johnson Samson A., Penny Matthew, Gaudi B. Scott, Kerins Eamonn, Rattenbury Nicholas J., Robin Annie C., Novati Sebastiano Calchi, Henderson Calen B. Predictions of the Nancy Grace Roman Space Telescope Galactic Exoplanet Survey. II. Free-floating Planet Detection Rates* // *AJ*. VIII 2020. 160, 3. 123.
- Kammerer Jens, Cooper Rachel A., Vandal Thomas, Thatte Deepashri, Martinache Frantz, Sivaramakrishnan Anand, Chaushev Alexander, Stolker Tomas, Lloyd James P., Albert Loïc, Doyon René, Sallum Steph, Perrin Marshall D., Pueyo Laurent, Mérand Antoine, Gallenne Alexandre, Greenbaum Alexandra, Sanchez-Bermudez Joel, Blakely Dori, Johnstone Doug, Volk Kevin, Martel Andre, Goudfrooij Paul, Meyer Michael R., Willott Chris J., Furio Matthew De, Dang Lisa, Radica Michael, Noirot Gaël. The Near Infrared Imager and Slitless Spectrograph for JWST. V. Kernel Phase Imaging and Data Analysis // *PASP*. II 2023. 135, 1043. 014502.
- Kaplan-Lipkin Avi, Macintosh Bruce, Madurowicz Alexander, Sowmya Krishnamurthy, Shapiro Alexander, Krivova Natalie, Solanki Sami K. Multiwavelength Mitigation of Stellar Activity in Astrometric Planet Detection // *The Astronomical Journal*. V 2022. 163. 205.
- Kasdin N. Jeremy, Bailey Vanessa, Mennesson Bertrand, Zellem Robert, Ygouf Marie, Rhodes Jason, Luchik Thomas, Zhao Feng, Riggs A J Eldorado, Seo Byoung-Joon, Krist John, Kern Brian D., Tang Hong, Nemati Bijan, Groff Tyler D., Zimmerman Neil T., Macintosh Bruce A., Turnbull Margaret, Debes John, Douglas Ewan S., Lupu Roxana E. The Nancy Grace Roman Space Telescope Coronagraph Instrument (CGI) Technology Demonstration // *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*. Online Only, United States: SPIE, XII 2020. 194.
- Kepler J. *Astronomia Nova* // G Voegelinus, Heidelberg. 1609.

Kepler Johannes. *Epitome Astronomiae Copernicanae*. Frankfurt, 1618.

Konopacky Q. M., Marois C., Macintosh B. A., Galicher R., Barman T. S., Metchev S. A., Zuckerman B. ASTROMETRIC MONITORING OF THE HR 8799 PLANETS: ORBIT CONSTRAINTS FROM SELF-CONSISTENT MEASUREMENTS // *AJ*. VII 2016. 152, 2. 28.

Konopacky Quinn M., Barman Travis S., Macintosh Bruce A., Marois Christian. Detection of Carbon Monoxide and Water Absorption Lines in an Exoplanet Atmosphere // *Science*. III 2013. 339, 6126. 1398–1401.

Koop Gary. *Bayesian Econometrics*. Chichester ; Hoboken, N.J: J. Wiley, 2003.

Lacour S., Wang J. J., Rodet L., Nowak M., Shangguan J., Beust H., Lagrange A.-M., Abuter R., Amorim A., Asensio-Torres R., Benisty M., Berger J.-P., Blunt S., Boccaletti A., Bohn A., Bolzer M.-L., Bonnefoy M., Bonnet H., Bourdarot G., Brandner W., Cantalloube F., Caselli P., Charnay B., Chauvin G., Choquet E., Christiaens V., Clénet Y., Coudé Du Foresto V., Cridland A., Dembet R., Dexter J., de Zeeuw P. T., Drescher A., Duvert G., Eckart A., Eisenhauer F., Gao F., Garcia P., Garcia Lopez R., Gendron E., Genzel R., Gillessen S., Girard J. H., Haubois X., Heißel G., Henning Th, Hinkley S., Hippler S., Horrobin M., Houllé M., Hubert Z., Jocu L., Kammerer J., Keppler M., Kervella P., Kreidberg L., Lapeyrière V., Le Bouquin J.-B., Léna P., Lutz D., Maire A.-L., Mérand A., Mollière P., Monnier J. D., Mouillet D., Nasedkin E., Ott T., Otten G. P. P. L., Paladini C., Paumard T., Perraut K., Perrin G., Pfuhl O., Rickman E., Pueyo L., Rameau J., Rousset G., Rustamkulov Z., Samland M., Shimizu T., Sing D., Stadler J., Stolker T., Straub O., Straubmeier C., Sturm E., Tacconi L. J., van Dishoeck E. F., Vigan A., Vincent F., von Fellenberg S. D., Ward-Duong K., Widmann F., Wieprecht E., Wiezorrek E., Woillez J., Yazici S., Young A., GRAVITY Collaboration . The Mass of β Pictoris c from β Pictoris b Orbital Motion // *Astronomy & Astrophysics*. X 2021. 654. L2.

Lafrenière David, Doyon René, Marois Christian, Nadeau Daniel, Oppenheimer Ben R., Roche Patrick F., Rigaut François, Graham James R., Jayawardhana Ray, Johnstone Doug, Kalas Paul G., Macintosh Bruce, Racine René. The Gemini Deep Planet Survey // *The Astrophysical Journal*. XII 2007a. 670. 1367–1390.

Lafrenière David, Marois Christian, Doyon René, Nadeau Daniel, Artigau Étienne. A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging // *The Astrophysical Journal*. V 2007b. 660. 770–780.

Lardière O., *al. et.* Optical design of SPIDERS, a Subaru Pathfinder Instrument for Detecting Exoplanets & Retrieving Spectra // *Adaptive Optics Systems VIII*. 12185-156. 2022. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series).

Lardière Olivier, Gerard Benjamin, Thompson William, Marois Christian, Véran Jean-Pierre, Blain Célia, Heidrich Wolfgang, Fu Qiang. Optical Design and Preliminary Results of NEW EARTH, First Canadian High-Contrast Imaging Laboratory Test Bench // *Adaptive Optics Systems VII*. 11448. XII 2020. 114486Y.

Le Coroller H., Nowak M., Delorme P., Chauvin G., Gratton R., Devinat M., Bec-Canet J., Schneeberger A., Estevez D., Arnold L., Beust H., Bonnefoy M., Boccaletti A., Desgrange C., Desidera S., Galicher R., Lagrange A. M., Langlois M., Maire A. L., Menard F., Vernazza P., Vigan A., Zurlo A., Fenouillet T., Lambert J. C., Bonavita M., Cheetham A., D'orazi V., Feldt M., Janson M., Ligi R., Mesa D., Meyer M., Samland M., Sissa E., Beuzit J. L.,

- Dohlen K., Fusco T., Le Mignant D., Mouillet D., Ramos J., Rochat S., Sauvage J. F. K-Stacker: An Algorithm to Hack the Orbital Parameters of Planets Hidden in High-Contrast Imaging. First Applications to VLT/SPHERE Multi-Epoch Observations // *Astronomy and Astrophysics*. VII 2020. 639. A113.
- Le Coroller H., Nowak M., Wagner K., Kasper M., Chauvin G., Desgrange C., *al et.* Efficiently Combining Alpha CenA Multi-Epoch High-Contrast Imaging Data: Application of K-Stacker to the 80 Hr NEAR Campaign // *A&A*. IX 2022.
- Liu Michael C., Dupuy Trent J., Allers Katelyn N. The Hawaii Infrared Parallax Program. II. Young Ultracool Field Dwarfs // *The Astrophysical Journal*. XII 2016. 833. 96.
- Llop-Sayson Jorge, Wang Jason J., Ruffio Jean-Baptiste, Mawet Dimitri, Blunt Sarah, Absil Olivier, Bond Charlotte, Brinkman Casey, Bowler Brendan P., Bottom Michael, Chontos Ashley, Dalba Paul A., Fulton B. J., Giacalone Steven, Hill Michelle, Hirsch Lea A., Howard Andrew W., Isaacson Howard, Karlsson Mikael, Lubin Jack, Madurowicz Alex, Matthews Keith, Morris Evan, Perrin Marshall, Ren Bin, Rice Malena, Rosenthal Lee J., Ruane Garreth, Rubenzahl Ryan, Sun He, Wallack Nicole, Xuan Jerry W., Ygouf Marie. Constraining the Orbit and Mass of Epsilon Eridani b with Radial Velocities, Hipparcos IAD-Gaia DR2 Astrometry, and Multiepoch Vortex Coronagraphy Upper Limits // *The Astronomical Journal*. XI 2021. 162. 181.
- Lustig-Yaeger Jacob, Meadows Victoria S., Mendoza Guadalupe Tovar, Schwieterman Edward W., Fujii Yuka, Luger Rodrigo, Robinson Tyler D. Detecting Ocean Glint on Exoplanets Using Multiphase Mapping // *AJ*. XII 2018. 156, 6. 301.
- Lyot Bernard. La Couronne Solaire Étudiée En Dehors Des Éclipses // *Comptes rendus*. 1930. 191. 834–837.
- Macintosh B., Graham J. R., Barman T., Rosa R. J. De, Konopacky Q., Marley M. S., Marois C., Nielsen E. L., Pueyo L., Rajan A., Rameau J., Saumon D., Wang J. J., Patience J., Ammons M., Arriaga P., Artigau E., Beckwith S., Brewster J., Bruzzone S., Bulger J., Burningham B., Burrows A. S., Chen C., Chiang E., Chilcote J. K., Dawson R. I., Dong R., Doyon R., Draper Z. H., Duchêne G., Esposito T. M., Fabrycky D., Fitzgerald M. P., Follette K. B., Fortney J. J., Gerard B., Goodsell S., Greenbaum A. Z., Hiben P., Hinkley S., Cotten T. H., Hung L.-W., Ingraham P., Johnson-Groh M., Kalas P., Lafreniere D., Larkin J. E., Lee J., Line M., Long D., Maire J., Marchis F., Matthews B. C., Max C. E., Metchev S., Millar-Blanchaer M. A., Mittal T., Morley C. V., Morzinski K. M., Murray-Clay R., Oppenheimer R., Palmer D. W., Patel R., Perrin M. D., Poyneer L. A., Rafikov R. R., Rantakyro F. T., Rice E. L., Rojo P., Rudy A. R., Ruffio J.-B., Ruiz M. T., Sadakuni N., Saddlemyer L., Salama M., Savransky D., Schneider A. C., Sivaramakrishnan A., Song I., Soummer R., Thomas S., Vasisht G., Wallace J. K., Ward-Duong K., Wiktorowicz S. J., Wolff S. G., Zuckerman B. Discovery and Spectroscopy of the Young Jovian Planet 51 Eri b with the Gemini Planet Imager // *Science*. X 2015. 350, 6256. 64–67.
- Macintosh Bruce, Graham James R., Ingraham Patrick, Konopacky Quinn, Marois Christian, Perrin Marshall, Poyneer Lisa, Bauman Brian, Barman Travis, Burrows Adam S., Cardwell Andrew, Chilcote Jeffrey, De Rosa Robert J., Dillon Daren, Doyon Rene, Dunn Jennifer, Erikson Darren, Fitzgerald Michael P., Gavel Donald, Goodsell Stephen, Hartung Markus, Hiben Pascale, Kalas Paul, Larkin James, Maire Jerome, Marchis Franck,

