

TOWARDS THE HABITABLE ZONE: DIRECT IMAGING OF
EXTRASOLAR PLANETS WITH THE MAGELLAN AO SYSTEM

by

Jared Robert Males

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF ASTRONOMY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2013

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Jared Robert Males entitled Towards the Habitable Zone: Direct Imaging of Extrasolar Planets with the Magellan AO System and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

Philip M. Hinz

Date: 8 May 2013

Dennis Zaritsky

Date: 8 May 2013

Olivier Guyon

Date: 8 May 2013

Glenn Schneider

Date: 8 May 2013

Date: 8 May 2013

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.
I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Dissertation Director: Laird M. Close

Date: 8 May 2013

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at the University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: Jared Robert Males

ACKNOWLEDGEMENTS

First and foremost I must give credit to my parents, *Jim* and *Marcia Males*. I'm eternally grateful for the opportunities that you've given me. I grew up in a house with Mom's calculus textbooks on the shelf and Dad on his way to one science experiment or another. *Of course* I ended up getting a Ph.D.

I also want to thank my advisor, *Laird Close*. You once told me that you thought we would do a lot of good work together, and you were right. Thanks for your patience, for your trust, and for giving me the freedom to explore and do it my way. I'm extremely proud of what we accomplished during my Ph.D.

I'm incredibly grateful to the ladies of the Phoenix ARCS foundation, who have helped fund my astronomical adventures for the last three years. ARCS is a truly great organization with a noble purpose, and they also know how to throw a party.

Thanks to the whole MagAO Team, including: my office-mate and fellow Ph.D. student *Derek Kopon*; our patient, responsive, flexible, and handy project manager/mechanical engineer/plumber/heavy lifter/etc., *Victor Gasho*; though she was late to the party, *Katie Morzinski* has been a huge part of our success and has helped make the VisAO blog a hit; *Alfio Puglisi* has been incredibly helpful and tolerant of the many changes and tweaks I made to his code; *Kate Follette* has contributed much in the way of science cases and establishing the VisAO blog; *Jason Lewis* did great work both in the lab with VisAO and in Chile; MagAO wouldn't have worked — wouldn't have even fit on the telescope — without *Alan Uomoto* and *Tyson Hare* of Carnegie; to the whole Arcetri team, lead by the incomparable *Simone Esposito*, the infallible *Armando Riccardi*, and their incredible crew including *Enrico*, *Marco*, *Runa*, *Fernando*, *Lorenzo*, *Javier*, *Paulo*, and many more; the Microgate team, especially *Roberto Biasi*; the LCO staff and crew, including *Povilas Polunas* and *Dave Osip*.

Thanks to *Andy Skemer* for letting me tag along on MMT runs and helping me get some good data. Thanks also to *Phil Hinz* and *Bill Hoffmann* for lots of advice, guidance, and encouragement.

MagAO itself received funding from several sources. Our adaptive secondary was supported by the NSF MRI program. Our pyramid wavefront sensor and telescope interfaces were developed with help from the TSIP program and the Magellan partners. The VisAO camera and system commissioning were supported with funds

from the NSF ATI program. We also received a great deal of support from Steward Observatory, and for that special thanks goes to *Peter Strittmatter*.

Much of the work presented here has been published previously. Parts of Chapter 2 and Chapter 3 appeared in SPIE proceedings as Males et al. (2010) and Males et al. (2012a), and/or in MagAO technical documents (MAOPs), which can be viewed on-line at <https://visao.as.arizona.edu/documentation/>. Chapter 6 was published in the Astrophysical Journal as Males et al. (2013). My second year project, which appears here as Appendix C, was published in the Astrophysical Journal as Males et al. (2012b).

DEDICATION

To Mom and Dad.

TABLE OF CONTENTS

LIST OF FIGURES	12
LIST OF TABLES	15
ABSTRACT	16
CHAPTER 1 INTRODUCTION	17
1.1 Adaptive Optics	17
1.2 LBTAO To Magellan	18
1.3 Visible AO	20
CHAPTER 2 THE MAGELLAN ADAPTIVE OPTICS SYSTEM AND VISAO CAMERA	22
2.1 Introduction	22
2.2 The VisAO Camera	22
2.2.1 Control Software	25
2.2.2 VisAO Components	26
2.3 CCD-47 Characterization	27
2.3.1 Measurements of Gain and RON	27
2.3.2 The 80 kHz Frame-Transfer Dark Current	30
2.3.3 CCD-47 Linearity	37
2.4 The VisAO Photometric System	40
2.4.1 Photometry in the $r'i'z'$ bandpasses	40
2.4.2 Photometry in Y_S	44
2.4.3 The Impact of Water Vapor on z' and Y_S	44
2.4.4 Exposure Time and Gain scalings	45
2.5 Performance Simulations	45
2.6 Tower Tests	47
2.6.1 Seeing Validation	48
2.6.2 Fitting Error	49
2.6.3 Tower Test Results	51
2.7 On-sky Results	54
2.7.1 Astrometric Calibration	54
2.7.2 Beamsplitter Ghost Calibration	56
2.7.3 Y_s Strehl Ratio	57
2.7.4 Throughput	61

TABLE OF CONTENTS – *Continued*

2.7.5	VisAO Images	62
CHAPTER 3 REAL TIME FRAME SELECTION		67
3.1	Introduction	67
3.2	A Generic Frame Selection Algorithm	69
3.2.1	WFS Telemetry Based RTFS	71
3.3	The Costs & Benefits of Frame Selection	72
3.3.1	Signal to Noise Ratio and Duty Cycle	72
3.3.2	Encircled Energy and Aperture Size	73
3.3.3	Effective Duty Cycle	75
3.3.4	The Speckle Limited Case	76
3.3.5	The Halo Limited Case	76
3.3.6	Background and Read Noise Limited	77
3.3.7	Simulated Faint Guide Star Strehl Selection	77
3.4	RTFS Implementation	80
3.4.1	Mechanical Shutter Performance	80
3.4.2	Telemetry	83
3.4.3	GPU Based Reconstruction	83
3.4.4	Digital Filter Design	85
3.4.5	Reconstructor Calibration	85
3.4.6	Strehl Classification Algorithms	86
3.5	Laboratory Demonstration	90
3.5.1	Experimental Setup	90
3.5.2	Results	92
3.6	Conclusion	94
CHAPTER 4 HIGH CONTRAST IMAGING WITH VISAO: OBSERVA- TIONS OF β PICTORIS B		97
4.1	Introduction	97
4.1.1	Disk Observations and Variability	97
4.1.2	Detection of β Pictoris b	98
4.2	High Contrast Observations of β Pictoris with VisAO	99
4.2.1	The VisAO Coronagraph	99
4.2.2	Observations and Data Reduction	101
4.2.3	VisAO Y_S Contrast Limits	112
4.3	Detection of β Pictoris b with VisAO	113
4.4	Prior Measurements of β Pic b photometry in J , H , and K_S	117
4.5	Prior Exoplanet Photometry in the Y Band	118
4.6	Discussion	121

TABLE OF CONTENTS – *Continued*

4.7	Conclusion and Future Work	126
CHAPTER 5 DETECTABILITY OF EGPS IN THE HZ		128
5.1	The Habitable Zone	128
5.2	Are there planets in the HZ?	130
5.3	The Radius and Temperature of a Giant Planet	132
5.4	Thermal Infrared Brightness of EGPs	134
5.5	Habitable Zone EGPs in the Visible	138
5.6	Blazing The Trail	139
CHAPTER 6 DIRECT IMAGING IN THE HABITABLE ZONE AND THE PROBLEM OF ORBITAL MOTION		142
6.1	Introduction	142
6.2	Motivation and Related Work	143
6.2.1	Nearby Habitable Zones	143
6.2.2	A Different Regime	144
6.2.3	Long Integration Times	145
6.2.4	Related Work	146
6.3	Quantifying The Problem	147
6.3.1	Basic Equations	148
6.3.2	Impact on Signal-to-Noise Ratio	151
6.3.3	Impact on Statistical Sensitivity	155
6.4	Blind Search: Recovering SNR after Orbital Motion	160
6.4.1	Limiting Trial Orbits	161
6.4.2	Choosing Orbital Elements	162
6.4.3	De-orbiting: Unique Sequences of Whole-Pixel Shifts	165
6.4.4	N_{orb} Scalings	165
6.4.5	Recovering SNR	169
6.4.6	Correlations And The True Impact On P_{FA}	172
6.4.7	Impact on Completeness of the Double Test	173
6.4.8	Tractability of a Blind Search	174
6.5	Cued Search: Using RV Priors	175
6.5.1	Constraints	176
6.5.2	Initial Detection	177
6.5.3	Calculating Orbits and Shifts	179
6.6	Conclusions	181
CHAPTER 7 CONCLUSION		184

TABLE OF CONTENTS – *Continued*

APPENDIX A POINT SPREAD FUNCTION RADIOMETRY AND PHOTOMETRY	186
A.1 PSF Modeling	186
A.2 Characterizing the PSF	187
A.3 Tip & Tilt Errors, FWHM, and Strehl Ratio	188
A.4 Gaussian PSF Photometry	189
A.5 Propagation of Errors	190
A.5.1 Absolute Magnitude	190
A.5.2 Physical Photometry	191
APPENDIX B SYNTHETIC PHOTOMETRY AND CONVERSIONS	193
B.1 Filters	193
B.1.1 The Y Band	193
B.1.2 The 2MASS System	195
B.1.3 The MKO System	195
B.1.4 The NACO System	195
B.1.5 The NICI System	196
B.2 Synthetic Photometry	198
B.3 Photometric Conversions	200
B.3.1 Converting 2MASS to MKO	200
B.3.2 Y Band Conversions	201
B.3.3 J Band Conversions	201
B.3.4 H Band Conversions	201
B.3.5 K Band Conversions	201
APPENDIX C FOUR DECADES OF IRC +10216: EVOLUTION OF A CARBON RICH DUST SHELL RESOLVED AT $10\mu\text{M}$ WITH MMT ADAPTIVE OPTICS AND MIRAC4	207
C.1 Introduction	207
C.1.1 The carbon star IRC +10216	207
C.1.2 SiC dust	208
C.1.3 IRC +10216 in the spatial domain	209
C.1.4 New results from the MMT	211
C.2 Observations and data reduction	211
C.2.1 2009 bandpass photometry	212
C.2.2 2010 grism spectroscopy	214
C.3 Archival data	218
C.3.1 Introduction to the spectral datasets	219
C.3.2 Bandpass photometry archives	221

TABLE OF CONTENTS – *Continued*

C.3.3	Comparison of Archival Data	223
C.4	Discussion	224
C.4.1	Changes in the 10 μm spectrum of IRC +10216	224
C.4.2	The spatial signature of SiC emission	227
C.5	Conclusion	229
C.6	Appendix: The possibly erroneous 8.7 μm photometry from 2009 . . .	230
REFERENCES	248

LIST OF FIGURES

2.1	The W-unit	23
2.2	The W-unit as built.	24
2.3	80 kHz dark frame	32
2.4	80 kHz dark frame standard deviation	33
2.5	80 kHz 41 sec frame	34
2.6	80 kHz dark current	35
2.7	80 kHz dark current at different temperatures	36
2.8	80 kHz dark current at histograms	36
2.9	80 kHz column standard deviation	38
2.10	CCD-47 Linearity	39
2.11	VisAO SDSS bandpasses	41
2.12	VisAO z' and Y_S bandpasses	42
2.13	Simulated 1x1 binning performance of MagAO	47
2.14	Seeing limited profile in the tower	49
2.15	MagAO fitting error	50
2.16	Test tower results	52
2.17	Test tower performance vs. predictions	53
2.18	Y_S PSF	58
2.19	θ^1 Ori C	63
2.20	θ^1 Ori B	64
2.21	HR 4796 with VisAO	66
3.1	Simulated Strehl selection results	78
3.2	Simulated Strehl selection FWHM and duty cycle	81
3.3	Shutter performance	82
3.4	Schematic of RTFS	84
3.5	Reconstructed Strehl ratio	87
3.6	Predicting Strehl ratio	88
3.7	Short exposure S statistics	91
3.8	RTFS images	93
3.9	RTFS laboratory results	95
3.10	RTFS duty cycle	96
4.1	Transmission of the VisAO Coronagraph (log scale)	100
4.2	Transmission of the VisAO Coronagraph	102
4.3	VisAO PSFs under the coronagraph	103

LIST OF FIGURES – *Continued*

4.4	Raw PSF FWHM under the mask	104
4.5	Ratio of FWHMs under the mask	105
4.6	VisAO Y_S Coronagraphic PSF	108
4.7	β Pic Median ADI	109
4.8	β Pic Reduction steps	111
4.9	VisAO Y_S Contrast Limits	114
4.10	β Pic Flux and S/N Map	115
4.11	β Pic b S/N	116
4.12	2M1207b and HR 8799b	120
4.13	β Pic b Colors	122
4.14	β Pic b brown dwarf close matches	123
4.15	β Pic b Spectral Type	124
4.16	β Pic b SpType without various filters	125
4.17	β Pic b HR Diagrams	127
5.1	Projected HZs of select nearby stars	131
5.2	EGPs around α Cen A	136
5.3	EGPs around Sirius	137
5.4	Visible reflected light contrast of EGPs	140
6.1	Orbital speed	150
6.2	SNR with orbital motion	152
6.3	Ground-based SNR with orbital motion	154
6.4	Statistics after orbital motion	157
6.5	Whole-pixel shift sequences	166
6.6	Trial orbits for α Cen A	167
6.7	Orbital motion scalings	170
6.8	Possible starting points for Gl 581d	178
6.9	Trial orbits for Gl 581d	180
6.10	Trial orbits for Gl 581d	182
B.1	Y-band filters	196
B.2	J-band filters	197
B.3	H-band filters	197
B.4	K-band filters	198
B.5	2MASS to MKO Conversions	202
B.6	Y Band Conversions	203
B.7	J Band Conversions	203
B.8	H Band Conversions	204
B.9	K Band Conversions	205

LIST OF FIGURES – *Continued*

C.1	MIRAC4 grism observations of IRC +10216	232
C.2	Spatial profiles of IRC +10216 and PSF	233
C.3	FWHM vs wavelength	234
C.4	Slit-loss	235
C.5	IRC+10216 spectrum	236
C.6	IRC +10216 $10\mu\text{m}$ spectra and photometry from 1973 to 1978	237
C.7	IRC +10216 $10\mu\text{m}$ spectra from 1983 to 1988	238
C.8	IRC +10216 $10\mu\text{m}$ spectra from 1993 to 1996	239
C.9	IRC+10216 N band spectra over four decades	240
C.10	$F_\nu(12.5\mu\text{m})/F_\nu(10.55\mu\text{m})$ vs. time	241
C.11	IRC +10216 on the “Carbon-Rich Dust Sequence”	242
C.12	The spectral and spatial signatures of SiC dust around IRC +10216	243
C.13	MIRAC4 detector linearity measurement	244

LIST OF TABLES

2.1	CCD-47 Gain and RON	28
2.2	VisAO Photometry	43
2.3	Predicted VisAO error budget	46
2.4	Tower Test Fitting Error	51
2.5	Clio astrometry of θ^1 Ori $B1$ and $B2$	55
2.6	VisAO Y_S platescale and rotator calibration	56
2.7	50/50 ghost calibration	57
2.8	Strehl measurements at Y_S ($0.984\mu\text{m}$)	61
2.9	VisAO Observation Log: LHS 14	61
2.10	VisAO Throughput	62
4.1	Observations of β Pictoris	106
4.2	Astrometry of β Pictoris b	114
4.3	Photometry of β Pictoris b	118
4.4	Estimated Y_S and NACO photometry of 2M1207b and HR 8799b . .	119
6.1	Trial orbit statistics	168
6.2	SNR recovery after orbital motion	171
6.3	False alarm probabilities after de-orbiting	173
6.4	G1 581d parameters	175
A.1	PSF Characterization	188
B.1	Atmospheres.	194
B.2	Synthetic Photometric System	199
B.3	Photometric conversion coefficients.	206
C.1	IRC+10216 Observations	245
C.2	Bandpass photometry of IRC +10216	246
C.3	Grism photometry of IRC +10216	247

ABSTRACT

One of the most compelling scientific quests ever undertaken is the quest to find life in our Universe somewhere other than Earth. An important piece to this puzzle is finding and characterizing extrasolar planets. This effort, particularly the characterization step, requires the ability to directly image such planets. This is a challenging task — such planets are much fainter than their host stars. One of the major solutions to this problem is Adaptive Optics (AO), which allows us to correct the turbulence in the Earth’s atmosphere, and thereby further the hunt for exoplanets with ground based telescopes. The Magellan Adaptive Optics system has recently obtained its first on-sky results at Las Campanas Observatory, marking a significant step forward in the development of high-resolution high-contrast ground-based direct imaging. MagAO includes a visible wavelength science camera, VisAO, which — for the first time — provides diffraction limited imaging, in long exposures, on a large filled-aperture (6.5 m) telescope. In this dissertation we report on the design, development, laboratory testing, and initial on-sky results of MagAO and VisAO, which include the first ground-based image of an exoplanet (β Pictoris b) with a CCD. We also discuss some of the exciting science planned for this system now that it is operational. We close with an analysis of a new problem in direct imaging: planets orbiting their stars move fast enough in the habitable zone to limit our ability to detect them.

CHAPTER 1

INTRODUCTION

1.1 Adaptive Optics

Everyone knows that stars appear to twinkle. While many have a fond attachment to this phenomenon, for astronomers it presents a significant problem. This twinkling is caused by turbulence in our own atmosphere, and sets a natural limit to the resolution achievable in ground based astronomical imaging. No matter how big we build a telescope, the twinkle-causing turbulence limits our resolution to about 1 arcsecond.

There is, of course, a solution. Using a technology called adaptive optics (AO), we can partially correct for atmospheric turbulence in real-time, effectively “de-twinkling” the stars. In what follows, we report on the development and first on-sky results with a new AO system — the Magellan AO system, or MagAO. MagAO was developed at the UA, with funds from the NSF MRI, TSIP, and ATI programs under principle investigator Laird Close. MagAO is now resident at Las Campanas Observatory (LCO), Chile, where it is installed on the 6.5 m Magellan Clay Telescope for dedicated observing runs.

MagAO uses an adaptive secondary mirror (ASM), a technology proven at the MMT telescope on Mt. Hopkins, Az (Wildi et al., 2003). By minimizing the number of warm surfaces in the optical train, the thermal performance of the system is improved compared to a conventional AO system (Lloyd-Hart, 2000). ASMs offer other benefits as well. Their larger size, relative to conventional deformable mirrors, makes it easier to achieve high actuator densities. ASMs also employ a contactless face-sheet. A significant benefit of this design over other technologies is that it allows failed actuators to be deactivated without creating surface discontinuities.

An important innovation developed for the LBT and now in use at Magellan is

the pyramid wavefront sensor (WFS). A significant improvement this provides over a conventional Shack-Hartmann WFS is the ability to re-bin the detector, which improves its performance on faint guide stars. Pyramid WFSs also suffer from much lower aliasing than other WFS implementations, allowing higher contrast to be reached.

All of these benefits have been taken advantage of in the latest generation of ASM AO systems, first at the Large Binocular Telescope (LBT), on Mt. Graham in Arizona, and now at Magellan.

1.2 LBTAO To Magellan

The MagAO system consists of a near clone of the LBT ASMs. One might initially think that this is a downgrade, as the diffraction-limited spatial resolution of a telescope is given by

$$FWHM = 0.2063 \frac{\lambda}{D} \text{ arcseconds}$$

where $FWHM$ means the full-width at half-maximum of the point spread function (PSF) in arcseconds, λ is the wavelength of light being observed in μm , and D is the diameter of the telescope in meters. The 8.4 m LBT primaries should have nearly 30% better resolution (smaller $FWHM$) than the 6.5 m Magellan Clay telescope¹. Furthermore, the sensitivity (“light-grasp”) of the telescope goes as at least D^2 , that is with the collecting area. With AO, in fact, we expect to see D^4 improvement in point-source sensitivity due to the complimentary effects of increased collecting area and smaller PSF reducing background noise.

But there is another consideration: actuator pitch. When projected over the pupil, that is the primary mirror, the actuators have an effective spacing of

$$d = \sqrt{\frac{\pi}{4} D^2 (1 - \epsilon^2) \left(\frac{1}{N_{act}} \right)} \text{ meters.}$$

where ϵ is the central obscuration ratio and N_{act} is the total number of illuminated actuators². At the LBT, with $D = 8.4\text{m}$ and $N_{act} = 666$, we find $d = 28.7 \text{ cm}$. At

¹Here we are ignoring details such as undersized cold stops

²Taking into account the central obscuration

Magellan, with $D = 6.5\text{m}$ and $N_{act} = 561$, we have $d = 23.3\text{ cm}$. Note that we have accounted for the central obscurations so N_{act} is not actually 672 and 585 for the two systems, respectively. The point to this calculation is that on the smaller primary mirror at Magellan, the same physical actuators are closer together when projected onto the sky.

To appreciate the impact actuator spacing has, we must consider the quantity called Strehl ratio, which we will denote as S . This is a measure of image quality defined as the ratio of the PSF peak obtained with an imaging system to that expected if that system were perfect (Hardy, 1998). Theoretical calculations of S typically proceed by determining the errors from different sources. The error concerning us here is the so-called “fitting-error”, which quantifies the fact that we can only correct atmospheric turbulence up to a certain spatial frequency. The fitting-error for a continuous face-sheet deformable mirror is given by (Hardy, 1998)

$$\sigma_{fit}^2 = 0.28 \left(\frac{d}{r_0} \right)^{5/3} \text{ rad}^2.$$

The quantity r_0 is the Fried coherence length, a measure of seeing. If we assume that all sources of error add in quadrature, we can employ the Marechal approximation (Born and Wolf, 1999) to calculate S due to fitting error:

$$S_{fit} = e^{-\sigma_{fit}^2}.$$

Finally, to fully apply these formulas, we need an estimate for r_0 . It can be expressed in terms of the seeing $FWHM$, θ , as

$$r_0 = 0.2022 \frac{\lambda}{\theta} \text{ meters}$$

where λ is in μm and θ is in arcseconds. On Mt. Graham, Arizona, the location of the LBT, median seeing at V band ($0.55\mu\text{m}$) is $0.8''$. So at the LBT median $r_0 = 13.9\text{ cm}$. At LCO median seeing at V is $0.625''$ (Floyd et al., 2010)³, so

³This is median DIMM seeing, measured seeing on the full 6.5m aperture is better than this due to outer scale effects.

median r_0 is 17.8 cm. Using the relationship $r_0 \propto \lambda^{6/5}$, we can make the following comparison:

At the LBT, observing at J band ($1.2\mu\text{m}$):

$$\begin{aligned} S_{fit} &= 81\% \\ FWHM &= 0.029'' \end{aligned}$$

At Magellan, observing at i' ($0.77\mu\text{m}$):

$$\begin{aligned} S_{fit} &= 80\% \\ FWHM &= 0.024'' \end{aligned}$$

So we see that the smaller effective actuator spacing, combined with the better seeing at LCO, means that the same AO technology should be able to provide the same level of spatial correction at i' at LCO as it does at J at the LBT — but actually realizing an improvement in resolution due to the shorter λ despite the smaller primary D . The full story is more complicated than this, as we have ignored such things as errors due to servo lag, but these simple arguments illustrate the advantage a smaller primary mirror provides for an AO system.

The high performance of the LBT ASMs has been demonstrated on sky at Mt. Graham (Esposito et al., 2010). Using the smaller Magellan primary, the relatively higher projected actuator pitch allows the same technology to provide the same excellent performance at shorter wavelengths. This motivated the development of a visible wavelength science camera for Magellan, which we call VisAO.

1.3 Visible AO

Other groups have implemented visible light AO in one form or another. ViLLaGEs is a MEMs-based visible wavelength AO testbed on the 1.0 m Nickel telescope at Lick Observatory, in California (Morzinski et al., 2010). Baranec et al. (2012)

have developed a visible AO capability on a 1.5 m telescope at Palomar observatory (Robo-AO). The USAF 3.6 m AEOS telescope has also worked in the visible (Roberts and Neyman, 2002), achieving moderate correction. The Palomar AO system had some capability to work in the visible, particularly when employing Lucky imaging (Law et al., 2009). The largest telescope with a visible AO capability is the 8.2m Subaru, with the AO188 system, which can achieve the diffraction limit in the visible utilizing Lucky imaging techniques in the Fourier domain (Garrel et al., 2010).

The Magellan VisAO camera represents a true step forward. In the following pages we report on the design, performance analysis, laboratory testing, and finally the on-sky demonstration of the world’s first truly diffraction limited visible light imager on a large ($> 6\text{m}$) telescope. The distinction setting VisAO apart from previous efforts is that it delivers filled-aperture, long exposure images, with diffraction limited cores, and Strehls greater than 20% — at visible wavelengths.

In Chapter 2 we present an overview of VisAO, and the simulation and laboratory characterization of the MagAO/VisAO system carried out prior to shipping to LCO. We also describe our initial on-sky characterization efforts. In Chapter 3 we discuss our version of the Lucky imaging technique, Real Time Frame Selection (RTFS), which uses a fast shutter to select images using WFS telemetry. In Chapter 4 we present observations of the exoplanet host star β Pictoris, demonstrating the first high-contrast exoplanet science with a CCD on the ground. Then in Chapter 5 we lay out some of the future exoplanet science goals for the system, namely a search for extrasolar giant planets (EGPs) in the habitable zones (HZs) of the nearest stars. Finally, in Chapter 6, we consider the far future, when the next generation of giant telescopes will enable HZ observations of many more stars, but the higher projected orbital speeds of these planets will degrade our sensitivity.

CHAPTER 2

THE MAGELLAN ADAPTIVE OPTICS SYSTEM AND VISAO CAMERA

2.1 Introduction

In this chapter we present an introduction to the VisAO camera, and an overview of the work done to prepare it and the MagAO system to go on-sky. We then present some of our first on-sky results, and show some of our early efforts to characterize the system.

Some of the work presented in this chapter has appeared in print in Males et al. (2010) and Males et al. (2012a). Much is also contained in one form or another in MagAO technical documents (MAOPs), which can be viewed on-line at <https://visao.as.arizona.edu/documentation/>

2.2 The VisAO Camera

The main focus of this dissertation is MagAO's visible light science camera, VisAO. VisAO is the world's first diffraction-limited imager on a large telescope capable of working at visible wavelengths. On the 6.5m Magellan Clay Telescope, VisAO is capable of 19 milliarcsecond resolution at $0.62\mu\text{m}$ (the r' central wavelength) (Close, et al., (2013, submitted)).

A very basic visible camera is part of the LBT W-unit baseline, where it is used almost exclusively as a wide-field seeing-limited acquisition camera. Here we provide an introduction to the VisAO camera and highlight the many upgrades and optimizations we have made to ready the system to record some of the highest resolution filled aperture images ever taken.

This 1024x1024 CCD camera provides $0.0079''$ pixels, Nyquist sampling the diffraction limited PSF down to $\sim 0.5\mu\text{m}$ with an $8.1''$ field of view (FOV). A

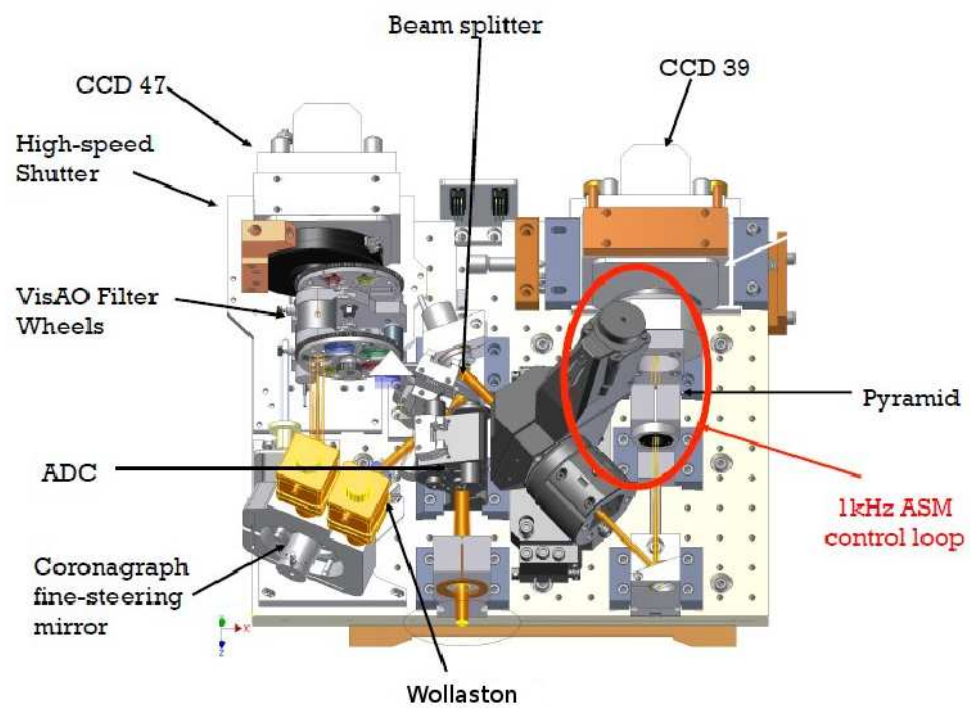


Figure 2.1 The Magellan AO system WFS and VisAO camera.

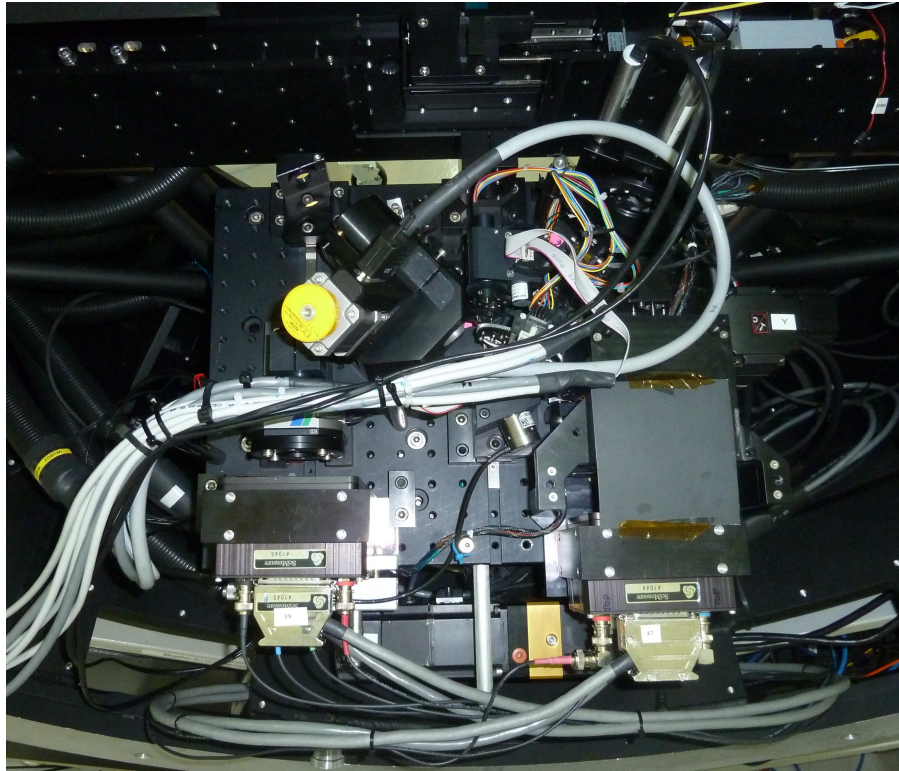


Figure 2.2 The Magellan AO system WFS and VisAO camera, as built. Shown here prior to installation of the CCD and shutter cooling system. Note that orientation is flipped with respect to Figure 2.1.

key feature of this design regarding VisAO performance is that the CCD47 is on a common mount with the WFS. Figure 2.1 illustrates the design of our Magellan AO WFS and VisAO optical board, and Figure 2.2 shows the as built system.

2.2.1 Control Software

Adding significant science observation capability to what is essentially a static acquisition camera offered many challenges. An important one identified early on was the need to not break the LBT AO system software, that is to make sure that the system operated at Magellan as similarly as possible to the LBT. In addition, the LBT makes use of a Microgate basic computational unit (BCU) as the framegrabber for the CCD-47. It became apparent that this architecture did not offer enough flexibility to fully exploit the capabilities of the VisAO camera. To compensate, we implemented our own framegrabber using a PCI card provided by Scimeasure, and developed a software system for emulating the BCU so that the AO system doesn't notice the missing device. All VisAO extensions to the LBT AO “adopt” software are seamlessly integrated, such that there are very few changes to the core of the AO control software. We took pains to ensure that process control (starting, stopping, state reporting, etc.) are identical to the “adopt” system. In short, VisAO functions as a native component of the LBT AO software used to control the MagAO system.

An important consideration was to ensure that the AO system, or rather the AO operator, could not inadvertently corrupt a science operation by, say, changing a filter wheel. Conversely, during AO acquisition the VisAO astronomer must be careful to not reconfigure the CCD-47 or other components lest the acquisition sequence fail. We implemented a hierarchical system of control, whereby all VisAO processes have a control state — REMOTE when under control of the AO system, LOCAL when under control of the VisAO astronomer, and SCRIPT when an observation script is running. For instance, the VisAO astronomer can not change filters unless she explicitly takes control of the filter wheel first.

We also developed several real-time components, mainly in support of the Real Time Frame Selection (RTFS) technique, which is described in full in Chapter 3.

Included in this system is real-time reconstruction of WFS slopes, which are used to calculate instantaneous Strehl ratio. We also use this system to record wavefront error (WFE) and write the average value of WFE during an exposure to the FITS headers. Details of this are also provided in Chapter 3, and we demonstrate the use of WFE to estimate Strehl ratio below.

The resulting software system, written almost entirely by the author, consists of over 300 files of source code in C, C++, idl, python, and BASH scripts. These files contain over 57000 lines of code, including code and comments but ignoring whitespace. The source, source documentation, and a user’s guide can be browsed at https://visao.as.arizona.edu/software_files/visao/html/index.html.

2.2.2 VisAO Components

VisAO has several custom components which optimize the camera for high-contrast diffraction-limited circumstellar science. These components are additions to the LBT W-unit baseline. Here we provide a brief overview of the major components. For more information on these and other components see the MagAO technical documents (MAOPs) provided at <https://visao.as.arizona.edu/documentation/>.

The Wollaston: A Wollaston beamsplitter prism is located just before the gimbal on a manually actuated elevator stage. When placed in the beam, the beam is split vertically allowing simultaneous differential imaging (SDI) using custom double filters located in filter wheel 3.

The Gimbal: The VisAO gimbal mirror is actuated, providing steering of the beam on the CCD-47 detector. This is broadly necessary as each beamsplitter has a different tilt. It is also necessary for coronagraph alignment, and optimizing the FOV for various targets. The mechanical FOV of the gimbal is $\sim 12 \times 17$ arcseconds, compared to $\sim 8 \times 8$ arcseconds on the detector.

The Focus Stage: The VisAO focus stage moves the CCD47, shutter, and VisAO filter wheel assembly in the z direction. The AO keeps the system focused on the tip of the pyramid at all times, so the VisAO focus stage only compensates for changes in the relative focus between the CCD-47 detector and the pyramid tip.

This depends on wavelength (filter selection), and whether or not the Wollaston prism is in the beam. Further information about the focus stage and its operation is provided in MAOP-706.

Filter Wheels: The VisAO camera has 2 filter wheels, instead of the one in the LBT baseline. The first wheel contains our broad bandpass filters: SDSS r' , i' , z' , and what we call Y short (Y_S) at $1\mu\text{m}$. The second wheel contains custom SDI filters (filters with two bandpasses on a common substrate), an ND 3, and a partially transmissive coronagraphic occulting mask (see Chapter 4). The filter curves and other characteristics of the filters are shown later in this chapter.

The CCD-47: The MagAO CCD 47 is the system acquisition camera and the main sensor of the visible wavelength science camera (VisAO). It is an EEV CCD-47 detector, with a Scimeasure Little Joe controller. Though the hardware is the same as those in a standard LBT W-unit, we acquired several unique operating modes (programs). Visible AO works best with a bright natural guide star (NGS), so one of our biggest challenges is avoiding saturation. As such, we frequently operate in high speed modes, up to 42 fps. For sensitive work, such as high-contrast circumstellar science, we take longer exposures in more sensitive gain settings. In the next section we describe our extensive characterization of these different modes.

2.3 CCD-47 Characterization

During laboratory integration, and now in on-sky testing, we have characterized the CCD-47 sensitivity and linearity. The results reported here are also provided in MAOP-702.

2.3.1 Measurements of Gain and RON

Each CCD-47 program defines a readout speed, window (or sub-array) size, binning, and gain. Each combination of readout speed, window, and binning has 4 possible choices of gain: high, medium-high, medium-low, and low. These correspond to most sensitive to least sensitive, respectively. Together, the readout speed and gain

set the readout noise (RON) of the camera. The CCD-47 is a 14 bit camera, which has important system sensitivity implications. The lowest sensitivity gain is ~ 13 electrons per ADU. That means that it takes 13 electrons to register a signal. At low flux levels, e.g. in the wings of the PSF halo, this makes the camera much less sensitive than from RON and photon noise alone. N.B. that when in this sensitivity setting, integrating longer will not reduce noise as expected (\sqrt{N}).

Gain and read-out noise were measured in the Magellan AO lab at Steward Observatory in February 2010, prior to being mounted on the W-Unit board. For these measurements the CCD head was wrapped in Al foil, place in a cardboard box, and had liquid cooling applied. The cardboard box had a hole cut in it, a paper placed over the hole to provide a somewhat flat illumination, and an LED flashlight was used as the source. The lab thermostat was set to minimum to provide a cool ambient temperature to minimize the impact of dark current. For these tests the Little Joe case temperature was 20C. CCD47 Head temperature was -36C, except in the 64x64 and 32x32 modes when it rose to -33C due to the high frame rate. At each pixel rate and gain setting we took 2 darks and 2 flats, which were then analyzed using the findgain task in IRAF. Two sets of data were taken at each setting, and typical variations between these sets was 0.01 for gain and 0.02 ADU for RON. The 64x64 and 32x32 modes had larger variations, and the numbers presented are the average of the two sets. Results are presented in Table 2.1. These data are also published in MAOP-702.