- Marley Mark S., McBride James, Millar-Blanchaer Max, Morzinski Katie, Norton Andrew, Oppenheimer B. R., Palmer David, Patience Jennifer, Pueyo Laurent, Rantakyro Fredrik, Sadakuni Naru, Saddlemeyer Leslie, Savransky Dmitry, Serio Andrew, Soummer Remi, Sivaramakrishnan Anand, Song Inseok, Thomas Sandrine, Wallace J. Kent, Wiktorowicz Sloane, Wolff Schuyler. First Light of the Gemini Planet Imager // Proceedings of the National Academy of Science. IX 2014. 111. 12661–12666.
- Maire A.-L., Skemer A. J., Hinz P. M., Desidera S., Esposito S., Gratton R., Marzari F., Skrutskie M. F., Biller B. A., Defrère D., Bailey V. P., Leisenring J. M., Apai D., Bonnefoy M., Brandner W., Buenzli E., Claudi R. U., Close L. M., Crepp J. R., De Rosa R. J., Eisner J. A., Fortney J. J., Henning T., Hofmann K.-H., Kopytova T. G., Males J. R., Mesa D., Morzinski K. M., Oza A., Patience J., Pinna E., Rajan A., Schertl D., Schlieder J. E., Su K. Y. L., Vaz A., Ward-Duong K., Weigelt G., Woodward C. E. The LEECH Exoplanet Imaging Survey. Further Constraints on the Planet Architecture of the HR 8799 System // Astronomy & Astrophysics. IV 2015. 576.
- Males Jared R., Belikov Ruslan, Bendek Eduardo. Orbital Differential Imaging: A New High-Contrast Post-Processing Technique for Direct Imaging of Exoplanets // Techniques and Instrumentation for Detection of Exoplanets VII. 9605. IX 2015. 414–422.
- Males Jared R., Skemer Andrew J., Close Laird M. Direct Imaging in the Habitable Zone and the Problem of Orbital Motion // The Astrophysical Journal. VII 2013. 771. 10.
- Markley F. Landis. Kepler Equation Solver // Celestial Mech Dyn Astr. III 1995. 63, 1. 101–111.
- Marley Mark, Saumon Didier, Morley Caroline, Fortney Jonathan, Visscher Channon, Freedman Richard, Lupu Roxana. Sonora Bobcat: Cloud-Free, Substellar Atmosphere Models, Spectra, Photometry, Evolution, and Chemistry. VII 2021.
- Marlow W., Carlton A., Yoon H., Clark J., Haughtwout C., Cahoy K., Males J., Close L., Morzinski K. Laser Guidestar Satellite for Ground-based Adaptive Optics Imaging of Geosynchronous Satellites // Advanced Maui Optical and Space Surveillance Technologies Conference. IX 2016. 99.
- Effects of Quasi-Static Aberrations in Faint Companion Searches. // . 2003. 8. 233–243.
- Marois C., Macintosh B., Barman T., Zuckerman B., Song I., Patience J., Lafreniere D., Doyon R. Direct Imaging of Multiple Planets Orbiting the Star HR 8799 // Science. XI 2008a. 322, 5906. 1348–1352.
- Marois C., *al. et.* Deployment of focal plane WFS technologies on 8-m telescopes: from the Subaru SPIDERS pathfinder, to the facility-class GPI 2.0 CAL2 system // Adaptive Optics Systems VIII. 12185-70. 2022. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series,).
- La recherche de naines brunes et d'exoplanètes : développement d'une technique d'imagerie multibande. // . 2004.
- Marois Christian, Correia Carlos, Galicher Raphael, Ingraham Patrick, Macintosh Bruce, Currie Thayne, Rosa Rob De. GPI PSF Subtraction with TLOCI: The next Evolution in Exoplanet/Disk High-Contrast Imaging // Adaptive Optics Systems IV. 9148. VII 2014. 287–299.

- Marois Christian, Doyon René, Nadeau Daniel, Racine René, Riopel Martin, Vallée Philippe, Lafrenière David.* TRIDENT: An Infrared Differential Imaging Camera Optimized for the Detection of Methanated Substellar Companions // Publications of the Astronomical Society of the Pacific. VII 2005. 117. 745–756.
- Marois Christian, Doyon René, Racine René, Nadeau Daniel.* Efficient Speckle Noise Attenuation in Faint Companion Imaging // PUBL ASTRON SOC PAC. I 2000. 112, 767. 91–96.
- Marois Christian, Gerard B., Lardière O., Anthony A., Bradley C., Dunn J., Fu Q., Hardy T., Heidrich W., Herriot G., Nielsen E., Sivanandam S., Sivransky D., Thibault S., Thompson W., Véran J.-P.* Upgrading the Gemini Planet Imager Calibration Unit with a Photon Counting Focal Plane Wavefront Sensor // Adaptive Optics Systems VII. 11448. XII 2020. 1144873.
- Marois Christian, Lafrenière David, Doyon René, Macintosh Bruce, Nadeau Daniel.* Angular Differential Imaging: A Powerful High-Contrast Imaging Technique* // ApJ. IV 2006a. 641, 1. 556.
- Marois Christian, Lafrenière David, Macintosh Bruce, Doyon René.* Accurate Astrometry and Photometry of Saturated and Coronagraphic Point Spread Functions // The Astrophysical Journal. VIII 2006b. 647. 612–619.
- Marois Christian, Lafrenière David, Macintosh Bruce, Doyon René.* Confidence Level and Sensitivity Limits in High-Contrast Imaging // ApJ. I 2008b. 673, 1. 647.
- Marois Christian, Lardière Olivier, Thompson William, Singh Garima, Johnson Adam, Hardy Tim, Fitzsimmons Joeleff, Gerard Benjamin L., Sivanandam Suresh, Thibault Simon, Savransky Dmitry, Bradley Colin, Jensen-Clem Rebecca, Demers Mathieu, Fu Qiang, Heidrich Wolfgang, N'Diaye Mamadou.* Deployment of Focal Plane WFS Technologies on 8-m Telescopes: From the Subaru SPIDERS Pathfinder, to the Facility-Class GPI 2.0 CAL2 System // Adaptive Optics Systems VIII. 12185. VIII 2022. 594–603.
- Marois Christian, Macintosh Bruce, Véran Jean-Pierre.* Exoplanet Imaging with LOCI Processing: Photometry and Astrometry with the New SOSIE Pipeline // Adaptive Optics Systems II. 7736. VII 2010a. 77361J.
- Exoplanet Detection with Simultaneous Spectral Differential Imaging: Effects of out-of-Pupil-Plane Optical Aberrations. // . VI 2006c. 6269. 62693M.
- A Fresnel Propagation Analysis of NFIRAOS/IRIS High-Contrast Exoplanet Imaging Capabilities. // . VII 2012. 8447. 844726.
- Marois Christian, Zuckerman B., Konopacky Quinn M., Macintosh Bruce, Barman Travis.* Images of a Fourth Planet Orbiting HR 8799 // Nature. XII 2010b. 468, 7327. 1080–1083.
- Matthews Brenda, Kennedy Grant, Sibthorpe Bruce, Booth Mark, Wyatt Mark, Broekhoven-Fiene Hannah, Macintosh Bruce, Marois Christian.* Resolved Imaging of the HR 8799 Debris Disk with Herschel // The Astrophysical Journal. I 2014. 780. 97.

- Mawet D., Milli J., Wahhaj Z., Pelat D., Absil O., Delacroix C., Boccaletti A., Kasper M., Kenworthy M., Marois C., Mennesson B., Pueyo L.* FUNDAMENTAL LIMITATIONS OF HIGH CONTRAST IMAGING SET BY SMALL SAMPLE STATISTICS // *ApJ*. VIII 2014. 792, 2. 97.
- Mawet Dimitri, Hirsch Lea, Lee Eve J., Ruffio Jean-Baptiste, Bottom Michael, Fulton Benjamin J., Absil Olivier, Beichman Charles, Bowler Brendan, Bryan Marta, Choquet Elodie, Ciardi David, Christiaens Valentin, Defrère Denis, Gomez Gonzalez Carlos Alberto, Howard Andrew W., Huby Elsa, Isaacson Howard, Jensen-Clem Rebecca, Kosiarek Molly, Marcy Geoff, Meshkat Tiffany, Petigura Erik, Reggiani Maddalena, Ruane Garreth, Serabyn Eugene, Sinukoff Evan, Wang Ji, Weiss Lauren, Ygouf Marie.* Deep Exploration of ϵ Eridani with Keck Ms-band Vortex Coronagraphy and Radial Velocities: Mass and Orbital Parameters of the Giant Exoplanet // *The Astronomical Journal*. I 2019. 157. 33.
- Mayor Michel, Queloz Didier.* A Jupiter-mass Companion to a Solar-Type Star // *Nature*. XI 1995. 378. 355–359.
- Scientific Design of a High Contrast Integral Field Spectrograph for the Subaru Telescope. // . IX 2012. 8446. 84469C.
- Mennesson Bertrand, Kasdin N. Jeremy, Macintosh Bruce, Turnbull Margaret, Douglas E., Stark C., Rizzo M., Krist J., Trauger John, Rhodes Jason, Moustakas Leonidas, Frerking Margaret A., Zhao Feng, Poberezhskiy Ilya Y., Demers Richard T., Roberge Aki, Debes John H., Nemati B., Zimmerman N., Cahoy K., Bailey V.* The WFIRST Coronagraph Instrument: A Major Step in the Exploration of Sun-like Planetary Systems via Direct Imaging // *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*. Austin, United States: SPIE, VIII 2018. 88.
- Meshkat Tiffany, Kenworthy Matthew A., Quanz Sascha P., Amara Adam.* Optimized Principal Component Analysis on Coronagraphic Images of the Fomalhaut System // *The Astrophysical Journal*. I 2014. 780. 17.
- Meyer L., Ghez A. M., Schödel R., Yelda S., Boehle A., Lu J. R., Do T., Morris M. R., Becklin E. E., Matthews K.* The Shortest-Known-Period Star Orbiting Our Galaxy's Supermassive Black Hole // *Science*. X 2012. 338, 6103. 84–87.
- Mogensen Patrick K., Riseth Asbjørn N.* Optim: A Mathematical Optimization Package for Julia // *Journal of Open Source Software*. IV 2018. 3, 24. 615.
- N'Diaye M., Vigan A., Dohlen K., Sauvage J.-F., Caillat A., Costille A., Girard J. H. V., Beuzit J.-L., Fusco T., Blanchard P., Le Merrer J., Le Mignant D., Madec F., Moreaux G., Mouillet D., Puget P., Zins G.* ZELDA, a Zernike Wavefront Sensor for the Fine Measurement of Quasi-Static Aberrations in Coronagraphic Systems: Concept Studies and Results with VLT/SPHERE // *SPIE Astronomical Telescopes + Instrumentation*. Edinburgh, United Kingdom, VII 2016. 99096S.
- Nagpal Vighnesh, Blunt Sarah, Bowler Brendan P., Dupuy Trent J., Nielsen Eric L., Wang Jason J.* The Impact of Bayesian Hyperpriors on the Population-level Eccentricity Distribution of Imaged Planets // *The Astronomical Journal*. II 2023. 165. 32.

Nielsen Eric L., De Rosa Robert J., Macintosh Bruce, Wang Jason J., Ruffio Jean-Baptiste, Chiang Eugene, Marley Mark S., Saumon Didier, Savransky Dmitry, Ammons S. Mark, Bailey Vanessa P., Barman Travis, Blain Célia, Bulger Joanna, Burrows Adam, Chilcote Jeffrey, Cotten Tara, Czekala Ian, Doyon Rene, Duchêne Gaspard, Esposito Thomas M., Fabrycky Daniel, Fitzgerald Michael P., Follette Katherine B., Fortney Jonathan J., Gerard Benjamin L., Goodsell Stephen J., Graham James R., Greenbaum Alexandra Z., Hibon Pascale, Hinkley Sasha, Hirsch Lea A., Hom Justin, Hung Li-Wei, Dawson Rebekah Ilene, Ingraham Patrick, Kalas Paul, Konopacky Quinn, Larkin James E., Lee Eve J., Lin Jonathan W., Maire Jérôme, Marchis Franck, Marois Christian, Metchev Stanimir, Millar-Blanchaer Maxwell A., Morzinski Katie M., Oppenheimer Rebecca, Palmer David, Patience Jennifer, Perrin Marshall, Poyneer Lisa, Pueyo Laurent, Rafikov Roman R., Rajan Abhijith, Rameau Julien, Rantakyö Fredrik T., Ren Bin, Schneider Adam C., Sivaramakrishnan Anand, Song Inseok, Soummer Remi, Tallis Melisa, Thomas Sandrine, Ward-Duong Kimberly, Wolff Schuyler. The Gemini Planet Imager Exoplanet Survey: Giant Planet and Brown Dwarf Demographics from 10 to 100 Au // *The Astronomical Journal*. VII 2019. 158. 13.

Nielsen Eric L., Rosa Robert J. De, Wang Jason J., Sahlmann Johannes, Kalas Paul, Duchêne Gaspard, Rameau Julien, Marley Mark S., Saumon Didier, Macintosh Bruce, Millar-Blanchaer Maxwell A., Nguyen Meiji M., Ammons S. Mark, Bailey Vanessa P., Barman Travis, Bulger Joanna, Chilcote Jeffrey, Cotten Tara, Doyon Rene, Esposito Thomas M., Fitzgerald Michael P., Follette Katherine B., Gerard Benjamin L., Goodsell Stephen J., Graham James R., Greenbaum Alexandra Z., Hibon Pascale, Hung Li-Wei, Ingraham Patrick, Konopacky Quinn, Larkin James E., Maire Jérôme, Marchis Franck, Marois Christian, Metchev Stanimir, Oppenheimer Rebecca, Palmer David, Patience Jennifer, Perrin Marshall, Poyneer Lisa, Pueyo Laurent, Rajan Abhijith, Rantakyö Fredrik T., Ruffio Jean-Baptiste, Savransky Dmitry, Schneider Adam C., Sivaramakrishnan Anand, Song Inseok, Soummer Remi, Thomas Sandrine, Wallace J. Kent, Ward-Duong Kimberly, Wiktorowicz Sloane, Wolff Schuyler. The Gemini Planet Imager Exoplanet Survey: Dynamical Mass of the Exoplanet \upbeta Pictoris b from Combined Direct Imaging and Astrometry // *AJ*. I 2020. 159, 2. 71.

Nise Norman S. Control Systems Engineering. VI 2020.

Nowak M., Coroller H. Le, Arnold L., Dohlen K., Estevez D., Fusco T., Sauvage J.-F., Vigan A. K-Stacker: Keplerian Image Recombination for the Direct Detection of Exoplanets // *A&A*. VII 2018. 615. A144.

O'Neil K. Kosmo, Martinez G. D., Hees A., Ghez A. M., Do T., Witzel G., Konopacky Q., Becklin E. E., Chu D. S., Lu J. R., Matthews K., Sakai S. Improving Orbit Estimates for Incomplete Orbits with a New Approach to Priors: With Applications from Black Holes to Planets // *AJ*. VI 2019. 158, 1. 4.

Perryman Michael, Hartman Joel, Bakos Gáspár Á, Lindegren Lennart. ASTROMETRIC EXOPLANET DETECTION WITH GAIA // *ApJ*. XI 2014. 797, 1. 14.

Philcox Oliver H E, Goodman Jeremy, Slepian Zachary. Kepler's Goat Herd: An Exact Solution to Kepler's Equation for Elliptical Orbits // *Monthly Notices of the Royal Astronomical Society*. X 2021. 506, 4. 6111–6116.

Pogorelyuk Leonid, Fitzgerald Riley, Vlahakis Sophia, Morgan Rhonda, Cahoy Kerri. Deconfusing Detections in Directly Imaged Multiplanet Systems* // *ApJ*. IX 2022. 937, 2. 66.

- Potier, *al. et.* Improving VLT/SPHERE without additional hardware: First on-sky calibration of the quasi-static aberrations with the dark hole technique // *Adaptive Optics Systems VIII*. 12185-236. 2022. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series).
- Potier A., Mazoyer J., Wahhaj Z., Baudoz P., Chauvin G., Galicher R., Ruane G. Increasing the Raw Contrast of VLT/SPHERE with the Dark Hole Technique. II. On-sky Wavefront Correction and Coherent Differential Imaging // *Astronomy and Astrophysics*. IX 2022. 665. A136.
- Pueyo L., Soummer R., Hoffmann J., Oppenheimer R., Graham J. R., Zimmerman N., Zhai C., Wallace J. K., Vesceles F., Veicht A., Vasisht G., Truong T., Sivaramakrishnan A., Shao M., Roberts L. C. Jr., Roberts J. E., Rice E., Parry I. R., Nilsson R., Lockhart T., Ligon E. R., King D., Hinkley S., Hillenbrand L., Hale D., Dekany R., Crepp J. R., Cady E., Burruss R., Brenner D., Beichman C., Baranec C. Reconnaissance of the HR 8799 Exosolar System. II. Astrometry and Orbital Motion // *The Astrophysical Journal*. IV 2015. 803. 31.
- Pueyo Laurent, Crepp Justin R., Vasisht Gautam, Brenner Douglas, Oppenheimer Ben R., Zimmerman Neil, Hinkley Sasha, Parry Ian, Beichman Charles, Hillenbrand Lynne, Roberts Lewis C., Dekany Richard, Shao Mike, Burruss Rick, Bouchez Antonin, Roberts Jenny, Soummer Rémi. Application of a Damped Locally Optimized Combination of Images Method to the Spectral Characterization of Faint Companions Using an Integral Field Spectrograph // *ApJS*. III 2012. 199, 1. 6.
- Racine René, Walker Gordon A. H., Nadeau Daniel, Doyon René, Marois Christian. Speckle Noise and the Detection of Faint Companions // *PASP*. V 1999. 111, 759. 587.
- Ragazzoni Roberto. Pupil Plane Wavefront Sensing with an Oscillating Prism // *Journal of Modern Optics*. II 1996. 43, 2. 289-293.
- Raymond Sean N., Morbidelli Alessandro. The Grand Tack Model: A Critical Review // *Proceedings of the International Astronomical Union*. VII 2014. 9, S310. 194-203.
- Rein Hanno, Tamayo Daniel. WHFAST: A Fast and Unbiased Implementation of a Symplectic Wisdom-Holman Integrator for Long-Term Gravitational Simulations // *Monthly Notices of the Royal Astronomical Society*. IX 2015. 452. 376-388.
- Revels Jarrett, Lubin Miles, Papamarkou Theodore. Forward-Mode Automatic Differentiation in Julia // *arXiv:1607.07892 [cs]*. VII 2016.
- Rhee Joseph H., Song Inseok, Zuckerman B., McElwain Michael. Characterization of Dusty Debris Disks: The IRAS and Hipparcos Catalogs // *ApJ*. V 2007. 660, 2. 1556.
- Ricker George R., Winn Joshua N., Vanderspek Roland, Latham David W., Bakos Gáspár Á, Bean Jacob L., Berta-Thompson Zachory K., Brown Timothy M., Buchhave Lars, Butler Nathaniel R., Butler R. Paul, Chaplin William J., Charbonneau David B., Christensen-Dalsgaard Jørgen, Clampin Mark, Deming Drake, Doty John P., Lee Nathan De, Dressing Courtney, Dunham Edward W., Endl Michael, Fressin François, Ge Jian, Henning Thomas, Holman Matthew J., Howard Andrew W., Ida Shigeru, Jenkins Jon M., Jernigan Garrett, Johnson John Asher, Kaltenegger Lisa, Kawai Nobuyuki, Kjeldsen Hans, Laughlin Gregory, Levine Alan M., Lin Douglas,

- Lissauer Jack J., MacQueen Phillip, Marcy Geoffrey, McCullough Peter R., Morton Timothy D., Narita Norio, Paegert Martin, Palle Enric, Pepe Francesco, Pepper Joshua, Quirrenbach Andreas, Rinehart Stephen A., Sasselov Dimitar, Sato Bun'ei, Seager Sara, Sozzetti Alessandro, Stassun Keivan G., Sullivan Peter, Szentgyorgyi Andrew, Torres Guillermo, Udry Stephane, Villaseñor Joel. Transiting Exoplanet Survey Satellite // JATIS. X 2014. 1, 1. 014003.
- Ruffio Jean-Baptiste, Horstman Katelyn, Mawet Dimitri, Rosenthal Lee J., Batygin Konstantin, Wang Jason J., Millar-Blanchaer Maxwell, Wang Ji, Fulton Benjamin J., Konopacky Quinn M., Agrawal Shubh, Hirsch Lea A., Howard Andrew W., Blunt Sarah, Nielsen Eric, Baker Ashley, Bartos Randall, Bond Charlotte Z., Calvin Benjamin, Cetre Sylvain, Delorme Jacques-Robert, Doppmann Greg, Echeverri Daniel, Finnerty Luke, Fitzgerald Michael P., Jovanovic Nemanja, López Ronald, Martin Emily C., Morris Evan, Pezzato Jacklyn, Ruane Garreth, Sappéy Ben, Schofield Tobias, Skemer Andrew, Venenciano Taylor, Wallace J. Kent, Wallack Nicole L., Wizinowich Peter, Xuan Jerry W. Detecting Exomoons from Radial Velocity Measurements of Self-luminous Planets: Application to Observations of HR 7672 B and Future Prospects // AJ. III 2023. 165, 3. 113.
- Ruffio Jean-Baptiste, Macintosh Bruce, Konopacky Quinn M., Barman Travis, De Rosa Robert J., Wang Jason J., Wilcomb Kielan K., Czekala Ian, Marois Christian. Radial Velocity Measurements of HR 8799 b and c with Medium Resolution Spectroscopy // AJ. X 2019. 158, 5. 200.
- Ruffio Jean-Baptiste, Macintosh Bruce, Wang Jason J., Pueyo Laurent, Nielsen Eric L., Rosa Robert J. De, Czekala Ian, Marley Mark S., Arriaga Pauline, Bailey Vanessa P., Barman Travis, Bulger Joanna, Chilcote Jeffrey, Cotten Tara, Doyon Rene, Duchêne Gaspard, Fitzgerald Michael P., Follette Katherine B., Gerard Benjamin L., Goodsell Stephen J., Graham James R., Greenbaum Alexandra Z., Hiben Pascale, Hung Li-Wei, Ingraham Patrick, Kalas Paul, Konopacky Quinn, Larkin James E., Maire Jérôme, Marchis Franck, Marois Christian, Metchev Stanimir, Millar-Blanchaer Maxwell A., Morzinski Katie M., Oppenheimer Rebecca, Palmer David, Patience Jennifer, Perrin Marshall, Poyneer Lisa, Rajan Abhijith, Rameau Julien, Rantakyö Fredrik T., Savransky Dmitry, Schneider Adam C., Sivaramakrishnan Anand, Song Inseok, Soummer Remi, Thomas Sandrine, Wallace J. Kent, Ward-Duong Kimberly, Wiktorowicz Sloane, Wolff Schuyler. Improving and Assessing Planet Sensitivity of the GPI Exoplanet Survey with a Forward Model Matched Filter // ApJ. VI 2017. 842, 1. 14.
- Ruffio Jean-Baptiste, Mawet Dimitri, Czekala Ian, Macintosh Bruce, De Rosa Robert J., Ruane Garreth, Bottom Michael, Pueyo Laurent, Wang Jason J., Hirsch Lea, Zhu Zhaohuan, Nielsen Eric L. A Bayesian Framework for Exoplanet Direct Detection and Non-detection // The Astronomical Journal. XI 2018. 156. 196.
- Ryu Tsuguru, Sato Bun'ei, Kuzuhara Masayuki, Narita Norio, Takahashi Yasuhiro H., Uyama Taichi, Kudo Tomoyuki, Kusakabe Nobuhiko, Hashimoto Jun, Omiya Masashi, Harakawa Hiroki, Abe Lyu, Ando Hiroyasu, Brandner Wolfgang, Brandt Timothy D., Carson Joseph C., Currie Thayne, Egner Sebastian, Feldt Markus, Goto Miwa, Grady Carol A., Guyon Olivier, Hayano Yutaka, Hayashi Masahiko, Hayashi Saeko S., Helminiak Krzysztof G., Henning Thomas, Hodapp Klaus W., Ida Shigeru, Ishii Miki, Itoh Yoichi, Iye Masanori, Izumiura Hideyuki, Janson Markus, Kambe Eiji, Kandori Ryo, Knapp Gillian R., Kokubo Eiichiro, Kwon Jungmi, Matsuo Taro, Mayama Satoshi, McElwain Michael W., Mede Kyle, Miyama Shoken, Morino Jun-Ichi, Moro-Martin Amaya, Nishimura Tetsuo, Pyo Tae-Soo, Serabyn Eugene, Suenaga Takuya, Suto Hiroshi, Suzuki Ryuji, Takami Michihiro, Takato Naruhisa, Takeda Yoichi, Terada Hiroshi, Thalmann Christian, Turner Edwin L., Watanabe Makoto, Wisniewski John, Yamada Toru, Yoshida Michitoshi, Takami Hideki, Usuda Tomonori, Tamura Motohide.