Table 2.1: CCD-47 Gain and readout noise measurements

		Measured		Scimeasure	
Mode	Gain Setting	Gain (e^- /ADU)	RON (e^-)	Gain (e^- /ADU)	RON (e^-)
2500 kHz	High	0.53	9.7	0.55	10.2
1024x1024	Med High	1.93	9.55	1.97	9.83
Bin 1x1	Med Low	3.58	10.74	3.62	10.4

continued on next page

Table2.1– *continued from previous page*

		Measured		Scimeasure	
Mode	Gain Setting	Gain (e^- /ADU)	RON (e^-)	Gain (e^- /ADU)	RON (e^-)
3.53 fps	Low	13.23	15.47	13.3	15.3
2500 kHz	High	0.54	9.62	-	-
64x64	Med High	1.93	9.58	-	-
Bin 1x1	Med Low	3.58	10.86	-	-
31.48 fps	Low	13.14	15.49	-	-
250 kHz	High	0.47	4.52	0.49	5.81
1024x1024	Med High ¹	1.77	4.67	1.71	5.66
Bin 1x1	Med Low ²	3.34	5.28	3.29	6.59
0.44 fps	Low ¹	12.3	11.11	12.1	10.8
80 kHz ³	High	0.48	7.35 / 3.54	0.48	3.37
1024x1024	Med High	1.78	6.3 / 3.69	1.79	3.53
Bin 1x1	Med Low ¹	3.33	6.23 / 4.38	3.31	4.28
0.143 fps	Low ¹	12.43	12.35 / 11.02	12.2	10.3
80 kHz ³	High	0.48	5.69 / 3.62	0.48	3.28
1024x1024	Med High	1.74	5.98 / 3.72	1.79	3.61
Bin 2x2	Med Low ¹	3.27	6.18 / 4.43	3.31	3.29
0.551 fps	Low	11.08	12.1 / 9.68	12.2	10.3
80 kHz	High ⁴	0.47	9.76 / 6.85	0.46	3.62
1024x1024	Med High ⁴	1.76	10.07 / 7.31	1.74	3.95
Bin 16x16	Med Low ⁴	3.25	10.42 / 7.43	3.31	4.63
10.42 fps	Low ⁴	12.38	14.43 / 12.13	11.6	10.3
2500 kHz	Low	0.53	9.59	-	-
512x512	Med Low	1.93	9.54	-	-
Bin 1x11	Med High	3.57	10.71	-	-
6.70 fps	High	13.26	15.55	-	-
2500 kHz	High	0.54	9.46	-	-
32x32	Med High	1.88	9.57	-	-

continued on next page

Table 2.1– *continued from previous page*

		Measured		Scimeasure	
Mode	Gain Setting	Gain (e^- /ADU)	RON (e^-)	Gain (e^- /ADU)	RON (e^-)
Bin 1x1 42.78 fps	Med Low	3.5	10.59	-	-
	Low	12.61	14.51	-	-
250 kHz 512x512	High	0.48	3.84	-	-
	Med High	1.77	4.25	-	-
Bin 1x1 1.49 fps	Med Low	3.32	4.88	-	-
	Low	12.36	10.52	-	-
80 kHz 512x512	High ⁵	0.47	9.06 / 8.66	-	-
	Med High	1.74	4.13 / 3.36	-	-
Bin 1x1 0.535 fps	Med Low	3.32	4.82 / 4.24	-	-
	Low	12.46	10.94 / 10.58	-	-

¹Used 1 pass of 5σ clipping

²One bad dark frame here gives odd results. Ignored.

³The 80kHz RON measurements require special handling due to excess frame-transfer dark current. The 2nd number is from the alternate 100 frame method described in 2.3.2.

⁴The bad results here are explainable by the excess dark current. Taking into account both the decreased frame time and the larger number of pixels in each bin, there is 3.5 as much dark current per pixel in these images

⁵This mode appears to be genuinely out of spec. We had to adjust black levels in this mode (a consequence of low Joe temperature) but it would be surprising if this affects RON.

2.3.2 The 80 kHz Frame-Transfer Dark Current

As noted in Table 2.1, the raw 80kHz RON was significantly worse than expected. The number one suspect is dark current since we did not measure RON with 0 exposure time. Upon investigating, we found that a dark current is the likely culprit,

however it appears that it is not simply a dark current which scales with exposure time.

To start our investigation we took 100 dark frames (cap on) at 80 kHz. Figure 2.3 shows the median of these frames. We next took the standard deviation of the 100 frames on a pixel by pixel basis, shown in Figure 2.4. It appears that the signal shown in Figure 2.3 is a source of Poisson noise, which is at the same level as expected to explain the high RON results. This dark signal is much higher than expected based on the E2V specifications for our CCD47.

In Figure 2.5 we show the median of 50 41 second dark frames. Here we see the first hint that the dark signal in Figure 2.3 is not scaling with time¹. A separate pattern is now becoming visible. In Figure 2.6 we show the median dark current, which was calculated by subtracting a 6.9 second exposure from a 94 second exposure. The short exposure was not scaled, so we see that the dark signal in Figure 2.3 is indeed not scaling with time, and once it is subtracted a dark signal more in line with that expected is evident.

Our best guess to explain these results is that the high dark signal found in Figure 2.3, which causes the high RON at 80kHz, is associated with the frame-transfer architecture and that it depends only on readout time which is a constant set by pixel-rate and is independent of exposure time. In other words, this is dark current on the transfer frame and its impact is controlled by how long charge sits on the transfer frame during the readout process, not on how long charge sits on the exposed science frame.

We found that this frame-transfer dark signal does scale with temperature. At a head temperature of -29.5C RON was 11.32 electrons, and at -33C RON was 8.41 electrons. To further test this, we added a second cold plate which got head temperature down to -36C. In Figure 2.7 we show a side-by-side comparison of this signal at -32C and -36C, demonstrating the reduction in the dark signal. In Figure 2.8 we show the change in the histogram of the RON of all pixels on the array with the reduced temperature.

¹We note that temperature did not change significantly during these measurements.

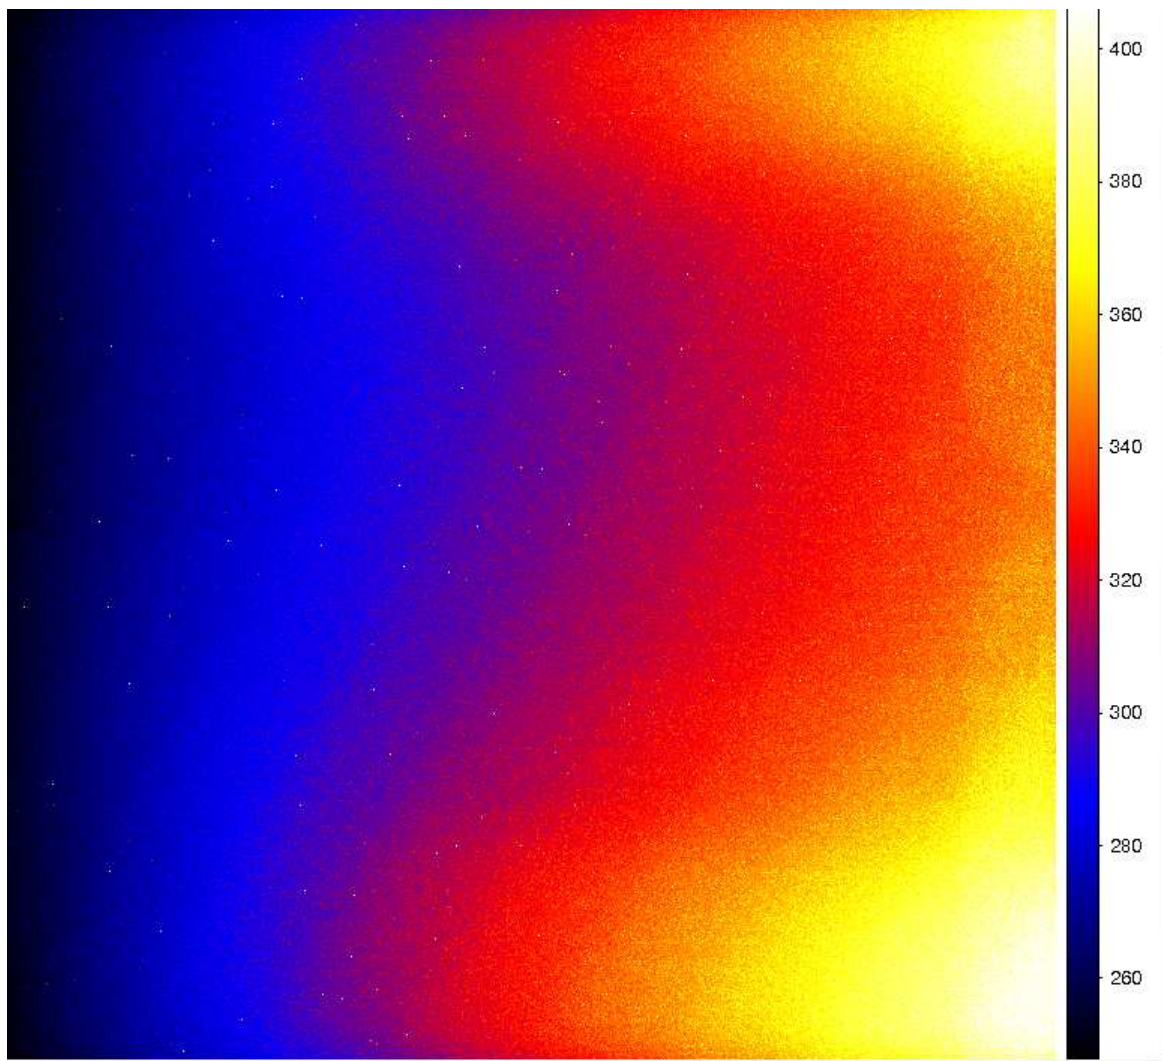


Figure 2.3 Median of 100 80 kHz dark frames, showing the structure of the excess dark current. Exposure time of individual frames was 6.9 secs.

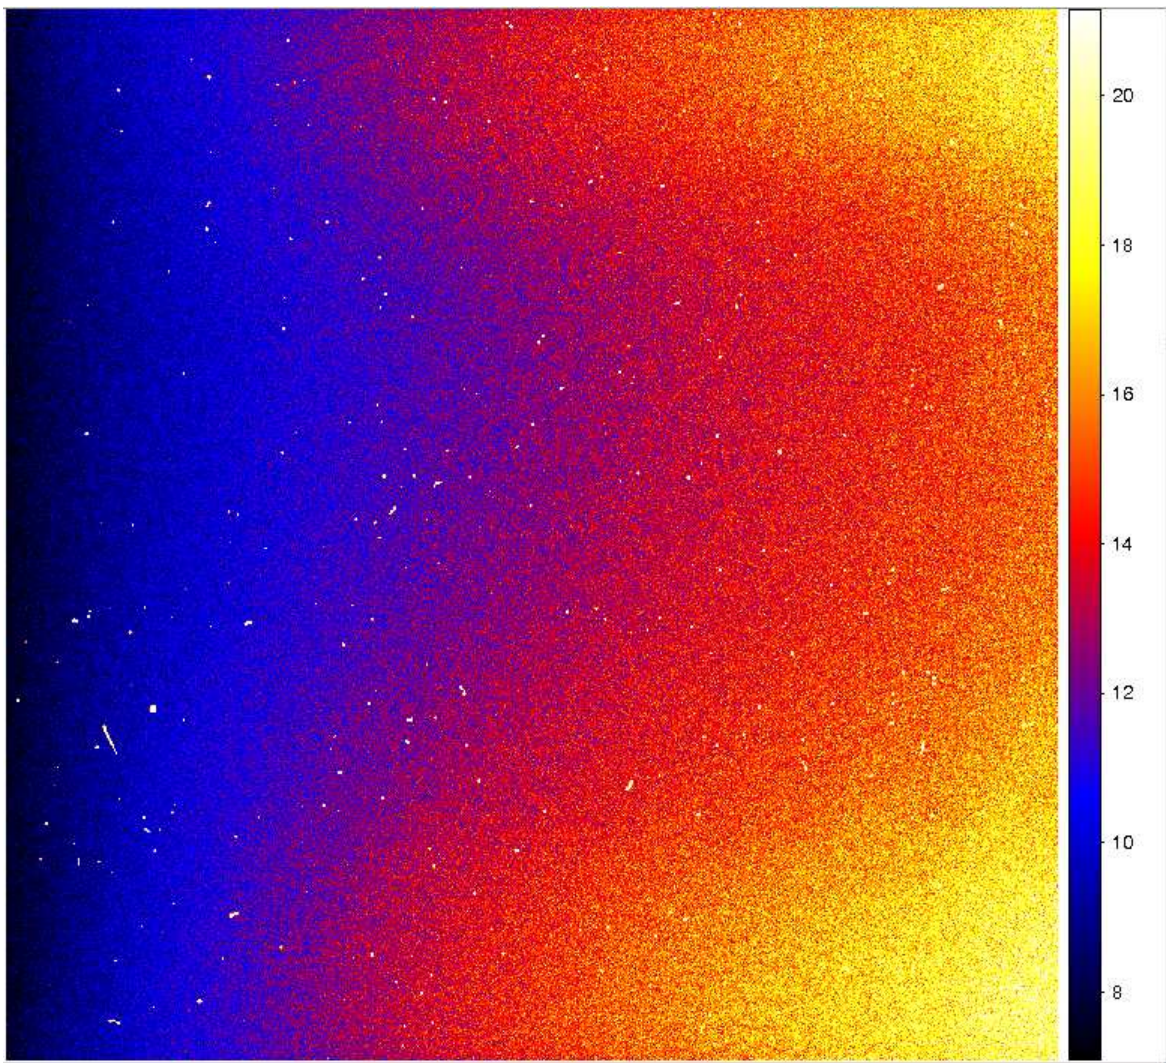


Figure 2.4 Pixel by pixel standard deviation of 100 80 kHz dark frames. The excess structure seen in Figure 2.3 appears to be a source of Poisson noise.

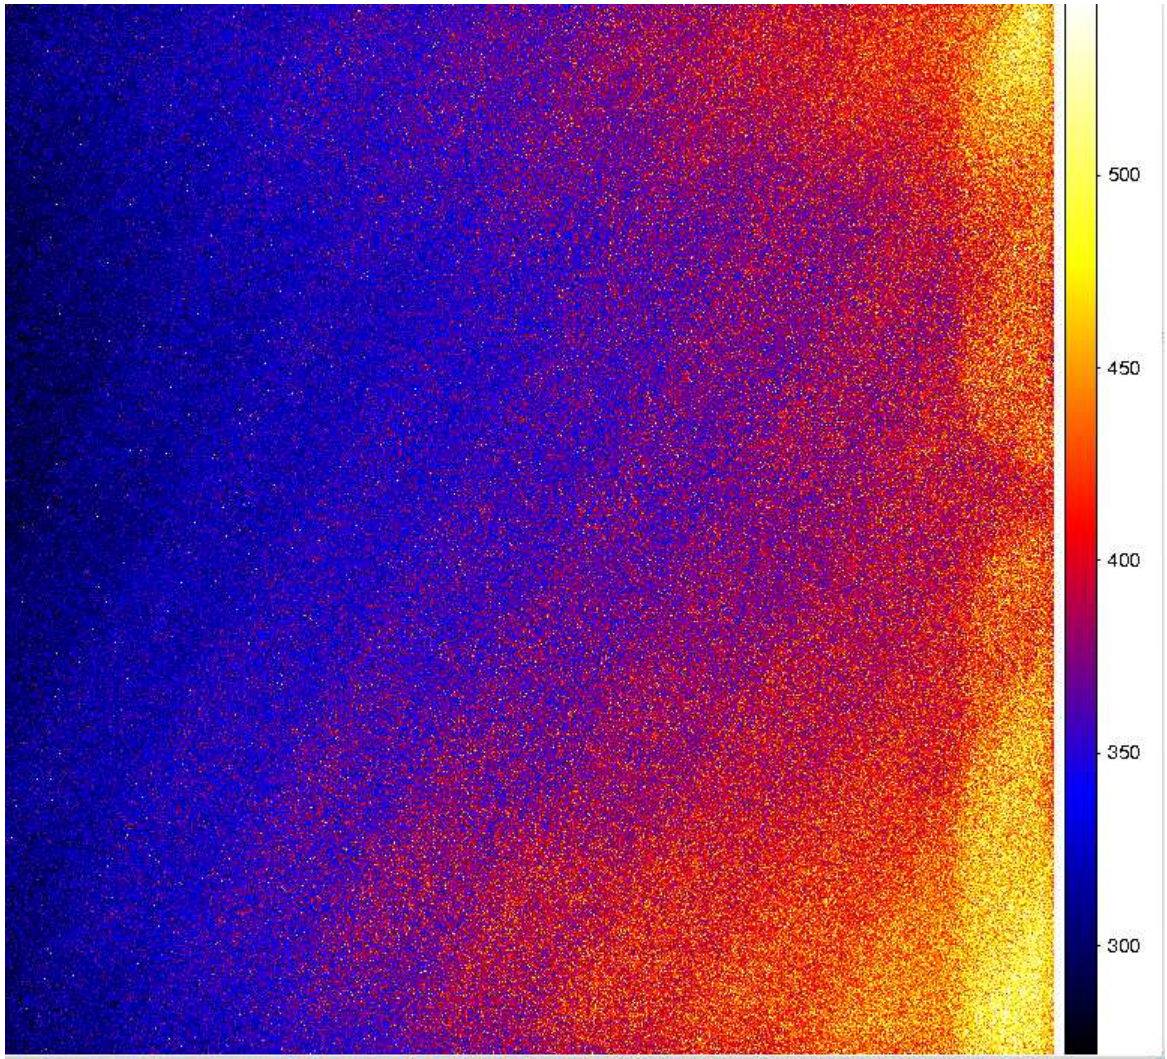


Figure 2.5 Median of 50 41 sec frames, at 80 kHz. The structure evident in Figure 2.3 does not scale with exposure time. This, and all images that follow, are full 1024x1024 frames.

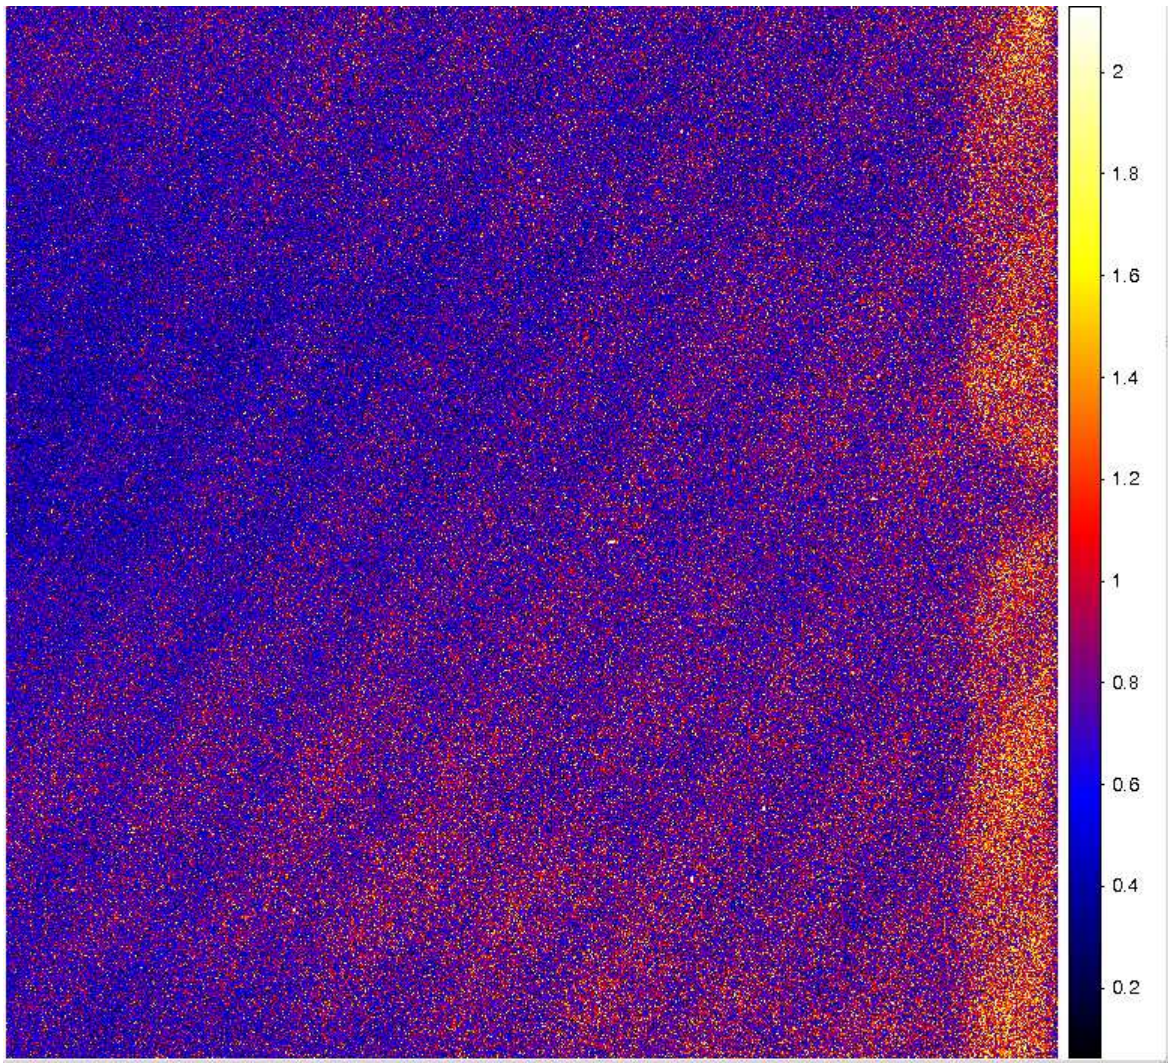


Figure 2.6 This is the dark current, calculated by subtracting a 6.9 second frame from a 94 second frame (80kHz) and dividing by exposure time. Note that the structure in the first image has almost completely subtracted out, but the pads and the waves are clearly visible.

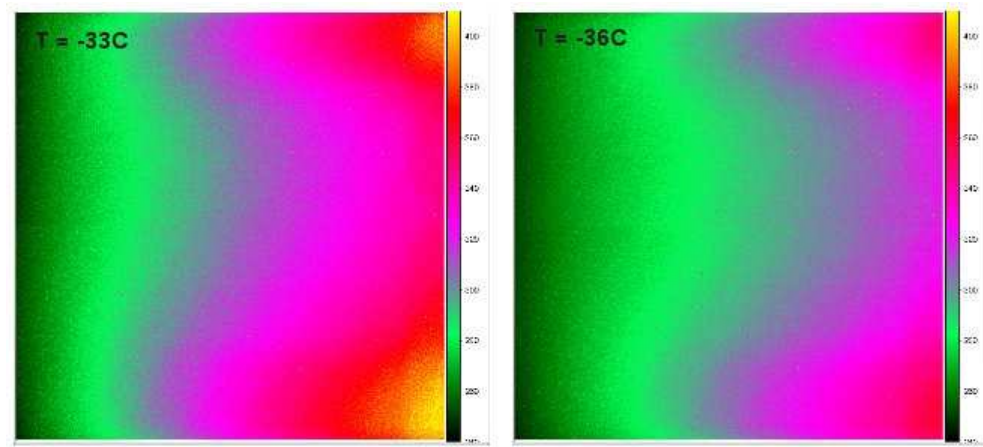


Figure 2.7 6.9 sec images at different temperatures, same stretch and colorbar. The dark current is lower at -36C.

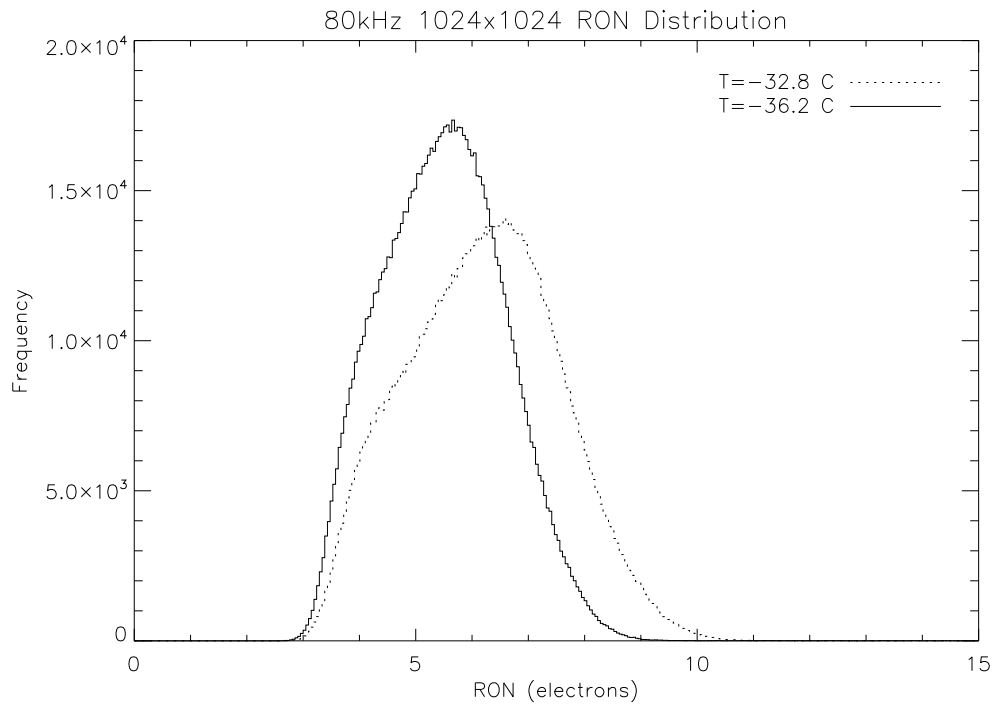


Figure 2.8 Histograms of the full array at two different temperatures, in the 80kHz readout speed. As temperature is lowered, it appears that we might approach the expected value of $\sim 3.5e^-$ RON.

In Figure 2.9 we plot the column-wise standard deviation, at -32C and -36C. We see that the signal is lowest at column 0, so we assume that this is the first column read out and column 1023 is the last column read out. To provide an estimate of our true RON measured in the lab, we calculate the value of the column 0 standard deviation by fitting a line to the first 100 columns and taking its intercept, which is 3.54 electrons in this case. This technique provides the second measurements in the 80kHz sections of Table 1.

If we achieved a -50C head temperature — which would require a 0-10C ambient temperature — linear extrapolation predicts that we will achieve the expected (based on Sciencemeasure’s measurements) value of 3.37 electrons. In practice we do not achieve this temperature very often at LCO, the coldest seen so far is roughly -45C. As such, the 80 kHz readout speed is not used as it has worse sensitivity than the 250kHz $3.8e^-$ RON mode due to this excess “frame-transfer” dark current.

2.3.3 CCD-47 Linearity

We measured the linearity of the VisAO CCD-47 using ambient light in the Auxiliary building at LCO, and varying the exposure time in the 2500 kHz full frame mode. This was done only in the LOW gain setting, as the higher gains will all digitally saturate at 16383 ADU before reaching non-linearity. Data was dark subtracted to remove the bias, and the median of a subarray was calculated at each exposure time. The subarray was chosen to correspond to the brightest part of the ambient light pattern. The results are shown in Figure 2.10. The CCD-47 is linear up to at least 9000 ADU, which corresponds to 119000 electrons. The manufacturer quoted typical value for well depth is 100000 electrons. In all other modes the CCD-47 is linear to 16383 ADU. In future work we plan to analyze the linearity of each pixel separately.

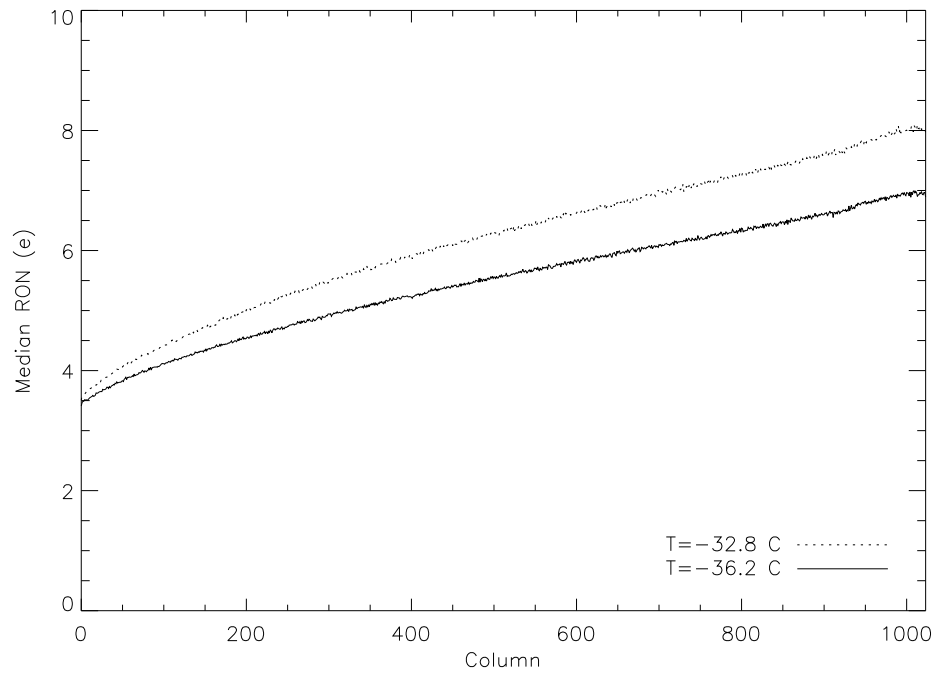


Figure 2.9 Column-wise standard deviation in the 80kHz mode. We hypothesize that column 0 is the first read out, and so has minimal frame-transfer dark current. Column 0 has the expected $\sim 3.5e^-$ RON.

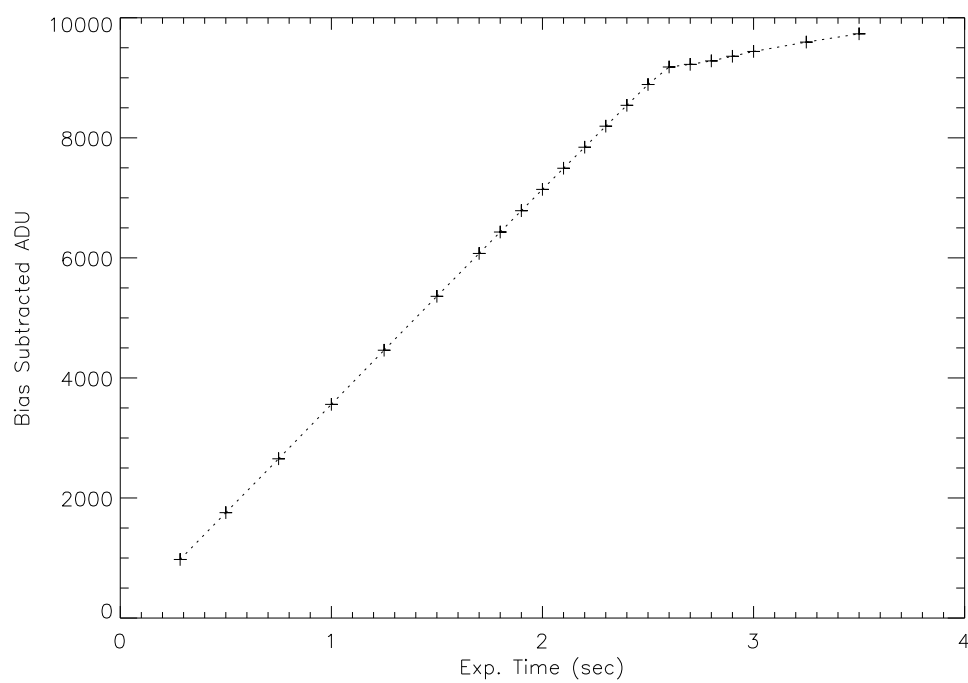


Figure 2.10 CCD-47 linearity measurement. The CCD-47 is linear up to 9000 bias subtracted ADU.

2.4 The VisAO Photometric System

The VisAO camera has 4 broadband filters: r' , i' , z' and Y_S (Y -short). The $r'i'z'$ filters are based on the Sloan Digital Sky Survey (SDSS) specifications (Fukugita et al., 1996), and were provided by Asahi Spectra. The Y_S bandpass is defined by a Melles-Griot long wavepass filter (LPF-950) which passes $\lambda \gtrsim 0.95\mu\text{m}$. We convolved the transmission curves provided by the respective filter manufacturers with the quantum efficiency (QE) for our EEV CCD47-20 with near-IR coating, and included the effects of 3 Al reflections. We also photon-weighted, as appropriate for a CCD, using the following equation (Bessell, 2000)

$$T(\lambda) = \frac{1}{hc} \lambda T_0(\lambda) \quad (2.1)$$

where T_0 is the raw energy-weighted profile.

The resultant VisAO filter profiles are shown in Figure 2.11 along with comparable Johnson and SDSS filter profiles. The 0 airmass (AM) transmission profiles for z' and Y_S are shown in Figure 2.12 along with examples of Y and J filters commonly used in exoplanet observations.

The SDSS system is an AB system (Fukugita et al., 1996), however to-date most, if not all, exoplanet direct-imaging observations have been reported in Vega based magnitudes (as are most galactic observations of any type). To facilitate comparisons of results from VisAO, we here define the VisAO photometric system such that for Vega $V = r' = i' = z' = 0.03$ mag. We integrated the filter profiles with the HST CALSPEC spectrum of Vega from Bohlin (2007) to determine the flux densities of a 0 mag star in each filter. The results, along with other relevant filter characteristics, are shown in Table 2.2.

2.4.1 Photometry in the $r'i'z'$ bandpasses

In the AB magnitude system the flux of a 0 magnitude star in any bandpass is defined as 3631 Jy. Using the 0 mag flux shown in Table 2.2 we find the transformations

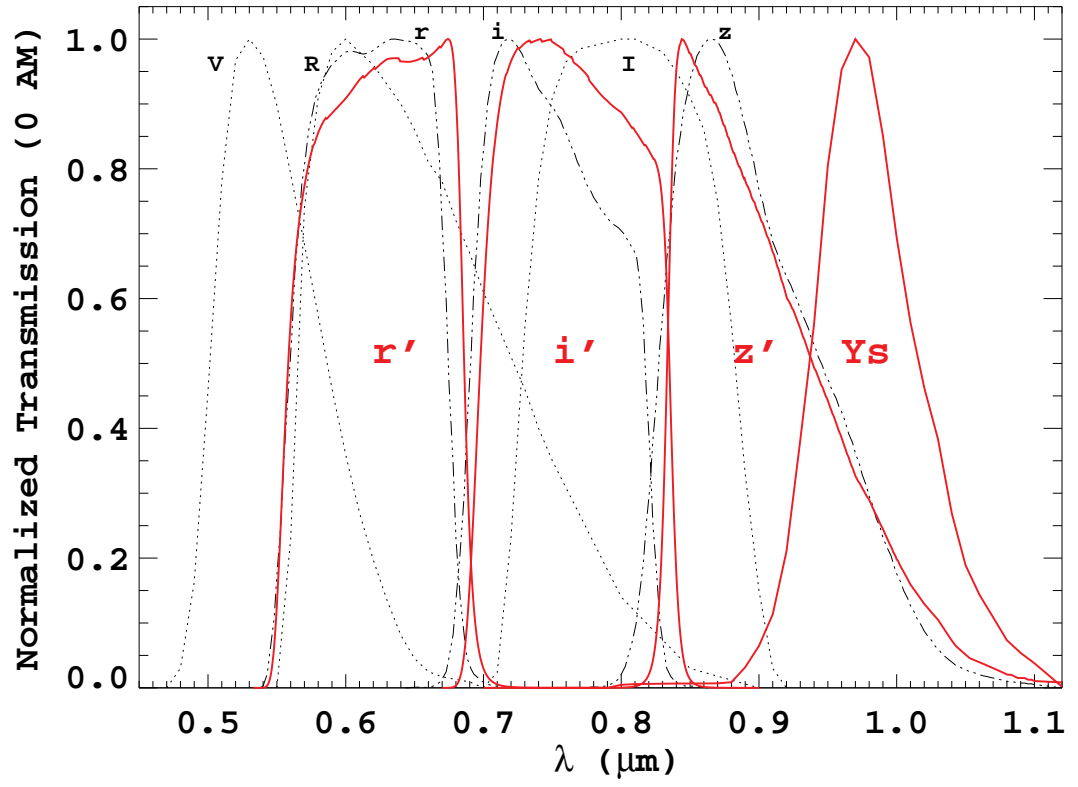


Figure 2.11 The VisAO broadband filters are shown in red. For comparison the SDSS riz filter profiles are shown. We also show the VRI profiles from Bessell (1990).

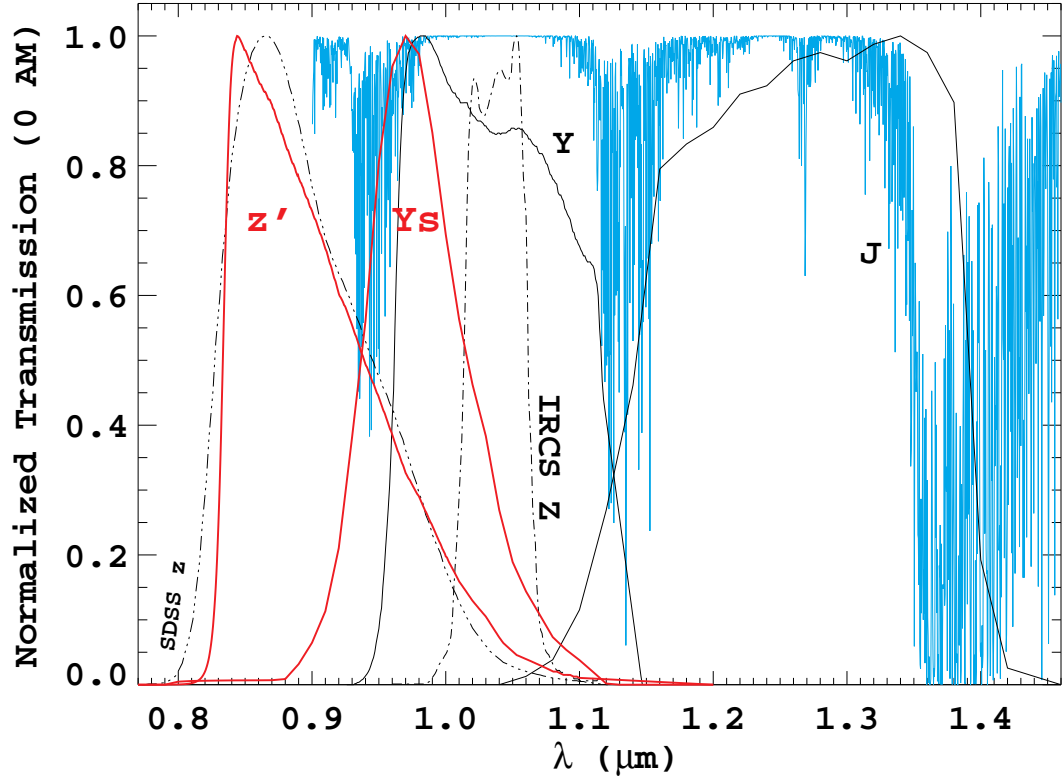


Figure 2.12 A comparison of the VisAO z' and Y_S bandpasses with near-IR filters used to observe exoplanets. The Y bandpass is from Hillenbrand et al. (2002). The Keck/NIRC2 “Z” bandpass was digitized from a plot obtained at <http://www2.keck.hawaii.edu/inst/nirc2/filters.html>, and appears identical to the Subaru/IRCS Z filter. The J profile is also from Keck/NIRC2. Atmospheric transmission for 2.3mm precipitable water vapor is from the ATRAN models (Lord, 1992) provided by Gemini Observatory.

Table 2.2 The VisAO Photometric system

Filter	AM0 λ_0 ¹ (μm)	FWHM (μm)	$\Delta\lambda$ ² (μm)	0 mag F_λ ³ ($\text{ergs/s/cm}^2/\mu\text{m}$)	0 mag F_ν ³ (Jy)	0 mag F_γ ³ ($\gamma/\text{s/m}^2/\mu\text{m}$)
r'	0.626	0.125	0.112	2.510×10^{-5}	3221	7.822×10^{10}
[OI]	0.630	0.004	0.005	2.435×10^{-5}	3228	7.727×10^{10}
[OI]/ $H\alpha$ Cont.	0.643	0.006	0.007	2.285×10^{-5}	3149	7.394×10^{10}
$H\alpha$	0.656	0.005	0.005	1.733×10^{-5}	2491	5.728×10^{10}
[SII]	0.673	0.005	0.006	1.997×10^{-5}	3012	6.760×10^{10}
[SII] Cont	0.700	0.005	0.006	1.768×10^{-5}	2889	6.228×10^{10}
i'	0.767	0.137	0.132	1.353×10^{-5}	2616	5.177×10^{10}
z'	0.910	0.116	0.121	8.453×10^{-6}	2302	3.839×10^{10}
Y_S	0.984	0.081	0.091	6.957×10^{-6}	2230	3.431×10^{10}

¹ $\lambda_0 = \int_0^\infty \lambda T(\lambda) d\lambda / \int_0^\infty T(\lambda) d\lambda$ where $T(\lambda)$ is filter transmission.