- HIGH-CONTRAST IMAGING OF INTERMEDIATE-MASS GIANTS WITH LONG-TERM RADIAL VELOCITY TRENDS // *ApJ*. VII 2016. 825, 2. 127.
- Sadakane K., Nishida M.* Twelve Additional "Vega-like" Stars. // *Publications of the Astronomical Society of the Pacific*. VII 1986. 98. 685–689.
- Sauvage J. F., Mugnier L., Paul B., Villecroze R.* Coronagraphic phase diversity: a simple focal plane sensor for high-contrast imaging // *Optics Letters*. XII 2012. 37, 23. 4808.
- Sepulveda Aldo G., Bowler Brendan P.* Dynamical Mass of the Exoplanet Host Star HR 8799 // *AJ*. II 2022. 163, 2. 52.
- Serabyn E., Huby E., Matthews K., Mawet D., Absil O., Femenia B., Wizinowich P., Karlsson M., Bottom M., Campbell R., Carlomagno B., Defrère D., Delacroix C., Forsberg P., Gonzalez C. Gomez, Habraken S., Jolivet A., Liewer K., Lilley S., Piron P., Reggiani M., Surdej J., Tran H., Catalán E. Vargas, Wertz O.* THE W. M. KECK OBSERVATORY INFRARED VORTEX CORONAGRAPH AND A FIRST IMAGE OF HIP 79124 B // *AJ*. I 2017. 153, 1. 43.
- Serabyn Eugene, Wallace J. Kent, Mawet Dimitri.* Speckle-phase measurement in a tandem-vortex coronagraph // *Appl. Opt.*. X 2011. 50, 28. 5453.
- Service M., Lu J. R., Campbell R., Sitarski B. N., Ghez A. M., Anderson J.* A New Distortion Solution for NIRC2 on the Keck II Telescope // *PASP*. IX 2016. 128, 967. 095004.
- Singh G., al. et.* Pupil-plane LLOWFS simulation and laboratory results from NEW-EARTH's high-contrast imaging testbed // *Adaptive Optics Systems VIII*. 12185-192. 2022. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series,).
- Singh Garima, Lozi Julien, Guyon Olivier, Baudoz Pierre, Jovanovic Nemanja, Martinache Frantz, Kudo Tomoyuki, Serabyn Eugene, Kuhn Jonas.* On-Sky Demonstration of Low-Order Wavefront Sensing and Control with Focal Plane Phase Mask Coronagraphs // *PASP*. IX 2015. 127, 955. 857.
- Singh Garima, Martinache Frantz, Baudoz Pierre, Guyon Olivier, Matsuo Taro, Jovanovic Nemanja, Clergeon Christophe.* Lyot-based Low Order Wavefront Sensor for Phase-mask Coronagraphs: Principle, Simulations and Laboratory Experiments // *PASP*. VI 2014. 126, 940. 586.
- Sivaramakrishnan Anand, Oppenheimer Ben R.* Astrometry and Photometry with Coronagraphs // *The Astrophysical Journal*. VIII 2006. 647. 620–629.
- Sivaramakrishnan Anand, Tuthill Peter, Lloyd James P., Greenbaum Alexandra Z., Thatte Deepashri, Cooper Rachel A., Vandal Thomas, Kammerer Jens, Sanchez-Bermudez Joel, Pope Benjamin J. S., Blakely Dori, Albert Loïc, Cook Neil J., Johnstone Doug, Martel André R., Volk Kevin, Soullain Anthony, Artigau Étienne, Lafrenière David, Willott Chris J., Parmentier Sébastien, Ford K. E. Saavik, McKernan Barry, Vila M. Begoña, Rowlands Neil, Doyon René, Beaulieu Mathilde, Desdoigts Louis, Fullerton Alexander W., Furio Matthew De, Goudfrooij Paul, Holfeltz Sherie T., LaMassa Stephanie, Maszkiewicz Michael, Meyer Michael R., Perrin Marshall D., Pueyo Laurent, Sahlmann Johannes, Sohn Sangmo Tony, Teixeira Paula S., Zheng Sheng-hai.* The Near

- Infrared Imager and Slitless Spectrograph for the James Webb Space Telescope. IV. Aperture Masking Interferometry // *PASP*. II 2023. 135, 1043. 015003.
- Skemer Andrew J., Close Laird M.* Sirius B Imaged in the Mid-infrared: No Evidence for a Remnant Planetary System // *The Astrophysical Journal*. III 2011. 730. 53.
- Skemer Andrew J., Hinz Philip M., Esposito Simone, Burrows Adam, Leisenring Jarron, Skrutskie Michael, Desidera Silvano, Mesa Dino, Arcidiacono Carmelo, Mannucci Filippo, Rodigas Timothy J., Close Laird, McCarthy Don, Kulesa Craig, Agapito Guido, Apai Daniel, Argomedo Javier, Bailey Vanessa, Boutsia Konstantina, Briguglio Runa, Brusa Guido, Busoni Lorenzo, Claudi Riccardo, Eisner Joshua, Fini Luca, Follette Katherine B., Garnavich Peter, Gratton Raffaele, Guerra Juan Carlos, Hill John M., Hoffmann William F., Jones Terry, Krejny Megan, Males Jared, Masciadri Elena, Meyer Michael R., Miller Douglas L., Morzinski Katie, Nelson Matthew, Pinna Enrico, Puglisi Alfio, Quanz Sascha P., Quiros-Pacheco Fernando, Riccardi Armando, Stefanini Paolo, Vaitheeswaran Vidhya, Wilson John C., Xompero Marco.* First Light LBT AO Images of HR 8799 Bcde at 1.6 and 3.3 Mm: New Discrepancies between Young Planets and Old Brown Dwarfs // *The Astrophysical Journal*. VII 2012. 753. 14.
- Skemer Andrew J., Marley Mark S., Hinz Philip M., Morzinski Katie M., Skrutskie Michael F., Leisenring Jarron M., Close Laird M., Saumon Didier, Bailey Vanessa P., Briguglio Runa, Defrere Denis, Esposito Simone, Follette Katherine B., Hill John M., Males Jared R., Puglisi Alfio, Rodigas Timothy J., Xompero Marco.* Directly Imaged L-T Transition Exoplanets in the Mid-infrared // *The Astrophysical Journal*. IX 2014. 792. 17.
- Skilling John.* Nested Sampling // *AIP Conference Proceedings*. XI 2004. 735, 1. 395–405.
- Skrutskie M. F., Jones T., Hinz P., Garnavich P., Wilson J., Nelson M., Solheid E., Durney O., Hoffmann W., Vaitheeswaran V., McMahon T., Leisenring J., Wong A.* The Large Binocular Telescope Mid-Infrared Camera (LMIRcam): Final Design and Status // *Ground-Based and Airborne Instrumentation for Astronomy III*. 7735. VII 2010. 77353H.
- Soulain A., Sivaramakrishnan A., Tuthill P., Thatte D., Volk K., Cooper R., Albert L., Artigau É., Cook N., Doyon R., Johnstone D., Lafrenière D., Martel A.* The James Webb Space Telescope aperture masking interferometer // *Optical and Infrared Interferometry and Imaging VII*. 11446. XII 2020. 1144611. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series).
- Soummer Rémi.* Apodized Pupil Lyot Coronagraphs for Arbitrary Telescope Apertures // *ApJ*. XII 2004. 618, 2. L161.
- Soummer Remi, Ferrari Andre, Aime Claude, Jolissaint Laurent.* Speckle Noise and Dynamic Range in Coronagraphic Images // *ApJ*. XI 2007. 669, 1. 642–656.
- Soummer Rémi, Hagan J. Brendan, Pueyo Laurent, Thormann Adrien, Rajan Abhijith, Marois Christian.* Orbital Motion of HR 8799 b, c, d Using Hubble Space Telescope Data from 1998: Constraints on Inclination, Eccentricity, and Stability // *The Astrophysical Journal*. XI 2011. 741. 55.
- Soummer Remi, Pueyo Laurent, Larkin James.* Detection and Characterization of Exoplanets and Disks Using Projections on Karhunen-Loeve Eigenimages // *ApJ*. VIII 2012. 755, 2. L28.

- Sozzetti A., Bonavita M., Desidera S., Gratton R., Lattanzi M. G. Gaia: The Astrometry Revolution // Proc. IAU. XI 2015. 10, S314. 264–269.
- Spiegel David S., Burrows Adam. SPECTRAL AND PHOTOMETRIC DIAGNOSTICS OF GIANT PLANET FORMATION SCENARIOS // ApJ. I 2012. 745, 2. 174.
- Su K. Y. L., Rieke G. H., Stapelfeldt K. R., Malhotra R., Bryden G., Smith P. S., Misselt K. A., Moro-Martin A., Williams J. P. The Debris Disk Around HR 8799 // The Astrophysical Journal. XI 2009. 705. 314–327.
- Sudol Jeffrey J., Haghighipour Nader. High-Mass, Four-planet Configurations for HR 8799: Constraining the Orbital Inclination and Age of the System // The Astrophysical Journal. VIII 2012. 755. 38.
- Talts Sean, Betancourt Michael, Simpson Daniel, Vehtari Aki, Gelman Andrew. Validating Bayesian Inference Algorithms with Simulation-Based Calibration // arXiv:1804.06788 [stat]. X 2020.
- Thiébaud Éric, Young John. Principles of image reconstruction in optical interferometry: tutorial // J. Opt. Soc. Am. A. Jun 2017. 34, 6. 904–923.
- Thompson William, Marois Christian. Extremely Bright Orbital Guide Beacons for Extremely Large Telescopes // Adaptive Optics for Extremely Large Telescopes 6. V 2019. 6.
- Thompson William, Marois Christian. Improved Contrast in Images of Exoplanets Using Direct Signal-to-noise Ratio Optimization // AJ. V 2021. 161, 5. 236.
- Thompson William, Marois Christian, Do Ó Clarissa R., Konopacky Quinn, Ruffio Jean-Baptiste, Wang Jason, Skemer Andy J., De Rosa Robert J., Macintosh Bruce. Deep Orbital Search for Additional Planets in the HR 8799 System // AJ. I 2023. 165, 1. 29.
- Trifonov T., Caballero J. A., Morales J. C., Seifahrt A., Ribas I., Reiners A., Bean J. L., Luque R., Parviainen H., Pallé E., Stock S., Zechmeister M., Amado P. J., Anglada-Escudé G., Azzaro M., Barclay T., Béjar V. J. S., Bluhm P., Casasayas-Barris N., Cifuentes C., Collins K. A., Collins K. I., Cortés-Contreras M., de Leon J., Dreizler S., Dressing C. D., Esparza-Borges E., Espinoza N., Fausnaugh M., Fukui A., Hatzes A. P., Hellier C., Henning Th., Henze C. E., Herrero E., Jeffers S. V., Jenkins J. M., Jensen E. L. N., Kaminski A., Kasper D., Kossakowski D., Kürster M., Lafarga M., Latham D. W., Mann A. W., Molaverdikhani K., Montes D., Montet B. T., Murgas F., Narita N., Oshagh M., Passegger V. M., Pollacco D., Quinn S. N., Quirrenbach A., Ricker G. R., Rodríguez López C., Sanz-Forcada J., Schwarz R. P., Schweitzer A., Seager S., Shporer A., Stangret M., Stürmer J., Tan T. G., Tenenbaum P., Twicken J. D., Vanderspek R., Winn J. N. A Nearby Transiting Rocky Exoplanet That Is Suitable for Atmospheric Investigation // Science. III 2021. 371. 1038–1041.
- Trifonov Trifon, Tal-Or Lev, Zechmeister Mathias, Kaminski Adrian, Zucker Shay, Mazeh Tsevi. Public HARPS Radial Velocity Database Corrected for Systematic Errors // A&A. IV 2020. 636. A74.
- The TOLIMAN Space Telescope. // . VII 2018. 10701. 107011J.
- Udell Madeleine, Mohan Karanveer, Zeng David, Hong Jenny, Diamond Steven, Boyd Stephen. Convex Optimization in Julia // arXiv:1410.4821 [cs, math, stat]. X 2014.

- Vigan A., N'Diaye M., Dohlen K., Sauvage J.-F., Milli J., Zins G., Petit C., Wahhaj Z., Cantalloube F., Caillat A., Costille A., Merrer J. Le, Carlotti A., Beuzit J.-L., Mouillet D. Calibration of Quasi-Static Aberrations in Exoplanet Direct-Imaging Instruments with a Zernike Phase-Mask Sensor - III. On-sky Validation in VLT/SPHERE // *A&A*. IX 2019. 629. A11.
- Vousden W. D., Farr W. M., Mandel I. Dynamic Temperature Selection for Parallel Tempering in Markov Chain Monte Carlo Simulations // *Monthly Notices of the Royal Astronomical Society*. I 2016. 455, 2. 1919–1937.
- Wagenmakers Eric-Jan, Lodewyckx Tom, Kuriyal Himanshu, Grasman Raoul. Bayesian Hypothesis Testing for Psychologists: A Tutorial on the Savage–Dickey Method // *Cognitive Psychology*. V 2010. 60, 3. 158–189.
- Wagner K., Boehle A., Pathak P., Kasper M., Arsenault R., Jakob G., Käufl U., Leveratto S., Maire A.-L., Pantin E., Siebenmorgen R., Zins G., Absil O., Ageorges N., Apai D., Carlotti A., Choquet É, Delacroix C., Dohlen K., Duhoux P., Forsberg P., Fuenteseca E., Gutruf S., Guyon O., Huby E., Kampf D., Karlsson M., Kervella P., Kirchbauer J.-P., Klupar P., Kolb J., Mawet D., N'Diaye M., Orban de Xivry G., Quanz S. P., Reutlinger A., Ruane G., Riquelme M., Soenke C., Sterzik M., Vigan A., de Zeeuw T. Imaging Low-Mass Planets within the Habitable Zone of α Centauri // *Nat Commun*. II 2021. 12, 1. 922.
- Wahhaj Z., Milli J., Romero C., Cieza L., Zurlo A., Vigan A., Peña E., Valdes G., Cantalloube F., Girard J., Pantoja B. A Search for a Fifth Planet around HR 8799 Using the Star-Hopping RDI Technique at VLT/SPHERE // *A&A*. IV 2021. 648. A26.
- Wahhaj Zahed, Cieza Lucas A., Mawet Dimitri, Yang Bin, Canovas Hector, de Boer Jozua, Casassus Simon, Ménard François, Schreiber Matthias R., Liu Michael C., Biller Beth A., Nielsen Eric L., Hayward Thomas L. Improving Signal-to-Noise in the Direct Imaging of Exoplanets and Circumstellar Disks with MLOC1 // *A&A*. IX 2015. 581. A24.
- Shades of Black: Searching for Brown Dwarfs and Giant Planets. // . I 1999. 56. 449.
- The Gemini Planet Imager Calibration Wavefront Sensor Instrument. // . VII 2010. 7736. 77365D.
- Wang Jason J., Graham James R., Dawson Rebekah, Fabrycky Daniel, De Rosa Robert J., Pueyo Laurent, Konopacky Quinn, Macintosh Bruce, Marois Christian, Chiang Eugene, Ammons S. Mark, Arriaga Pauline, Bailey Vanessa P., Barman Travis, Bulger Joanna, Chilcote Jeffrey, Cotten Tara, Doyon René, Duchêne Gaspard, Esposito Thomas M., Fitzgerald Michael P., Follette Katherine B., Gerard Benjamin L., Goodsell Stephen J., Greenbaum Alexandra Z., Hibon Pascale, Hung Li-Wei, Ingraham Patrick, Kalas Paul, Larkin James E., Maire Jérôme, Marchis Franck, Marley Mark S., Metchev Stanimir, Millar-Blanchaer Maxwell A., Nielsen Eric L., Oppenheimer Rebecca, Palmer David, Patience Jennifer, Perrin Marshall, Poyneer Lisa, Rajan Abhijith, Rameau Julien, Rantakyro Fredrik T., Ruffio Jean-Baptiste, Savransky Dmitry, Schneider Adam C., Sivaramakrishnan Anand, Song Inseok, Soummer Remi, Thomas Sandrine, Wallace J. Kent, Ward-Duong Kimberly, Wiktorowicz Sloane, Wolff Schuyler. Dynamical Constraints on the HR 8799 Planets with GPI // *AJ*. X 2018. 156, 5. 192.

- Wang Jason J., Ruffio Jean-Baptiste, Morris Evan, Delorme Jacques-Robert, Jovanovic Nemanja, Pezzato Jacklyn, Echeverri Daniel, Finnerty Luke, Hood Callie, Zanazzi J. J., Bryan Marta L., Bond Charlotte Z., Cetre Sylvain, Martin Emily C., Mawet Dimitri, Skemer Andy, Baker Ashley, Xuan Jerry W., Wallace J. Kent, Wang Ji, Bartos Randall, Blake Geoffrey A., Boden Andy, Buzard Cam, Calvin Benjamin, Chun Mark, Doppmann Greg, Dupuy Trent J., Duchêne Gaspard, Feng Y. Katherina, Fitzgerald Michael P., Fortney Jonathan, Freedman Richard S., Knutson Heather, Konopacky Quinn, Lilley Scott, Liu Michael C., Lopez Ronald, Lupu Roxana, Marley Mark S., Meshkat Tiffany, Miles Brittany, Millar-Blanchaer Maxwell, Ragland Sam, Roy Arpita, Ruane Garreth, Sappéy Ben, Schofield Tobias, Weiss Lauren, Wetherell Edward, Wizinowich Peter, Ygouf Marie. Detection and Bulk Properties of the HR 8799 Planets with High Resolution Spectroscopy // *AJ*. X 2021. 162, 4. 148.
- Wang Ji, Mawet Dimitri, Ruane Garreth, Hu Renyu, Benneke Björn. Observing Exoplanets with High Dispersion Coronagraphy. I. The Scientific Potential of Current and Next-generation Large Ground and Space Telescopes // *The Astronomical Journal*. IV 2017. 153. 183.
- Wang Ji, Wang Jason J., Ma Bo, Chilcote Jeffrey, Ertel Steve, Guyon Olivier, Ilyin Ilya, Jovanovic Nemanja, Kalas Paul, Lozi Julien, Macintosh Bruce, Strassmeier Klaus G., Stone Jordan. On the Chemical Abundance of HR 8799 and the Planet c // *AJ*. IX 2020. 160, 3. 150.
- Wertz O., Absil O., Gómez González C. A., Milli J., Girard J. H., Mawet D., Pueyo L. VLT/SPHERE Robust Astrometry of the HR8799 Planets at Milliarcsecond-Level Accuracy: Orbital Architecture Analysis with PyAstrOFit★ // *A&A*. II 2017. 598. A83.
- Wilner David J., MacGregor Meredith A., Andrews Sean M., Hughes A. Meredith, Matthews Brenda, Su Kate. Resolved Millimeter Observations of the HR 8799 Debris Disk // *The Astrophysical Journal*. III 2018. 855. 56.
- Wooten Alwyn, Thompson A. Richard. The Atacama Large Millimeter/Submillimeter Array // *Proceedings of the IEEE*. 2009. 97, 8. 1463–1471.
- Wright J. T., Howard A. W. Efficient Fitting of Multiplanet Keplerian Models to Radial Velocity and Astrometry Data // *The Astrophysical Journal Supplement Series*. V 2009. 182. 205–215.
- Wunderlich Fabian, Godolt Mareike, Grenfell John Lee, Städt Steffen, Smith Alexis M. S., Gebauer Stefanie, Schreier Franz, Hedelt Pascal, Rauer Heike. Detectability of Atmospheric Features of Earth-like Planets in the Habitable Zone around M Dwarfs // *A&A*. IV 2019. 624. A49.
- Xu Kai, Ge Hong, Tebbutt Will, Tarek Mohamed, Trapp Martin, Ghahramani Zoubin. AdvancedHMC.JI: A Robust, Modular and Efficient Implementation of Advanced HMC Algorithms // *Proceedings of The 2nd Symposium on Advances in Approximate Bayesian Inference*. II 2020. 1–10.
- Yelda S., Lu J. R., Ghez A. M., Clarkson W., Anderson J., Do T., Matthews K. IMPROVING GALACTIC CENTER ASTROMETRY BY REDUCING THE EFFECTS OF GEOMETRIC DISTORTION // *ApJ*. XII 2010. 725, 1. 331–352.
- Zahnle K., Marley M. S., Morley C. V., Moses J. I. PHOTOLYTIC HAZES IN THE ATMOSPHERE OF 51 ERI B // *ApJ*. VI 2016. 824, 2. 137.