²Effective width, such that $F_\lambda(\lambda_0)\Delta\lambda = \int_0^\infty F_\lambda(\lambda)T(\lambda)d\lambda$.

³Using the STIS calibration spectrum of Vega from Bohlin (2007), which has an uncertainty of 1.5%.

$$\begin{aligned}
r'(\text{Vega-mag}) &= r'(\text{AB-mag}) - (0.130 \pm 0.016) \\
i'(\text{Vega-mag}) &= i'(\text{AB-mag}) - (0.356 \pm 0.016) \\
z'(\text{Vega-mag}) &= z'(\text{AB-mag}) - (0.495 \pm 0.016)
\end{aligned} \tag{2.2}$$

Rodgers et al. (2006) gives transformations from $UBVRcIc$ to $u'g'r'i'z'$ (AB). As a consistency check we can transform the photometry of Vega from Bessel (1990) which yields the alternative transformations

$$\begin{aligned}
r'(\text{Vega-mag}) &= r'(\text{AB-mag}) - (0.141 \pm 0.034) \\
i'(\text{Vega-mag}) &= i'(\text{AB-mag}) - (0.349 \pm 0.039) \\
z'(\text{Vega-mag}) &= z'(\text{AB-mag}) - (0.499 \pm 0.046).
\end{aligned}$$

Comparing results, we see that all three bandpasses agree within the 1σ errors.

2.4.2 Photometry in Y_S

The observations reported here were taken during commissioning, and we have not yet fully calibrated the Y_S filter. This filter has a central wavelength very close to the 99 filter of the 13-color photometric system (Johnson and Mitchell, 1975). For stars with published photometry in this system we can just use the 99 magnitude as the Y_S magnitude. For other stars, one approach is to use the Stellar Spectral Flux Library of Pickles (1998), the calibrated Vega spectrum from Bohlin (2007), and our 0 AM filter profile to calculate $V - Y_S$ colors for main sequence (M.S.) stars, setting $Y_S = 0.03$ for Vega. To assess the accuracy of this technique, we applied it to the synthetic VRI (Johnson) profiles of Bessell (1990) and compared the results to the intrinsic M.S. colors compiled by Ducati et al. (2001). Based on these calculations we estimate the uncertainty of $V - Y_S$ colors determined in this fashion as $\sigma_{V-Y_S} = 0.1$ mag.

2.4.3 The Impact of Water Vapor on z' and Y_S

Both our z' and Y_S filters are affected by telluric water vapor. Using the ATRAN models (Lord, 1992) for Cerro Pachon provided by Gemini Observatory² we assessed the impact changes in both AM and precipitable water vapor (PWV) have on these filters. For $AM \geq 1$ both the mean wavelength λ_0 and the mean total transmission of the filter changes. In the z' filter mean transmission changes by $\pm 2\%$ over the ranges $1.0 \leq AM \leq 1.5$ and $2.3 \leq PWV \leq 10.0$ mm. In the same range Y_S transmission changes by $\pm 3\%$. This change in transmission has little impact on differential photometry so long as PWV does not change between measurements, and the overall effect of extinction changes due to airmass can be removed.

AM has almost no effect on λ_0 but changes in PWV do change it by 2 to 4 nm. This is relatively small and since we have no contemporaneous PWV measurements for the observations reported here we neglect this effect.

²<http://www.gemini.edu/?q=node/10789>

2.4.4 Exposure Time and Gain scalings

It is useful to convert from ADU to electrons per second when comparing images taken with different exposure times. To convert we calculate the scaling factor (SF) to multiply each pixel by. The formula is:

$$SF = \frac{GAIN}{EXPTIME} 10^{ND} \quad (2.3)$$

Where:

SF = scale factor which converts ADU to electrons/seconds

GAIN = the gain factor, in electrons/ADU. This depends on the gain setting (LOW, MLOW, MHIGH, HIGH) which is given in the fits header as V47GAIN, and the pixel rate which is V47PIXRT, and very weakly on the window size and binning. See Table 2.1 for the measured gains.

EXPTIME = the exposure time of the image, in VisAO fits headers it is given by the standard EXPTIME keyword.

ND = value of the neutral density filter if used, 0 otherwise.

2.5 Performance Simulations

To predict the performance of the Magellan AO system we made use of the Code for Adaptive Optics System (CAOS) package (Carbillet et al., 2005). This IDL based “problem solving environment” provides good off-the-shelf functionality and flexibility, and has been used to simulate the LBT AO system (Carbillet et al., 2003).

Our atmosphere model is derived from the GMT site survey of LCO (Thomas-Osip et al., 2008). We use 6 turbulent layers with C_n^2 and wind speed and direction as determined by the survey. We also made use of recent work establishing $L_0 = 25m$ at LCO (Floyd et al., 2010). Based on these data we use von Karman turbulence with $r_0 = 14cm$ as our performance baseline, which corresponds to the $\sim 75^{th}$ percentile at LCO (Thomas-Osip et al., 2008).

The CAOS calibration procedure allows us to calculate interaction matrices for various pyramid sensor configurations. A typical simulation for a bright ($R \sim 7$ mag)

Table 2.3 The predicted VisAO error budget for a spectral type G2 R=7 magnitude guide star. Assumptions include $r_0 = 14cm$ at 550nm, science wavelength of $0.7\mu m$, and the LCO site atmosphere layer model. Estimates using standard AO thumb-rules agree very well with our simulation results. For this analysis the assumed performance of our VisAO tip-tilt loop is 5mas r.m.s. At 10mas r.m.s. tip-tilt control our long exposure Strehl would degrade to 0.2.

VisAO $0.7\mu m$ Error Budget for an R=7 G2 Guide Star			
Error Term	Est. (nm)	Sim. (nm)	notes
Fitting	77.2	...	Estimates from standard thumb-rules Roddier, F. (1999). These are added in quadrature for loop total.
Servo	47.4	...	
Recon.	47.0	...	
Loop Total	102.1	102.4	CAOS simulations as described in the text.
Static.	30	30	Based on LBT design specifications.
Non-Com. Path	30	30	Based on 4D interferometer measurements.
Resid. T/T	52.6	52.6	For 5 mas residual. (Sandler et al., 1994)
Total	122.4	122.7	Sum in quadrature.
$0.7\mu m$ Strehl	0.3	0.3	Using extended Marechal approximation.

guide star uses 392 modes with 1 khz sampling, a gain of 0.4, pyramid modulation of $2\lambda/D$, and pyramid sensor CCD39 parameters based on the manufacturer specification. We simulate with 1ms time steps and apply a 2ms delay to each update to account for WFS readout, calculations, and mirror motion. For each setup (guide star magnitude, etc.) we allow 100ms for loop closing, and then run the simulation for 2 seconds of loop time. At each time step we save a simulated science image at various wavelengths. These images are stored with no sources of noise and we use a 1nm wide bandpass, which allows us to make a very accurate Strehl measurement. The Strehl ratio at each point is measured by comparison to a perfect Airy pattern for a 6.5m telescope with a 29% central obscuration. We typically use the mean short exposure Strehl from these time series, since this value does not include tip-tilt (which we add in quadrature from our error budget).

Table 2.3 shows the error budget for the Magellan AO system at $0.7\mu m$ band on a spectral type G2 R=7 magnitude guide star. We compare estimates for the fitting, servo, and reconstruction errors with our CAOS simulation results and find good

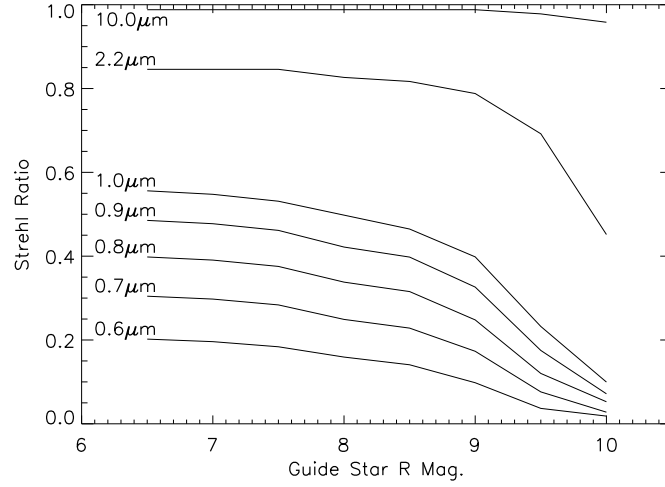


Figure 2.13 Simulated performance of the Magellan AO system vs. guide star R magnitude for 75th percentile seeing at LCO. Based on CAOS simulations as described in the text, each curve includes the error terms listed Table 2.3. Since the primary focus of this dissertation is performance at visible wavelengths, we do not present results for fainter guide stars where VisAO will not perform as well. We expected VisAO to consistently provide Strehl ratios > 0.2 for bright guide stars, and usable correction out to at least $R=9.5$.

agreement. The error calculated from simulations is based on the mean short exposure Strehl ratio. We then add (in quadrature) the unsimulated errors from static mirror aberrations, non-common path aberrations, and finally the long exposure degradation of Strehl due to tip-tilt. In Figure 2.13 we show our simulation-based performance predictions vs. guide star magnitude at various wavelengths. These curves are calculated in similar fashion to the Strehl in Table 2.3, with appropriate differences for wavelength.

2.6 Tower Tests

MagAO was integrated and tested in Arcetri, Italy, between March 2011 and March 2012. This period culminated with a successful pre-ship review (PSR) by an external panel in late February 2012. Here we provide a brief overview of our results from this

testing and describe our attempts to validate the simulated atmospheric turbulence and how we correct our results to produce estimates of on-sky performance.

2.6.1 Seeing Validation

In the test tower, atmospheric turbulence is generated using the ASM itself. A pre-calculated phase screen is applied to the mirror in parallel to the AO corrections. A full description of how turbulence is generated using the ASM was provided by Esposito et al. (2010). To provide a baseline for evaluating performance we took data with AO off but the phase screen propagating across the ASM, that is we took simulated seeing limited data. We then used this data to test whether the seeing produced by the ASM matches expectations from theory. We typically used a phase screen generated to have a seeing limited full-width at half-maximum ($FWHM$) of $0.8''$ at $0.55\mu\text{m}$, or $r_0 = 0.14\text{m}$. For SDSS i', with central wavelength $\lambda_0 = 0.765\mu\text{m}$, we have $r_0(0.765\mu\text{m}) = 0.21\text{m}$. So in the SDSS i' bandpass, assuming Kolmogorov turbulence, we expect the seeing limited PSF to have $FWHM = 0.75''$.

We must also consider that simulated turbulence outer scale was set to $L_0 = 40\text{m}$ (which does not depend on wavelength). Assuming von Karmen statistics on a large aperture, this causes a reduction in $FWHM$ by a factor 0.8159 at $0.765\mu\text{m}$ (Tokovinin, 2002). So our expected $FWHM(0.765\mu\text{m}) = 0.611''$. In Figure 2.14 we show a cut through the seeing limited PSF generated by this phase screen, recorded at SDSS i' with the CCD 47. The best Moffat profile fit to the seeing limited data is $FWHM = 0.617''$, assuming a plate scale of $0.0080''$, corresponding to f/52.6. Of note, the Moffat index of the fit was $\beta = 3.9$. It has been reported that $\beta = 4$ provides a good match to a telescope seeing limited PSF using on-sky data (Racine, 1996). We conclude that the seeing generated in the test tower using the ASM does a very good job of producing the expected image at the CCD 47.

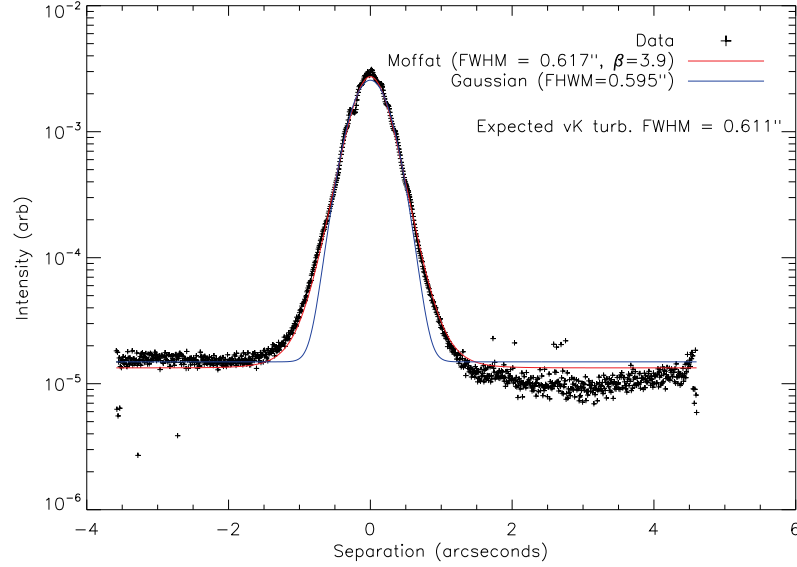


Figure 2.14 A cut through a seeing limited image, and a cut through the best fit Moffat profile. We also show the best fit Gaussian for comparison. The fits were conducted in two dimensions.

2.6.2 Fitting Error

The MagAO ASM influence functions were measured in the Arcetri test tower using an interferometer. The best fit projection of these into a Karhunen-Loeve (KL) basis set was then computed. As is done at the LBT, these KL modes are used during on-sky closed-loop operations at Magellan. To determine the fitting error of our modal basis, 500 independent Kolmogorov phase screens were generated and fit with progressive numbers of our KL modes ³. The residuals for each number of compensated modes were computed, and these points were fit with a function of the standard form:

$$\sigma^2 = A(j_{max})^B (D/r_0)^{(5/3)} \quad (2.4)$$

³these calculations were carried out by Fernando Quiros-Pacheco at Arcetri

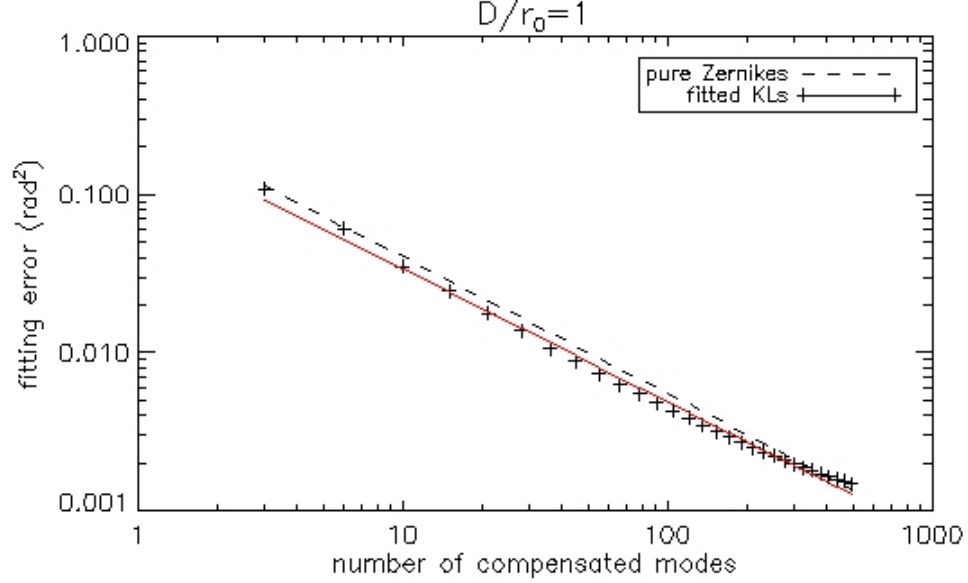


Figure 2.15 Fitting error of the MagAO ASM after correcting j_{max} modes. We show the measured residuals after fitting with the KL modes that will be used on-sky, the residuals expected using Zernike polynomials, and the best fit of Equation 2.4 to the KL mode residuals in red. Note that our KL modes become less efficient than Zernikes after about mode 400.

with

$$A = 0.232555$$

$$B = -0.840466$$

A comparison of this function with the fitting error expected from a pure Zernike polynomial basis is shown in Figure 2.15. The MagAO basis is less efficient than Zernikes for modes greater than number 400, a different result than obtained for the LBT ASMs with an identical procedure. We speculate that this is due in part to the asymmetry caused by the machined slot at the outer edge of the shell.

As discussed above, in the test tower atmospheric turbulence is simulated using the ASM itself, so the phase screen contains only a limited number of spatial frequencies corresponding to the maximum degrees of freedom of the mirror. In the case of MagAO this means that only the first 495 modes of turbulence are simu-

Table 2.4 Tower Test fitting error corrections for 0.8'' seeing. Any S measurement made in the tower can be multiplied by the appropriate S_{corr} to determine an estimate of the on-sky S .

Filter	$\lambda(\mu\text{m})$	$\sigma^2(\text{rad}^2)$	S_{corr}
r'	0.626	0.60	0.55
i'	0.767	0.40	0.67
z'	0.910	0.28	0.75
Y_S	0.984	0.24	0.78

lated, so we must correct our laboratory results for the wavefront variance caused by modes $496 - \infty$ which will be present on-sky. We can use Equation 2.4 to estimate the correction factors to apply to our results. Table 2.4 lists the correction factors for our standard VisAO filters.

2.6.3 Tower Test Results

Over the course of the Arcetri tower testing we took data in many different system configurations, including different magnitude guide stars and different VisAO filter selections. A typical experiment involved taking measurements without simulated turbulence to capture the small amount of turbulence present in the test tower tube, due to internal convection and tip/tilt from flexure caused by the outside wind. We then took an identically configured data set with the ASM simulating turbulence as described above. Finally we nearly always took seeing limited data. Figure 2.16 shows an example of results from such an experiment conducted on a bright star with the loop operating at 800Hz. There we compare the three measurements to a theoretical Airy pattern. Figure 2.17 compares the same experiment to the simulation based performance predictions made above, and also shows the magnitude of the fitting error correction which we apply to form an on-sky performance prediction. See Close et al. (2012a) for additional tower test results .

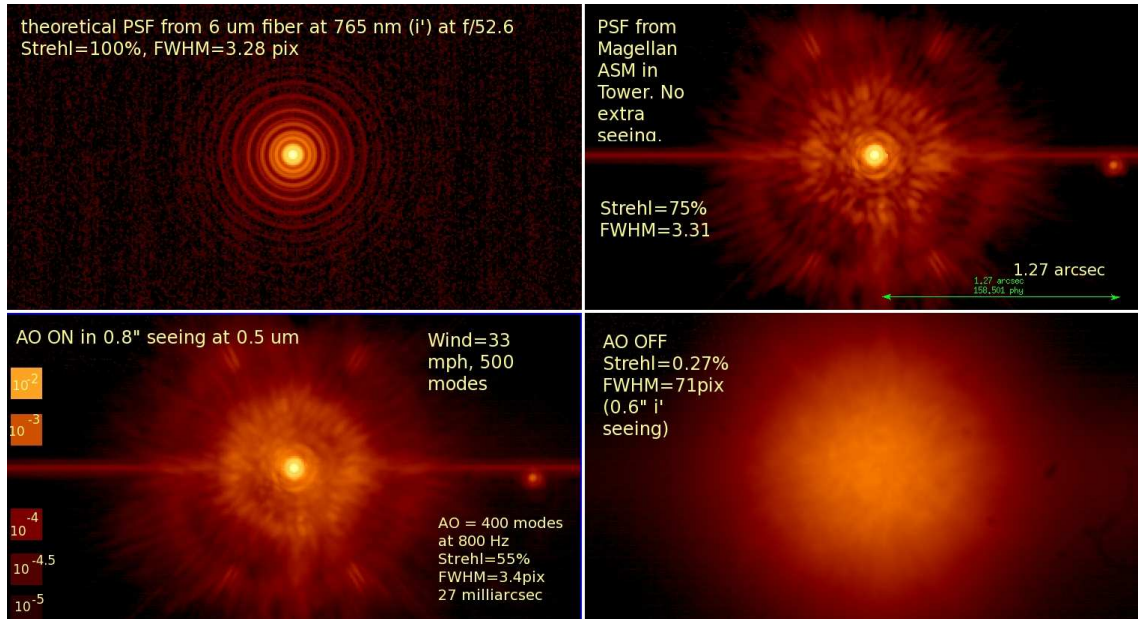


Figure 2.16 Example tower test results. In this case the system was correcting 400 modes at 800Hz, and data were taken in the SDSS i' bandpass. At upper left is a theoretical Airy pattern. At upper right is the MagAO PSF with no simulated turbulence applied, so that the system was correcting only the small amount of turbulence present in the test tower tube. At lower left is the PSF with 0.8" simulated turbulence applied. At lower right is the result with the AO correction off, showing the seeing limited PSF in the same simulated atmosphere. Note that S values quoted in the figure do not have a fitting error correction applied.

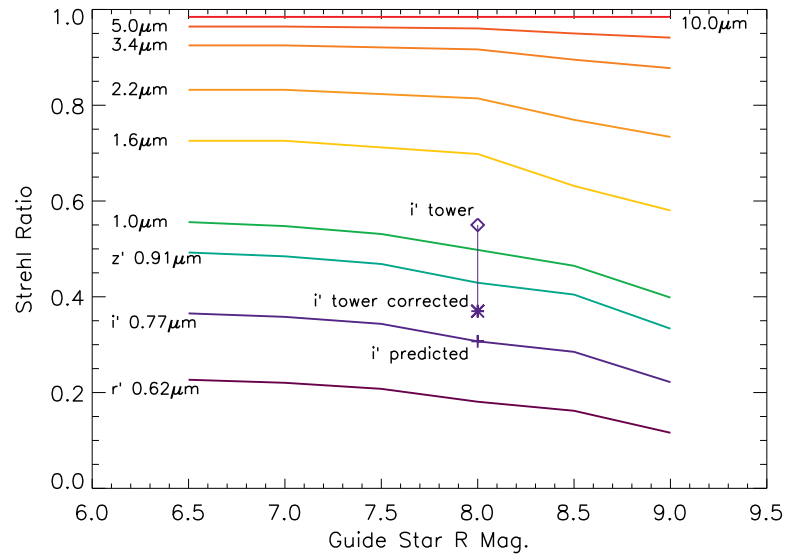


Figure 2.17 A comparison of predicted performance and our SDSS i' test tower results. The solid curves are the same as in Figure 2.13. The data points connected by a vertical line show the raw measured SDSS i' S from our results shown in Figure 2.16 and the value obtained after applying the correction for fitting error from Table 2.4. We also highlight the predicted value for an 8th magnitude guide star. This plot shows that MagAO is performing as expected, if not a bit better.

2.7 On-sky Results

MagAO saw first light in December, 2012, and completed commissioning in May, 2013. Here we present some of the on-sky calibration and characterization, and close with a brief survey of some of images obtained with VisAO so far.

2.7.1 Astrometric Calibration

To calibrate the VisAO platescale and rotator orientation, we used the Orion Trapezium cluster stars. The primary stars used for VisAO calibration were θ^1 Ori $B1$ and $B2$. These stars were observed repeatedly throughout the commissioning run, and with a separation of $\sim 0.94''$ $B2$ is well within the isoplanatic path when guiding on $B1$. This separation is also convenient for dithering the stars around the chip to test for distortions. Other asterisms in Trapezium are challenging to work with due to anisoplanatism and the FOV of VisAO. The drawback to using these two stars is that they have been shown to be in orbit around each other, albeit slowly (see Close, et al., (2013, submitted)).

To avoid uncertainty from the orbits of $B1$ and $B2$, we boot-strapped their current astrometry using the wider FOV Clio camera. With Clio we were able to first use combinations of Trapezium stars to measure the platescale and orientation of Clio. This was done by distortion correcting Clio to match the astrometry given in Close et al. (2012c) which used LBTAO/Pisces. We then measured the separation and position angle of $B1 - B2$. The results are shown in Table 2.5. We track the contribution of errors in the Clio astrometry from measurement and from LBTAO/Pisces separately so that we can compare results between the cameras as well as with other measurements.

The θ^1 Ori $B1 - B2 - B3 - B4$ “mini-cluster” was dithered around the CCD-47 using the gimbal in the Y_S filter on 2012 Dec 3 UT with the rotator tracking. Each dither position was reduced separately: frames were selected for good correction, registered, and median combined. The images were not de-rotated. At each dither position, the separation and PA of $B1 - B2$ were measured using the *starfinder.pro*

Table 2.5 Clio astrometry of θ^1 Ori $B1$ and $B2$ (Dec 2012)

	Value	Measurement Uncertainty	Astrometric Uncertainty	Total Uncertainty
sep	0.9400"	0.0011"	0.0023"	0.0025"
PA	254.87 deg	0.10 deg	0.3 deg	0.32 deg

idl program (Diolaiti et al., 2000), making use of $B1$ itself as the PSF. We also measured the position of the optical beam splitter ghost of $B1$, as well as its relative flux. Since the beam splitter is very near a pupil, we expect this ghost to be very stable.

The relative x and y positions of $B1 - B2$ were found to depend on position on the chip, consistent with a focal plane tilt. The primary symptom of this was a ~ 0.25 degree scatter in PA measurements, which was well outside the formal errors from the PSF fitting astrometry. We found that fitting a plane through the Δx and Δy measurements of $B1 - B2$ reduced this scatter to be consistent with the formal errors. The equations to correct measurements to the center of the chip (512x512) are

$$\begin{aligned}\delta x &= -0.0003892(x - 512) + 0.0008432(y - 512) \\ \delta y &= -0.0002576(x - 512) - 0.0024045(y - 512).\end{aligned}\tag{2.5}$$

These corrections are then added to the measured Δx and Δy to get the center-of-chip value. After applying these corrections we measured the center-of-chip platescale and the value of $NORTH_{VisAO}$, which is used to de-rotate images by

$$DEROT_{VisAO} = ROTOFF + 90 + NORTH_{VisAO}\tag{2.6}$$

These results are presented in Table 2.6.

Table 2.6 VisAO Y_S platescale and rotator calibration. Measurement uncertainty includes both Clio and VisAO scatter. Astrometric uncertainty is propagated from LBTAO/Pisces

	Value	Measurement Uncertainty	Astrometric Uncertainty	Total Uncertainty
Platescale	0.007910"	0.000009"	0.000019"	0.000021"
$NORTH_{VisAO}$	-0.59 deg	0.10 deg	0.3 deg	0.32 deg

2.7.2 Beamsplitter Ghost Calibration

The WFS beamsplitters control how much light is sent to the pyramid sensor vs. the VisAO camera. They are selectable in a filter wheel, and the choice of beamsplitter is an integral part of AO system setup based on what the science goal of the observation is. Each beamsplitter has an optical ghost, which is slightly out of focus. We use these ghosts for both registration and photometry when the central star is either saturated, or under the coronagraph. As noted above, we extracted both the relative flux and position of the ghost in the 50/50 beamsplitter as part of our analysis of the Trapezium astrometric field. The relative flux of the ghost did not depend on position, but the position of the ghost does. In Table 2.7 we present the ghost parameters for the 50/50 beamsplitter, in the Y_S filter with filter wheel 3 open, when the guide star is at 512x512 (counting from 0).

In our second commissioning run, we dithered HIP 56004 around the detector in the same setup as above, but this time the coronagraph was in the beam. The results are also shown in Table 2.7. It appears that the astrometric properties of the ghost change slightly when the coronagraph is in. This change is on the order of 0.5 pixels, which can be significant for techniques based on angular differential imaging (ADI).

The position of the ghost depends on the position of the guide star. As before we fit a plane to the $B1$ -ghost data and found the following correction for F/W 3 open:

Table 2.7 Photometric and astrometric parameters of the 50/50 beamsplitter ghost in the Y_s filter.

	Value	Notes
F/W 3 Open		
Rel. Flux.	0.00718 ± 0.00013	independent of position
Δx	160.369	at 512, 512
Δy	-9.040	at 512, 512
F/W 3 Coronagraph		
Δx	159.745	at 512, 512
Δy	-9.483	at 512, 512

$$\delta x_g = -0.0004939(x - 512) - 0.0003123(y - 512) \quad (2.7)$$

$$\delta y_g = -0.0005014(x - 512) - 0.0039757(y - 512)$$

and when the coronagraph is in the correction equations are:

$$\delta x_g = -0.0004812(x - 512) - 0.0007174(y - 512) \quad (2.8)$$

$$\delta y_g = -0.0003305(x - 512) - 0.0037534(y - 512)$$

2.7.3 Y_s Strehl Ratio

Of particular interest to this dissertation, we observed the known exoplanet host β Pictoris several times during the first MagAO commissioning run. Our best correction was obtained on 2012 Dec 4, when we were observing in Y_s using the VisAO coronagraph. The high contrast data reduction and analysis of our detection of β Pic b are discussed later. Here we present our off-coronagraph PSF and photometric calibration image.

To avoid saturation, the off-coronagraph data was taken in the 2500 kHz pixel-rate, full frame mode, resulting in 0.283 individual exposure times. The camera was in the LOW gain setting, allowing access to the full well-depth. Even with

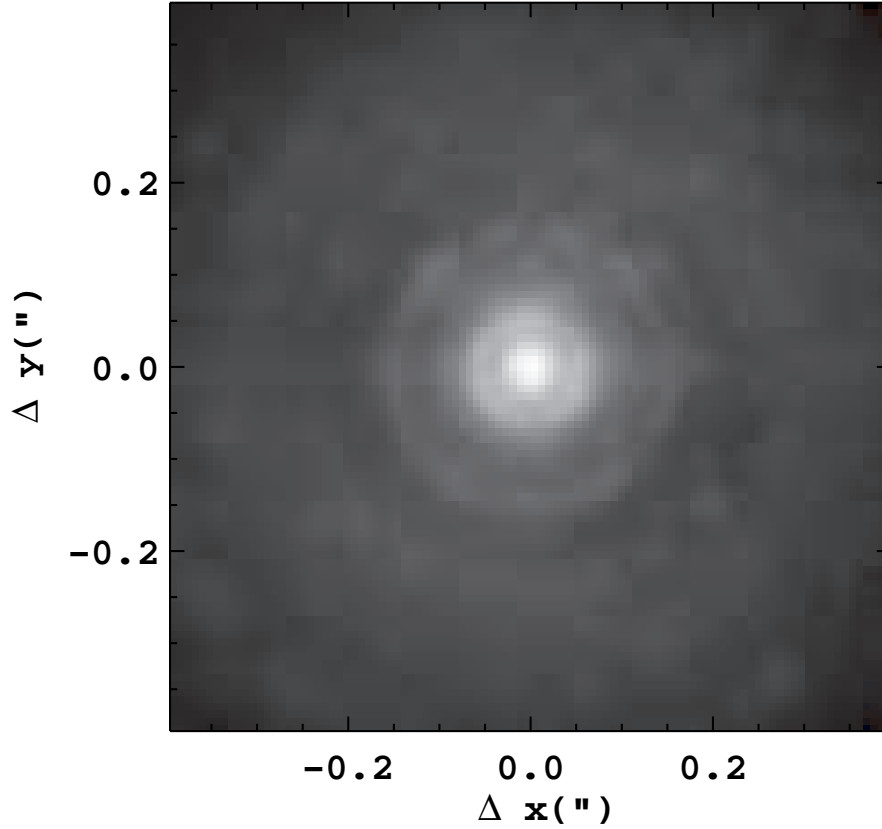


Figure 2.18 Image of β Pictoris, a $Y_s = 3.5$ mag A5V star, made by shifting and adding 0.283 sec exposures. This is a log stretch. This image has a raw Strehl of $\sim 32\%$, and an optical Strehl (corrected for PRF) of $\sim 40\%$. Instantaneous Strehl ratio was $> 50\%$, but there was a ~ 0.9 pixel RMS jitter.

these settings, we saturated the peak pixel in roughly a third of the exposures. To compensate, and avoid any possible non-linearity, we selected only frames with peak pixel between 8000 and 9000 ADU. This cut out roughly half the frames, and has the effect that we are using data between approximately the 75th and 25th percentiles — so not the very best. We also applied a WFE cut at 130 nm RMS phase, using the VisAO real time telemetry stream. The 491 selected frames were then registered, and median combined. The result is shown in Figure 2.18.

The resultant PSF core has a FWHM of 4.73 pixels (37 mas). The diffraction

limit at Y_S is 3.87 pixels (31 mas). We expect some broadening due to the pixel response function (PRF) - mainly from charge diffusion. When a photon is detected by a CCD pixel, the resulting photo-electron can diffuse into a neighboring pixel. This causes a blurring effect, which has been well documented in the HST ACS and WFPC cameras. See Krist (2003), Anderson and King (2006), and the ACS handbook for more on charge diffusion PRFs.

We measured the PRF of the CCD-47 in the lab, by switching between it and a well over-sampled CCD. The effects of charge diffusion lessen as more pixels are placed across the PSF, so this allowed us to compare the true optical PSF to the PSF measured on the CCD-47. Using a blurring kernel developed from these measurements, we find that a perfectly diffraction limited PSF should measure $FWHM = 4.18$ pixels. So we have ~ 0.5 pixels of broadening. This implies an RMS tip-tilt residual error of 0.94 pixels (see Equation (A.4)).

Strehl from WFE and PRF: We can use the real-time reconstructed WFE recorded in the VisAO fits headers to form an initial estimate of SR in the PSF shown in Figure 2.18. The mean WFE for the images included in the final SAA was 123.37 nm RMS phase. This gives us

$$S_i = 0.54$$

for the instantaneous Strehl ratio. Using Equation A.5 and our estimate of jitter from $FWHM$ broadening we estimate

$$S_{tt} = 0.75$$

so

$$S_{wfe} = S_i S_{tt} = 0.40$$

The PRF also lowers measured Strehl ratio. After broadening a theoretical Airy pattern by the PRF kernel, we find that

$$S_{prf} = 0.80.$$

This means the expected focal plane Strehl ratio due to WFE reconstruction and PRF is

$$S_{rec} = S_i S_{tt} S_{prf} = 0.32.$$

Core-to-Halo Strehl: Another way to measure Strehl is to use the ratio of the flux enclosed in the core (to the first Airy minimum) to the total flux. This will tend to be robust against effects like the PRF and T/T broadening. For MagAO, with a 29% central obscuration, the first Airy minimum occurs at $1.12\lambda/D$ and encloses 74.7% of the total flux at $S = 1$ (see Section A.2 in the Appendix). We measured the total flux in the PSF with IRAF imexam, with a photometric aperture of 239 pixels, a sky radius of 240 pixels, and a sky width of 5 pixels. This should enclose 99.6% of the flux. We then changed the aperture radius to 4.42 pixels and measured the core flux. We estimate the core-to-halo S from these measurements by

$$S_{c/h} = \frac{\mathcal{E}(1.12\lambda/D)}{0.747\mathcal{E}(\infty)} = 0.34$$

where $\mathcal{E}(r)$ denotes encircled flux at radius r .

Peak-to-Halo Strehl: We can use the total flux in the image to calculate the expected peak height for $S = 1$

$$I_o = \frac{\pi P_{tot}(1 - \epsilon^2)}{4} \frac{0.249}{\lambda^2}.$$

See Appendix A for the derivation of this expression. We fit a Gaussian to the PSF, and comparing the resultant peak height I_{pk} we find

$$S_{p/h} = \frac{I_{pk}}{I_o} S_{prf} = 0.30$$

is the peak-to-halo estimate of S .

Theory Strehl: Finally, we can compare the PSF shown in Figure 2.18 to a theoretical Airy pattern. We first broaden the theoretical PSF using the PRF kernel. Then normalizing the theory PSF with the same 239 pixel photometric aperture, we find

$$S_{theory} = 0.30$$

Table 2.8 Strehl measurements at Y_S ($0.984\mu\text{m}$)

	Strehl	Notes
S_{rec}	0.32	reconstructed WFE
$S_{c/h}$	0.34	core-to-halo
$S_{p/h}$	0.30	peak-to-halo
S_{theory}	0.30	PRF broadened Airy
Average	0.32 ± 0.02	
S_{wfe}	0.40	before PRF

Table 2.9 Observation log for LHS 14, observed on 2012 Dec 3 UT

Filter	Gain	Exp Time (sec)	No. Exp.	Tot. Exp (sec)
r'	Med. High	0.283	199	56.3
i'	Med. High	0.283	271	76.7
z'	Med. High	0.283	263	74.4
Y_S	Med. High	0.283	225	63.7

We summarize our Strehl measurements in Table 2.8. These four different ways of estimating Strehl ratio agree well. An interesting quantity is the optical Strehl, that is the calculated S_{wfe} . We can estimate this by dividing by S_{prf} . This quantity is noteworthy as PRF is easily removed by deconvolution, allowing us to recover the optical resolution achieved by the AO system.

2.7.4 Throughput

On 2012 Dec 3 UT we observed LHS 14, an M2.5V star with published $u'g'r'i'z'$ photometry, in all four VisAO broad bandpasses. These observations, summarized in Table 2.9, were used to measure system throughput. The 50/50 WFS beamsplitter was selected. Data were taken in 5 point dither patterns to minimize corruption from dust spots, and images were manually selected, dark subtracted, registered, and median combined.

Next the IRAF task daophot was used to conduct aperture photometry on each

Table 2.10 VisAO total throughput measurements on LHS 14 in 50/50 beam-splitter.

Filter	Mag (AB)	Mag (Vega)	0AM Flux (phot/sec)	Meas. Flux (phot/sec)	Expected Throughput	Measured Throughput
r'	9.481	9.356	5.248×10^7	7.990×10^6	15.9%	15.2%
i'	8.547	8.194	3.430×10^7	1.177×10^7	17.0%	34.3%
z'	8.104	7.611	3.806×10^7	8.348×10^7	12.5%	21.3%
Y_S	—	7.084	4.401×10^7	1.201×10^6	3.6%	2.7%

final image. The photometry was converted from ADU to e^-/s using Equation (2.3). We then compared these results to the photometry of LHS 14 from Smith et al. (2002), converting from AB to Vega magnitudes using Equation (2.3) and making use of the parameters presented in Table 2.2. The Y_S magnitude of LHS 14 was estimated using the Pickles spectral library as described in Section 2.4.2 above. The resulting throughput measurements are shown in Table 2.10.

We also calculated our expected throughput given the beamsplitter, atmosphere, and the filter curves calculated as described above. These numbers are also presented in Table 2.10. Of note, our measured throughput is roughly a factor of two higher than expected in i' and z' . Two effects may account for this. In i' especially, our filters are redder than the standard SDSS bandpasses, due to our IR-coated CCD QE being higher in this region. It is also possible that our CCD QE is somewhat better in this region than assumed in our filter profiles, as the QE is based only on a catalog plot and is not a measurement of the actual device. The lower than expected Y_S throughput is conversely possibly explained by over-estimating the tail of Silicon QE for $\lambda > 1\mu\text{m}$, and also possibly on a poor quality catalog transmission curve for the filter itself.