- Zink Jon K, Hansen Bradley M S. Accounting for Multiplicity in Calculating Eta Earth // Monthly Notices of the Royal Astronomical Society. VII 2019. 487, 1. 246–252.
- Zuckerman B., Song Inseok. Dusty Debris Disks as Signposts of Planets: Implications for Spitzer Space Telescope // ApJ. III 2004. 603, 2. 738.
- Zurlo A., Vigan A., Galicher R., Maire A.-L., Mesa D., Gratton R., Chauvin G., Kasper M., Moutou C., Bonnefoy M., Desidera S., Abe L., Apai D., Baruffolo A., Baudoz P., Baudrand J., Beuzit J.-L., Blancard P., Boccaletti A., Cantalloube F., Carle M., Cascone E., Charton J., Claudi R. U., Costille A., de Caprio V., Dohlen K., Dominik C., Fantinel D., Feautrier P., Feldt M., Fusco T., Gigan P., Girard J. H., Gisler D., Gluck L., Gry C., Henning T., Hugot E., Janson M., Jaquet M., Lagrange A.-M., Langlois M., Llored M., Madec F., Magnard Y., Martinez P., Maurel D., Mawet D., Meyer M. R., Milli J., Moeller-Nilsson O., Mouillet D., Origné A., Pavlov A., Petit C., Puget P., Quanz S. P., Rabou P., Ramos J., Rousset G., Roux A., Salasnich B., Salter G., Sauvage J.-F., Schmid H. M., Soenke C., Stadler E., Suarez M., Turatto M., Udry S., Vakili F., Wahhaj Z., Wildi F., Antichi J. First Light of the VLT Planet Finder SPHERE. III. New Spectrophotometry and Astrometry of the HR 8799 Exoplanetary System // Astronomy and Astrophysics. III 2016. 587. A57.
- van Leeuwen F. Validation of the New Hipparcos Reduction // Astronomy and Astrophysics. XI 2007. 474. 653–664.

Appendix A

Derivation of Cartesian and Celestial Position, Velocity, and Acceleration

This appendix is a companion to Chapter 3 and based on the same material submitted to the *Astronomical Journal*. *The authorship statement at the start of Chapter 3 applies here as well. This text was contributed to that paper by Jensen Lawrence and is included here for completeness.*

To properly fit to relative astrometry, proper motion anomaly, and radial velocity data, expressions for the position, velocity, and acceleration of a secondary companion (such as a planet) about a primary mass (such as a star) in Cartesian and celestial coordinates are required. It is well-known that the solution to the Keplerian two-body problem (for example [Goldstein et al., 2008](#)) is

$$r(t) = \frac{a(1 - e^2)}{1 + e \cos(\nu(t))}, \quad (\text{A.1})$$

where r is the radial separation of the two bodies, a is the orbital semi-major axis, e is the orbital eccentricity, and $\nu(t)$ is the true anomaly. The true anomaly describes the angular position of the orbit with respect to periastron, and is the solution to the equation

$$\tan\left(\frac{\nu(t)}{2}\right) = \sqrt{\frac{1+e}{1-e}} \tan\left(\frac{E(t)}{2}\right). \quad (\text{A.2})$$

The eccentric anomaly $E(t)$ is in turn the solution to

$$M(t) = E(t) - e \sin(E(t)), \quad (\text{A.3})$$

with $M(t)$ being the mean anomaly

$$M(t) = \frac{2\pi(t - t_0)}{T}. \quad (\text{A.4})$$

Here, t_0 is the time of periastron passage and T is the orbital period. Since Eq. A.3 is a transcendental equation, $E(t)$ must be determined numerically. Once calculated, though, the derivatives of the eccentric anomaly with respect to the mean anomaly and eccentricity have analytic expressions:

$$\frac{\partial E}{\partial M} = \frac{1}{1 - e \cos E(t)} \quad (\text{A.5})$$

and

$$\frac{\partial E}{\partial e} = \frac{\sin E(t)}{1 - e \cos E(t)}. \quad (\text{A.6})$$

These expressions can be determined using implicit differentiation.

A.1 Cartesian Coordinates

Using the well-known relation between polar coordinates (r, ν) and Cartesian coordinates (x, y, z) , we see that in the frame of the secondary, its position with respect to the primary is

$$\mathbf{r}(t) = \begin{bmatrix} r(t) \cos(\nu(t)) \\ r(t) \sin(\nu(t)) \\ 0 \end{bmatrix} = \frac{a(1 - e^2)}{1 + e \cos(\nu(t))} \begin{bmatrix} \cos(\nu(t)) \\ \sin(\nu(t)) \\ 0 \end{bmatrix}. \quad (\text{A.7})$$

However, the frame of the secondary and the frame of an external observer are unlikely to be coincident. Instead, an external observer is likely to be rotated with respect to the frame of the secondary. This rotation is quantified by ω (the argument of periastron of the secondary), i (the inclination of the orbit), and Ω (the longitude of the orbit's ascending node). The rotations associated with these angles are given by the rotation

operators

$$R_z(\omega) = \begin{bmatrix} \cos(\omega) & -\sin(\omega) & 0 \\ \sin(\omega) & \cos(\omega) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (\text{A.8})$$

$$R_x(i) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(i) & -\sin(i) \\ 0 & \sin(i) & \cos(i) \end{bmatrix} \quad (\text{A.9})$$

$$R_z(\Omega) = \begin{bmatrix} \cos(\Omega) & -\sin(\Omega) & 0 \\ \sin(\Omega) & \cos(\Omega) & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (\text{A.10})$$

Thus, the position of the secondary with respect to its primary as seen by an external observer is

$$\mathbf{r}_{\text{obs}}(t) = R_z(\Omega)R_x(i)R_z(\omega)\mathbf{r}(t) \quad (\text{A.11})$$

$$\mathbf{r}_{\text{obs}}(t) = r(t) \begin{bmatrix} \cos(\nu(t) + \omega) \cos(\Omega) - \sin(\nu(t) + \omega) \cos(i) \sin(\Omega) \\ \cos(\nu(t) + \omega) \sin(\Omega) + \sin(\nu(t) + \omega) \cos(i) \cos(\Omega) \\ \sin(\nu(t) + \omega) \sin(i) \end{bmatrix}. \quad (\text{A.12})$$

Explicitly, we have

$$\mathbf{r}_{\text{obs}}(t) = \frac{a(1 - e^2)}{1 + e \cos(\nu(t))} \begin{bmatrix} \cos(\nu(t) + \omega) \cos(\Omega) - \sin(\nu(t) + \omega) \cos(i) \sin(\Omega) \\ \cos(\nu(t) + \omega) \sin(\Omega) + \sin(\nu(t) + \omega) \cos(i) \cos(\Omega) \\ \sin(\nu(t) + \omega) \sin(i) \end{bmatrix} \quad (\text{A.13})$$

Inspecting this equation, we see that $\mathbf{r}_{\text{obs}}(t)$ is dependent on a , e , i , ω , and Ω . In addition, there is dependence on t_0 and T because of Eqs. A.2 and A.3. However, a and T are related by Kepler's third law, meaning dependence on T can be reformulated as dependence on a . Thus, a Keplerian orbit is uniquely described by the parameters $(a, e, i, \omega, \Omega, t_0)$; these are referred to as the Campbell elements. Note that in this paper, we replace t_0 with either τ or θ . τ is given by

$$\tau = \left(\frac{t_0 - t_{\text{ref}}}{T} \right) \text{ mod } 1, \quad (\text{A.14})$$

where we choose $t_0 = 58849.0$ MJD (January 1, 2020) (Blunt et al., 2020), and θ is discussed in Appendix C. We now differentiate $\mathbf{r}_{\text{obs}}(t)$ to obtain $\mathbf{v}_{\text{obs}}(t)$. Kepler's second law implies that

$$\frac{r(t)^2 \dot{v}(t)}{2} = \frac{\pi a^2 \sqrt{1 - e^2}}{T}. \quad (\text{A.15})$$

From this, we obtain

$$\dot{r}(t) = \frac{2\pi a e \sin(v(t))}{T \sqrt{1 - e^2}}, \quad r(t) \dot{v}(t) = \frac{2\pi a}{T} \frac{1 + e \cos(v(t))}{\sqrt{1 - e^2}}. \quad (\text{A.16})$$

Using these identities to differentiate $\mathbf{r}_{\text{obs}}(t)$ produces

$$\mathbf{v}_{\text{obs}}(t) = \frac{2\pi a}{T} \frac{1}{\sqrt{1 - e^2}} \begin{bmatrix} -\cos(i) \sin(\Omega) [\cos(v(t) + \omega) + e \cos(\omega)] + \cos(\Omega) (\sin(v(t) + \omega) + e \sin(\omega)) \\ \cos(i) \cos(\Omega) [\cos(v(t) + \omega) + e \cos(\omega)] - \sin(\Omega) (\sin(v(t) + \omega) + e \sin(\omega)) \\ \sin(i) [\cos(v(t) + \omega) + e \cos(\omega)] \end{bmatrix}. \quad (\text{A.17})$$

$$(\text{A.18})$$

We determine $\mathbf{a}_{\text{obs}}(t)$ in an identical manner. Observing that

$$\dot{v}(t) = \frac{2\pi}{T} \frac{(1 + e \cos(v(t)))^2}{(1 - e^2)^{3/2}}, \quad (\text{A.19})$$

differentiating $\mathbf{v}_{\text{obs}}(t)$ yields

$$\mathbf{a}_{\text{obs}}(t) = \frac{4\pi^2 a}{T^2} \frac{1}{(1 - e^2)^2} \begin{bmatrix} (1 + e \cos(v(t)))^2 [\cos(i) \sin(\Omega) \sin(v(t) + \omega) - \cos(\Omega) \cos(v(t) + \omega)] \\ -(1 + e \cos(v(t)))^2 [\cos(i) \cos(\Omega) \sin(v(t) + \omega) + \sin(\Omega) \cos(v(t) + \omega)] \\ -(1 + e \cos(v(t)))^2 \sin(i) \sin(v(t) + \omega) \end{bmatrix}. \quad (\text{A.20})$$

$$(\text{A.21})$$

Having expressions for $\mathbf{r}_{\text{obs}}(t)$, $\mathbf{v}_{\text{obs}}(t)$, and $\mathbf{a}_{\text{obs}}(t)$ is quite useful, since they allow us to determine the corresponding quantities in celestial coordinates with ease. Furthermore, the radial velocity of the secondary with respect to its primary can be read off directly as the z-component of Eq. A.18, giving us

$$v_{z,\text{obs}} = \frac{2\pi a}{T} \frac{\sin(i) [\cos(v(t) + \omega) + e \cos(\omega)]}{\sqrt{1 - e^2}}. \quad (\text{A.22})$$

To get the primary's radial velocity with respect to the secondary, we simply employ conservation of momentum and multiply the previous expression by $-M_s/(M_s + M_p)$, with M_p and M_s being the masses of the primary and secondary, respectively. To determine the radial velocities measured by an external observer, we add the barycentre radial velocity γ . Additional linear trends, quadratic trends, or other trends can be accounted for in Octofitter by making the instrument zero-points a function of other parameters in the fit. Overall, we get that the observed radial velocities of the primary and secondary are

$$v_{r,s} = \frac{2\pi a \sin(i)[\cos(\nu(t) + \omega) + e \cos(\omega)]}{T \sqrt{1 - e^2}} + \gamma + \xi, \quad (\text{A.23})$$

$$v_{r,p} = -\frac{M_s}{M_s + M_p} \frac{2\pi a \sin(i)[\cos(\nu(t) + \omega) + e \cos(\omega)]}{T \sqrt{1 - e^2}} + \gamma + \xi, \quad (\text{A.24})$$

where ξ is the combined effect of all relevant trends. Using these expressions, orbital and physical parameters can be fit to radial velocity data.