2.7.5 VisAO Images

Here we very briefly present some other on-sky results from the MagAO commissioning periods. We do no astrophysical analysis here, in essence we are just offering pretty pictures to establish that MagAO and VisAO are working well on-sky. Later

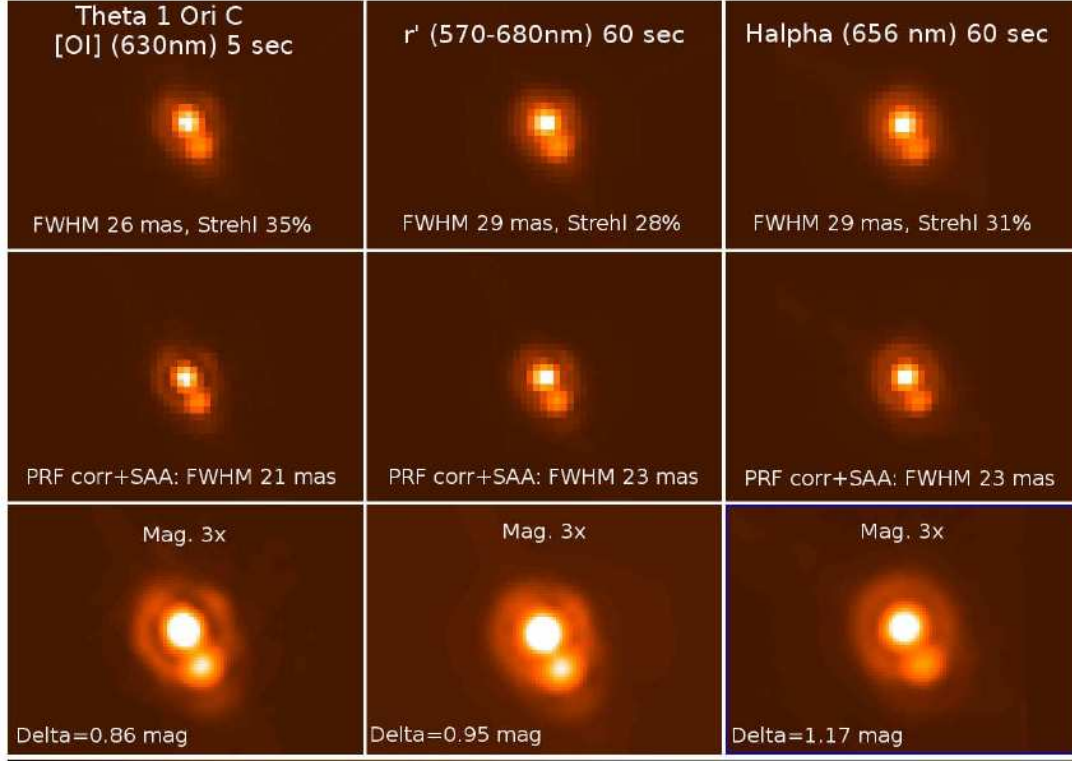


Figure 2.19 Image of the θ^1 Ori C binary, obtained with VisAO in December, 2012. We resolved this binary at only 31 mas separation — the first time this has been done with a filled aperture long exposure. These images, from Close, et al., (2013, submitted), show data taken in [OI] (630 nm), r' (centered at 624 nm), and $H\alpha$, at (656 nm). These reductions are by L. Close. SAA is shift-and-add, and PRF was corrected where indicated by deconvolution.

we will present a detailed analysis of our observations of an exoplanet host star.

As discussed above we used the Trapezium cluster in Orion for astrometric calibrations. We have also been able to do some science with these data. In Figure 2.19 we show the first filled-aperture long exposure images to resolve the θ^1 Ori C binary. We also demonstrate how, using the high speed readout modes and our well characterized PRF we can achieve 21 mas resolution with VisAO. In Figure 2.20 we show a z' image of the θ^1 Ori B “mini-cluster”. These data have been used in Close, et al., (2013, submitted) to demonstrate that the B2-B3 barycenter is orbiting B1, along with B4.

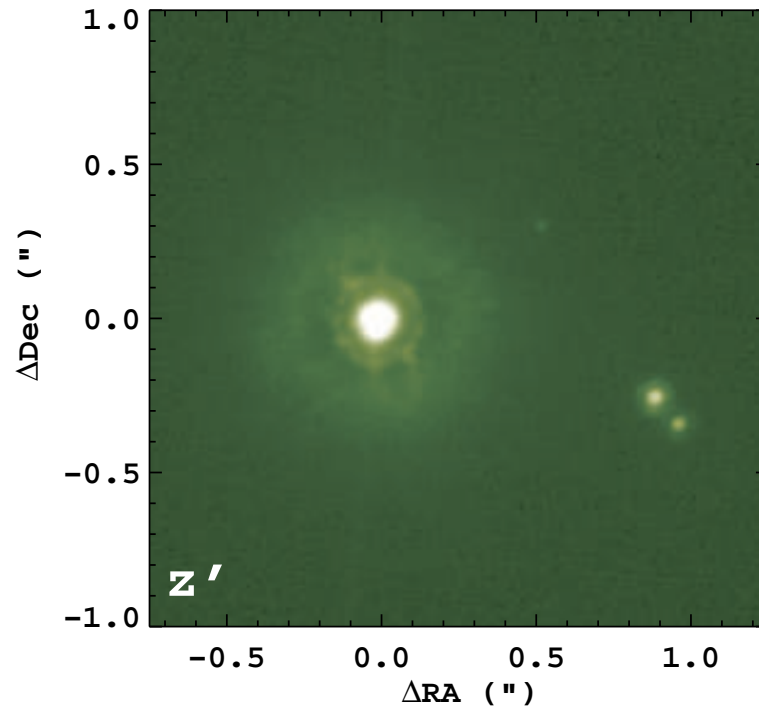


Figure 2.20 An image of the θ^1 Ori B cluster, obtained with VisAO in December, 2012. VisAO observations of this system have shown, for the first time, conclusive evidence for orbital motion of the B2-B3 barycenter around B1. See Close, et al., (2013, submitted) for details.

During the May 2013 commissioning 2 run we observed HR 4796A, a star with a well known circumstellar disk, as part of a program with TJ Rodigas and Alycia Weinberger. Data on this star was taken simultaneously on Clio and VisAO, covering 6 filters in total. We present a very quick reduction of the VisAO data in Figure 2.21.

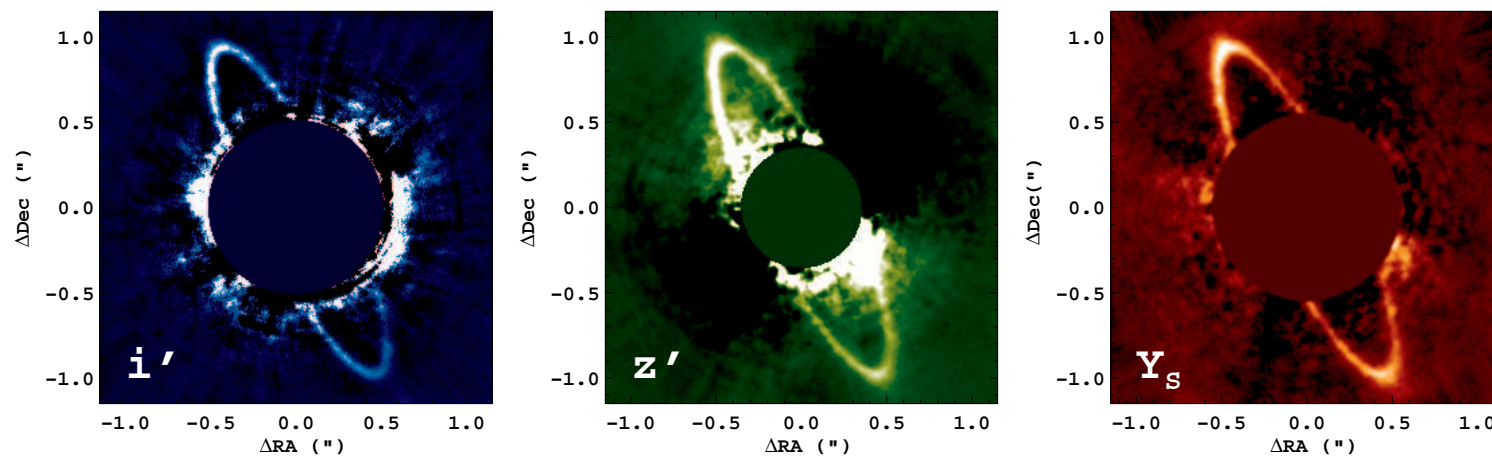


Figure 2.21 Images of HR 4796A and its circumstellar disk, obtained with VisAO in May, 2013. Only basic ADI processing, with a radial profile subtraction step, and high-pass filtering (unsharp mask) was used to produce these images. Look for higher fidelity reductions and analysis of these data in Rodigas, et al., (in prep).

CHAPTER 3

REAL TIME FRAME SELECTION

3.1 Introduction

Lucky imaging is a technique which selects the best images from a series of short exposures, then shifts and adds them to produce a final image with higher spatial resolution than a single long exposure. First proposed by Fried (1978) to counter the effects of atmospheric seeing, it is now in common use at several telescopes (Law et al., 2006). It has also been adapted for use with an AO system, where the correction quality in the visible was¹ typically low, but has short periods of high Strehl (Law et al., 2009, 2008). By selecting images based on Strehl ratio, it has been shown that both resolution (Law et al., 2009) (measured by full width at half maximum (FWHM)) and sensitivity (Gladysz et al., 2008a) (measured by signal-to-noise ratio (S/N)) can be improved.

To be effective, lucky imaging typically uses very short exposures, requiring cameras that operate faster than ~ 10 frames-per-second (fps). For normal astronomical CCDs this imposes a significant readout noise (RON) penalty, as each read will produce a few electrons of noise which then quickly adds to overwhelm faint signals. This has been overcome to great effect using EMCCDs, which offer very low RON - typically $\sim 0.1e^-$ per read (Daigle et al., 2009) - when operated in the photon counting mode. Photon counting EMCCDs have some (small) drawbacks though. If flux is higher than 1 photon/pixel/read the device effectively has its quantum efficiency (QE) lowered by 50% (Mackay et al., 2004). On VisAO, however, we almost always have a very bright guide star in our FOV, and if we had low to moderate Strehl ratios we will have a bright uncorrected halo from the star spread over the detector, making this QE penalty impossible to ignore.

¹prior to the MagAO era

An additional issue we identified with any lucky imaging system is the trade-off between FOV and camera speed. The isoplanatic patch at visible wavelengths at LCO will typically be $\sim 4''$ in radius. To provide good sampling across this FOV we need a 1024×1024 array. When we began the investigations detailed here the fastest EMCCD cameras of this size could only be operated at ~ 10 fps (EMCCDs have since become a little faster, now achieving ~ 30 fps over arrays of this size). As we will show 10 fps is not quite fast enough to fully take advantage of the peaks in Strehl ratio, which are typically shorter than 100ms in our simulations. The common solution with an EMCCD is to window such a device and operate at 50fps or faster, which also helps to mitigate the high-flux penalty. This carries its own drawback in that FOV is cut by 25% in area, which for many observations is itself equivalent to a QE penalty.

A final consideration, and perhaps most important, is that the Magellan VisAO system was largely based on an already designed instrument. Since the VisAO CCD47 is used as an acquisition camera integral to AO system operation, changing detectors was judged too risky to overall performance. Changing detectors would significantly increase the cost of this system, but we still desire to take advantage of frame selection.

Given the bright guide star specific and FOV vs. speed drawbacks of EMCCD based lucky imaging, and the pre-existing system designs, we developed a new imaging concept which we call real time frame selection (RTFS). In this mode of operation, we make use of a high speed mechanical shutter and telemetry from the AO system to only expose our CCD47 when Strehl is high. The shutter is both fast and responsive enough to provide the equivalent temporal resolution of a 100fps camera, and can do this over the entire $8.7''$ FOV of our 1024×1024 array. As we will show, this technique can improve resolution (when compared with doing nothing) by nearly 100% of λ/D .

We first present a generic frame selection algorithm, providing a formal definition of frame selection. Then we present a model of S/N in AO imaging, which we use to analyze the costs and benefits of frame selection and to compare different imaging

techniques. Using this model and the output of performance simulations we calculate the performance of an ideal RTFS system both in terms of resolution and sensitivity. We then describe our implementation of RTFS, including characterization of the shutter and development of real time telemetry processing. An area of ongoing development is Strehl prediction, necessary because of a short but unavoidable delay in shutter actuation time. We then show the results of laboratory testing of RTFS, conducted with realistic simulations in the test tower at Arcetri, Italy.

Much of this chapter has appeared in print in (Males et al., 2010) and (Males et al., 2012a). An important caveat to all of this work, taking advantage of perfect hindsight, is that it was almost entirely unnecessary for MagAO and VisAO. The need for Lucky-style imaging was anticipated because, simply put, we just did not think VisAO would work very well. As shown in detail below, any frame selection technique involves a trade-off between sensitivity and resolution. Above a certain image quality (Strehl ratio of, say, about 20%), it is not worth it (at least for point sources). Despite the effort that went into developing this system, I am perfectly happy to never need it on-sky — VisAO works really well!

3.2 A Generic Frame Selection Algorithm

Before we can analyze its benefits, we first state what we mean by frame selection. To do this we will develop a general description of a selection algorithm, leaving specific details for later.

We begin by collecting a stream of raw data at a time t_i , such as AO control loop telemetry or short exposure science image pixels, with n elements

$$\vec{X}(t_i) = \begin{bmatrix} x_1(t_i) \\ \vdots \\ x_n(t_i) \end{bmatrix}.$$

Next the data is converted to a set of m attributes² by some operation F

$$\vec{Y}(t_i) = \begin{bmatrix} y_1(t_i) \\ \vdots \\ y_m(t_i) \end{bmatrix} = F\left(\vec{X}(t_i)\right).$$

For instance, F may include the calculation of slopes given the raw WFS pixels or sub-aperture counts. Finally we use a classifier G to determine whether some image quality metric, say Strehl ratio S , is above some threshold value, say S_T . The classifier uses the previous l samples of the m attributes, possibly with a delay of k time steps (meaning the classifier is also a predictor):

$$G\left(\vec{Y}(t_i), \dots, \vec{Y}(t_{i-l+1})\right) = \begin{cases} 0 & \text{if } S(t_{i+k}) < S_T \\ 1 & \text{if } S(t_{i+k}) \geq S_T. \end{cases} \quad (3.1)$$

The value of G represents the decision whether to include the data at t_{i+k} in the final image. We explicitly allow for prediction since this will be necessary for the real time implementation we discuss later. Finally we note that this formulation does not require that the actual value of the image quality metric (e.g. S) be calculated. This opens the door to using, for instance, machine learning classification techniques without *a priori* knowledge of relationships between the raw data and image quality.

The standard lucky imaging technique can be described using this algorithm. In this case, the data vector $\vec{X}(t_i)$ is made up of the pixel values of a short exposure image at time t_i . The corresponding attribute is just the Strehl Ratio $S(t_i)$, and the operation F is the reduction pipeline which results in the Strehl ratio measurement. The classifier G is a simple comparison between the measured $S(t_i)$ and S_T . In standard Lucky imaging only the current time step is used and no prediction is performed, i.e. $k = 0$ and $l = 1$.

Another implementation of this algorithm is RTFS (discussed above), developed for the Magellan VisAO system, which is used to control a camera shutter in real time. The primary goal of this technique is to minimize the number of detector reads, while gaining the benefits of frame selection.

²It is not necessary that $m = n$

3.2.1 WFS Telemetry Based RTFS

RTFS uses a fast shutter to block moments of bad correction, causing only periods of high S to be recorded by the science camera. This prevents us from using direct measurements of S to trigger the shutter. Above we developed a notation to describe a generic frame selection algorithm, including conventional Lucky imaging. Here we adapt that algorithm and notation to the specific case of using only WFS telemetry to reconstruct S .

We record a slope vector at time t_i with n elements, $\vec{X}(t_i)$. The wavefront is reconstructed by multiplying the slope vector by the reconstructor matrix \mathbf{R} (typically the same one in use in the main AO loop).

$$\vec{A}(t_i) = \mathbf{R}\vec{X}(t_i)$$

where $\vec{A}(t_i)$ is the vector of reconstructed mode amplitudes at time t_i . The orthogonal KL modal basis is normalized such that each mode has unit variance, so we can calculate the wavefront variance by summing the amplitudes in quadrature

$$\sigma^2 = (4 \times 10^9)^2 \vec{A}(t_i) \cdot \vec{A}(t_i)$$

where the factor of 4 accounts for the double pass of the ASM in the test tower, and 10^9 converts to nanometers. We then calculate the reconstructed S using the extended Marechal approximation

$$S_{rec}(t_i) = e^{-\left(\frac{2\pi}{\lambda}\right)^2 \sigma^2}.$$

Next we apply an empirical calibration, using two parameters.

$$S_{cal}(t_i) = aS_{rec}(t_i) + b.$$

See below for further discussion of this calibration step and the interpretation of these parameters.

Finally, we apply a finite impulse response (FIR) low-pass filter of order N to prevent high frequencies from over-driving the shutter.

$$S_{filt}(t_i) = \sum_{k=0}^{k=N} f_k S_{cal}(t_{i-k})$$

where the f_k are the filter coefficients. The design of appropriate digital filters is described below.

We use the reconstructed filtered S to classify each moment as a good or bad according to whether it is above or below a threshold S_T .

$$G = \begin{cases} 0 & \text{if } S_{filt}(t_i) < S_T \\ 1 & \text{if } S_{filt}(t_i) \geq S_T. \end{cases} \quad (3.2)$$

The value of G represents the decision whether to open ($G = 1$) or close ($G = 0$) the shutter.

3.3 The Costs & Benefits of Frame Selection

To assess the benefits of frame selection, we use a simple model of the AO imaging process to calculate S/N and the output of our CAOS simulations to determine the resulting resolution. The following development relies heavily on the work of Racine et al. (1999), and benefits from the work of Law et al. (2009, 2008, 2006) and Gladysz et al. (2008a,b, 2006).

3.3.1 Signal to Noise Ratio and Duty Cycle

The obvious drawback to frame selection is that only a fraction of the telescope time allotted to the observation is used in the final result. This fraction can be thought of as the duty cycle \mathcal{DC} . The act of throwing away the fraction $(1 - \mathcal{DC})$ of the signal can be expected to negatively effect the sensitivity of the observation. It might be true, however, that by keeping only the “good” frames we can overcome this loss in signal by reducing the noise in our final image, and this has been demonstrated on-sky (Gladysz et al., 2008a). In conventional lucky imaging, one always has all the data available from an observation and so has lost no telescope time. In RTFS, however, we will irretrievably lose the time when the shutter is closed and so we must understand the trade-offs with sensitivity for this technique.

To that end, we use a simple model of the S/N in AO imaging, based heavily on that developed by Racine et al. (1999) with only slight modifications,

to determine the net efficiency cost of frame selection. Here we skip most of the derivation and present the results for the limiting cases most likely to be encountered in natural guide star (NGS) AO.

Our S/N model is

$$S/N = \frac{\mathcal{E} f_c \sqrt{\mathcal{DC} t}}{(\mathcal{E} f_c + \bar{S} \bar{\mathcal{P}}(\theta) f_* n + (1 - \bar{S}) H(\theta) f_* n + 0.53 \tau_0 [(1 - \bar{S}) H(\theta) f_* n]^2 + (N_{sky} + N_{det}) n)^{\frac{1}{2}}} \quad (3.3)$$

where f_* and f_c are the flux (photons sec^{-1}) at the telescope of the central star and the companion (located at separation θ); \mathcal{E} is the total flux enclosed in the photometric aperture; \mathcal{DC} is the duty cycle (discussed above); $\bar{\mathcal{P}}$ is the point spread function (PSF), ideally an obscured Airy pattern, averaged by tip-tilt; n is the number of pixels contained in the photometric aperture; $H(\theta)$ is the uncorrected halo flux per pixel; τ_0 is the speckle lifetime (Racine et al., 1999); N_{sky} is the per-pixel flux due to sky background (BG); and N_{det} is the per-pixel flux due to detector noise. The quantity \bar{S} is the mean short exposure Strehl, which we use rather than the long exposure (tip-tilt degraded) Strehl ratio to account for the guide star's halo contribution to the noise as it is what quantifies the relative fraction of flux in the halo. Using the long exposure Strehl here would make the halo noise too large.

3.3.2 Encircled Energy and Aperture Size

The fraction of incident photons contained in a circular aperture assuming a perfect obscured-Airy PSF is

$$\mathcal{E}_{\mathcal{P}}(\rho) = \frac{1}{1 - \epsilon^2} \left((1 + \epsilon^2)(1 - J_0^2(\rho) - J_1^2(\rho)) - 4\epsilon \int_0^\rho \frac{J_1(t) J_1(\epsilon t)}{t} dt \right) \quad (3.4)$$

where ϵ is the telescope central obscuration and ρ is the aperture size (see Appendix A and Mahajan (1986)). Now in an AO corrected image only a fraction \bar{S} of the flux is contained in the diffraction limited component of the PSF, and $(1 - \bar{S})$ is contained in the halo. Following Racine et al. (1999) we adopt the function

$$H(\theta) = \frac{0.488}{W_h^2} \left[1 + \frac{11}{6} \left(\frac{\theta}{W_h} \right)^2 \right]^{-11/6}. \quad (3.5)$$

to describe the uncorrected halo, where W_h is a width parameter.

We can integrate equation (3.5) to calculate the fraction of the incident photons in the halo encircled by an aperture of size ρ :

$$\mathcal{E}_H(\rho) = 1 - \left(1 + \frac{11}{6} \left(\frac{\rho}{W_h} \right)^2 \right)^{-5/6}. \quad (3.6)$$

Racine et al. (1999) argued that the halo contribution to flux could safely be ignored since its S/N would be comparatively low. While this is likely true for high Strehls, we expect to employ frame selection with only low to moderate Strehls. At Magellan, with $\epsilon = 0.29$, the enclosed fraction is $\mathcal{E}_P = 0.747$ for $\rho = 1.12\lambda/D$ (the first Airy minimum). From simulations we derive a value of $W_h = 0.23''$ at $\lambda = 0.7\mu\text{m}$, so $\mathcal{E}_H = 0.018$. For $S = 0.1$ then

$$\frac{(1 - S)\mathcal{E}_H}{S\mathcal{E}_P + (1 - S)\mathcal{E}_H} = 0.18$$

is a non-negligible fraction of the total signal collected by the photometric aperture defined by the first Airy minimum. This means that we must account for the halo component of the PSF of the science object and so cannot assume that \mathcal{E} is simply proportional to S .

In our simulations we have found that the usual model just employed of a diffraction limited core on top of a partially corrected halo has limited ability to describe the location of the photons in our image. At low Strehls, imperfect correction causes the core to broaden, and the PSF tends to elongate in the direction of the prevailing winds. Even at higher Strehls residual telescope jitter can have the same effect. We also consider it desirable to avoid relying on an analytic model for the PSF, especially the halo component as this is likely to depend greatly on atmospheric seeing and guide star brightness. In practice, we find that using elliptical apertures fit to contours of constant flux will give consistent results from our simulated data.

Now let $\mathcal{A} = \mathcal{E}/n$ be the average fractional photon flux per pixel in the ellipse described by $\vec{x} = (a, b, \phi)$, which are the semi-major axis, semi-minor axis, and orientation of the ellipse. If we then define \vec{x}^* as the aperture which maximizes

S/N, then the optimum faint companion detection S/N is

$$S/N = \frac{\mathcal{A}(\vec{x}^*)f_c\sqrt{\mathcal{DC}tn^*}}{(\mathcal{A}(\vec{x}^*)f_c + \bar{S}\bar{\mathcal{P}}(\theta;\epsilon)f_* + (1 - \bar{S})H(\theta)f_* + 0.53\tau_0[(1 - \bar{S})H(\theta)f_*]^2n^* + N_{sky} + N_{det})^{\frac{1}{2}}}. \quad (3.7)$$

where $n^* = \pi a^* b^*$. We have gone to this effort because we find that using other algorithms to estimate S/N, such as peak-pixel or a fixed aperture size, tends to incorrectly analyze the benefits of frame selection in our simulations in various cases - especially when the contribution of the companion's halo component is ignored.

3.3.3 Effective Duty Cycle

Now we can solve Equation (3.7) for the time t it takes to reach a desired S/N:

$$t = \frac{(S/N)^2}{\mathcal{A}^2(\vec{x}^*)n^*f_c^2\mathcal{DC}}(\mathcal{A}(\vec{x}^*)f_c + \bar{S}\bar{\mathcal{P}}(\theta;\epsilon)f_* + (1 - \bar{S})H(\theta)f_* + 0.53\tau_0[(1 - \bar{S})H(\theta)f_*]^2n^* + N_{sky} + N_{det}).$$

We can then compare two data taking techniques, e.g. frame selection to simple integration. In order for a technique to provide a S/N advantage then the effective duty cycle $\mathcal{DC}_{\text{eff}}$ must satisfy the inequality

$$\mathcal{DC}_{\text{eff}} = \frac{t_o}{t_1} \geq 1 \quad (3.8)$$

where t_o is the the time needed to reach a S/N goal with simple integration (i.e. doing nothing), and t_1 is the time needed with a particular frame selection technique.

The effective duty cycle concept allows us to compare the trade-offs between resolution, encircled energy, and efficiency, and then decide the optimal imaging technique for our AO system and science goals. Now we consider the limiting cases of equations (3.7) and (3.8) which we expect to routinely encounter with Magellan AO. By choosing cases where specific sources of noise dominate we can compare imaging techniques without specifying details such as companion brightness and separation, or the desired S/N.

3.3.4 The Speckle Limited Case

Several differential imaging techniques are in common use to reduce the impact of the coherent speckle noise in high contrast imaging (cf. ADI, SDI, ASDI (Oppenheimer and Hinkley, 2009)). These techniques each have weaknesses, typically being less effective close to the guide star and when used on extended objects such as a circumstellar disk. SDI also requires a strong spectral feature in the companion when compared to the guide star. In cases where speckle suppression cannot be achieved and the term $0.53\tau_0[(1 - \bar{S})H(\theta)f_*]^2n^*$ dominates in Equation (3.7), then Equation (3.8) reduces to

$$\mathcal{DC}_{eff}^{sp} = \mathcal{DC}_1 \left(\frac{\mathcal{A}_1^*}{1 - \bar{S}_1} \right)^2 \left(\frac{1 - \bar{S}_o}{\mathcal{A}_o^*} \right)^2. \quad (3.9)$$

It is apparent from this expression that if, through frame selection, we can increase \bar{S}_1 and/or \mathcal{A}_1^* we will have at least some compensation for the loss of efficiency represented by $\mathcal{DC}_1 < 1$. Furthermore, given the right conditions, frame selection has the potential to maintain or even improve sensitivity while delivering the higher resolution represented by increased \bar{S} and \mathcal{A}_1^* .

3.3.5 The Halo Limited Case

If we are able to suppress the speckles, then when the term $(1 - \bar{S})H(\theta)f_*$ dominates in equation (3.7) we are in the halo photon-noise limited regime. Equation (3.8) then becomes

$$\mathcal{DC}_{eff}^h = \mathcal{DC}_1 \left(\frac{(\mathcal{A}_1^*)^2 n_1^*}{1 - \bar{S}_1} \right) \left(\frac{1 - \bar{S}_o}{(\mathcal{A}_o^*)^2 n_o^*} \right) \quad (3.10)$$

Once again we see that if we can increase \bar{S}_1 and/or \mathcal{A}_1^* we have some leeway with lower \mathcal{DC}_1 . It should also be noted that the S/N maximizing aperture will be different between this case and the speckle limited case above due to the different dependence on n^* .

3.3.6 Background and Read Noise Limited

The next limiting case that should be considered occurs when the dominate noise terms are due to sky BG and detector RON. However, the present work is focused on bright NGS AO. In this regime, we will almost always be limited by halo noise (photon and speckle) within the FOV of our camera. As such, we will only state here, without proof, that because RTFS allows arbitrarily long integrations the detector read-noise performance can be competitive with the current generation of EMCCDs. When we also consider that RTFS allows us to do this over our entire detector FOV without windowing, RTFS retains its competitiveness even for wider separations from faint stars.

3.3.7 Simulated Faint Guide Star Strehl Selection

Now that we have a framework for comparing imaging techniques, we investigate the performance of an ideal RTFS technique on an observation simulated with CAOS. The setup of this simulation was nearly identical to that described in Section 2.5, except seeing was set to the median $r_0 = 18cm$. For an R=10 mag A5V guide star the loop is stable, but undergoes significant fluctuation in correction quality. We ran the simulation with this guide star for 5 seconds of observation time, and extracted the simulated image at the CCD47 at 1 ms intervals and measured Strehl on each of these short frames. We then applied the same corrections for static and non-common-path aberrations to the 1ms measurements as in Table 2.3, but do not use the tip-tilt correction on the 1ms frames.

Next, we establish a threshold S_T and stack each 1ms frame which is above this value. On the combined frame, we then fit elliptical contours at various flux levels (isophots). For the reasons described in Section 3.3.2 we use these contours as apertures to calculate the enclosed flux and number of pixels, and then choose the S/N maximizing aperture for the speckle and halo limited cases. Finally we measure the FWHM resolution using the 50% peak flux contour. The results of this algorithm for various thresholds are shown in Figure 3.1.

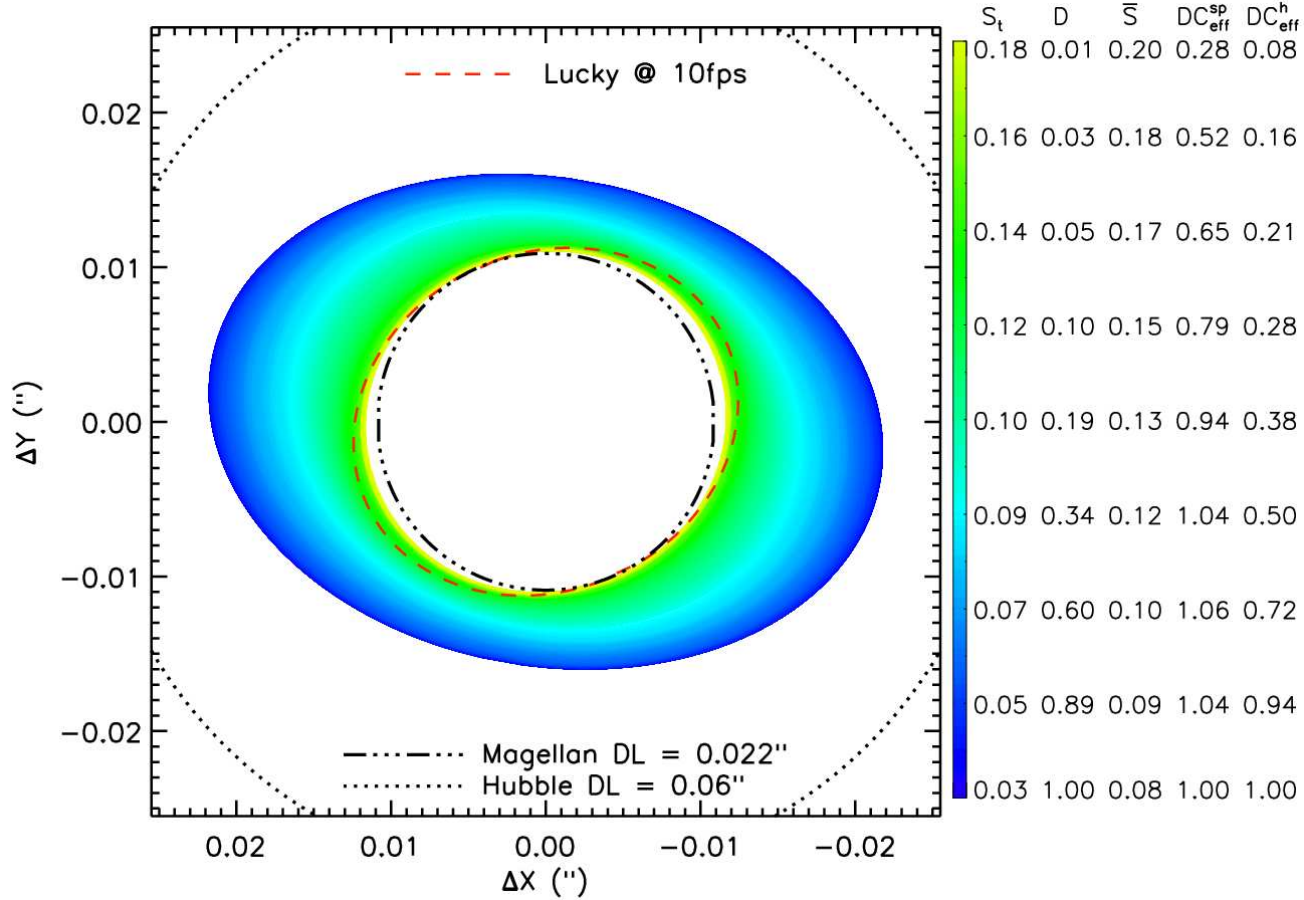


Figure 3.1 Results from Strehl selection on a simulated $R = 10$ guide star, showing FWHM contours corresponding to various selection thresholds. The colorbar encodes the threshold S_t and resultant gross duty cycle \mathcal{DC} for each contour, as well as the resultant mean Strehl \bar{S} and \mathcal{DC}_{eff} for the speckle and halo limited cases. We see that significant gains in resolution can be achieved, and compare these gains to those possible with 10 fps lucky imaging. For comparison the diffraction limits (DL) of the Magellan VisAO system and the Hubble Space Telescope (HST) are plotted.

The very competitive design choice for our system would be to use an EMCCD as the science camera. To study the trade-offs with RTFS, we assume a 1024x1024 array which can be operated at 10fps with negligible read noise. We also assume a QE penalty of 50% when flux is greater than 1 electron/pixel/read. The same simulated frames used for RTFS are combined in 100ms exposures, and we then apply the typical Lucky imaging method of shifting these longer frames before adding based on S_T . In Figure 3.2 we compare this technique to RTFS. When flux is low, EMCCD based lucky provides a large \mathcal{DC}_{eff} advantage due to the resolution boost from shift-and-add, however the ultimate resolution achieved is $\sim 10\%$ worse due to the lower temporal resolution. For brighter objects, or those close to a bright guide star, the EMCCD based lucky performs worse due to the QE reduction.

An important caveat to this discussion is that we have assumed that the EMCCD “QE reduction”, which is actually due to an increase in photon noise, applies identically to the speckle noise in the speckle limited case. This is almost certainly not strictly true, but rather depends on subtle details such as the plate scale and speckle lifetime. As such, the lower lucky-imaging curves in Figure 3.2 should be considered merely an illustration of the point, and actual performance in this case could be better or worse.

A further consideration is the impact of our S/N maximizing apertures. In the speckle-limited case this is almost always the peak pixel due to the strong dependence on the number of pixels in the aperture. Photometry is seldom conducted on a single pixel, however, and so the performance of frame selection (whether real-time or conventional lucky) is understated here if one uses a larger aperture. This is somewhat true in the halo limited case as well, as the optimum aperture tends to be $< 1\lambda/D$ in radius, which is probably also smaller than normally used.

Finally we assume a VisAO tip-tilt loop with $\sim 5\text{mas}$ rms control, simulated by shifting the 1ms images before applying the RTFS algorithm. The correction of residual atmospheric tip-tilt ($\sim 15\text{mas}$ rms) provides a huge improvement in \mathcal{DC}_{eff} . Interestingly, the achievable resolution is the same without tip-tilt control. This is because at very high thresholds, only a small fraction of the simulated data is

used. This implies that by accepting only the best images, the dominant source of resolution degradation (Strehl loss) is rejected and residual tip-tilt has a smaller effect.

The main conclusion of this effort is that RTFS has potential to provide significant gains in resolution and sensitivity - with some trade-offs between the two - similar to conventional Lucky imaging. The benefit of RTFS is that these can be realized over the full FOV of a camera, and with its full QE, using already installed detectors.

3.4 RTFS Implementation

Having established that RTFS offers significant performance enhancement for the Magellan VisAO system in a low correction regime, we next report on the implementation of RTFS. We first test the performance limitations of the mechanical shutter, and then develop algorithms to provide real-time control of the shutter.

3.4.1 Mechanical Shutter Performance

The Magellan VisAO camera uses a Uniblitz VS-25 mechanical shutter. This shutter has a 25mm aperture, and is capable of operation at up to 40hz. Here we follow the manufacturer and discuss shutter speed in terms of a complete open and shut cycle, so 40hz implies 12.5ms exposures if we use a symmetric square wave pulse. This is equivalent to 80 fps with $\mathcal{DC} = 0.5$, over the full FOV of our 1024x1024 detector. The minimum exposure time of the shutter is ~ 10 ms, and it can be operated asynchronously. This gives us the time-resolution equivalent to a 100 fps camera.

We have performed a series of bench tests to determine the accuracy and stability of the VS-25. Our device has an LED-photosensor synchronization circuit, which is interrupted by one of the two shutter blades. The state change of this circuit occurs at 80% shut and 20% open. Figure 3.3 shows the results of one of these tests at 25hz, which is equivalent to 50 fps. We have found the timing of the shutter motion

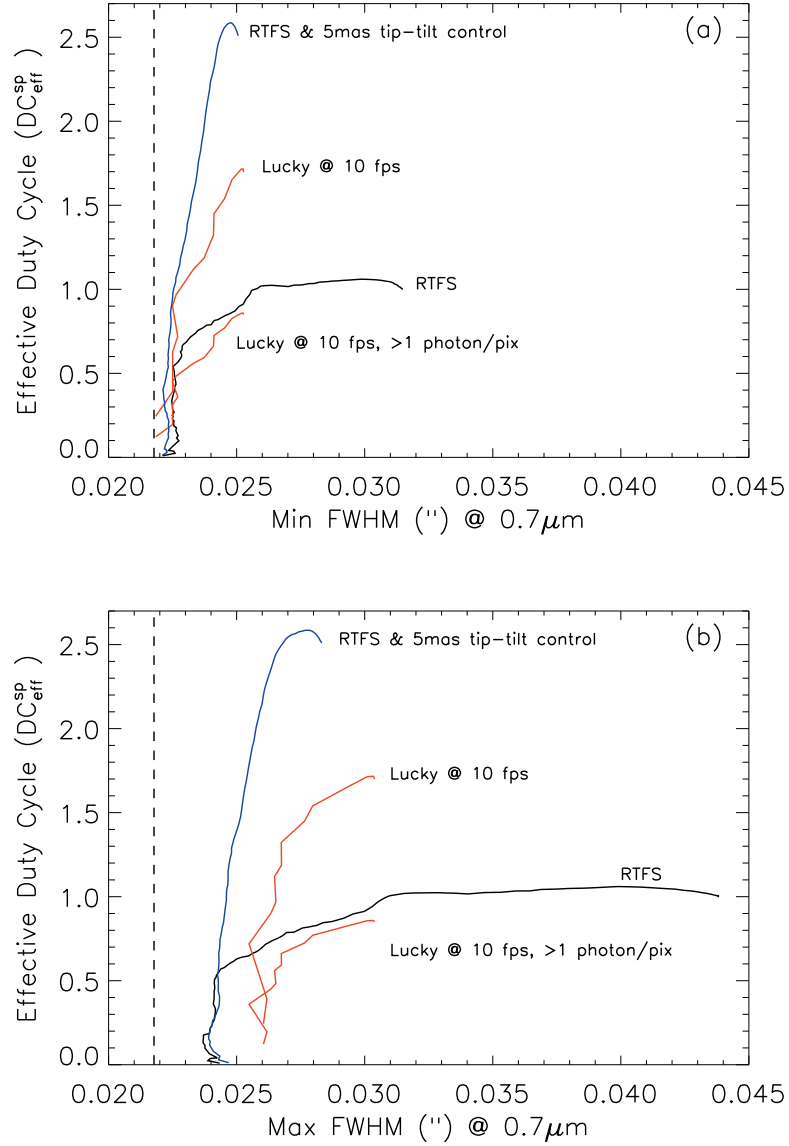


Figure 3.2 Results from Strehl selection on an $R=10$ guide star, showing the effective duty cycle vs the resultant resolution for the speckle limited case. Minimum FWHM (a) corresponds to the semi-minor axis of the elliptical contours shown in Figure 3.1, and maximum FWHM (b) likewise to the semi-major axis. We compare the results with those possible with a 10 fps EMCCD (which has a 50% QE penalty for bright targets). We also show the significant improvement possible with the combination of a fast tip-tilt loop and RTFS - this system could potentially more than double observing efficiency in the speckle limited case.

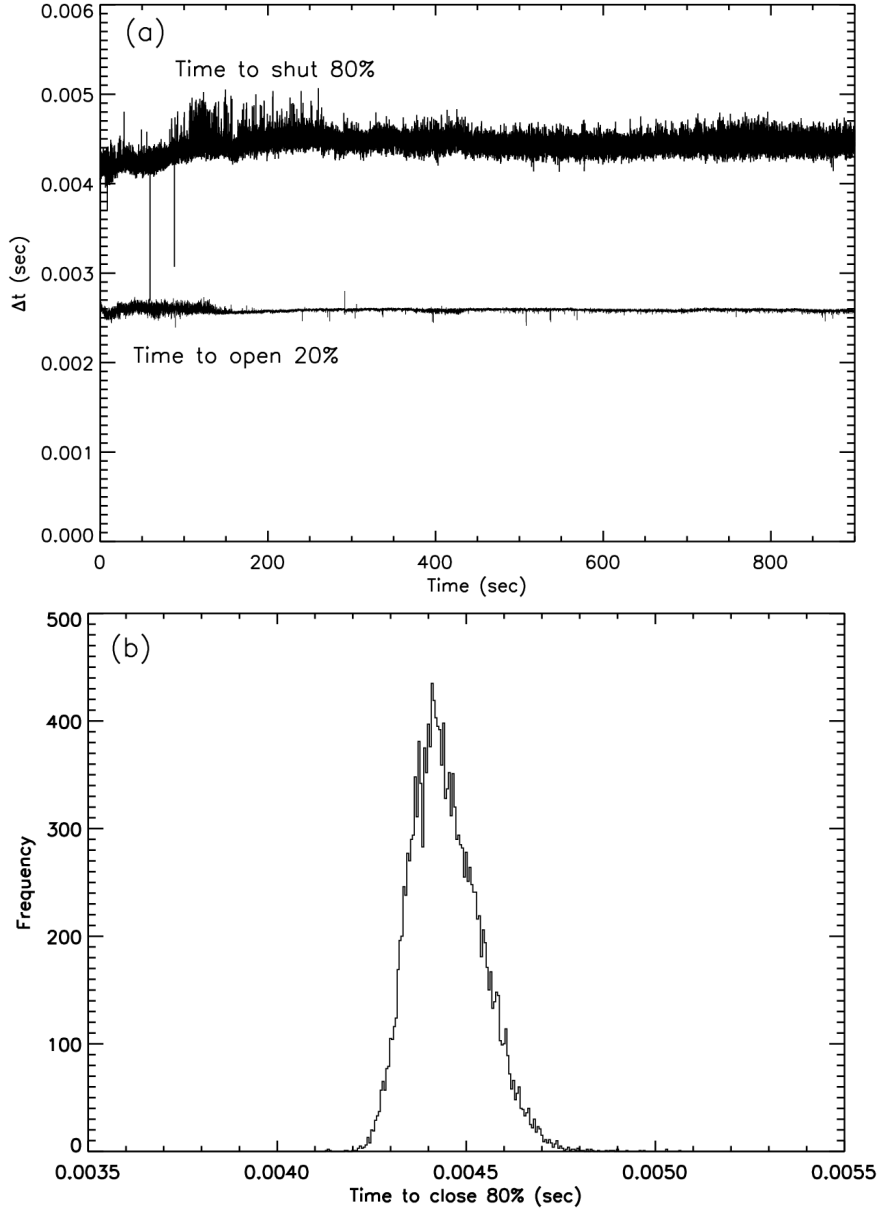


Figure 3.3 Performance of the VS-25 shutter. In (a) we present the time to open and shut while the VS-25 shutter was operated at 25hz continuously (50 fps equivalent, with $\mathcal{DC} = 50\%$). As proxies for full open and shut we use the time to change an LED-photosensor circuit, which corresponds to 80% open and 20% shut. After a few minutes of warm up time the device is very stable. We show the histogram of the 80% shut times after 200 s has elapsed in (b). The standard deviation of this distribution is $\sigma = 92.5\mu\text{s}$, and the open distribution has $\sigma = 11.2\mu\text{s}$. For 180,000 cycles, a total exposure of 1 hour, the uncertainty in exposure time will be 55ms, or 0.002%.

to be very reliable and stable. If we operated the shutter at 25hz for 2 hours, to obtain a total exposure of 1 hour, the resulting uncertainty in exposure time would only be 55ms, or 0.002%.

3.4.2 Telemetry

Our system, based on the LBT architecture, does not have a dedicated real-time telemetry system built in. For RTFS, WFS slopes are taken from an auxiliary output of the slope computer (Microgate BCU 39) so as to leave the main AO loop unaltered. This auxiliary output, a UDP broadcast over standard ethernet, is normally used to send WFS frames to the AO operator's workstation. To capture the slope output in near real-time without requiring any changes to the AO control software, we use an ethernet bridge. A bridge consists of two ethernet adapters, and a kernel software module which passes packets from one adapter to the other so as to be transparent to other devices on the network but allowing one to capture packets. We thus can transparently intercept slope computer diagnostic frames and extract the slope vector. The slope vector is then used to reconstruct the wavefront modal amplitudes as described above. See the schematic in Figure 3.4 which outlines how data flows from the WFS detector, through the bridge, to the GPU (discussed next), and finally becomes a command to the shutter.

3.4.3 GPU Based Reconstruction

We have implemented the matrix-vector multiplication step on a GPU. Our current device is an NVIDIA GeForce GTX 465, which has 352 cores. We used the NVIDIA CUDA³ basic linear algebra (BLAS) library, cuBLAS, to perform the multiplication with the SGEMV routine. For comparison, we also implemented the reconstruction on the CPU using the automatically tuned linear algebra software (ATLAS) package (Whaley and Petit, 2005), which was compiled from source to fully optimize for speed. On the GPU, time to reconstruct a single frame averages 206 μ sec, including

³http://www.nvidia.com/object/cuda_home_new.html

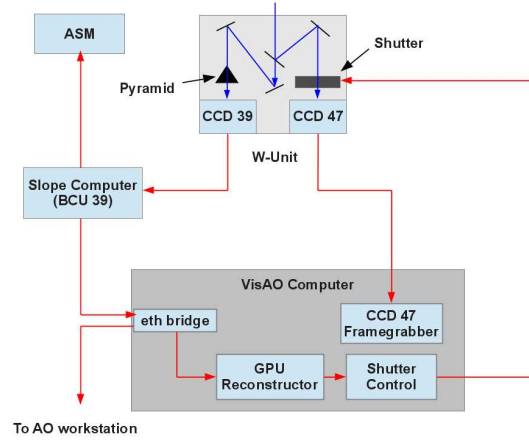


Figure 3.4 Schematic showing the flow of data in our system. The blue lines show the path of visible wavelength light from the telescope. A selectable beam-splitter sends light to both the PWFS (left) and the VisAO camera (right). The red lines show the flow of data. The CCD 39 records the pyramid pupil images, which are processed by the BCU 39 slope computer. The slopes are sent to the ASM via a custom fast fiber link. A secondary output, for telemetry monitoring, sends the slopes over standard ethernet which we transparently intercept for use in RTFS. From there reconstruction occurs on an off the shelf GPU, which produces a real-time measure of Strehl ratio. This is used to determine whether to open or close the shutter.

slope transfer overhead. On the CPU with ATLAS the average time was 314 μsec , so using the GPU results in a 34% improvement in reconstruction speed.

3.4.4 Digital Filter Design

After reconstruction the Strehl time series is low-pass filtered using a finite impulse response (FIR) filter. An FIR is used because of its simplicity and guaranteed stability. Filters appropriate for each loop speed (and telemetry rate) were designed in Matlab using the filter design and analysis tool, fdatool. A major concern is keeping the phase lag low to prevent inaccuracy due to filter delay. Good results have been obtained with a pass frequency of 10 Hz and a stop frequency of 50Hz, and 20 dB attenuation, using the generalized equiripple minimum order technique.

3.4.5 Reconstructor Calibration

Due to fitting error, which is caused by the finite number of spatial frequencies sampled by the WFS, and non-common path (NCP) errors (for VisAO primarily caused by one beam-splitter) we expect our reconstructed Strehl (S_{rec}) to under-predict true focal plane S . Using the extended Marechal approximation, we expect the combined fitting and NCP errors to be a multiplicative correction to S_{rec} :

$$S = e^{(-\sigma_{fit}^2 - \sigma_{NCP}^2)} S_{rec}.$$

In our tests so far, we have found that this simple assumption is insufficient to fully describe the S time series. The logical way to proceed using the above relation is to set the combined fitting and NCP errors to match the mean S_{rec} to the mean S_{foc} measured in the CCD 47 focal plane. However, this technique underestimates the peak-to-valley variability of the true S significantly, which is unacceptable for frame-selection. To match both mean and variance, we use a simple two parameter model

$$S_{cal} = aS_{rec} + b$$

where we estimate the parameters by

$$\begin{aligned} a &= \frac{\text{stdev}(S_{foc})}{\text{stdev}(S_{rec})} \\ b &= \text{mean}(S_{rec})a - \text{mean}(S_{foc}). \end{aligned}$$

We do not use a fitting procedure as the sampling frequencies are different between S_{foc} and S_{rec} . S_{foc} was measured in 32x32 pixel frames taken at 42fps on the CCD 47. The frames were averaged without shifting, and the ratio of this long exposure peak height to the mean short exposure peak height provided an estimate of S loss due to image motion. We then calibrated S_{foc} so that the long exposure peak height is equal to the S measured in a long exposure full frame image, and then the mean short exposure S_{foc} was set to match the image-motion corrected mean S_{foc} . This process avoids the difficulties of accurately normalizing S measurements in small format images. The results of this calibration can be judged in Figure 3.5. Both the mean value and the peak-to-valley variations are well fit by our calibrated reconstructed S .

The parameter a retains a simple interpretation as

$$-\ln(a) = \sigma_{fit}^2 + \sigma_{NCP}^2.$$

Interpreting the parameter b is more challenging. We believe it is related to a phenomenon which has been dubbed the “optical loop gain”, whereby the sensitivity of the PWFS depends on the size of the spot on the pyramid tip and hence on the instantaneous quality of correction. That is, the calibration of S_{rec} depends on the value of S itself. Work is ongoing to understand and calibrate the effect of this optical loop gain. We also note that performing these calibrations in terms of WFE instead of S requires a similar parameter.

3.4.6 Strehl Classification Algorithms

We have previously discussed the use of wavefront sensor (WFS) telemetry exclusively to estimate Strehl ratio. We also considered adding a fast tip-tilt control loop to the VisAO camera. This would have included a small FOV EMCCD providing

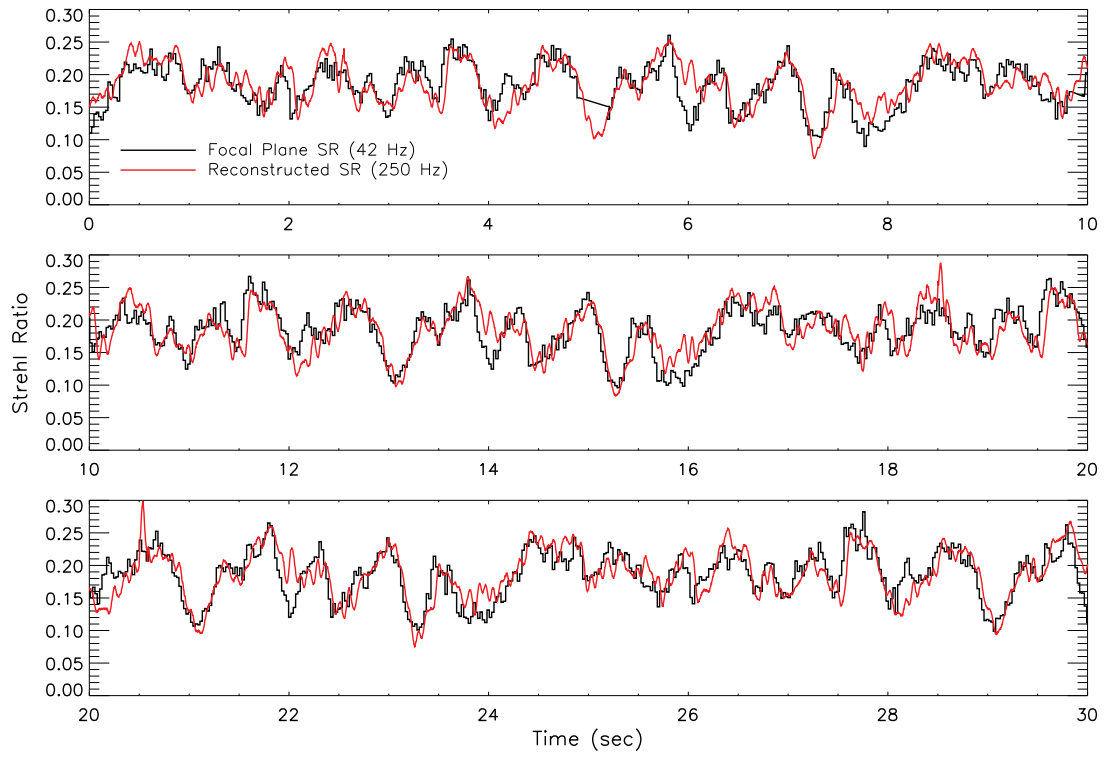


Figure 3.5 Reconstructed S_{rec} compared to S_{foc} measured in the VisAO focal plane, for our SDSS i' bandpass.

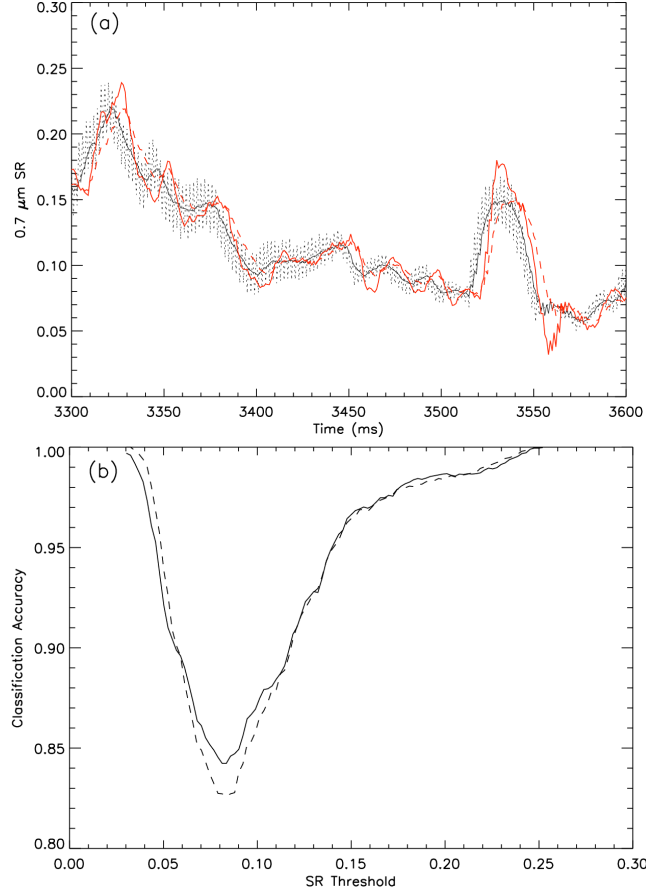


Figure 3.6 Strehl predictions. In part (a) we show a raw time series from the simulation described in Section 3.3.7, zoomed into a representative section. The solid black line is the output of the FIR digital filter designed to have a minimal phase lag. The dashed red line is the “prediction” result of applying the current filtered Strehl measurement with a 6ms shutter actuation delay. The solid red line is the prediction output of using an average of two different linear extrapolations. In (b) we present the classification accuracy of these two prediction strategies vs. Strehl selection threshold.

direct measurements of Strehl. Though the tip-tilt loop has not yet been implemented, here we describe investigations into using such direct measurements. In future work we will investigate the benefits of combining WFS telemetry with the direct measurements.

We can place the RTFS problem in the formalism established in section 3.2. In the simplest case, where we are using the high speed output of our tip-tilt sensor, the data vector $\vec{X}(t_i)$ consists of Strehl measurements, or simply the value of the brightest pixel. To start we can simply use these measurements as input to the classifier G , with a delay of k time steps set by the mechanical performance of our shutter. The main problem with this approach is that we suffer errors due to the delay. We also must deal with noise in a real system, and desire to reject moments of classification change which are shorter than our shutter actuation times (e.g. a 1ms long dip barely below the threshold should not close the shutter).

As discussed above, it is necessary to condition the Strehl time series. The tip-tilt loop will typically be operated at very high frame rates (1-3 khz), faster than the main AO loop. Simulations indicate that this produces a sawtooth pattern in the signal, with sharp increases when the latest correction is applied followed by decreases between corrections. At such short exposure times photon noise may become important on fainter stars as well. Due to the mechanical limitations of a shutter (i.e. minimum cycle time and finite response times) we must filter out these high frequency components of the signal. For this analysis we designed a filter with a pass frequency of 50 hz and a stop frequency of 75 hz, optimized for a minimum phase lag. In terms of the generic selection algorithm, this filter serves as the operator F and the filtered-signal output becomes the input to the classifier G . Figure 3.6 illustrates the performance of this filter on a simulated Strehl time series.

To accomplish prediction, we have found good performance with simple linear extrapolation. Our experiments with various fitting intervals indicate that combining the results of various intervals can improve performance, especially near peaks and valleys in the signal. By interval we mean the number of previous data points used for fitting. Based on the manufacturer specifications for the VS-25 shutter,

we extrapolate 6ms in the future. Figure 3.6 compares the prediction results of averaging 4 step and 10 step fitting intervals with no prediction.

Figure 3.6 also shows the classification accuracy of these two approaches. From Equation (3.1) we know that this is a binary problem, which makes it significantly easier to analyze compared to attempting to calculate the exact value of Strehl ratio. In the bottom panel of Figure 3.6 it is apparent that with a real time source of Strehl measurements we can achieve very good classification accuracy. The simple algorithms we are considering here struggle a little at lower Strehl thresholds (0.05 – 0.10). Considering Figure 3.1, where we find that our best resolution is achieved at thresholds $S_t > \sim 0.15$, these simple techniques appear quite sufficient as accuracy is $> 95\%$.

3.5 Laboratory Demonstration

We conducted a series of experiments in the Arcetri test tower to test our RTFS architecture, at the conclusion of our integration and testing period there.

3.5.1 Experimental Setup

On bright guide stars the MagAO system performs very well, producing high and stable S down to at least $\lambda \sim 0.7\mu\text{m}$. In this regime, RTFS will be counterproductive due to the small variation in S , and the relatively small improvement in signal-to-noise ratio for any corresponding reduction in total exposure time. We therefore expect RTFS to be most useful on fainter stars, where the MagAO system no longer produces such good correction. Due to the way guide star magnitude was controlled in the test tower, both with a variable brightness lamp and by changing beam-splitters, it was challenging to create a fully realistic faint guide star simulation for these tests which also allowed enough light to reach the CCD 47 for accurate measurement of short exposure S . To compensate, we instead used a 9.4 mag guide star but ran the loop at 500 Hz and intentionally set gains to produce a large amount of variability in S . As usual, $0.8''$ turbulence was simulated, with a 15 m/s

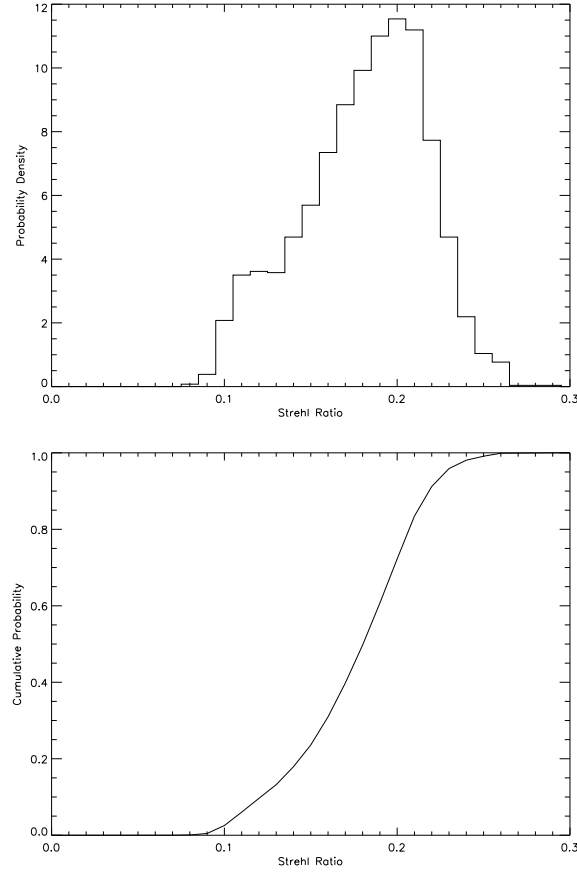


Figure 3.7 Short exposure S statistics.

wind. Though artificial, this setup provided a good test of our RTFS architecture due to the large temporal variability of S while allowing accurate short exposure measurements, and produced the time series shown in Figure 3.5.

The statistics of our S time series are shown in Figure 3.7. Of note, there is negative skewness in the S probability density function (PDF). As other authors have demonstrated, we actually expect positive skewness when the mean value of S is so low (Gładysz et al., 2008b). This discrepancy is likely due to the artificiality of our test setup.

3.5.2 Results

Using the loop setup described above, we set thresholds at the following selection fractions: 100%, 95%, 90%, 75%, 50%, 25%, 10%, and 5%. For example, the 5% threshold means we are attempting to select the best 5% of S . We use the term selection fraction to make it clear that we mean the amount of time the shutter will be open if it is following S with complete accuracy, or in other words the telescope duty cycle.

Once a threshold was selected, the RTFS system was activated in closed-loop while taking data on the CCD 47 in the SDSS i' filter. To avoid saturation, we took short exposures which were then summed (after dark subtraction) without shifting to form a long exposure image. Figure 3.8 shows the resultant images. As we selected higher and higher thresholds, there was a clear improvement in the resultant S , and $FWHM$. Figure 3.9 plots the improvement in these values vs. threshold.

We measured the flux in each long exposure image, and compared the result to the flux in the 100% selection fraction image. This provides a measure of the net telescope duty cycle, or shutter open time, for each threshold. This also serves as a quick-look proxy for total accuracy of our RTFS system, based on how closely our system follows the $y=x$ line. The system tends to under-select at the higher S thresholds. The shutter has an actuation delay of a few msec, and we enforced a re-actuation time of 35 msec in software as a mechanical safety measure. These delays should cause the system to miss many peaks which occur close together, though low pass filtering mitigates this to some extent.

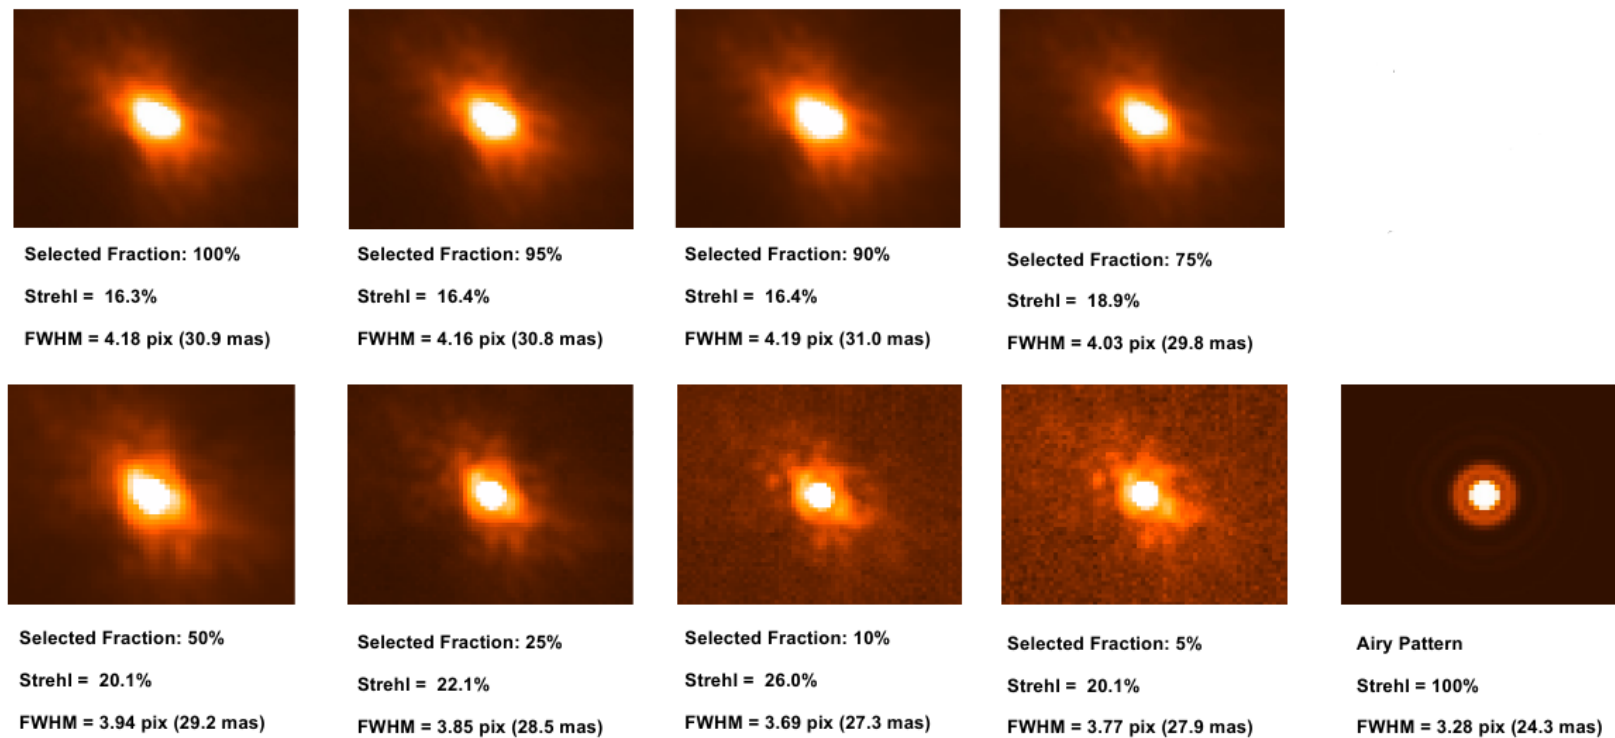


Figure 3.8 Long exposure images obtained with RTFS. Both long exposure S and $FWHM$ improved as the selection fraction is changed to higher values of S .

3.6 Conclusion

When we began simulating the performance of MagAO and VisAO, the simulations indicated that we would have low to moderate Strehl ratios for wavelengths less than $1.0\mu m$. As is typical for such conditions, we saw large fluctuations in correction quality over short periods of time. To compensate for this we planned to utilize a novel frame selection technique, which could provide the ability to reliably achieve the diffraction limit while offering sensitivity and efficiency improvements.

To understand the costs and potential benefits of frame selection, which unavoidably involves discarding valuable telescope time, we have employed a simple model of S/N in AO imaging. Using this model we developed a simple framework to compare various imaging strategies, and showed that our new RTFS technique is very competitive when compared to short-exposure based lucky imaging. This is especially true in NGS visible AO where the presence of a bright star in our FOV makes the use of low-RON EMCCDs challenging.

We demonstrated that commercially available mechanical shutters provide the timing accuracy and precision to support RTFS. It was also shown that with a real time source of short exposure Strehl ratio measurements we can obtain very accurate frame selection. The combination of a fast VisAO specific tip-tilt loop and RTFS would more than double our observing efficiency and consistently provide diffraction limited images at the 6.5m Magellan Clay telescope.

Finally we demonstrated RTFS in closed-loop for the first time. Using WFS telemetry, we were able to reconstruct S in real time and use this knowledge to select moments of good AO correction using a mechanical shutter. This allowed us to improve S from 16% to 26% under (somewhat artificial) laboratory conditions. There is room for further improvement in our system, as higher accuracy would permit more efficient use of telescope time.

As noted in the introduction, MagAO and VisAO are performing much better than our initial pessimistic estimates. One might say that they are performing as they should - something not always achieved by new AO systems. In any case, we

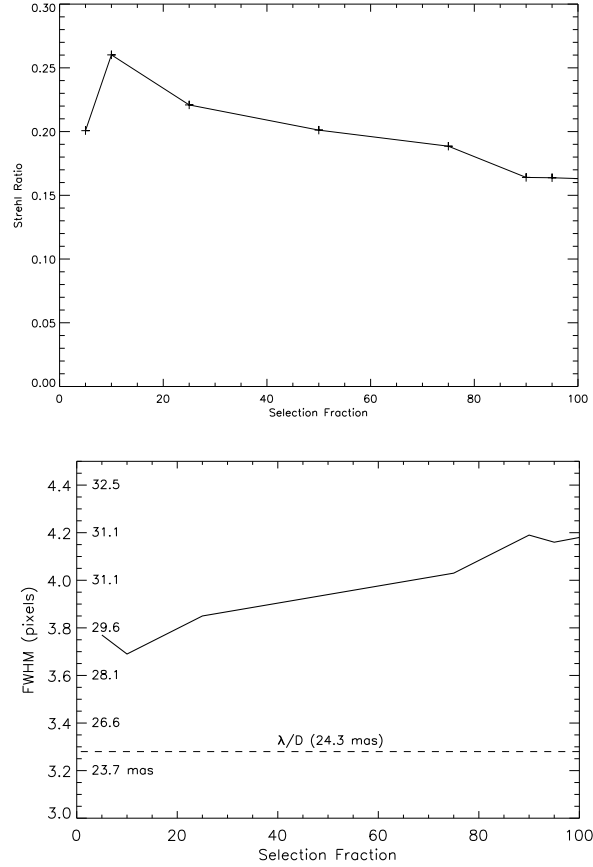


Figure 3.9 At left we plot achieved S for various selection thresholds, and at right we plot $FWHM$. The selection thresholds correspond to the S distribution, that is a 10% threshold means that the shutter opens only when S is in the best 10%. The achieved S was lower when at the 5% threshold, compared to the 10% threshold. The shutter has a finite actuation time of about 10 msec and this drop in S at the highest threshold is possibly due to this delay compared to the width of peaks of S above the threshold. If the shutter actuates too slowly to catch the highest S , it will instead be open during lower S periods.

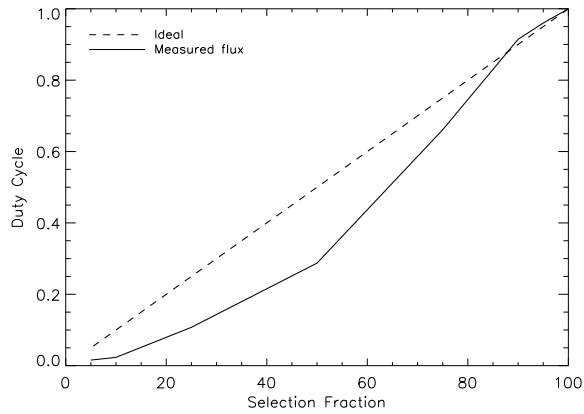


Figure 3.10 Duty cycle for the selection thresholds. Duty cycle was measured photometrically, and serves as a proxy for accuracy. A perfect system would follow the $y=x$ line. The system is less accurate, and tends to under-select, at the most selective thresholds. This is likely due to the shutter actuation delay of 10 msec, and a software safety re-actuation time of 35 msec. These delays cause the system to miss many peaks which occur close together. Low pass filtering mitigates this to some extent.

have so far not needed RTFS on sky. I have conducted a few short engineering trials - mainly to make sure the shutter still works - during our commissioning runs, but have not obtained a full RTFS data set. For now we use the shutter mainly for taking darks.

CHAPTER 4

HIGH CONTRAST IMAGING WITH VISAO: OBSERVATIONS OF β PICTORIS B

4.1 Introduction

In 1984, Astronomers from the University of Arizona, observing at LCO, spatially resolved an extended edge-on debris disk around β Pictoris. Smith and Terrile (1984) hypothesized that the disk, which had been inferred from observations of the IRAS satellite, indicated the presence of planets around β Pic, either forming or recently formed. Since this discovery, the star has been observed repeatedly at many wavelengths. Here we review a small portion of the literature, paying particular attention to the line of evidence pointing to the presence of at least one planet orbiting β Pictoris.

4.1.1 Disk Observations and Variability

The β Pic disk, being visible in seeing limited imaging, was subject to many investigations from ground-based telescopes. For instance, it was observed in R band by Golimowski et al. (1993), using a coronagraph and image stabilizer, also at the 100" du Pont at LCO. They detected wide asymmetries in the disk. Kalas and Jewitt (1995) also observed in R band with a coronagraph, and also detected asymmetry. The asymmetries found were on too large a spatial scale to be associated with planets closer to the star. Using an anti-blooming CCD, rather than a coronagraph, Lecavelier Des Etangs et al. (1993) obtained $BVRI_c$ images with information as close as 2" from the star and noted that the nature of the disk particles changed closer to the star, possibly related to planetary formation, and that there were hints of inhomogeneous dust distributions in the inner regions of the disk.

The disk was observed with HST as well, and a warp in the inner disk was detected (Burrows et al., 1995; Heap et al., 2000; Golimowski et al., 2006). Ground-based AO observations also show the warped inner disk (Mouillet et al., 1997). Analysis of this warp indicated that it should be unstable, and required the presence of a perturbing planet to be maintained (Augereau et al., 2001) .

Lagrange-Henri et al. (1988) interpreted spectroscopic variability in metallic absorption lines as the infall of solid bodies — comets — of km size. This led to the interpretation that these comets were being disturbed by a planet (Beust et al., 1990, 1991), possibly at high inclination. A further interesting incident of variability occurred on 1981 Nov 10, when β Pic was observed to grow rapidly fainter consistent with a transit by a Jupiter sized object (Lecavelier Des Etangs et al., 1995).

4.1.2 Detection of β Pictoris b

Lagrange et al. (2009) reported the detection of a point source 8 AU from β Pic, as many expected based on the several lines of circumstantial evidence pointing to the presence of a planet around the star. These first observations, from 2003, were followed up and confirmed in the fall of 2009 by Lagrange et al. (2010). In between the 2003 initial detection and the 2009 confirmation, the planet had orbited to the opposite side of the star. The planet has since been detected throughout the near and mid-IR (Quanz et al., 2010; Bonnefoy et al., 2011; Currie et al., 2011b; Chauvin et al., 2012a; Bonnefoy et al., 2013; Boccaletti et al., 2013).

Though it was reported to be mis-aligned with the warp by Currie et al. (2011b), dynamical analysis indicates b is most likely the cause of the warp (Dawson et al., 2011) and later observations of the planet and disk together in the same image indicate that b is in fact aligned with the warp (Lagrange et al., 2012). In addition to matching expectations based on disk morphology, the interesting possibility exists that a transit of β Pic b also explains the Nov 1981 brightness transient (Lecavelier Des Etangs and Vidal-Madjar, 2009).

4.2 High Contrast Observations of β Pictoris with VisAO

Here we present observations of β Pictoris with MagAO/VisAO. To perform these observations, we used the VisAO near-focal-plane occulting mask coronagraph. We expect β Pic b to be located under the mask, so we here present the extensive characterization of the coronagraph we performed in the lab and on-sky. Next we describe our observations and data reduction, and show the contrast obtained with VisAO.

4.2.1 The VisAO Coronagraph

The VisAO camera contains a partially transmissive occulting mask, used to prevent saturation of the CCD when observing bright stars. The mask has a radius equivalent to $0.1''$, were it in the focal plane, but it is approximately 60mm out of focus in an f/52 beam. Consequently it attenuates flux out to $\sim 0.8''$ in radius.

We calibrated the mask transmission and PSF by scanning an artificial test source across the mask with the W-unit off the telescope. The source was scanned along 12 different lines, spaced roughly 30 degrees apart, in the Y_S filter. For each scan, we found the best fit center assuming symmetry, and then combined data from all scans. Transmission is reported here using the ratio of the maximum pixel. As we will discuss shortly, the PSF changes shape somewhat dramatically under the mask so peak fitting is unreliable. The results of our laboratory scans are shown in Figure 4.1. Though the mask itself is not apodized, it is out of focus resulting in the smooth roll-off of attenuation. The maximum attenuation of the max pixel is 0.0015, that is $ND = 2.8$.

The profile of the mask is not well described by convenient functions (Gaussian, Moffatt, polynomials), so we merely re-binned the raw lab data (median) and interpolate between points as needed for analysis. This binned profile is shown in Figure 4.2.

We tested our laboratory calibration on-sky as well. We scanned a star, with the AO loop closed and the Y_S filter, across the mask using the same script as used

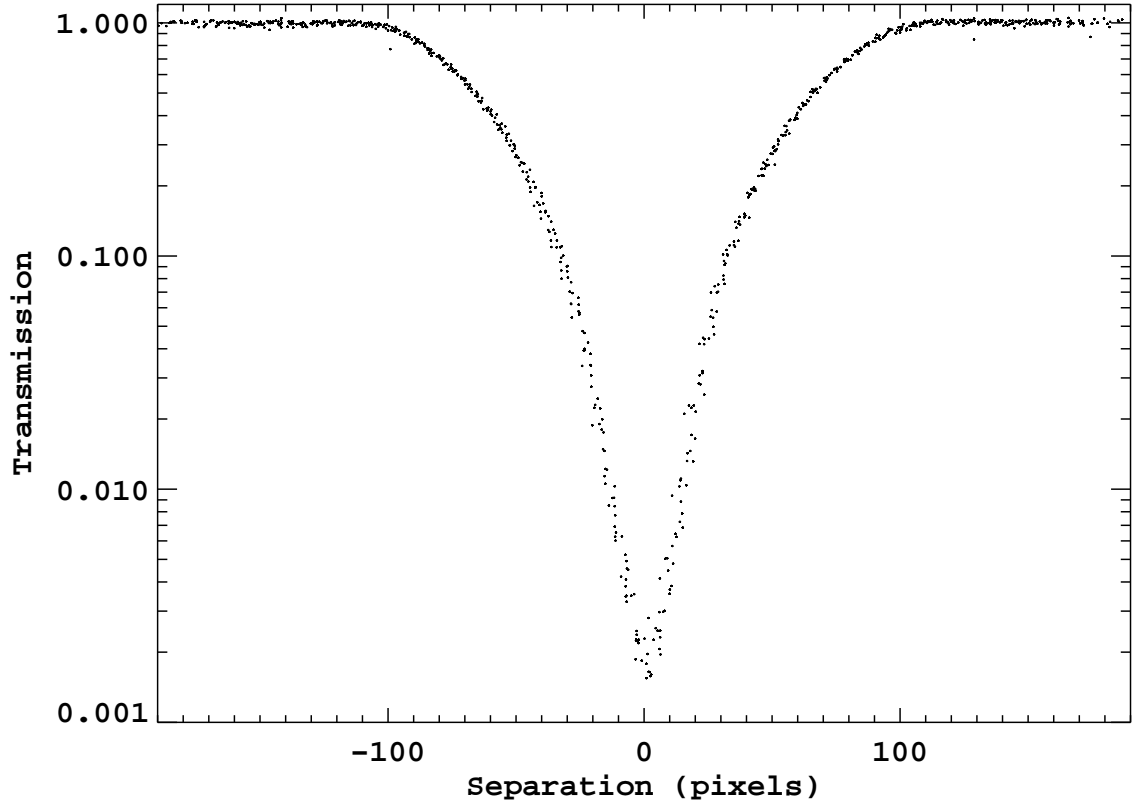


Figure 4.1 Transmission of the VisAO occulting mask, on a log scale. Shown here is the ratio of maximum pixel - no peak fitting was applied due to the change in shaped caused by the mask. Though the chrome mask is not apodized, it is out of focus resulting in this apodized transmission profile. Maximum attenuation at the center is 1.5×10^{-3} , or ND= 2.8. The separation of β Pic b is 59 pixels.

for the test source along two lines 90 degrees apart. We compare these on-sky scans to the lab scans in Figure 4.2. We also tested the mask using binary stars on-sky. The results of these tests are also shown in Figure 4.2.

As we noted, the apodized transmission profile changes the shape of the PSF under the mask. Relevant example PSFs are shown in Figure 4.3. The test source uses optics designed for the LBT, and so has a smaller PSF. We magnified these images to match the diffraction limited PSF we measured on-sky, which has a FWHM=4.7 pixels at $0.98\mu\text{m}$.