A.2 Celestial Coordinates

Now that we know $\mathbf{r}_{\text{obs}}(t)$, $\mathbf{v}_{\text{obs}}(t)$, and $\mathbf{a}_{\text{obs}}(t)$, we can calculate the corresponding expressions in celestial coordinates. Before doing so, a brief discussion of coordinate conventions is required. For this paper, we measure with reference to the celestial north pole. This means that, in reference to the previous section, the positive x -axis points in the direction of positive declination (upwards on the sky), while the positive y -axis points in the direction of positive right ascension (leftwards on the sky). This choice of coordinates is important to note, since it differs from the right/up orientation of the x/y axes commonly seen elsewhere. Furthermore, the positive z -axis points away from the observer.

Keeping this coordinate convention in mind, we can determine position, velocity, and acceleration in celestial coordinates. Normally, an observer a distance d away from an object of size r measures the object to have an angular size of

$$\Delta\theta = \arctan\left(\frac{r}{d}\right). \quad (\text{A.25})$$

However, for scenarios relevant to this work, r is on the order of astronomical units and d is on the order of parsecs, meaning the small angle approximation $\Delta\theta \approx r/d$

can be used with negligible error. Applying this approximation, we get that the right ascension and declination offsets of the secondary from its primary are

$$\Delta\alpha = \frac{y_{\text{obs}}}{d}, \quad \Delta\delta = \frac{x_{\text{obs}}}{d}, \quad (\text{A.26})$$

where d is the distance to the system, and x_{obs} and y_{obs} are the x - and y -components of Eq. A.13. From this, it immediately follows that

$$\Delta\dot{\alpha} = \frac{v_{y,\text{obs}}}{d}, \quad \Delta\dot{\delta} = \frac{v_{x,\text{obs}}}{d}, \quad (\text{A.27})$$

where $v_{x,\text{obs}}$ and $v_{y,\text{obs}}$ are the x - and y -components of Eq. A.18, and that

$$\Delta\ddot{\alpha} = \frac{a_{y,\text{obs}}}{d}, \quad \Delta\ddot{\delta} = \frac{a_{x,\text{obs}}}{d}, \quad (\text{A.28})$$

where $a_{x,\text{obs}}$ and $a_{y,\text{obs}}$ are the x - and y -components of Eq. A.21. The equations for $\Delta\alpha$ and $\Delta\delta$ can be used to fit orbital and physical parameters to relative astrometry data, while the equations for $\Delta\dot{\alpha}$ and $\Delta\dot{\delta}$ can be used to fit orbital and physical parameters to proper motion anomaly data. $\Delta\ddot{\alpha}$ and $\Delta\ddot{\delta}$ are used for calculating model derivatives for higher-order samplers, and may be applied to angular acceleration data in the future.

Appendix B

Thiele-Innes Elements

This appendix is a companion to Chapter 3 and based on the same material submitted to the *Astronomical Journal*. *The authorship statement at the start of Chapter 3 applies here as well. This text was contributed to that paper by Jensen Lawrence and is included here for completeness.*

In Appendix A.1, we saw that Keplerian orbits can be uniquely characterized by the Campbell elements $(a, e, i, \omega, \Omega)$, in addition to an orbit position parameter (t_0 , τ , or θ). However, Keplerian orbits can be equivalently characterized by the Thiele-Innes elements (e, A, B, F, G) , where e is the orbital eccentricity and A, B, F , and G are the Thiele-Innes constants, along with an orbit position parameter. The Thiele-Innes constants are defined as

$$A = a\bar{\omega}(\cos(\Omega)\cos(\omega) - \sin(\Omega)\sin(\omega)\cos(i)), \quad (\text{B.1})$$

$$B = a\bar{\omega}(\sin(\Omega)\cos(\omega) + \cos(\Omega)\sin(\omega)\cos(i)), \quad (\text{B.2})$$

$$F = a\bar{\omega}(-\cos(\Omega)\sin(\omega) - \sin(\Omega)\cos(\omega)\cos(i)), \quad (\text{B.3})$$

$$G = a\bar{\omega}(-\sin(\Omega)\sin(\omega) + \cos(\Omega)\cos(\omega)\cos(i)), \quad (\text{B.4})$$

where $\bar{\omega}$ is the parallax distance to the system. As we can see, it is straightforward to calculate A, B, F , and G from a, i, ω , and Ω . Inverting this relationship, although less straightforward, is still doable. We begin by defining

$$u = \frac{1}{2}(A^2 + B^2 + C^2 + D^2), \quad (\text{B.5})$$

$$v = AG - BF. \quad (\text{B.6})$$

We can recover a using u and v like so:

$$a = \frac{\sqrt{u + \sqrt{(u+v)(u-v)}}}{\bar{\omega}}. \quad (\text{B.7})$$

Once we have a , i immediately follows from

$$i = \arccos\left(\frac{v}{a^2 \bar{\omega}^2}\right). \quad (\text{B.8})$$

Since Eqs. B.1 – B.4 involve only $\cos(i)$, and $\cos(\theta) = \cos(-\theta)$, this is technically an equation for $|i|$. Next, we define

$$j = \arctan(A + G, B - F), \quad (\text{B.9})$$

$$k = \arctan(A - G, B + F), \quad (\text{B.10})$$

where $\arctan(y, x)$ is a quadrant-sensitive version of $\arctan(y/x)$ (sometimes written as $\arctan2(y, x)$ or $\text{atan2}(y, x)$). From these quantities, ω and Ω are given by

$$\omega = \frac{1}{2}(j - k), \quad (\text{B.11})$$

$$\Omega = \frac{1}{2}(j + k). \quad (\text{B.12})$$

Thus, we have successfully recovered a , i , ω , and Ω from the Thiele-Innes constants.

Appendix C

Derivation of the Parameter τ from Position Angle θ

This appendix is a companion to Chapter 3 and based on the same material submitted to the *Astronomical Journal*. *The authorship statement at the start of Chapter 3 applies here as well. I wrote the complete text of this appendix.*

The parameter τ used in Orbitize! and Octofitter is the fraction of the orbit completed by a planet after periastron, at some reference epoch t_{ref} . This is a useful parameterization since it lies between 0 and 1 for all orbits; however, if the reference epoch is not similar to the epochs of the observations, it can exhibit pathological behaviour as the other orbital parameters change. One solution is to adjust the reference epoch for each dataset, but a further improvement for many cases is to instead adopt the observed position angle θ at epoch t as an independent variable. θ is an improvement over τ because it is insensitive to changes in the other parameters and has a straightforward interpretation. A derivation of τ from θ is given here.

We begin with the parallax distance to the system $\bar{\omega}$, the host mass M , the eccentricity e , and the Thiele-Innes constants A, B, F , and G given in Eqs. B.1 – B.4. From these, we write the transformation matrix

$$T = \begin{bmatrix} A & B \\ F & G \end{bmatrix}, \quad (\text{C.1})$$

from which we find the unit vectors

$$\begin{bmatrix} \hat{x} \\ \hat{y} \end{bmatrix} = T^{-1} \begin{bmatrix} \sin(\theta) \\ \cos(\theta) \end{bmatrix}. \quad (\text{C.2})$$

We can then calculate true anomaly ν as

$$\nu = \arctan(\hat{y}, \hat{x}), \quad (\text{C.3})$$

and mean anomaly as

$$\text{MA} = \arctan(-\sin(\nu)\sqrt{1-e^2}, -\cos(\nu) - e) + \pi - \frac{e\sqrt{1-e^2}\sin(\nu)}{1+e\cos(\nu)}. \quad (\text{C.4})$$

From this, the parameter τ can be calculated using the period as follows. First, calculate the semi-major axis a using Eq. B.7. Next, combining Kepler's third law and the value of a gives us the period $P = 2\pi\sqrt{a^3/(GM)}$. The mean motion is then $n = 2\pi/P$, and the epoch of periastron passage follows from this as $t_{\text{peri}} = t - \text{MA}/n - t_{\text{ref}}$. Finally,

$$\tau = \frac{t_{\text{peri}}}{P}. \quad (\text{C.5})$$

The corner plot in Figure C.1 demonstrates the benefit of this parameterization. The bottom rows show the same orbits parameterized with τ and with θ . The θ marginal posterior is nearly Gaussian, and has simple relationships with all other orbital parameters. By contrast, the τ marginal posterior is less informative and has complex relationships to ω and Ω .

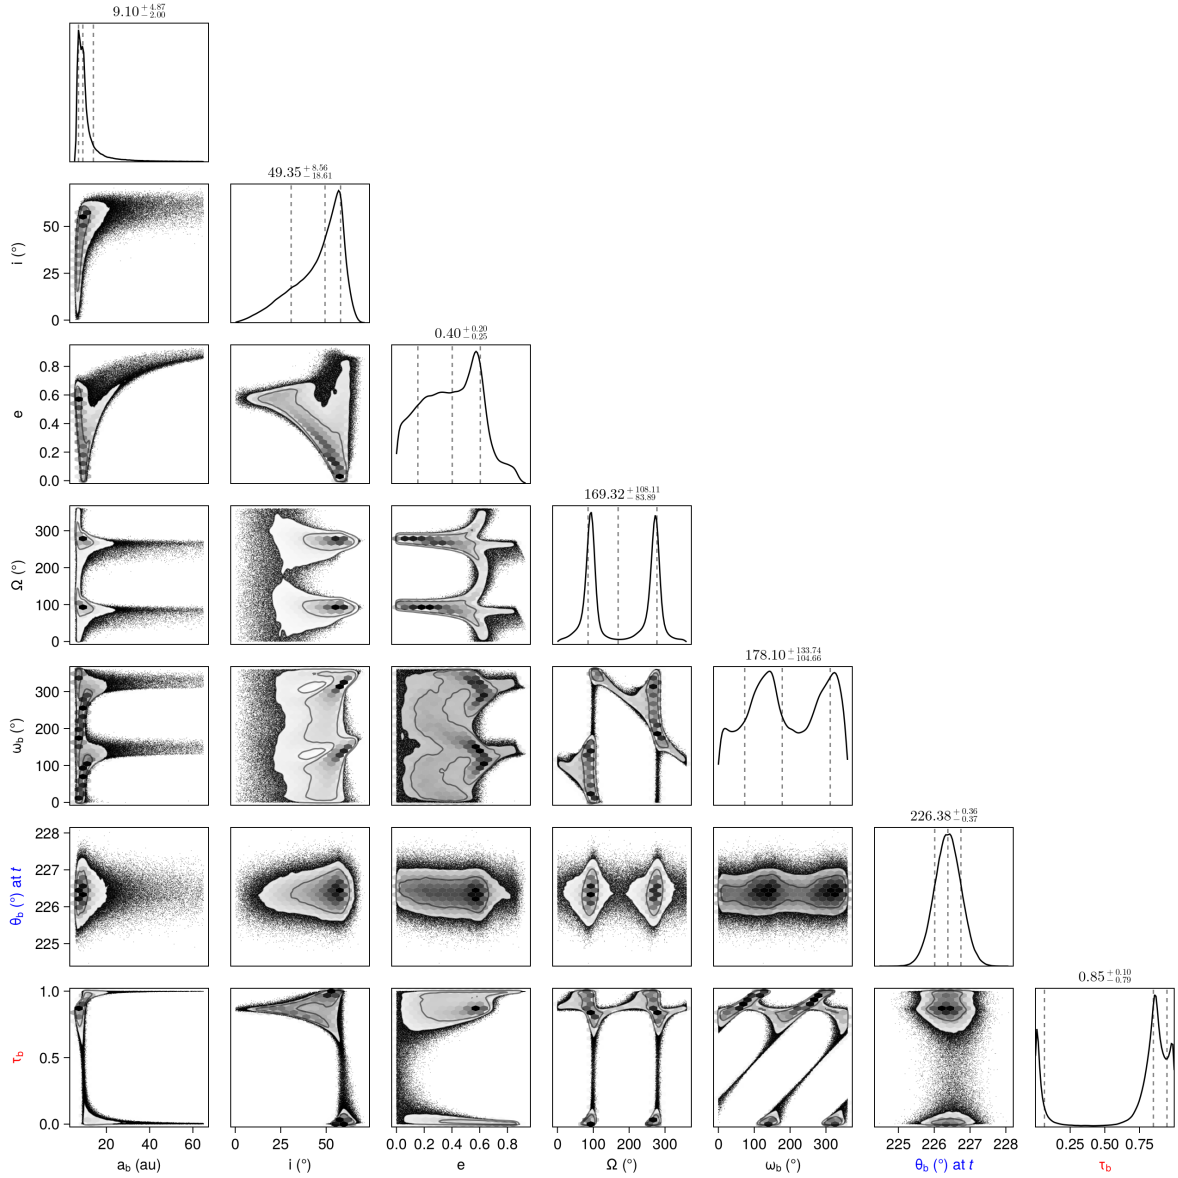


Figure C.1 Comparison of θ (blue, second last row) and τ (red, last row) parameterizations of the same orbital posterior. The θ marginal posterior is nearly Gaussian, and has simple relationships with all other orbital parameters. By contrast, the τ marginal posterior is less informative and has complex relationships to ω and Ω . Complex structures in a posterior decrease sampling efficiency and could in some cases lead to biased results.