In Figure 4.4 we quantify this change in shape, comparing the FWHM of an elliptical Gaussian fit to the PSF vs. position, along the semi-major and semi-minor axes. The PSF is elongated radially, but maintains the diffraction limited FWHM in the azimuthal direction. In Figure 4.5 we show the ratio of FWHMs and also show measurements from our on-sky scans.

4.2.2 Observations and Data Reduction

We observed β Pic on the night of 2012 Dec 04 UT, in the Y_S filter using the coronagraph. We used the 50/50 beamsplitter. Conditions were photometric, with variable seeing at the beginning of the observation, but settling down to $\sim 0.5''$ by the end. Towards the end of the observation we took off-coronagraph calibration data, which was presented earlier in Section 2.7.3, where we found that Strehl ratio was $32 \pm 2\%$ and that true optical Strehl was 40%. The latter number is important for high contrast imaging as it sets the flux in the halo as $(1 - S)$.

We used the WFE telemetry recorded in the Fits headers to select images, using a cut of 130 nm RMS phase. This resulted in a little over 2hrs of data, with an elapsed time of 4.17 hrs and 116 degrees of rotation. The complete details of the observation are given in Table 4.1.

Images were bias and dark subtracted, using shutter-closed darks taken at 15 min intervals throughout the observation. To better facilitate data processing, the 3399 selected and dark-subtracted individual exposures were median coadded in 30 second chunks, with a rotation limit of 0.5 degrees. That is, the images were coadded

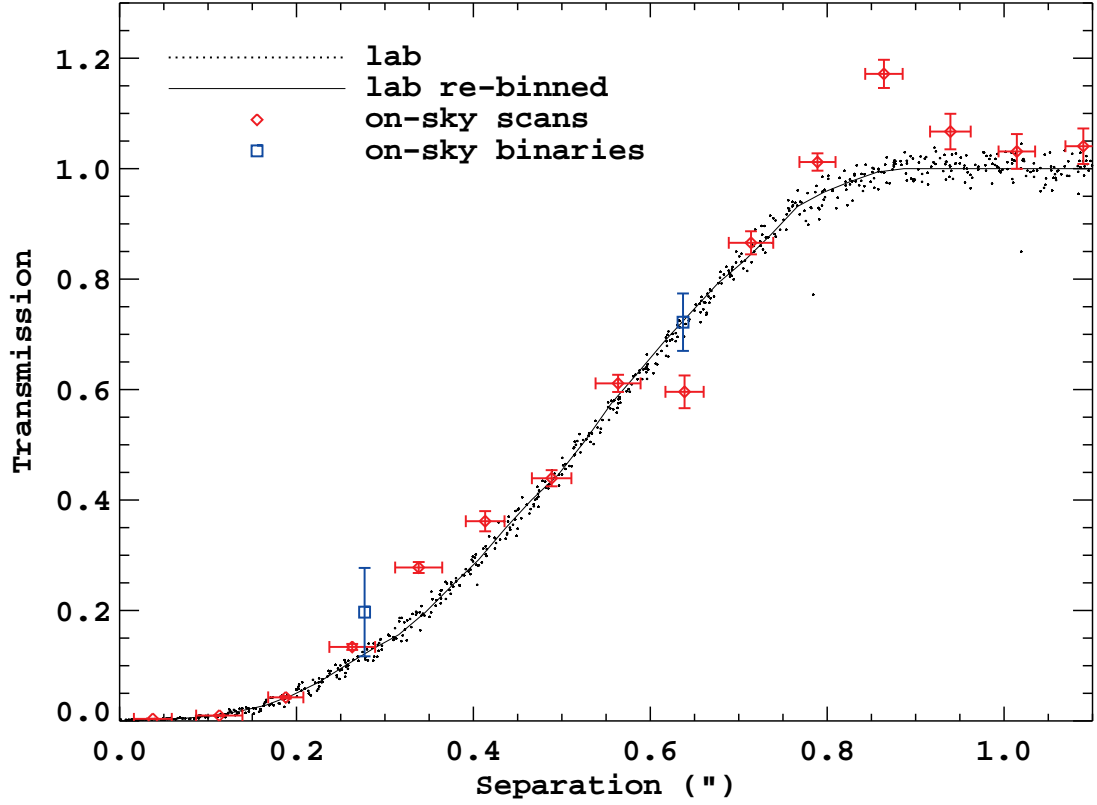


Figure 4.2 Transmission of the VisAO occulting mask. As before this is the ratio of maximum pixels. Here we also show the on-sky calibration checks conducted, including scans of a star across the mask and two binaries compared on and off the coronagraph. The on-sky scans were binned: the separation error bars are the standard deviation of the separations in each bin, and the transmission error bars are the standard deviation of the mean in that bin. The separation of β Pic b was $0.47''$.

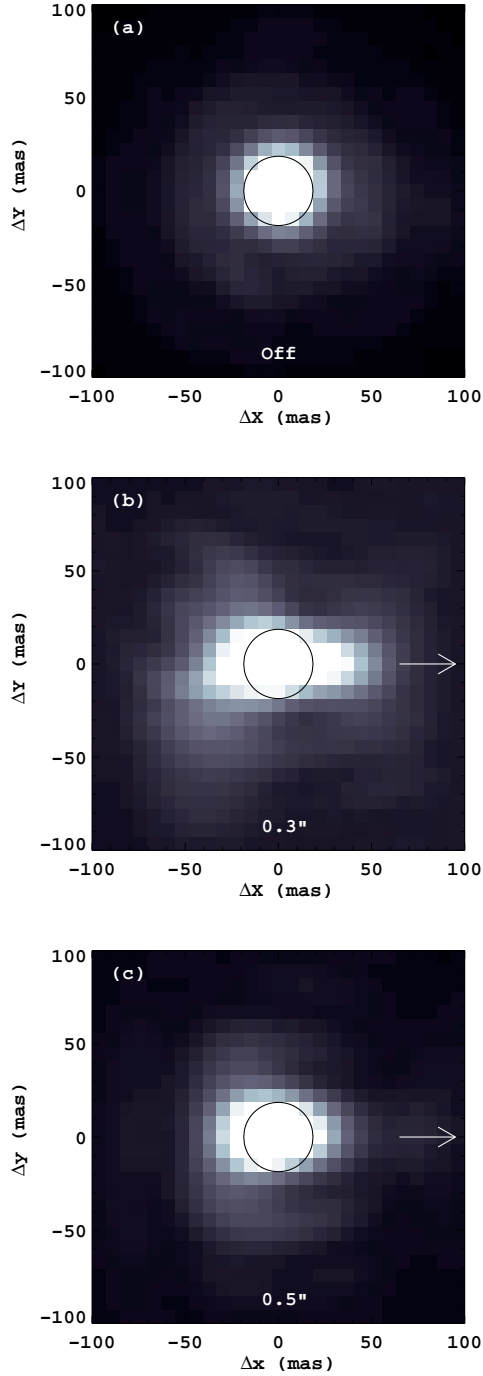


Figure 4.3 Test source calibrations of the PSF through the occulting mask. (a) shows the PSF of the test source off the mask. (b) shows the PSF at $\sim 0.3''$ and (c) shows the PSF at $\sim 0.5''$ separation from the mask center. The black circle has a diameter of 4.7 pixels, the FWHM of the un-occulted PSF measured on-sky. Images are in a linear stretch. The white arrow indicates the direction of the coronagraph center. The separation of β Pic b was $0.47''$.

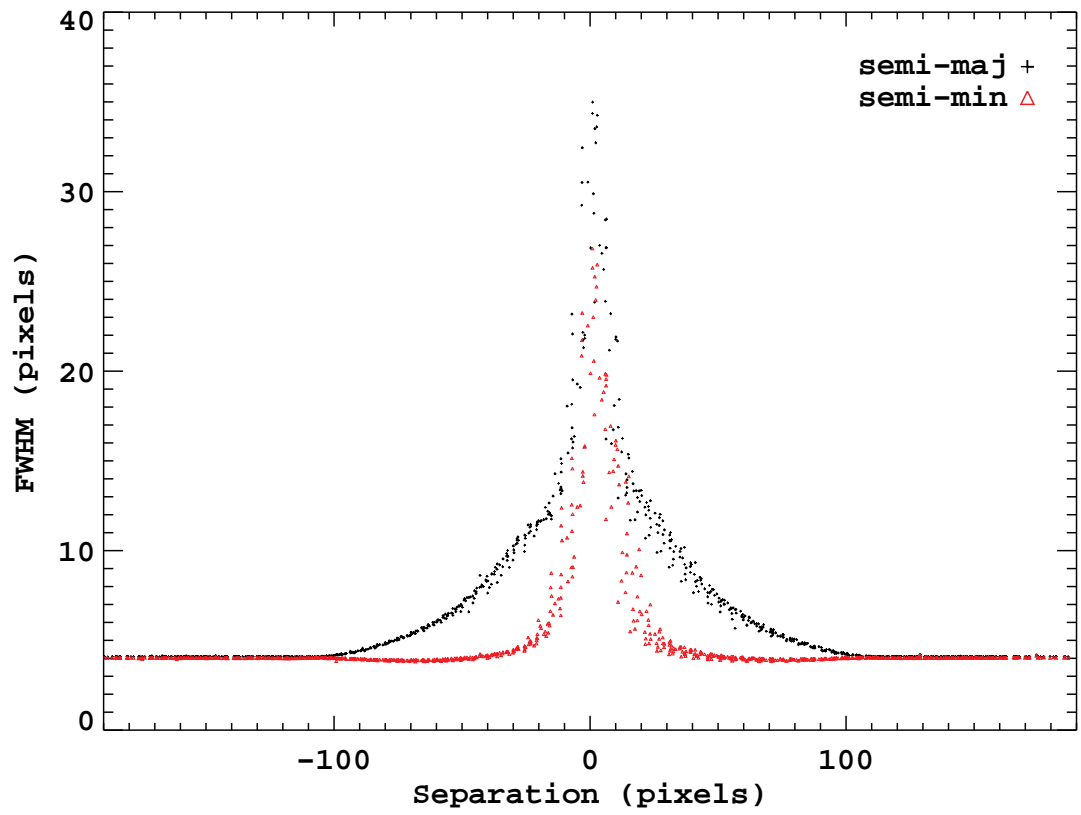


Figure 4.4 FWHMs of an elliptical Gaussian fit to the PSF vs. separation from the mask center. The separation of β Pic b was 59 pixels.

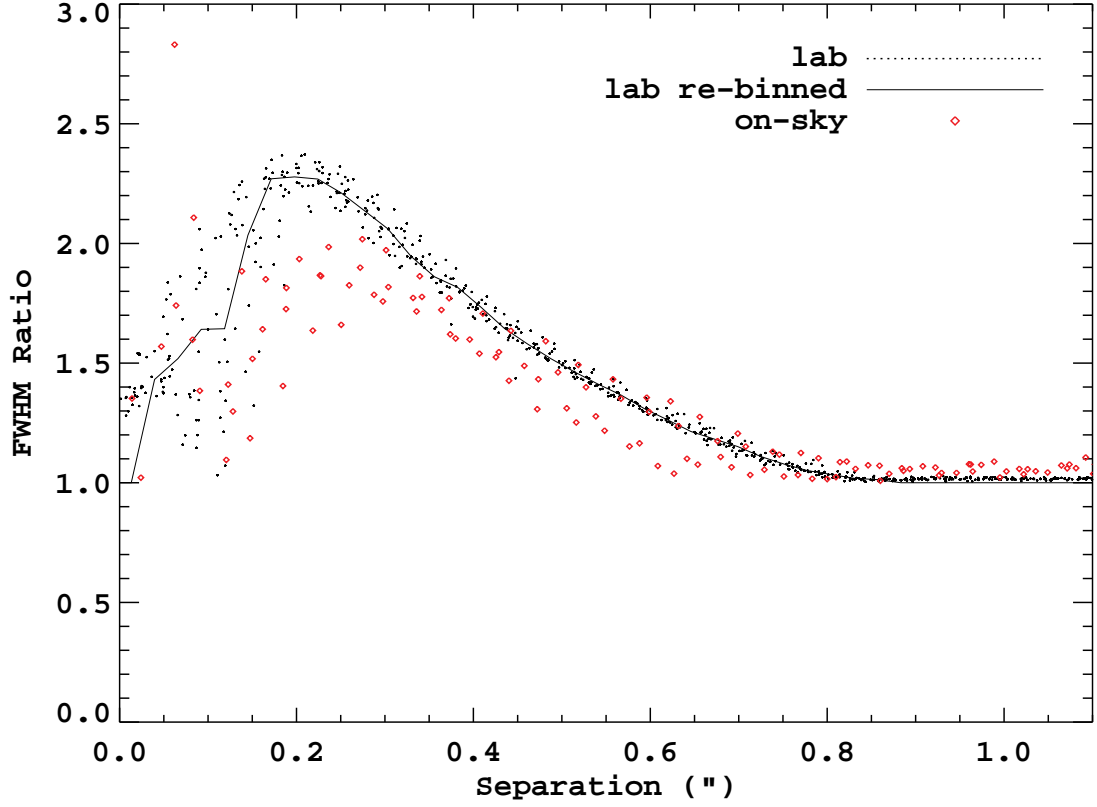


Figure 4.5 Ratio of Gaussian FWHM along the radial direction to the FWHM along the azimuthal direction under the mask, vs. separation from the mask center. The occulting mask changes the shape of the PSF, elongating it radially away from the mask center. The on-sky results are biased due to the slight elongation present in the PSF due to the wind. This causes the ratio to start > 1 , and to not reach the same peak value. The separation of β Pic b was $0.47''$.

Table 4.1 Observations of β Pictoris in Y_S with the coronagraph.

Date	2012 Dec 04 UT
Filter	Y_S ($0.984\mu\text{m}$)
CCD Speed	250 kHz
CCD Window	1024x1024
Exp. Time	2.273 s
Elapsed Time	4.17 hrs
Tot. Rotation	115.97 deg
WFE ¹ Threshold	130.0 nm rms
No. Images	3399
Tot. Exposure	2.15 hrs
Seeing FWHM	0.5'' - 0.7''
AO Speed	990 Hz
AO Binning	1x1
AO Modes	200
Mean WFE ¹	125.91 nm rms
Median WFE ¹	128.9 nm rms
Long Exp. Strehl	$32 \pm 2\%$
Optical Strehl ²	40%
PSF FWHM	4.7 pix=37.2 mas

¹Instantaneous WFE, average over a single exposure.²Corrected for PRF (charge diffusion)

until 30 seconds elapsed time or 0.5 degrees of rotation had occurred.

The coadded images were first registered and centered using the beamsplitter ghost, which has a good astrometric calibration to the center of the star (cf. Section 2.7.2). The coronagraph is partially transmissive ($ND \approx 2.8$) in the core, so on a bright star we are able to use the star itself for centering. After registering with the ghost, we then located the center of rotational symmetry of the attenuated star using cross correlation. This procedure finds the center of rotational symmetry to better than 0.05 pixels.

We then median combined the registered images, forming our master PSF under the mask. This is shown in Figure 4.6. Using this master PSF, we then carried out a standard median ADI (MADI) data reduction: the master PSF was subtracted from individual coadds, and then each was derotated and the images were median combined. No radial profile subtraction or subtraction of nearby images was applied. The resulting image is shown in Figure 4.7. We masked the central 200 pixels in this image, where the residuals are high. Outside the mask, the well known disk is clearly visible extending to the NE and SW. Next we describe our reduction of the central region with principle component analysis.

Closer to the star, within $\sim 1.5''$, MADI does not provide good speckle subtraction. Inside $1''$ the mask modulates speckles in a position dependent way, and combined with the small drift that occurred due to flexure, the speckle pattern changed over the course of the observations. We therefore employed principle component analysis (PCA), implementing the Karhunen-Loève Image Processing (KLIP) algorithm of Soummer et al. (2012).

We employed search areas, applying KLIP to these regions individually. This is similar to the optimization and subtraction regions employed in LOCI (Lafrenière et al., 2007b), though in this reduction our optimization and subtraction regions were the same. For each image, the reference PSFs were taken to be any of the remaining images where at least one $FWHM = 4.7$ pixels of rotation had occurred at the inner edge of the region being reduced. We then applied PCA to the reference PSFs, forming a KL basis set. The optimum PSF was then calculated by determining

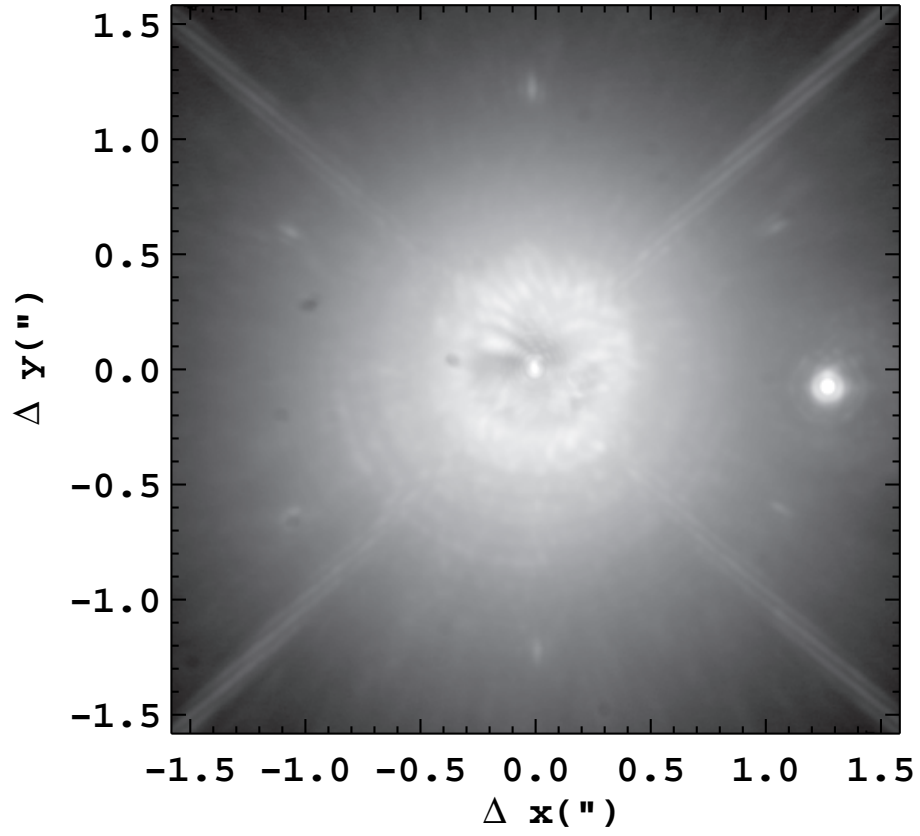


Figure 4.6 VisAO Y_S master coronagraphic PSF. Noteworthy features include the beamsplitter ghost to the right, which has Airy rings, and the six short diffraction spikes at $\sim 1.2''$ due to the six-fold symmetry of the ASM actuators. The flare-like structure extending to the upper left from the central star is characteristic of slight misalignment with the center of the mask.

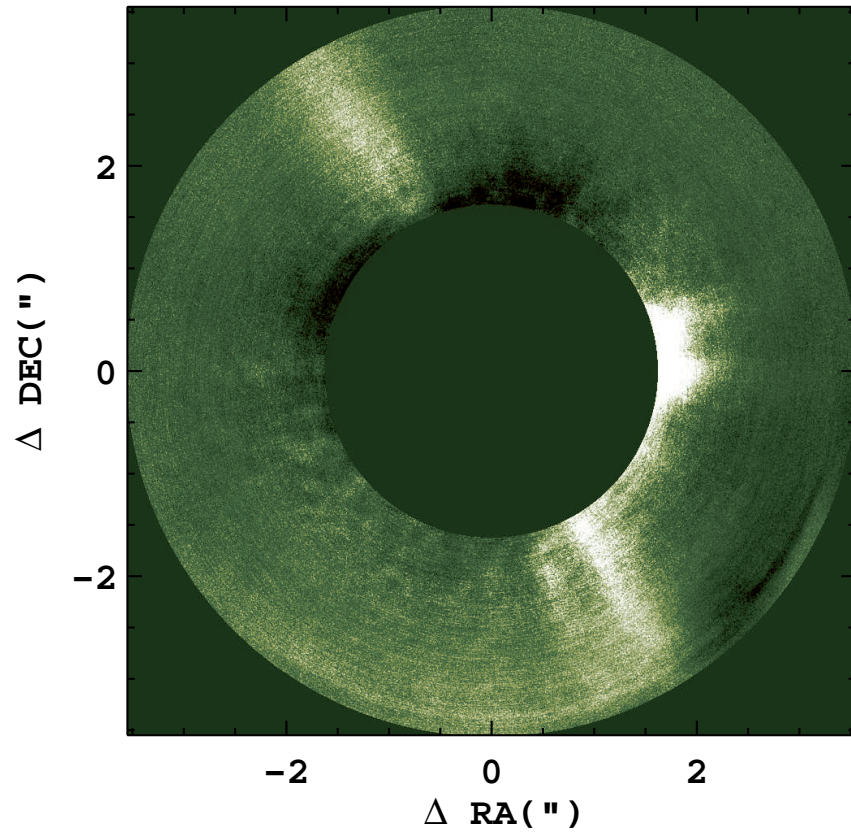


Figure 4.7 Median-ADI reduction of β Pictoris. The large central mask corresponds to the region reduced with PCA. The disk is detected, though we do no analysis of it here.

the coefficients to apply to each basis image, and summing. The optimum PSF for each image was subtracted, and the resulting images were de-rotated and median-combined. The result is shown in Figure 4.8(a).

Next we apply three spatial filters, the parameters of which were determined using fake planet injection described later. The images reduced with PCA have azimuthal structure, as is common in ADI reductions. To remove this, we apply an azimuthal unsharp mask whereby a smoothed image is constructed using a kernel with a Gaussian profile in azimuth and a boxcar profile in radius. The Gaussian has angular width $FWHM_\theta = 10$ degrees, and the boxcar has radial width $FW_r = 0.5$ pixels. This azimuthally smoothed image is subtracted to remove the residual azimuthal structure, the result of which is shown in Figure 4.8(b). We then apply a symmetric Gaussian unsharp mask to high-pass filter the image, with a $FWHM = 4.7$ pixels (Figure 4.8(c)). This removes most features that are wider than the PSF. Finally, we smooth the images (low-pass filter) with a Gaussian kernel of $FWHM = 2$ pixels, which effectively Nyquist samples the 4.7 pixel PSF. The final result is shown in Figure 4.8(d).

To aid in determining the optimal values for various parameters (including the number of KLIP modes, the radial and azimuthal size of the search regions, and the smoothing filter parameters), we injected fake planets of various contrasts using the measured coronagraphic PSF at $0.47''$ (the expected separation of β Pic b) at 21 position angles spaced 9 degrees apart, in the 180 degrees opposite the location of the planet. Images with the fake planets were reduced with PCA and filtered as described. We tried various combinations of region sizes and spatial filter parameters. For each set of parameters, we conducted aperture photometry on these fake planets. To estimate the noise in the aperture photometry we also tested apertures at the same radius where there was no fake or (known) real planet. This allowed an estimate of S/N at each contrast level, which we used along with visual inspection of the results to determine the best parameters.

We found that choice of region size, in both radial extent and angular width, affected throughput. In general, having fewer pixels in the search region decreased

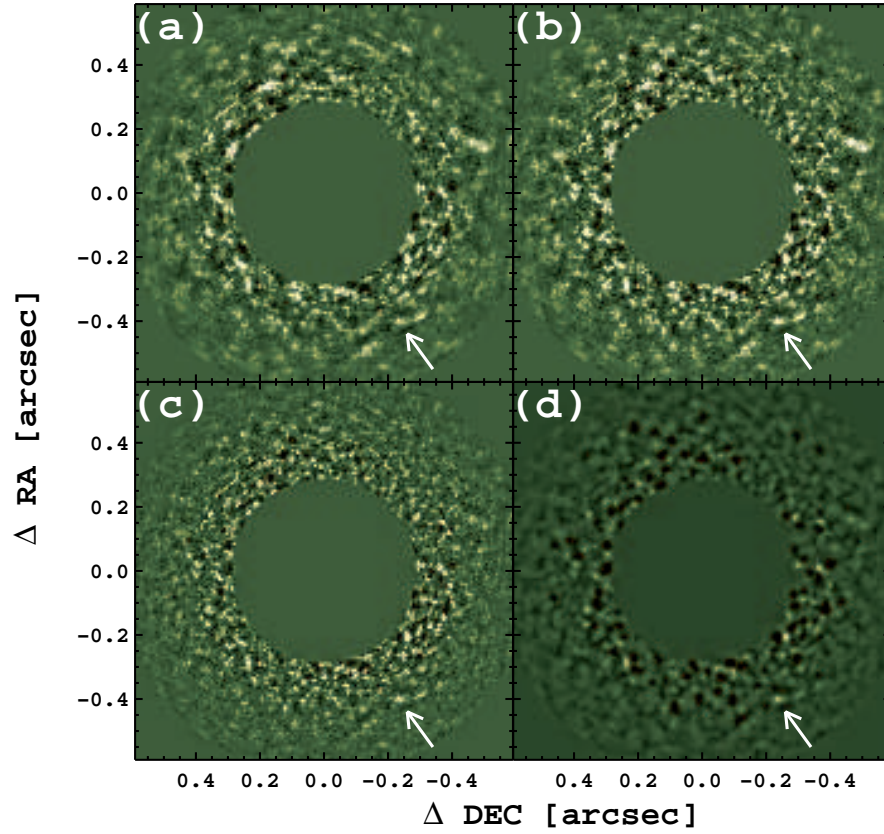


Figure 4.8 Here we show the step-by-step results of the 3 spatial filters we applied. In (a) we show the raw results of our KLIP/PCA reduction. (b) shows the image after applying the azimuthal unsharp mask. (c) shows the data after the symmetric unsharp mask, and (d) shows the final image after gaussian-smoothing by 2 pixels. In each panel the arrow highlights the detection of β Pic b.

throughput. Conversely, minimizing the size of the region allowed better speckle subtraction. These competing effects caused us to vary the search region size with radius. KLIP has the advantage that throughput can be modeled at any location in the image, even if a real astronomical source is there. This is done by projecting a model of the signal onto the same KL modes, and reducing the model as before. This “forward-modeling” (Soummer et al., 2012) is subtly distinct from fake planet injection, in that we are not injecting the model images into the data. After applying PCA we apply the same spatial filters as is done to the actual images. Throughput is estimated by comparing the heights of fitted peaks. Though photometry of β Pic b was calibrated with the fake planet injection, we used such throughput measurements to determine our detection limits as described next.

4.2.3 VisAO Y_S Contrast Limits

To determine the detection limits of this observation, we first quantified the fundamental noise floor. We use the master PSF from the MADI reduction to calculate a radial profile of the PSF. This was converted from ADU/image to total photons by multiplying by gain and the total exposure time, and the halo photon noise at each pixel is then the square root of this quantity. We next divided the 1σ noise by the coronagraph transmission profile. This limit in photons is converted into contrast by using the beamsplitter ghost and the calibration of ghost-peak to PSF-peak described in Section 2.7.2. The RON limit was calculated similarly, assuming $4.5e^-$ per readout (see Table 2.1). Note that the RON contrast limit is not constant, as it too is divided by the coronagraph transmission. Next, the halo photon noise and RON were added in quadrature, and then normalized for coronagraph and PSF peak. This then gives the fundamental detection limit of our observation. We do not expect to be photon noise limited, however, instead speckles will set our floor.

Next we calculated the standard deviation in 4 pixel annuli for both the MADI and the PCA reductions. This was chosen as it corresponds to the 2 pixel radius aperture we used for photometry. These noise profiles were converted to electrons, contrast relative to the peak, and corrected for coronagraph transmission. The PCA

limit was also corrected for throughput determined by forward modeling. Since the disk is detected in the MADI image, we mask out the disk $\text{PA} \pm 15$ degrees.

All of these contrast curves are shown in Figure 4.9, as the 3σ detectability contrast limits. We use 3σ in this analysis because we have a very strong prior on the location of the planet.

4.3 Detection of β Pictoris b with VisAO

As noted hinted above, there is a point-source-like signal in our PCA reduction at the expected location of β Pic b. In Figure 4.10(a) we show the reduced data. In Figure 4.10(b) we show the S/N map, which is the image divided by the radial noise profile. In both images there is a clear point source at or near the expected position of the planet (separation $\approx 0.47''$, $\text{PA} \approx 211$).

We observed this target simultaneously with Clio, where we had much higher S/N from $3\text{--}5\mu\text{m}$, thus providing us a position for the planet. On the same night as these VisAO observations, we obtained a high S/N detection of the planet in the M' bandpass, at separation $0.496 \pm 0.025''$ and $\text{PA } 212.0 \pm 2.0$ degrees (Morzinski et al., 2013, in prep), noting that these measurements are preliminary pending a refined distortion solution for Clio. We used the fake planets to calibrate our VisAO astrometry by measuring their position after reduction. We found that even at this low SNR, Gaussian centroiding produced unbiased results on the fake planets. The estimated position uncertainty from this analysis is 0.65 pixels. In tests on a binary in ADI mode using the coronagraph, we found an additional 1 degree systematic error in PA which is possibly due to the coronagraphic mask. We add this error directly to the statistical uncertainty estimated from fake planets. We did not find a similar systematic error in separation. We measured a separation of $0.4676 \pm 0.0059''$ and PA of 210.4 ± 1.65 degrees. Our astrometry is summarized in Table 4.3.

To further determine the significance of our detection, we conducted aperture photometry along circles of constant radius, with a spacing of 1 FWHM, using apertures 2 pixels in radius. We formed an estimate of the local sky at 5 pixels

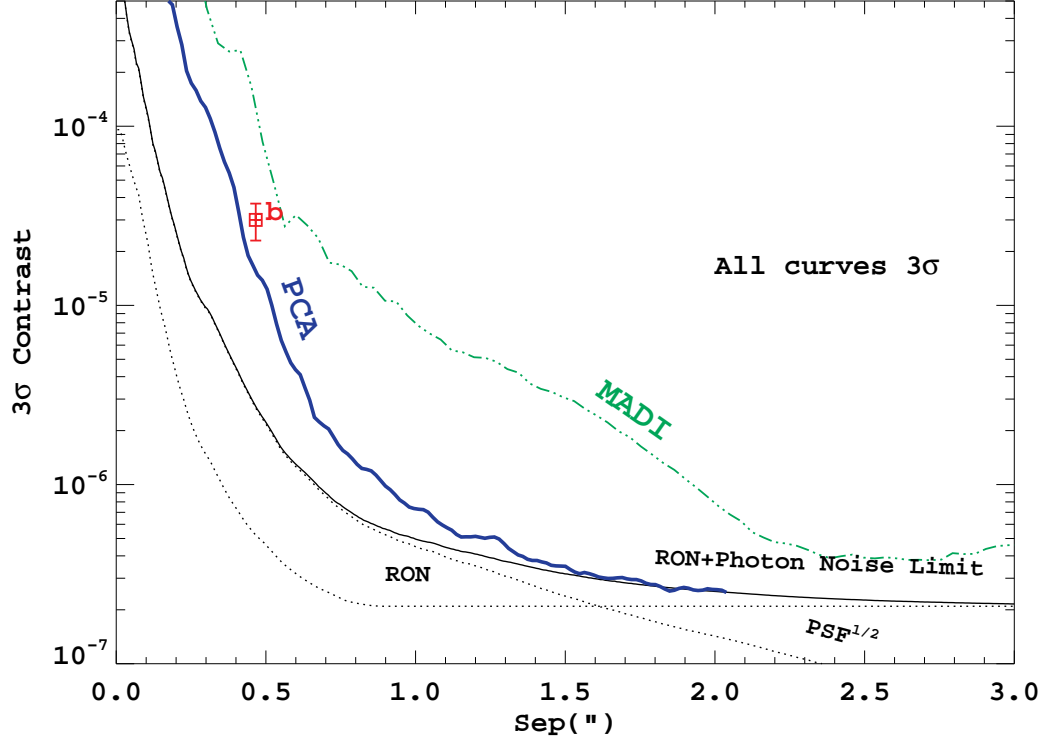


Figure 4.9 Here we show the 3σ detection limits, in terms of contrast ratio, for ~ 2 hrs of integration on β Pic at Y_S . The readout noise (RON) assumes $4.5 e^-$ per pixel. The PSF limit is derived from the median-PSF calculated as part of ADI processing. The radial profile of the median-PSF was multiplied by gain and divided by the coronagraph transmission profile shown in Figure 4.2. RON and PSF limits were added in quadrature to determine the basic \sqrt{N} limit. This sets the fundamental noise limit of this observation. We then show the 3σ contrast curves, calculated as standard deviation in 4 pixel annuli, for median ADI (MADI) with the disk masked, and for principle component analysis (PCA). PCA shows significant improvement over MADI, and past $\sim 1.0''$ is rapidly converging to the \sqrt{N} photon-noise limit. We also show our detection of β Pic b, at $0.47''$.

Table 4.2 Astrometry of β Pictoris b

Clio $M' \mu\text{m}$	Sep = $0.496 \pm 0.025''$	Morzinski et al, in prep
	PA = 212.0 ± 2.0 deg	(preliminary)
VisAO Y_S :	Sep = $0.4676 \pm 0.0059''$	This work
	PA = 210.4 ± 1.65 deg	

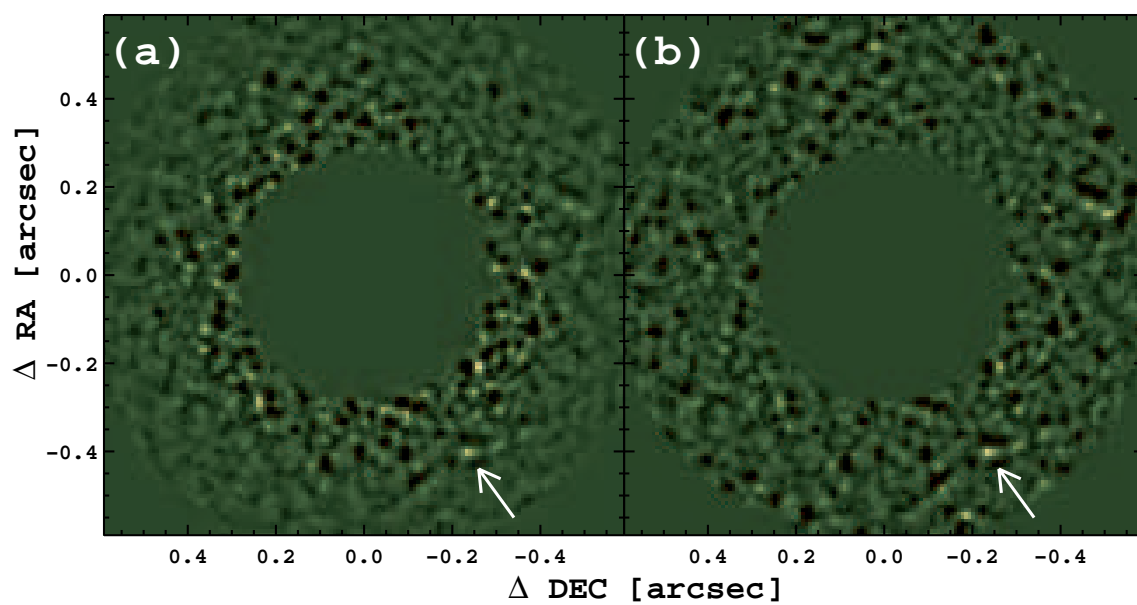


Figure 4.10 (a) Final reduction. (b) S/N map.

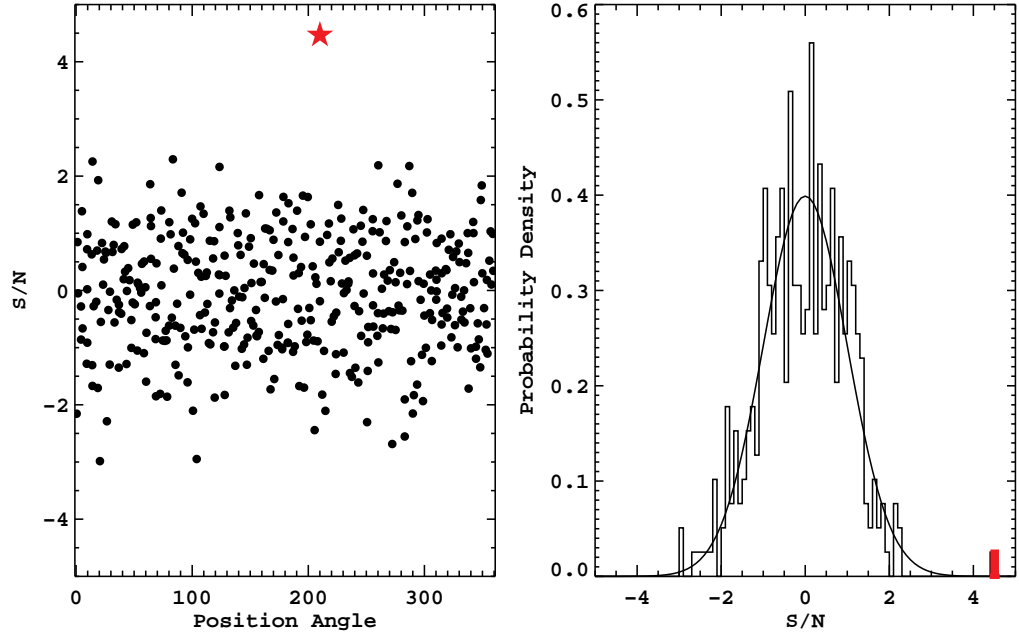


Figure 4.11 (a) Aperture S/N vs. position angle. The aperture containing β Pic b, which was excluded from the standard deviation, is shown as the red star. (b) Histogram of the results. The location of the planet is highlighted in red. The solid curve is a normal distribution with $\sigma = 1$.

radius, with a width of 1 pixel. We did this at 4 radii, separated by 1 FWHM each. The results for each radius were normalized by the standard deviation of aperture fluxes on that radius, and we excluded the aperture containing the known planet from the calculation of standard deviation. We show the results of this test in Figure 4.11. The statistics appear to be Gaussian, and β Pictoris b is detected with $S/N \sim 4.4$. This gives a false alarm probability of $\sim 5 \times 10^{-6}$.

We also used the fake planets to calibrate our photometry in the reduced image. This results in an estimate of Y_S contrast of $(3.0 \pm 0.7) \times 10^{-5}$, or $\Delta Y_S = 11.30 \pm 0.25$

mags. Mitchell and Johnson (1969) observed β Pic in the 99 filter of the 13-color system. The central wavelength of our Y_S filter and the 99 filter are nearly identical, so we use their measurement of 3.561 ± 0.035 mag as the brightness of β Pic A in Y_S . We take into account the following sources of uncertainty:

- Detection S/N: 23%
- Coronagraph centering (± 5 pixels): 16%
- Ghost calibration: 2.4%
- β Pic A Photometry: 3.5%
- β Pic A distance modulus: 0.005%

Our estimate for the apparent magnitude of β Pic b is then 14.87 ± 0.31 , and 13.43 ± 0.31 for the absolute magnitude.

4.4 Prior Measurements of β Pic b photometry in J , H , and K_S

β Pic b has been detected at other wavelengths by other observers. At J band, it was detected with NACO by Bonnefoy et al. (2013). This data was re-reduced by Currie et al. (2013) with slightly higher reported S/N. As these are reductions of the same data with different algorithms, we do not treat them as statistically independent results, rather here we adopt the higher S/N measurement of Currie et al. (2013). For the H band detection of NACO, we likewise choose the reported higher S/N measurement of Currie et al. (2013) over Bonnefoy et al. (2013). Currie et al. (2013) did conduct new observations of β Pic b with NICI at H band. Since there are slight differences in these two bandpasses, we do not average these results. In K_S band we have the NACO measurement of Bonnefoy et al. (2011) and the NICI measurement of Currie et al. (2013). There is also an unpublished NICI measurement provided by Mike Liu (private communication). We average the two NICI measurements. The $Y_S J H K_S$ photometry of β Pic b to-date is collected in Table 4.3.