Appendix D

Loading, Manipulating, and Visualizing Astronomical Images

This appendix describes a Julia ([Bezanson et al., 2012](#)) package for retrieving, processing, and visualizing astronomical FITS images developed in the course of completing this dissertation.

This material is based on work published in the proceedings of the JuliaCon Conferences (2022). Authors: William Thompson and Mose` Giordano. *My contribution to this work as lead author and developer of the version of the AstroImages.jl package described herein with comments and some initial programming work by M. Giordano*

D.1 Introduction

To study the cosmos, astronomers use data cubes with many dimensions representing images with axes for sky position, time, wavelength, polarization, and more. Since these large datasets often span many orders of magnitude in intensity and typically include colours invisible to humans, astronomers like to visualize their images using a variety of non-linear stretching and contrast adjustments. Additionally, images may contain metadata specifying arbitrary mappings of pixel positions to multiple celestial coordinate systems. Julia ([Bezanson et al., 2012](#)) is a powerful language for processing astronomical data, but these visualization tasks are a challenge for any tool.

D.2 Background

One of the most ubiquitous data formats used in astronomy is FITS, or the Flexible Image Transport System. Compared to a traditional raster image formats, FITS files are more like containers or small filesystems. Each file contains one or more header data units (HDUs) that pair a dataset with a header. The FITS format and header conventions have developed organically over several decades, first by convention, and then in a series of papers proposing new extensions and standardizing existing behaviour now summarized in [FITS Working Group \(2016\)](#).

In each HDU, headers are stored in a plain text ASCII format and are followed by data. The data may be in the format of an N-dimensional binary array (an “image” HDU), a binary table, or an ASCII table (“table HDUs”).

The metadata described by FITS headers are quite rich. A header consists of `KEY = VALUE / COMMENT` entries, long form `COMMENT` sections, and long form `HISTORY` sections used to describe the sequence of transformations used to generate the data. FITS headers often contain several hundred header entries most of which are specific to the instrument and software packages used to record and process the data. A subset of these are standardized and used to describe the physical units and coordinates of pixels in the image, called world coordinate system (WCS) header entries.

These WCS entries record pixel locations in one or more celestial coordinate systems including, right ascension and declination, galactic coordinates, velocity, frequency, wavelength, polarization, and more. For most coordinate schemes, pixel coordinates are encoded using an affine transformation matrix combined with a `CTYPE` specifier. Images are typically sampled regularly in the plane of a detector; however, coordinate projections are in general non-linear, and the physical coordinates of each pixel, spacings between pixels, etc. may vary across the image. In many instances, image HDUs contain 3, 4, or even 5 dimensional data cubes, and thanks to the affine transformation matrix, moving along any one of these dimensions can shift the coordinates along the other dimensions. This point is important since it means that the full coordinates of a given pixel are needed to calculate its physical position along any one axis.

For image data itself, FITS files do not in general store colour information that can be displayed directly. In contrast to digital cameras, astronomical data is almost invariably captured outside the human visible range or at least with filters that are not well-matched to human colour perception. Instead, astronomers use visualizations that map raw numerical data to a false colour image. These steps include a linear




Figure D.1 Conceptual schematic of a FITS file containing a 2D image HDU, a 3D hyperspectral image HDU, and a table HDU.

transformation from arbitrary numerical ranges to a standard 0 to 255 display range that produces a desired contrast and brightness at an intensity level of interest. The comparatively low dynamic range of electronic displays is often compensated for by applying non-linear mapping, or “stretching” such as log scales or arcsinh scales. Once the data is normalized to a reasonable intensity level, one typically presents a greyscale image or applies a false colour map. Finally, one may also create a colour composite image combining multiple information layers, especially when preparing data for public consumption.

D.3 Existing Tools

Julia has a robust package ecosystem for loading FITS files (`FITSIO.jl`), manipulating images (`Images.jl` ecosystem), calculating world coordinates (`WCS.jl`), and plotting (e.g. `Plots.jl`). For instance, the `FITSIO` package is capable of reading and writing FITS files, headers, image HDUs, and table HDUs; the `Images` packages support generating and displaying RGB images; and the `WCS` package supports converting between pixel and world coordinates given a FITS header. The Julia package ecosystem lacks, however, a package that combines these tools into a high-level interface for loading and visualizing astronomical image data.

D.4 The `AstroImages.jl` package

The aim of `AstroImages.jl` is to tie together relevant utility packages from the Julia ecosystem to facilitate the loading, manipulation, analysis, and visualization of




Figure D.2 AstroImage automatically rendered as a PNG. Image of the Carina nebula from JWST commissioning data (NASA, ESA, and CSA).

astronomical data.

`AstroImages.jl` is registered with `FileIO.jl` so that images and tables can be loaded from FITS files with the `load` and `save` functions.

```
using AstroImages
img = load("jwst-carina.fits")
```

By default, this will load the first image HDU from the file. If a different HDU or a table HDU is desired, the HDU number can be passed explicitly as a second argument. Multiple HDUs can be loaded at once by passing a tuple of HDU numbers or `:`.

The values returned from `load` are `AstroImage` objects which combine a parent array, a FITS header, and implements the `DimensionalData AbstractDimArray` interface. Any array type can be wrapped in an `AstroImage` by passing it to the constructor.

By implementing the `AbstractDimArray` interface, the dimensions of each axis of the array can be named and paired with coordinates to enable concise indexing of multidimensional cubes.

Once loaded, the user should be able to pass the `AstroImage` through any analysis code that accepts an `AbstractArray`. Indexing and slicing are tracked using `DimensionalData.jl` so that pixel locations in the originating parent data cube are recoverable.

FITS header keys can be accessed by indexing the image with a string key, the `Comment` type, and or the `History` type to read and write to the header. `WCS.jl` `Transformation` objects can be generated using the `wcs(img, wcsnum)` function and are cached unless/until the user modifies a relevant header field.

Finally, when the user wishes to display an `AstroImage`, they may use either `imview` to create a lazy mapping to RGB data or the `implot` `Plots.jl` recipe. `imview` accepts keyword arguments for controlling the image stretching, colour limits, colour scheme, bias, and contrast:

```
imview(img; clim=Percent(95), cmap=:magma,
       stretch=asinhstretch, contrast=1.2, bias=0.4)
```

`imview` is also called automatically with configurable default values whenever an `AstroImage` is displayed. In this case, the output is automatically downscaled to a reasonable size using `ImageTransformations.restrict`. For example, all that's needed to display a 2D image HDU from a FITS file is to type `load("filename.fits")` at an interactive prompt provided that prompt supports rich outputs.

The `implot` `Plots.jl` recipe supports additional functionality. It uses `imview` to first render the image and then uses `Plots.jl` to present an image series.

The plot recipe automatically applies WCS grid lines (which may be arbitrarily tilted and warped) and labels the axes with their coordinate types. It also creates a custom colour bar labelled with ticks at the correct, possibly non-linear spacing caused by chosen stretch. When applied to slices of multi-dimensional cubes, `implot` further titles the plot with the slice's location along the other axes.

```
using AstroImages, Plots
HIcube = load("HI.dat.fits")
slice = HIcube[Z=228]
implot(slice; cmap=:viridis, stretch=asinh)
```

`AstroImages.jl` is a complete, high-level package for working with astronomical

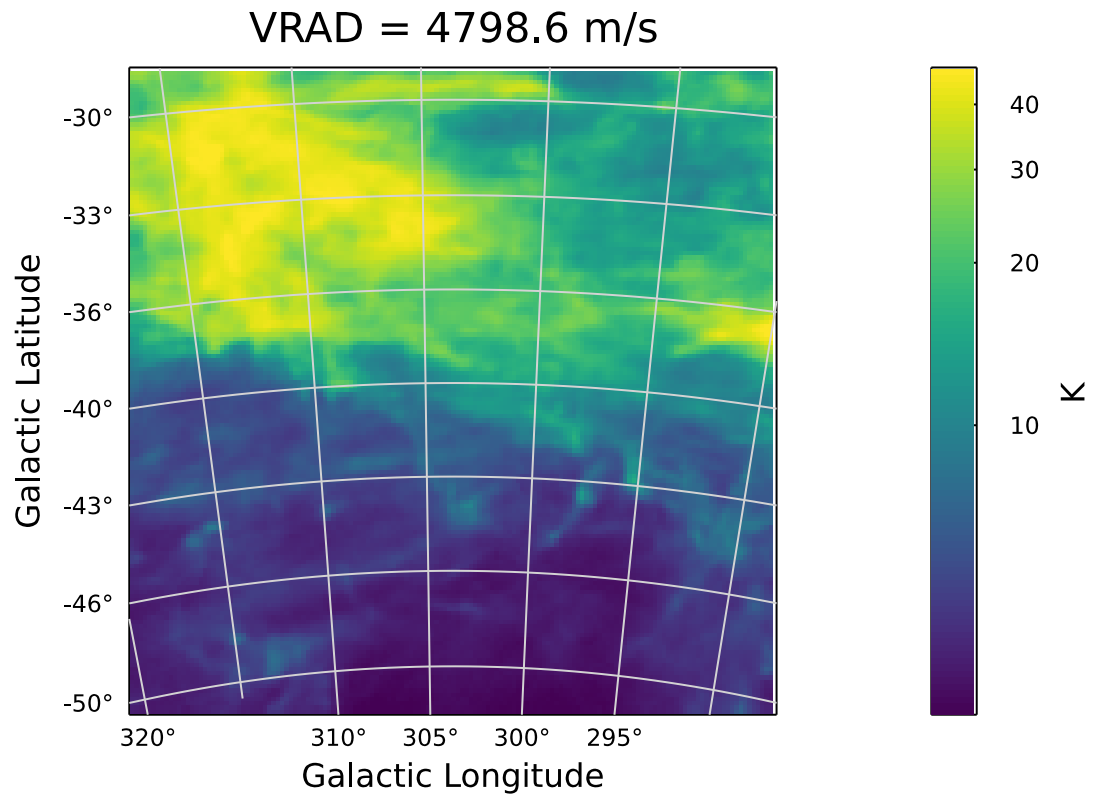


Figure D.3 Implot augments an image with world coordinates and automatically labelled axes, titles, and colourbar. Image is a slice through galactic latitude, longitude, and velocity space of neutral Hydrogen density ([HI4PI Collaboration et al., 2016](#), HI4PI survey;).

data in Julia.