Table 4.3 Photometry of β Pictoris b

Filter	Instrument	Date	Apparent Magnitude	Absolute Magnitude	Notes
Y_S	VisAO	2012/12/04	14.87 ± 0.31	13.43 ± 0.31	[1]
J	NACO	2011/12/16	14.11 ± 0.21	12.68 ± 0.21	[2]
H	NACO	2012/01/01	13.32 ± 0.14	11.89 ± 0.14	[2]
	NICI	2013/01/09	13.25 ± 0.18	11.82 ± 0.18	[2]
K_S	NACO	2010/04/10	12.6 ± 0.1	11.2 ± 0.1	[3]
	NICI	2012/12/15	12.37 ± 0.13	10.93 ± 0.13	[4]
	NICI	2013/01/09	12.47 ± 0.13	11.04 ± 0.13	[2]
	NICI mean		12.42 ± 0.09	10.98 ± 0.09	[1]

Notes: [1] this work, [2] Currie et al. (2013), [3] Bonnefoy et al. (2011)
[4] Mike Liu, private communication.

4.5 Prior Exoplanet Photometry in the Y Band

This is the first ground-based direct-detection of an exoplanet in the optical. Other efforts have pushed into the Y band on the HR8799 planets. Currie et al. (2011a) detected HR8799 b at $1.04\mu\text{m}$ in the unfortunately labeled z filter. From here on we, follow Liu et al. (2012) and refer to this filter as $z_{1.1}$. See further discussion in Appendix B.

Oppenheimer et al. (2013) have the ability to work down to $0.995\mu\text{m}$ with Project 1640, and reported low significance detections of HR8799 b and HR8799 c at $1.05\mu\text{m}$. At $0.986\mu\text{m}$, our detection of β Pic b further pushes our ability to probe exoplanet atmospheres blue-ward.

Another planet-mass object which has been observed at similar wavelengths is 2M1207b, which was observed with the Hubble Space Telescope by Mohanty et al. (2007). For both 2M1207b and HR8799b we would like to directly compare our results in the optical and near-IR. To do so, we must convert these measurements into our Y_S bandpass and the NACO filters. Though we show how this can be done

Table 4.4 Estimated Y_S and NACO photometry of 2M1207b and HR 8799b

Filter	λ_0 μm	Abs. Magnitude Measured	Estimated	References
2M1207b				
F090M	0.905	18.86 ± 0.25	—	[1]
Y_S	0.986	—	18.2 ± 0.26	[2]
Y	1.032	—	17.79 ± 0.26	[2]
$z_{1.1}$	1.039	—	17.73 ± 0.26	[2]
F110M	1.102	17.01 ± 0.16	—	[1]
J_{2MASS}	1.241	16.40 ± 0.21	—	[3]
J_{NACO}	1.256	—	16.34 ± 0.21	[2]
H_{2MASS}	1.651	14.49 ± 0.21	—	[3]
H_{NACO}	1.656	—	14.52 ± 0.22	[2]
$K_{S,2MASS}$	2.166	13.33 ± 0.12	—	[3]
$K_{S,NACO}$	2.160	—	13.38 ± 0.13	[2]
HR 8799b				
Y_S	0.986	—	18.84 ± 0.29	[2]
Y	1.032	—	18.31 ± 0.29	[2]
$z_{1.1}$	1.039	18.24 ± 0.29	—	[4]
J_{MKO}	1.249	16.30 ± 0.16	—	[5]
J_{NACO}	1.256	—	16.52 ± 0.17	[2]
H_{MKO}	1.634	14.87 ± 0.17	—	[5]
H_{NACO}	1.656	—	14.92 ± 0.17	[2]
$K_{S,MKO}$	2.156	14.05 ± 0.08	—	[5]
$K_{S,NACO}$	2.160	—	14.17 ± 0.09	[2]
Notes: [1] Song et al. (2006), [2] this work, [3] Mohanty et al. (2007), [4] Currie et al. (2011a), [5] Marois et al. (2008b)				

for normal brown dwarf spectral types in Appendix B, these two objects do not correspond to any known brown dwarf. Instead we turn to published best-fit model spectra. For 2M1207b we use the model of Barman et al. (2011), and for HR8799b we use the best-fit model from Madhusudhan et al. (2011). We use the models to calculate a Δmag between the filters with actual photometry and our desired bandpasses. The estimates for the photometry in other bands are then formed by applying this Δmag to the actual measurement. We illustrate this in Figure 4.12 and summarize the results in Table 4.4.

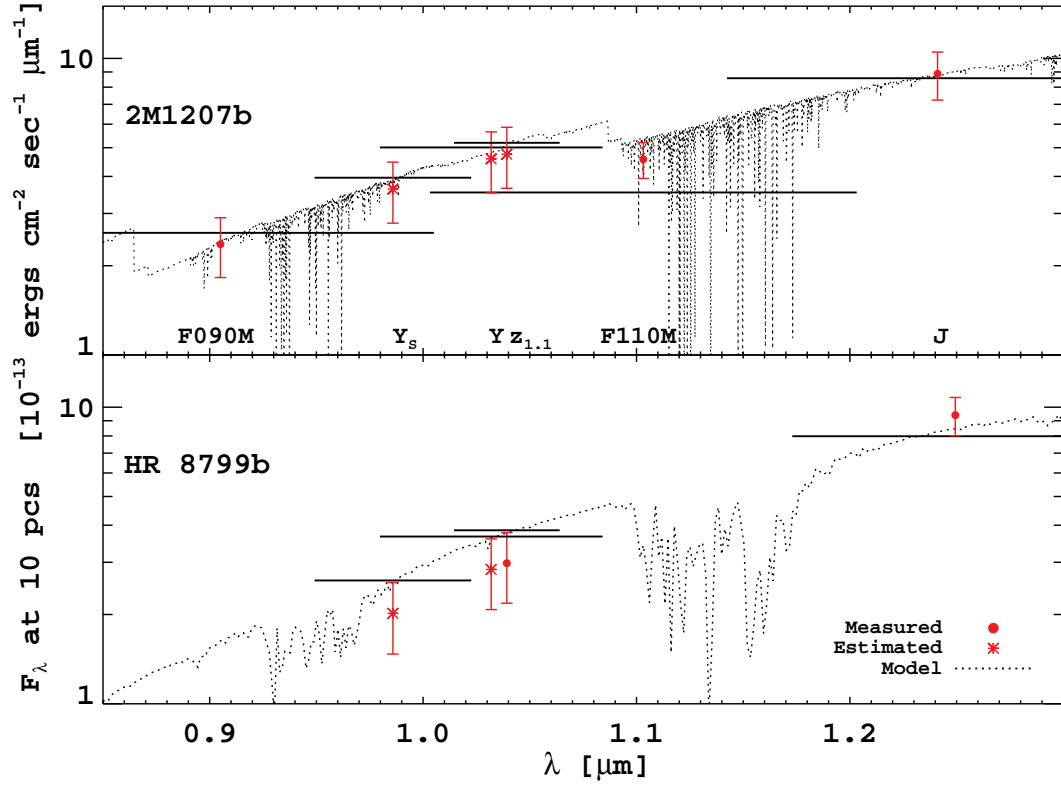


Figure 4.12 Optical and near-IR photometry and models of 2M1207b and HR 8799b. To estimate photometry in the Y atmospheric window and in the NACO filters we used the models to extrapolate or interpolate from the measured photometry. Measured photometry is indicated by filled circles, and estimated photometry is indicated by asterisks. The 2M1207b model is from Barman et al. (2011), and the HR 8799b model is from Madhusudhan et al. (2011). See also table 4.4

4.6 Discussion

It is interesting to compare β Pictoris b with other planetary mass objects, as well as with field brown dwarfs. We discussed 2M1207b and HR8799b above. For the field dwarfs, we collected YJHK spectra of 497 objects. 441 of these are from the SpeX Prism Spectral libraries maintained by Adam Burgasser¹ (from various sources), 23 are WISE brown dwarfs from Kirkpatrick et al. (2011), and 33 are from Allers and Liu (2013). We correlated 91 of these with parallax measurements, either listed in the SpeX Library (from various sources) or from Dupuy and Liu (2012). We conducted synthetic photometry on these spectra as described in Appendix B.

In Figure 4.13 we compare β Pic b to the field dwarfs in color-color plots, and include HR 8799b and 2M1207b. We also highlight the location of 2M0355, a young low-gravity object identified by Faherty et al. (2013) as having colors similar to the young, dusty planets. Two things are evident from the color-color plots. First, β Pic b is very different from the other planets. The other HR 8799 planets (c,d,e) can, based on their near-IR colors, be expected to fall in regions similar to the three comparison planets and planet-like objects we show. Furthermore, though Y_S appears somewhat brighter than expected, β Pictoris b has colors consistent with those of early L dwarfs. We note that age alone can not explain the differences here, as 2M1207 (~ 8 Myr) is younger than β Pic (~ 12 Myr), and both are likely younger than HR8799 (~ 30 Myr).

Based on this observation, we next attempted to fit the YJHK photometry of β Pic b to our collection of brown dwarf spectra. We conducted synthetic photometry on the spectra, as described in Appendix B, in each of the filters: VisAO Y_S , NACO J , NACO H , NICI H , NACO K_S , and NICI K_S . For each spectrum, we then found the magnitude offset that minimized χ^2 between these synthetic measurements and the measurements of β Pic b (see Table 4.3). The best five fitting objects are shown in Figure 4.14. Of note, the very best fitting spectra is a medium surface gravity (β) L1 dwarf.

¹<http://pono.ucsd.edu/~adam/browndwarfs/spexprism/>

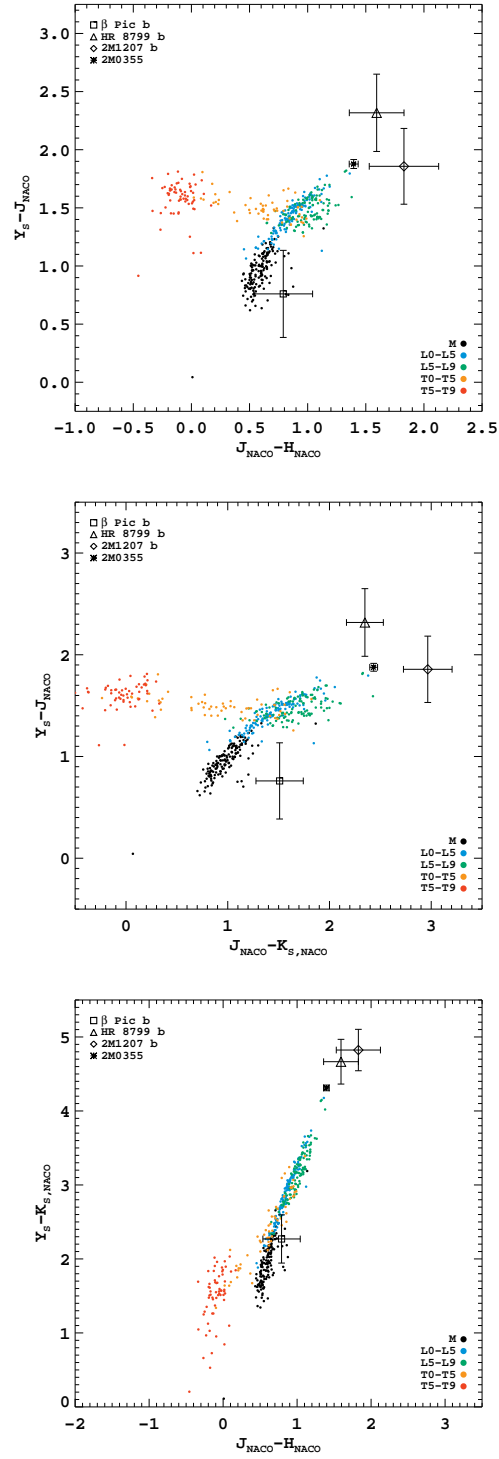


Figure 4.13 β Pic b color-color plots, comparing β Pic b to the field brown dwarfs, as well as the planets HR 8799b and 2M1207b and the low-gravity dusty brown dwarf 2M0355. β Pic b has colors consistent with a field L dwarf, or perhaps an early T dwarf. This is quite different from other directly imaged young giant planets.

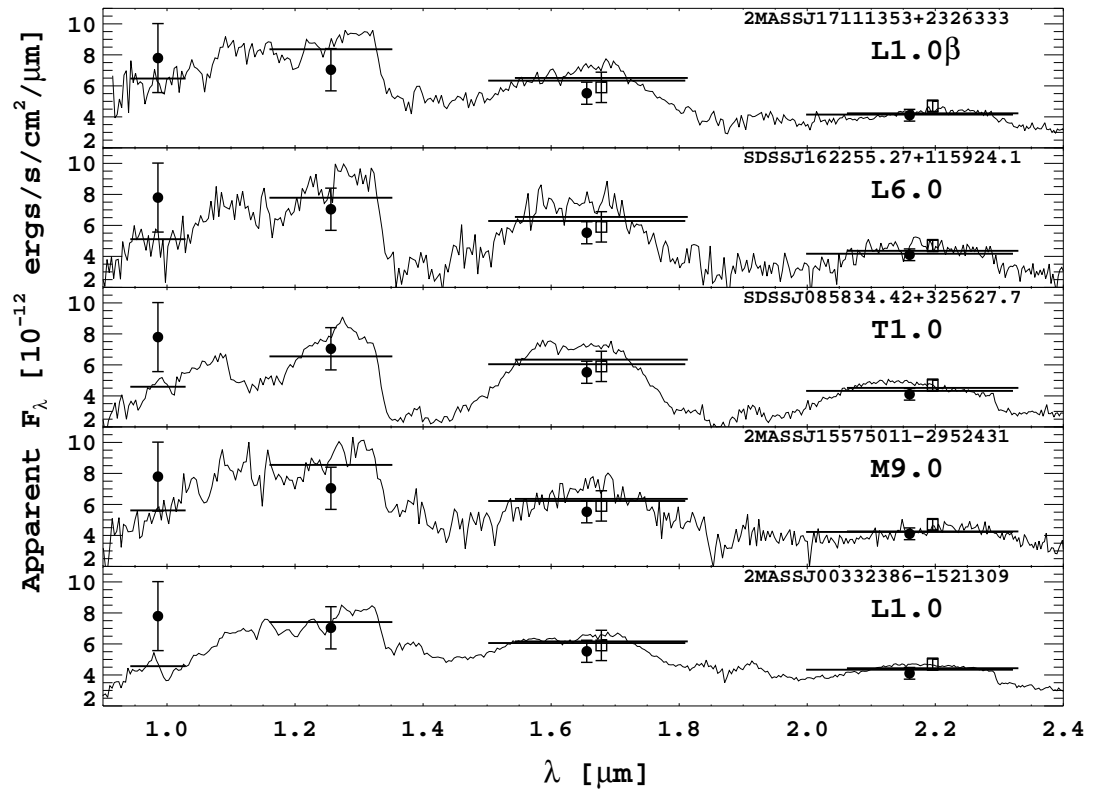


Figure 4.14 The field brown dwarfs which best match the photometry of β Pic b.

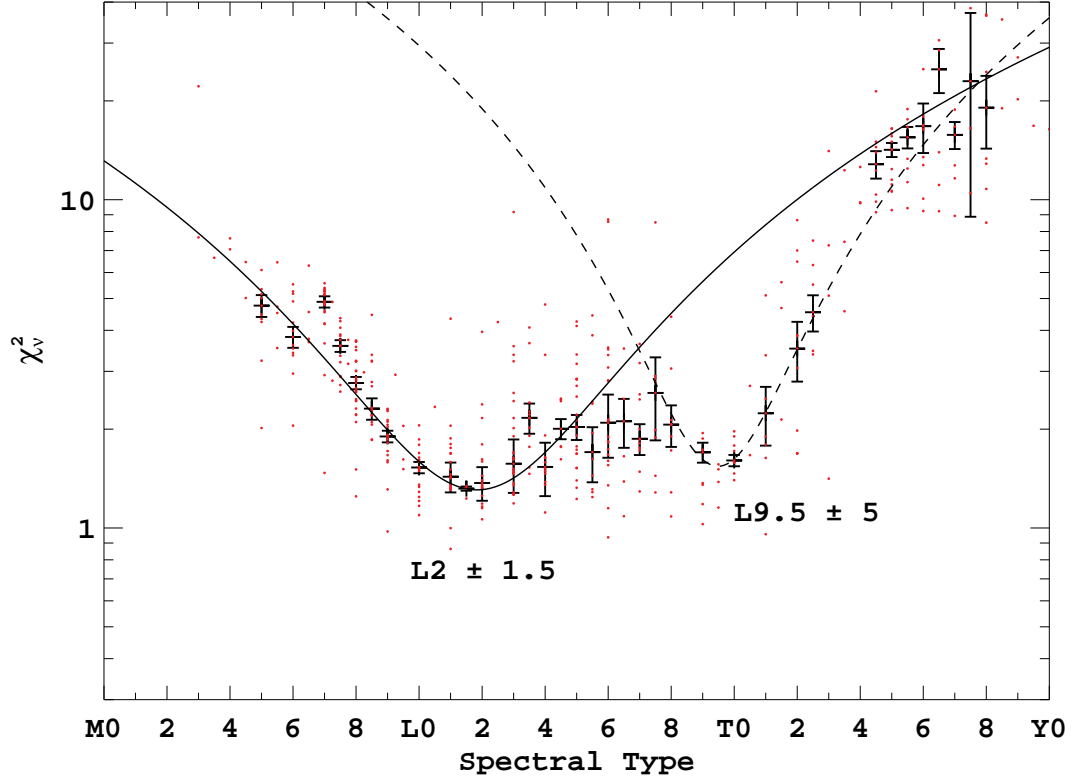


Figure 4.15 β Pic b spectral type. The red points are individual field dwarfs, and the black crosses with error bars show the median and standard deviation in each half spectral type. We fit parabolas to the two apparent minima to find the best-fit spectral type.

We collated the results of these fits by spectral type (to 0.5 types), shown in Figure 4.15. The median χ^2_ν for each half spectral type is also shown, with error bars determined by the standard deviation in each type. There are two minimums, one for early L and one for late L/early T. We fit each of these with a parabola. The results are $L1.8 \pm 1.2$ and $L9.5 \pm 4.8$. The L2 minimum is a better fit, so we adopt $L2 \pm 1.5$ as the best fit spectral type. This is consistent with $L2\gamma \pm 2$ adopted by Bonnefoy et al. (2013).

To test the effect that each filter has on the spectral type fit, we repeated the fit with one bandpass excluded. These results are shown in Figure 4.16. It is evident

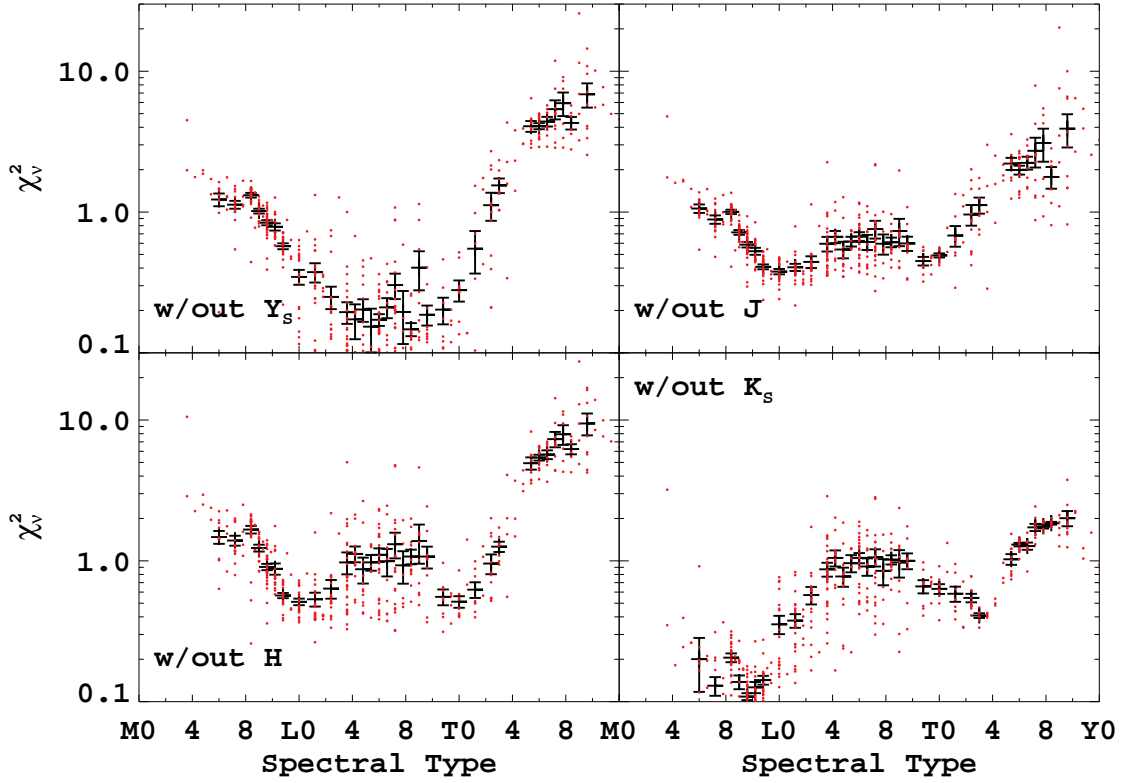


Figure 4.16 β Pic b spectral type without various filters. This shows that our Y_S measurement pushes the fit to earlier L types.

that our Y_S result is pushing the fit to earlier spectral types. Dropping either Y_S or K_S results in a dramatically lower χ^2_ν . Excluding either J or H does not significantly change the results. So we conclude that the best-fit spectral type is primarily being determined by Y_S and K_S .

Finally, we use our best-fit spectral type to place β Pic b on Y_S , J , and H color-magnitude diagrams. We show absolute magnitude in these filters vs. both $X - K_S$ color spectral type in Figure 4.17. As in the color-color plots, β Pic b is consistent with an early L dwarf. We also show the same three planet/planet-like comparison objects. We placed HR 8799b and 2M1207b on the spectral type diagram using their best fit spectral temperatures from Madhusudhan et al. (2011) and Barman

et al. (2011), respectively, and the spectral type to temperature relationships given by Stephens et al. (2009). These plots once again illustrate the wide diversity in the directly imaged exoplanets. Despite similar ages these objects have very different colors and luminosities

4.7 Conclusion and Future Work

We have presented the first ground-based detection of an exoplanet with a CCD. With its groundbreaking VisAO camera, the MagAO system detected β Pictoris b at $0.986\mu\text{m}$, with a contrast of 3.0×10^{-5} . This is certainly the shortest wavelength direct image of an exoplanet from the ground, and given the latest results indicating Fomalhaut b may not be a true planet, this could be the bluest image of an exoplanet to date. This results represents a new frontier in exoplanet science as we push into the visible.

We compared the $Y_S J H K_S$ photometry of β Pictoris b with field brown dwarfs. In color-color and color-magnitude plots, β Pic b is consistent with being an early type L dwarf. Using nearly five hundred objects with Y through K spectroscopy, we found a best-fit spectral type of $L2 \pm 1.5$. We also compared these results to the other directly imaged exoplanets with Y photometry HR 8799 b, 2M1207b. These objects are remarkably different from β Pictoris b at these wavelengths.

We made full use of MagAO's unique O/IR wavelength coverage, observing this exoplanet from $0.98\mu\text{m}$ to $5\mu\text{m}$. Our Y_S photometry is brighter than expected for a typical L dwarf given prior measurements. There are a few more verifications that we plan to conduct, but this result appears to be robust. Though we do not address it here, there are other, soon-to-be published, measurements in the near-IR using the NICI camera which are also bright compared to the others (Mike Liu, private communication). There may be more to this story, such as variability. We leave a comprehensive multi-spectral modeling effort to Morzinski et al. (2013, in prep). For now, we have shown the power of blue-wavelength imaging of exoplanets.

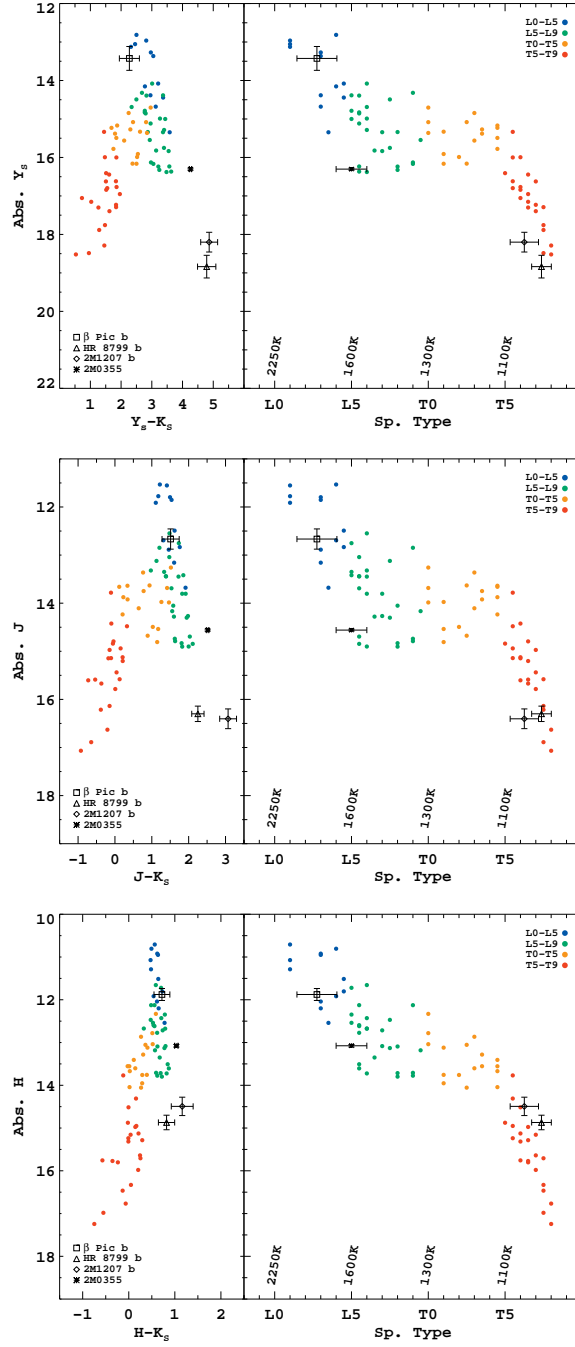


Figure 4.17 Color-magnitude diagrams. Y_s photometry, NACO photometry, and temperature equivalent spectral type for HR 8799b and 2M1207b estimated as described in the text.

CHAPTER 5

DETECTABILITY OF EGPS IN THE HZ

In this chapter we begin to consider what comes next. The unique thermal-IR and visible wavelength coverage of the high order Magellan and LBT adaptive optics systems allows the first direct imaging search of the habitable zone (HZ) of very nearby stars, where *Kepler* hints that planets are common.

Typical direct imaging searches can be characterized as searching for young, self-luminous, giant planets at wide separations. With the unique wavelength coverage offered by MagAO and the LBT AO systems, we can search nearby old stars at close separations, moving the hunt for EGPs into the HZ. Extrasolar gas-giant planets (EGPs) are roughly the same radius regardless of mass, and in the HZ irradiation will set the minimum planet temperature, so minimum HZ EGP thermal luminosity is nearly independent of mass and age. Similar logic applies to reflected visible light. In short, EGPs should be very detectable in the HZs of nearby stars. A search for such planets directly addresses the goals of Astro2010/New Worlds (Blandford, et al., 2010), and developing the ability to spectrophotometrically characterize HZ planets is of paramount importance to the search for life—the primary goal of NASA’s Exoplanet Exploration Program. The next few years should be very exciting.

5.1 The Habitable Zone

The HZ is generally defined as the region around a star where liquid water can exist on a planet’s surface. Though this is a simple statement, determining the location of this region is difficult as it depends on many factors such as the composition of the atmosphere, the level of plate tectonics and volcanism, and the time history of the star’s luminosity. In any case, it seems obvious that the most important quantity defining whether water is liquid will be the temperature of the planet. As a simplified

starting point we can assume that this is set by the equilibrium temperature

$$T_{eq} = \left[\frac{1 - A_B}{f} \right]^{1/4} 278.5 \text{K} \left(\frac{L}{L_\odot} \right)^{1/4} \left(\frac{r}{1 \text{AU}} \right)^{-1/2} \quad (5.1)$$

where A_B is the bond albedo and f is the fraction of the surface area where the heat is uniformly distributed. A_B is defined as the amount of incoming radiation reflected by the planet, over all wavelengths. L is the luminosity of the star in solar units, and r is the distance of the planet from the star. Rearranging we have a convenient scaling of distance for a given temperature and stellar luminosity

$$\frac{r(T_{eq})}{1 \text{AU}} = \left[\frac{1 - A_B}{f} \right]^{1/2} \left(\frac{278.5 \text{K}}{T_{eq}} \right)^2 \left(\sqrt{\frac{L_*}{L_\odot}} \right)^{1/2} \quad (5.2)$$

and most simply

$$r_{T_{eq}} \propto \sqrt{L_*}. \quad (5.3)$$

As a first guess at defining the HZ we can calculate the distance from a $1L_\odot$ star at which water boils at one Earth atmosphere of pressure is ($T = 373\text{K}$)

$$r_{boil} = 0.56 \text{ AU}$$

for $A_B = 0$ and $f = 1$. The distance from a $1L_\odot$ star at which water freezes ($T = 273\text{K}$) is

$$r_{freeze} = 1.04 \text{ AU}$$

It is of course more complicated than this. The seminal work on defining the location of the HZ is Kasting et al. (1993), who modeled an Earth-like atmosphere orbiting a Sun-like star. They argued that the inner edge of the HZ is set by water loss through the process of photodissociation of H_2O and the subsequent escape of H_2 to space. In this model the outer edge was set by the formation of CO_2 clouds in the planet atmosphere, which have a cooling effect. A major consideration within the HZ is the presence of CO_2 and its impact on temperature through the greenhouse effect. CO_2 concentration is modulated by weather, volcanism, and by temperature itself. Kasting et al. (1993) give a “conservative” HZ of 0.95-1.37 AU, quite different from the blackbody boiling-freezing guess above.

A further consideration is that stellar luminosity is not constant with time. Taking this into account, and considering conditions on Venus in the distant past (when it may have been wet), and Mars in the distant past (when it may also have had standing water), Kasting et al. (1993) give the “early-Venus” and “early-Mars” HZ estimate of 0.75-1.77 AU.

The Kasting et al. (1993) HZ model has been updated and adapted by many authors, but was most recently updated in Kopparapu et al. (2013), using essentially the same model but with a new H_2O and CO_2 line list. The new conservative HZ estimate is 0.99 - 1.70 AU, placing Earth at the inner edge. Another recent analysis by Zsom et al. (2013) argues that for a planet with low relative humidity, what they call a desert world, the inner edge of the HZ could be as close as 0.5 AU. Zsom et al. (2013) also provocatively argue (as have others before) that the outer edge of the HZ is effectively infinite, as other processes such as tidal heating can generate enough heat to keep water liquid even without a main sequence star.

For this analysis we will assume the “wide HZ” of (Traub, 2012), with separation $r = 0.72\text{-}2.0\text{AU}$ for stellar luminosity $L_* = 1L_\odot$. This very optimistic HZ definition, which is actually due to Kasting, was designed to ensure that NASA’s *Kepler* mission does not miss any potentially habitable worlds (Traub, 2012). We show this “wide HZ” for a sample of 12 nearby stars in Figure 5.1. These stars were chosen because part or all of their HZs project to $\gtrsim 300$ mas, making them accessible to MagAO and LBTAO.

5.2 Are there planets in the HZ?

Having agreed on a definition of the HZ, we might then ask the question “but are there planets in the HZs of other stars?”. One way to answer this question is to calculate the fraction of stars in the sky with a planet in the HZ. If this question is narrowed to consider only Earth-like planets, this quantity is frequently called eta-Earth, η_\oplus . Traub (2012) extrapolated from the first 136 days of *Kepler* data

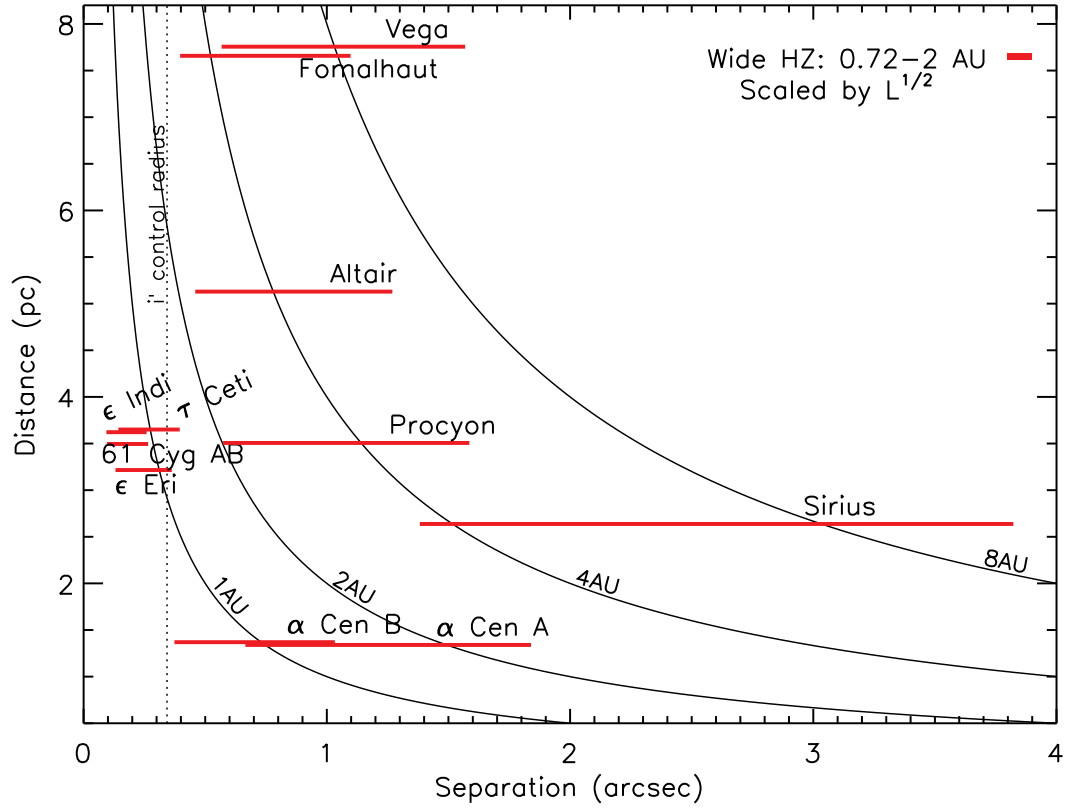


Figure 5.1 Projected HZs of select nearby stars. For these stars, the projected separation of the HZ makes it accessible with the high order LBT & MagAO systems in the thermal-IR and the visible. MagAO/Clio2/VisAO and LBTI can credibly search nearby HZs for EGPs with mass as low as $\sim 0.1M_{Jup}$.

that the frequency of earth-like planets around sun-like stars is

$$\eta_{\oplus} = 0.34.$$

Though this should be treated with caution since it is such a wide extrapolation, it stands as one of our best estimates of η_{\oplus} to date. More recent updates to the *Kepler* candidate catalog have been generally consistent with this analysis.

Now if we extend Traub (2012)’s analysis to *all* planet sizes, we find that

$$\eta_{planet} = 1.2.$$

In other words most stars have at least 1 planet in the HZ¹. For EGPs, with $R_P > 8R_{\oplus}$, $\eta_{EGP} = 0.11$, so ~ 1 in 10 stars will have an EGP in the HZ. This result is broadly consistent with radial velocity (RV) studies (Wittenmyer et al., 2011a,b; Mayor et al., 2011). Star-by-star RV completeness was given by Wittenmyer et al. (2011b), e.g., for α Cen A they are 50% complete to $0.12M_{Jup}$ and 90% complete to $0.26M_{Jup}$ – nearly the mass of Saturn – at $1.5\sqrt{L}$ AU. Similar survey-wide limits were given by Mayor et al. (2011). Planets should be common in the HZ of nearby stars, yet RV has not ruled out planets of nearly the mass of Saturn at 1.5AU around nearby stars.

The recently discovered $\sim 1M_{\oplus}$ planet α Cen Bb (Dumusque et al., 2012), though not in the HZ, has exciting implications—it begins to address remaining questions about the formation and stability of planets around binary stars (cf. Eggl et al. (2013) and references), and it is likely that there are multiple planets in the system (Lissauer et al., 2011). α Cen Bb makes a search of the HZ of α Cen A&B and other nearby stars highly compelling.

5.3 The Radius and Temperature of a Giant Planet

Old EGPs should be easily detectable in the HZ in the thermal-IR (3–5 μ m). The key result leading to this assertion is that EGPs with masses ranging from $\sim 0.1M_{Jup}$ to

¹For a completely different way to arrive at this estimate, Bovaird and Lineweaver (2013) claim that the Generalized Titius-Bode rule predicts that the average number of planets in the HZ should be 1 to 2

$10M_{Jup}$ are the same size to within 20%. At higher planetary masses, $M \gtrsim 3M_{Jup}$, electron degeneracy pressure supports the internal structure of the planet. Below $3M_{Jup}$, coulomb pressure supports the planet's atmosphere. The combination of these two effects means that the radius of the planet depends only weakly on the planet's mass. A useful scaling for $M \lesssim 3M_{Jup}$ is $R_P \propto M_P^{1/10}$, where R_P is planet radius and M_P is mass (Fortney et al., 2011).

The total luminosity of an EGP will closely vary as

$$L_P \propto R_P^2 T_P^4.$$

Where T_P is the planet temperature. Extrasolar planet direct-imaging search we have typically looked for young planets, still cooling. Their effective temperature, T_{eff} , is then set by their age and mass as they radiate the gravitational potential energy from formation. In addition to age, then, mass also plays an important role in setting temperature. Combined with the constraints imposed by the limitations of current generation high-contrast imagers, the planets imaged to-date have been young, massive ($\sim 10M_{Jup}$), and at wide physical separations.

Closer to the star, however, the situation is somewhat different. Here the radiation from the star also plays a role in setting the planet temperature. The minimum temperature is set by Equation (5.1) regardless of mass and age, so once the planet cools to the point that $T_{eff} \leq T_{eq}$, its luminosity will be set by T_{eq} and R_P . Now since R_P only weakly depends on mass, in this regime L_P itself only weakly depends on mass. **In the HZ the minimum EGP thermal luminosity is nearly independent of mass and age.**

Armed with this result, and an imaging system capable of high enough contrast in the HZ of a star, we should be able to detect EGPs of very low masses. It is then interesting to consider how low the mass-radius relationship for giant planets extends. This was recently addressed by Batygin and Stevenson (2013), who analyzed the possible composition of *Kepler*-30d, an $8.8 \pm 0.5 R_{\oplus}$ planet orbiting at a distance such that $T_{eq} = 364\text{K}$. They conclude that with a sufficiently large core ($3 - 5M_{\oplus}$), EGPs with H/He dominated atmospheres are stable under moderate irradiation

(as in the HZ) down to even a few Earth masses. Their models indicate EGPs will have $R_P > 8R_\oplus$ down to $M_P \sim 10M_\oplus = 0.03M_{Jup}$. An important implication of this result is that planets with Radii in the $R_P \sim 8R_\oplus$ regime could have very low masses, and such planets — based on current *Kepler* results — might be common.

5.4 Thermal Infrared Brightness of EGPs

To determine how bright an EGP is in the HZ, we can use the non-irradiated COND models (few, if any, other model grids extend to low enough temperatures). Our goal is to calculate the flux in a certain bandpass as a function of distance from the star, allowing temperature to vary with separation until $T_{eq} < T_{eff}$. Now the typically cited “cooling curves” or isochrones of Baraffe et al. (2003) can’t be used as is. At young ages, the planets are still contracting so their radii are inflated compared to a planet of the same mass at an older age. This means that we can’t simply interpolate magnitudes across T_{eff} on the grid of isochrones at a given mass, as the brightness will be too high due to the larger radius.

We downloaded the COND spectra, corresponding to the AMES-2000 line list². These are parameterized by T_{eff} and surface gravity $\log(g)$. The first step is then to use the isochrones of Baraffe et al. (2003) to find the expected radius of a planet of a given mass and age. For planets with mass lower than the lowest mass given in the isochrone, we extrapolate $R_P \propto M_P^{1/10}$. We can then use radius to find the surface gravity of the planet through

$$g = 2478.50 \left(\frac{M_P}{1M_{Jup}} \right) \left(\frac{1R_{Jup}}{R_P} \right)^2 \text{ [cgs]}.$$

Next, at a given separation from the star we calculate T_{eq} , assuming $A_B = 0.34$ (typical of planets in our solar system) and $f = 1$. We interpolate on the grid of spectra with $\log(g)$ and setting $T_{eff} = T_{eq}$, and apply the dilution factor correction to determine the flux measured at 10 pc. We then convolve the resulting spectrum with the M' bandpass of Clio (MKO M'). We also convolve the Vega spectrum

²<http://phoenix.ens-lyon.fr/Grids/AMES-Cond/>

of Bohlin (2007) with M' , and so determine the absolute magnitude in M' for the planet.

The final step is to apply the distance modulus for the star being considered. We show the results of these calculations for α Cen A in Figure 5.2 and for Sirius in Figure 5.3. Though we have very recently observed these stars at M' with MagAO, we have not yet derived a contrast curve. Instead, we use a contrast curve measured at the LBT with LMIRCAM³, and scale it for the smaller diameter Magellan Clay primary. The resulting 1, 5, and 10 hour detection limits are over-plotted in the figures.

With MagAO, and with LBTAO, we should be sensitive to HZ EGPs with $M \gtrsim 0.1 M_{Jup}$. It is worth noting that JHK imaging (e.g. GPI, Project 1640) cannot detect such planets. The major caveat is that we assumed photon-noise \sqrt{N} scaling for both D and exposure time. Due to speckle-noise, achieving \sqrt{N} scaling over 10hrs is challenging. It is one of the main goals of my future work to test the limits of \sqrt{N} scaling, and to develop active techniques to suppress speckles in this regime.

An additional question that should be addressed regards the choice of M' vs L' . An important consideration is the impact that Methane absorption has on L' brightness of EGPs. Exoplanets detected to date have generally been brighter than expected in bandpasses where methane should make them faint, but these results are for much warmer atmospheres than we consider here. The COND models predict (roughly) that these low-mass EGPs will be approximately 2 magnitudes fainter at L' than at M' . At the contrasts needed, around such bright stars, we do not expect to be background limited at these tight separations. That means that we should seek to maximize Strehl ratio to minimize speckle noise, and should maximize the number of photons collected to maximize signal-to-noise ratio. Taken together, these arguments favor M' despite the higher thermal background. For an example of AO imaging with high thermal backgrounds, see Appendix C.

³provided by Vanessa Bailey

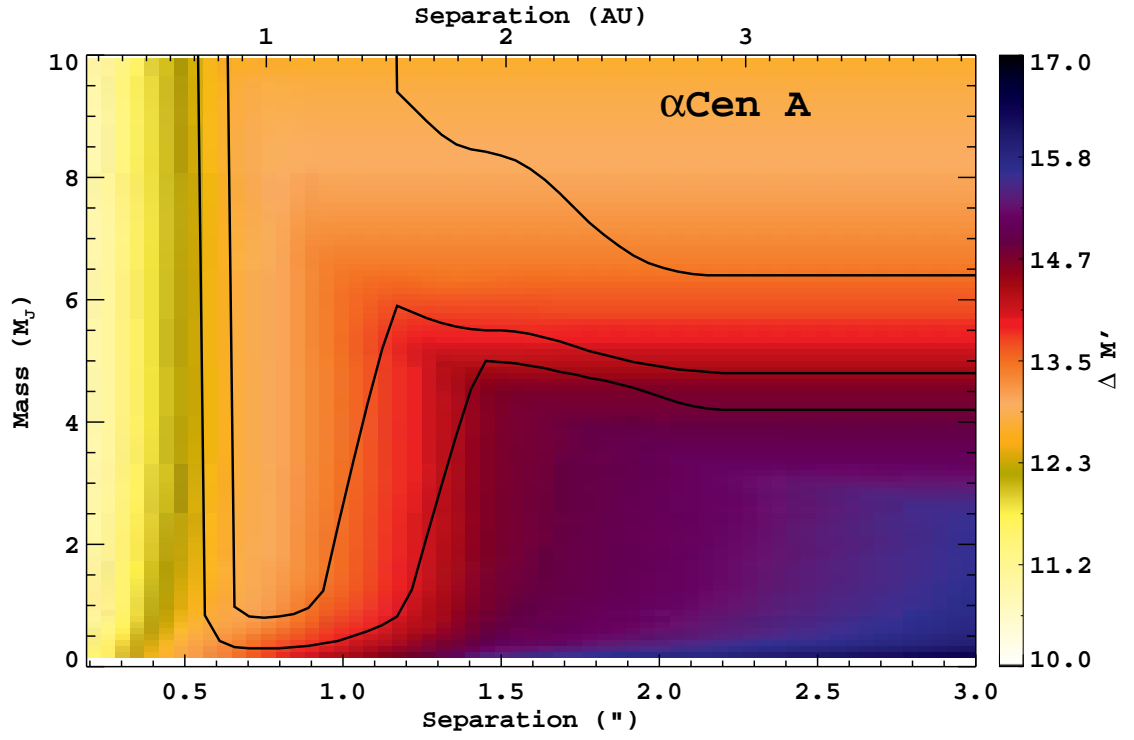


Figure 5.2 The contrast of EGPs around α Cen A using the COND spectra and our assumptions about T_{eq} . The calculations were cut off at an arbitrary $0.1M_{Jup}$. The contrast curves are for 1, 5, and 10 hrs with MagAO/Clio. The required $\Delta M' \approx 14$ required to detect very low mass planets is very achievable.

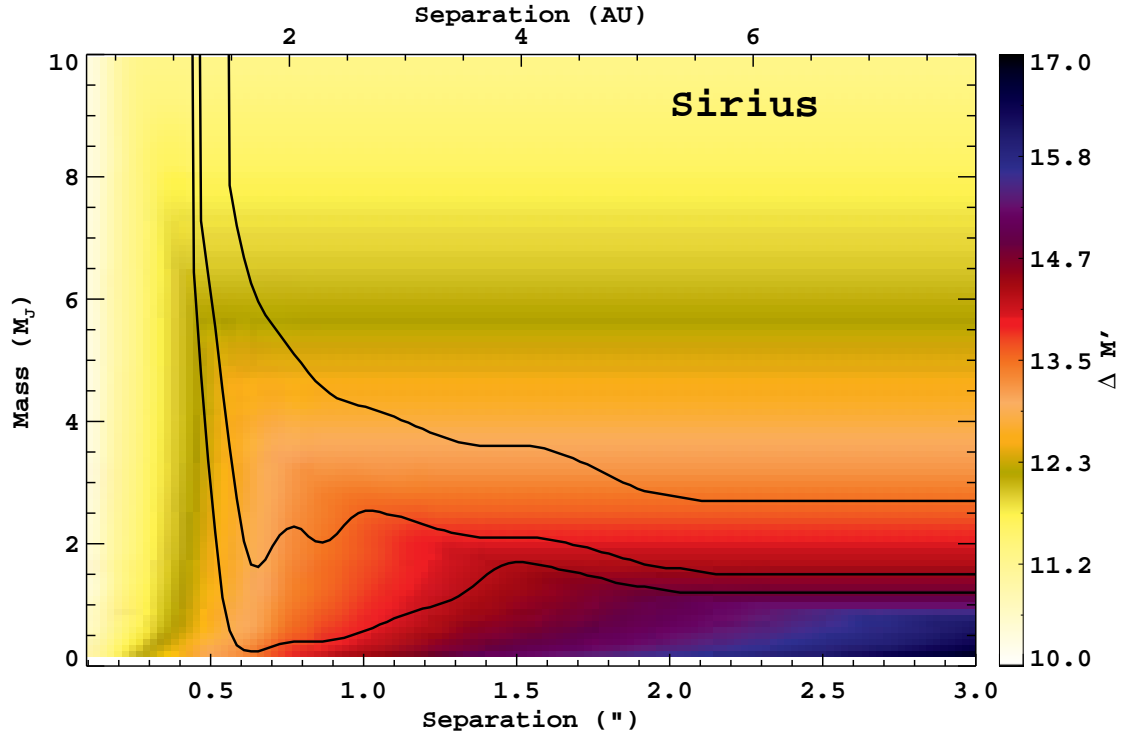


Figure 5.3 The contrast of EGPs around Sirius using the COND spectra and our assumptions about T_{eq} . The calculations were cut off at an arbitrary $0.1M_{Jup}$. The 5σ contrast curves are for 1, 5, and 10 hrs with MagAO/Clio.

5.5 Habitable Zone EGPs in the Visible

Next we consider the brightness of EGPs in the visible ($0.5\text{--}1.0\mu\text{m}$), where reflected starlight will be the dominant source of photons. The reflected flux from a planet, F_p , relative to the flux from its star, F_* , is

$$\frac{F_p}{F_*} = A_g \left(\frac{R_p}{r} \right)^2 \Phi(\alpha)$$

where A_g is the geometric albedo, R_p is the planet radius, r is the distance of the planet from the star, Φ is the phase function, and α is the phase angle. Putting in the appropriate constants gives

$$\frac{F_p}{F_*} = 1.818 \times 10^{-9} A_g \left(\frac{R_p}{1R_\oplus} \times \frac{1 \text{ AU}}{r_p} \right)^2 \Phi(\alpha) \quad (5.4)$$

Cahoy et al. (2010) modeled the atmospheres of Jupiter and Neptune-like planets, and present the resulting geometric albedos for such planets at a range of separations from a star. EGPs should have high geometric albedo near $r \approx 2\text{AU}$ due to water clouds, but these will evaporate closer to the star. The competing effects in Equation (5.4) of $1/r^2$ and A_g increasing with r cause a peak in EGP contrast to occur between 1 and 2 AU. EGP reflectivity peaks in the HZ, with only weak mass dependence. That is, the same low-mass EGPs we discussed above should be at their brightest in the HZ.

Assuming that the main driver of changes in A_g as separation changes is the change in T_{eq} , we can adapt the models of Cahoy et al. (2010) to stars of various luminosities. We show the contrast of EGPs orbiting α Cen A&B and τ Cet in Fig. 5.4, based on A_g from Cahoy et al. (2010), and predicted contrast for ϵ Eri b using its astrometric orbit (Benedict et al., 2006) (periastron in early 2014). These calculations fully include the effects of orbital phase modeled by Cahoy et al. (2010), rather than the simplified Lambert phase function.

In Figure 5.4 we plot the halo-photon-noise (\sqrt{N}) contrast limit for MagAO/VisAO at i' derived from our test-tower PSF measurements. We also show the \sqrt{N} limit derived from the PCA reduction of the β Pic observations described

earlier. In both of these cases we have made the following assumptions: 10 hours of open-shutter time on a -0.5 mag star, with a broadband $r'i'z'$ filter, in the 950 beamsplitter (meaning all blue photons go to the CCD-47). The scaling of the β Pic contrast curve assumed a factor of 10 increase in system throughput compared to Y_S . In all cases these assumptions were applied to the measurements assuming \sqrt{N} scaling.

Achieving such high contrasts will be very difficult, even at wide ($> 20\lambda/D$) separations. These observations will not be photon-noise limited! Rather, we expect to reach a speckle noise limit before reaching the \sqrt{N} limit. However, this regime (30-50% visible Strehl ratio (SR) on a 6.5m telescope) is completely new and unexplored: the nearest sun-like stars have not been imaged at close separations in visible light. Testing the limits of AO corrected high-contrast imaging in the visible is a major focus of my future research.

5.6 Blazing The Trail

In the next few years will use MagAO and LBTI to perform the first credible search for EGPs in the HZ of nearby stars. Our thermal-IR observations may lead to the first images of an exoplanet in the HZ. Even with no detections, we will provide constraints on HZ planets in parameter space not yet reached by RV. Our VisAO observations will be an important step towards detecting exoplanets in reflected light. This search has already begun: during MagAO commissioning run 2 we observed both Sirius and α Cen A in i' and M' .

The observations discussed here are limited to a handful of very nearby, bright stars, and will be sensitive only to EGPs with $R_P \sim 8R_\oplus$. However, on the next generation of giant telescopes D^2 will improve our sensitivity from $R_P \approx 8R_\oplus$ to $R_P \leq 2R_\oplus$, allowing detection and characterization of rocky, potentially habitable, planets. While probing the HZ of these nearby stars we will develop and test ExAO control strategies for both the thermal-IR and visible wavelength regimes, using MagAO and LBTAO. This work is an important precursor to future ground-based

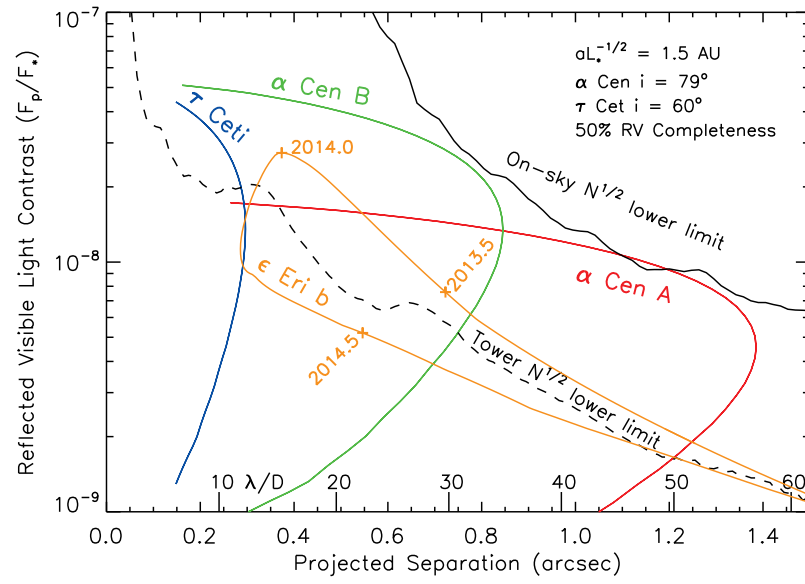


Figure 5.4 Visible reflected light contrasts for hypothetical EGPs on circular orbits around α Cen A&B and τ Ceti, and ϵ Eri b contrast based on the astrometric orbit (Benedict et al., 2006). Predictions include effects of separation and orbital phase (Cahoy et al., 2010). Radius was scaled by $R_P \propto M_P^{1/10}$, and mass was set by RV 50% completeness.

searches for rocky habitable planets such as that proposed by Guyon et al. (2012).

CHAPTER 6

DIRECT IMAGING IN THE HABITABLE ZONE AND THE PROBLEM OF ORBITAL MOTION

6.1 Introduction

In this chapter we turn to the far future, and consider the problem of detecting planets orbiting within the HZ using the next generation of giant telescopes. This chapter has previously been published in the *Astrophysical Journal* as Males et al. (2013).

Orbital motion (Kepler, 1609) has been used in one fashion or another to detect planets around stars other than our Sun in large numbers. The radial velocity (RV) technique monitors the Doppler shift of a stellar spectrum as the star itself orbits the planet-star center of mass, thus allowing us to infer the presence of a planet. Similarly, the astrometry technique monitors the motion of the star on the sky and likewise infers the presence of a planet. The transit technique monitors the reduction in brightness of the star as the orbiting planet temporarily crosses the line of sight between the telescope and the star.

Unlike these indirect techniques, direct imaging detects light from the planet itself and spatially resolves it from the light of the star (Traub and Oppenheimer, 2011). The extreme difference in brightness between star and planet at small projected separations has generally limited direct imaging efforts to wide separations where orbital motion is ignorable. The next generation of large telescopes will move us into a new regime of direct imaging, moving closer to the star. We will even be able to begin probing the liquid water habitable zone (HZ). Here we point out that at these tight separations orbital motion will no longer be negligible in direct imaging. As we will show the motion of planets in the HZ (and closer), during the required integration times, will be large enough to limit our sensitivity unless we

take action to correct it.

In Section 6.2 we present our motivation for this study and briefly review some of the related prior work. In Section 6.3 we develop the basic tools needed to analyze this problem, including the expected speed of orbital motion in the focal plane and the effect it has on signal-to-noise ratio (SNR). In Section 6.4 we analyze the impact orbital motion will have on a search of α Cen A by the Giant Magellan Telescope (GMT) working at $10\mu\text{m}$, and propose a method to mitigate this impact by de-orbiting a sequence of observations. Then in Section 6.5 we treat the more favorable case of a cued search, where we have prior information from an RV detection. To do so we analyze the case of the potentially habitable planet Gl 581d being observed by the planned European Extremely Large Telescope (E-ELT). Finally, in Section 6.6, we present our conclusions and prospects for future work.

6.2 Motivation and Related Work

Moving the hunt for exoplanets into the HZ of nearby stars marks a departure from prior efforts. Here we briefly discuss the definition of the HZ, review direct imaging results to date, discuss the differences between them and future efforts, and finally review some closely related prior work.

6.2.1 Nearby Habitable Zones

The HZ is generally agreed to be the region around a star where a planet can have liquid water on its surface. This is far from simply related to the blackbody equilibrium temperature, as it depends on atmospheric composition and the action of the greenhouse effect (Kasting et al., 1993; Kopparapu et al., 2013), among other factors. For our purposes it is enough to assume that the HZ is generally located at about one AU from a star, scaled by the star’s luminosity

$$a_{HZ} \approx \sqrt{L_*/L_\odot} \text{ AU}. \quad (6.1)$$

Traub (2012) provided three widths for the HZ based on various considerations, and then used the first 136 days of data from the *Kepler* mission to estimate that

the fraction of sun-like stars (spectral types FGK) with an earth-like planet in the HZ is $\eta_{\oplus} \approx 0.34$. More generally, this analysis indicates that $\eta_{planet} \approx 1.2$, implying that *every sun-like star is likely to have a planet in its HZ, and some will have more than one*. While this exciting result is based on a very large extrapolation from the earliest *Kepler* results, it is currently one of our best estimates of planet frequency in the HZ.

This topic was recently brought to the fore with the announcement of α Cen Bb by Dumusque et al. (2012). Discovered using the RV technique, α Cen Bb is an $m \sin i = 1.13 M_{\oplus}$ planet orbiting a K1 star at 0.04 AU. While certainly not in the HZ, this discovery has exciting implications for the presence of planets in the HZ of the nearest two sun-like stars.

The above arguments hint that planets will be common in the HZ of sun-like stars. We are about to enter a new era of exoplanet direct imaging. With the next generation of giant telescopes and high-performance spaced-based coronagraphs we will be searching for planets in this scientifically important region around nearby stars.

6.2.2 A Different Regime

The typical search for exoplanets with direct imaging has used 2.4m (Hubble Space Telescope, HST) to 10m (Keck) telescopes. These surveys have mostly concentrated on young giant planets, which are expected to be self-luminous as they dissipate heat from their formation. This allows them to be detected at wider separations from their host stars, where reflected starlight would be too faint. This has also caused planet searches to typically work at H band ($\sim 1.6\mu\text{m}$), with exposure times of ~ 1 hr. Examples conforming to these survey archetypes include Lowrance et al. (2005) using HST/NICMOS; the Gemini Deep Planet Search (Lafrenière et al., 2007a); the Simultaneous Differential Imaging survey using the Very Large Telescope and MMT (Biller et al., 2007); the Lyot Project at the Advanced Electro-Optical System telescope (Leconte et al., 2010); the International Deep Planet Survey (Vigan et al., 2012); and the Near Infrared Coronagraphic Imager at Gemini South (Liu et al.,

2010).

These searches have had some success. Examples include the 4 planets orbiting the A5V star HR 8799 (Marois et al., 2008b, 2010), with projected separations of 68, 38, 24, and ~ 15 AU. These correspond to orbital periods of ~ 460 , ~ 190 , ~ 100 , and ~ 50 years, respectively. The A5V star β Pic also has a planet (Lagrange et al., 2010) orbiting at ~ 8.5 AU with a period of ~ 20 years (Chauvin et al., 2012b). Another A star, Fomalhaut, has a candidate planet on an 872 year (115 AU) orbit (Kalas et al., 2008). At these wide separations it takes months, or even years, to notice orbital motion.

In the much closer HZ, however, orbital periods will be on the order of one year. We show in some detail that this is fast enough to yield projected motions of significant fractions of the point spread function (PSF) full width at half maximum (FWHM) over the course of an integration. The resulting smeared out image of the planet will have a lower SNR , making our observations less sensitive.

6.2.3 Long Integration Times

In addition to HZ planets having higher orbital speeds than the current generation of imaged exoplanets, integration times required to detect them will be much longer. Direct imaging surveys to date have mostly worked in the infrared while attempting to detect young planets still cooling after formation. The coming campaigns to image planets in the HZ of nearby stars will focus on older planets, which will be less luminous in the near infrared. In the HZ, starlight reflected from the planet will be more important. The result is integration times required to detect such planets will be tens of hours, rather than the ~ 1 hour characteristic of current campaigns.

Consider the Exoplanet Imaging Camera and Spectrograph (EPICS), an instrument proposed for the E-ELT. Kasper et al. (2010) predicted that EPICS will be able to image the RV detected planet Gl 581d, which has a semi-major axis of 0.22 AU with a period of ~ 67 days (Forveille et al., 2011; Vogt et al., 2012). This orbit places it on the outer edge of the HZ of its M2.5V star (von Braun et al., 2011). EPICS will be able to detect Gl 581d, at a planet/star contrast of 2.5×10^{-8} , in 20

hrs with $SNR = 5$ (Kasper et al., 2010). Since this is a ground based instrument, a 20 hour integration will be broken up over at least 2 nights. Plausible observing scenarios could extend this to several nights, taking into account such things as the need for sky rotation. As we will show, the planet will move several FWHM on the EPICS detector during a multi-day observation.

More generally, Cavarroc et al. (2006) showed that when realistic non-common path wavefront errors are taken into account, the integration times required to achieve the 10^{-9} to 10^{-10} contrast necessary to detect an earth-like planet around a sun-like star approach 100 hours on the ground, even on a 100m telescope with extreme-AO and a perfect coronagraph. One of several concerns about the feasibility of a 100 hour observation from the ground is that such a long observation will be broken up over many nights.

With net exposure times of 20 to 100 hrs, and total elapsed times for ground based observations of several to tens of days, HZ planets will move significantly over the course of a detection attempt. The focus of this investigation is the impact of the orbital motion of a potentially detectable planet on sensitivity.

6.2.4 Related Work

Though it has not yet been a significant issue in direct imaging of exoplanets, orbital motion has been considered in several closely related contexts. Here we briefly review a select portion of the literature. A very similar problem has been addressed in the context of searching for objects in our solar system, such as Kuiper Belt objects (KBOs), which can have proper motions on the order of $1''$ to $6''$ per hour (Chiang and Brown, 1999). Blinking images to look for moving objects by eye is a well established technique. A more computationally intensive form of blinking images proceeds by shifting-and-adding a series of short exposures along trial paths, usually assumed to be linear. This “digital tracking” makes it possible to detect KBOs too faint to appear in a single exposure. This has been done both from the ground (Chiang and Brown, 1999; Yamamoto et al., 2008) and from space with HST (Bernstein et al., 2004). More recently Parker and Kavelaars (2010) have taken into

account nonlinear motion and optimized selection of the search space, especially important given the large data sets that facilities such as the Large Synoptic Survey Telescope will produce.

Orbital motion is an important consideration when planning coronagraphic surveys of the HZs of nearby stars. Brown (2005) treats the problem of completeness extensively. Large parts of the HZ will be within the inner working angle of a Terrestrial Planet Finder-Coronagraph (TPF-C) type mission, and so undetectable during a single observation. Also discussed in Brown (2005) is photometric completeness - that is how long the TPF-C must integrate on a given star to detect an earth-like planet in the HZ. Other work on this topic includes Brown and Soummer (2010) and Brown (2004). These analyses consider orbital motion only between observations, not during a single observation as we do here. In general, the scenarios considered for these studies involved space-based high-performance coronagraphs on medium to large telescopes. In such cases exposure times were short enough and continuous so that orbital motion should be negligible during a single observation.

The work most similar to our analysis here is the detection of Sirius B at $10\mu\text{m}$ by Skemer and Close (2011), in fact, it was part of our motivation for the present study. Skemer and Close (2011) used the well known orbit of the white dwarf companion to Sirius to de-orbit 4 years worth of images. Before accounting for orbital motion, Sirius B appeared as only a low SNR streak, but after shifting based on its orbit it appears as a higher SNR point source from which photometry can be extracted. Similar to this method, we will analyze the prospects for de-orbiting sequences of images, only we consider the case with no prior information at all, and with orbital elements with significant uncertainties.

6.3 Quantifying The Problem

In this section we will quantify the effects of orbital motion on an attempt to detect an exoplanet. Our first step will be to determine how fast planets move when projected on the focal plane of a telescope. Then we will illustrate the impact this

motion will have on the SNR and the statistical sensitivity of an observation.

6.3.1 Basic Equations

We begin by considering a focal plane detector working at a wavelength λ in μm . The FWHM of the PSF for a telescope of diameter D in m, neglecting the central obscuration, is

$$\text{FWHM} = 0.2063 \frac{\lambda}{D} \text{ arcsec}. \quad (6.2)$$

If we are observing a planet in a face-on circular (FOC) orbit with a semi-major axis of a in AU at distance d in pc, its angular separation will be a/d arcsec. At the focal plane the projected separation will then be

$$\rho = 4.847 \frac{aD}{\lambda d} \text{ in FWHM}. \quad (6.3)$$

We note that it will occasionally be convenient to specify ρ in AU instead of FWHM. When it is not clear from the context we will use the notation ρ_{au} to denote this.

The orbital period is $P = 365.25 \sqrt{a^3/M_*}$ days around a star of mass M_* in M_\odot . In one period, the planet will move a distance equal to the circumference of its orbit, $2\pi\rho$, so the speed of the motion in a FOC orbit will be¹

$$v_{FOC} = 0.0834 \left(\frac{D}{1\text{m}} \right) \left(\frac{1\mu\text{m}}{\lambda} \right) \left(\frac{1\text{pc}}{d} \right) \sqrt{\left(\frac{M_*}{1M_\odot} \right) \left(\frac{1\text{AU}}{a} \right)} \text{ in FWHM day}^{-1}. \quad (6.4)$$

In the general case, the equations of motion in the focal plane are

$$\begin{aligned} \dot{x} &= v_{FOC} \sqrt{\frac{1}{1-e^2}} \left[e \sin(f) (\cos(\Omega) \cos(\omega + f) - \sin(\Omega) \sin(\omega + f) \cos(i)) \right. \\ &\quad \left. - (1 + e \cos(f)) (\cos(\Omega) \sin(\omega + f) + \sin(\Omega) \cos(\omega + f) \cos(i)) \right] \\ \dot{y} &= v_{FOC} \sqrt{\frac{1}{1-e^2}} \left[e \sin(f) (\sin(\Omega) \cos(\omega + f) + \cos(\Omega) \sin(\omega + f) \cos(i)) \right. \\ &\quad \left. - (1 + e \cos(f)) (\sin(\Omega) \sin(\omega + f) - \cos(\Omega) \cos(\omega + f) \cos(i)) \right] \\ v_{om} &= \sqrt{\dot{x}^2 + \dot{y}^2} \end{aligned} \quad (6.5)$$

¹This result is equivalent to defining the gravitational constant in the focal plane as $G = (0.0834D/(\lambda d))^2$ and using the equation for speed in a circular orbit $v_{circ} = \sqrt{\frac{GM_*}{a}}$.

where Ω is the longitude of the ascending node, ω is the argument of pericenter, i is the inclination, and the true anomaly f depends on a , e , and the time of pericenter passage τ through Kepler's equation (Murray and Correia, 2010).

In Figure 6.1 we show the variation in projected orbital speed for both circular orbits at several inclinations, and face-on eccentric orbits ($i = 0$), for a planet orbiting a $1M_{\odot}$ star at 1 AU. In the plots we normalized speed to 1, and provide v_{FOC} for several interesting cases. These various scenarios produce projected orbital speeds of appreciable fractions of a FWHM per day. We will later show that, especially for ground based imaging, this causes a significant degradation in our sensitivity.

Our main focus here is on planets in the HZ. Our simple definition of the HZ results in $a_{HZ} \propto \sqrt{L_*}$. Now, on the main sequence mass and luminosity approximately follow scaling laws of the form $L_* \propto M_*^b$, where $b > 2$ except for very massive stars. So according to Equation (6.4) we expect v_{FOC} in the HZ to increase as M_* decreases, i.e. M stars will have faster HZ planets than G stars. For example, a planet in the HZ of α Cen B ($M_* = 0.9M_{\odot}$, $L_* = 0.5L_{\odot}$) will be moving roughly 20% faster than a planet in the HZ of α Cen A ($M_* = 1.1M_{\odot}$, $L_* = 1.5L_{\odot}$) (stellar parameters from Bruntt et al. (2010)).

To provide a more concrete example we return to the 20 hour observation of Gl 581d by the E-ELT/EPICS proposed by Kasper et al. (2010). Using a wavelength of $0.75\mu m$ with Equation (6.4) we find $v_{FOC} = 0.82$ FWHM per day, or a total of 0.68 FWHM for a continuous 20 hour observation. Since this is a ground based observation the actual amount of motion to consider is ~ 1.15 FWHM over the ~ 1.4 days minimum it would take to integrate for 20 hours. Were this a face-on orbit, an eccentricity of 0.25 (Forveille et al., 2011) would increase the maximum orbital speed to as much as 1.05 FWHM per day, or 1.47 FWHM minimum for a 20 hour ground based observation.

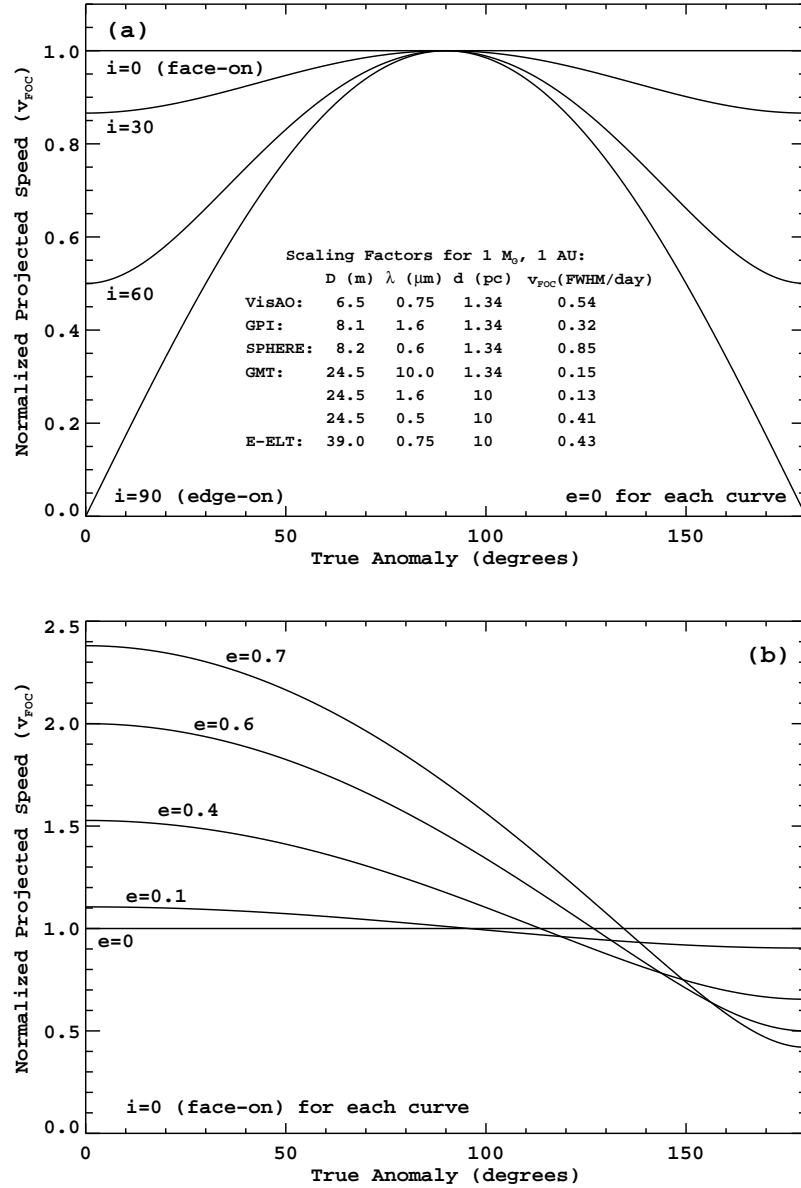


Figure 6.1 Magnitude of projected orbital speed, normalized to 1 FWHM day⁻¹, for 1 AU orbits around a $1M_{\odot}$ star. In (a) we show the orbital speeds for circular orbits at various inclinations, and in (b) we show the speeds for face-on orbits at various eccentricities. We give scaling factors in (a) for MagAO/VisAO (Close et al., 2012b), GPI (Macintosh et al., 2012), SPHERE/ZIMPOL (Roelfsema et al., 2010), GMT (Johns et al., 2012), and E-ELT/EPICS (Kasper et al., 2010). These scalings can be applied to the y-axis of either plot for various scenarios. These cases can also be scaled for different semi-major axes, telescopes, wavelengths, star masses and distances, by $v_{FOC} \propto \frac{D}{\lambda d} \sqrt{\frac{M_{*}}{a}}$. See the text for the general equations of motion for arbitrarily oriented eccentric orbits.

6.3.2 Impact on Signal-to-Noise Ratio

So what does the orbital motion calculated above do to our observations? To find out we consider a simple model of aperture photometry. Let us assume that we are conducting aperture photometry with a fixed radius r_{ap} , that the PSF is Gaussian, and that we are limited by Poisson noise from a photon flux N per unit area. With these assumptions, the optimum r_{ap} is 0.7 FWHM, but taking into account centroiding uncertainty $r_{ap} \approx 1$ FWHM is typical. We will approximate orbital motion at speed v_{om} by substituting $x \rightarrow x - v_{om}t - x_0$. Orbits are of course not linear, but this will be approximately valid over short periods of time. The parameter x_0 allows us to optimize the placement of the aperture to obtain the maximum signal, i.e. centering the aperture in the planet's smeared out flux. Note that with the exception of this centering parameter, this model appears quite naive in that we are not adapting the aperture radius and are pretending that we won't notice a smeared out streak in our images.

Now the SNR in the fixed-size aperture after time Δt will be

$$SNR_{fix} = \frac{\int_0^{r_{ap}} \int_0^{2\pi} \int_0^{\Delta t} I_0 e^{(-4 \ln 2 ((r \cos \theta - v_{om}t - x_0)^2 + r^2 \sin^2 \theta))} dt d\theta dr}{\sqrt{N \pi r_{ap}^2 \Delta t_{int}}} \quad (6.6)$$

where I_0 is the peak value of the PSF. In the case of no orbital motion $v_{om} = 0$ and aperture $r_{ap} = 1$ FWHM, so we have

$$SNR_o = \frac{0.6 I_0 \sqrt{\Delta t}}{\sqrt{N}}. \quad (6.7)$$

As a simple alternative to a fixed size aperture, we also consider allowing our photometric aperture to expand along with the motion of the planet. This aperture will collect the same signal as in SNR_o , but the noise increases with the area as $2r_{ap}v_{om}\Delta t$, so we have

$$SNR_{exp} = \frac{0.6 I_0 \sqrt{\Delta t}}{\sqrt{N (1 + (2/\pi)v_{om}\Delta t)}}. \quad (6.8)$$

A convenient scaling is to multiply top and bottom by $\sqrt{v_{om}}$ and work in normalized SNR units of $I_o/\sqrt{N v_{om}}$. This puts time in terms of FWHM of motion, $\epsilon = v_{om}\Delta t$, and allows comparisons without specifying v_{om} .

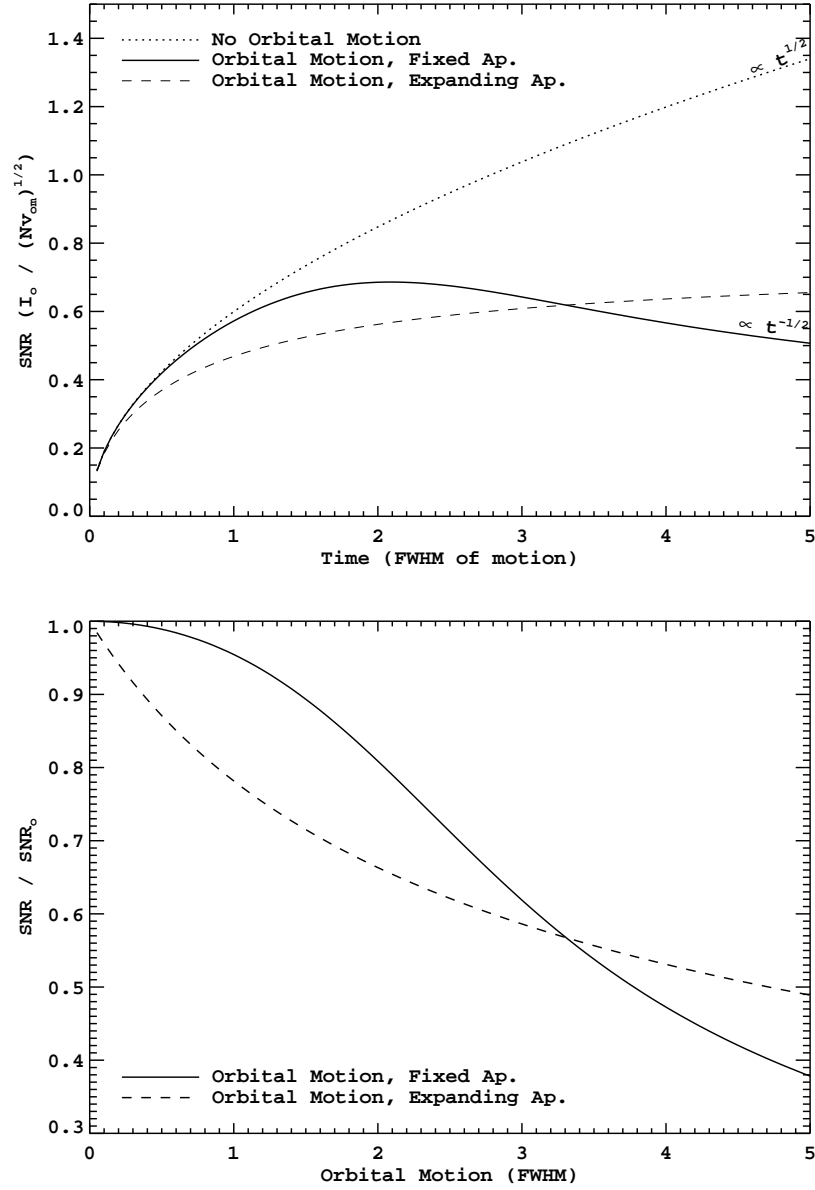


Figure 6.2 Top panel: SNR of a Gaussian PSF with and without orbital motion, in normalized units with time given as FWHM of motion. With no orbital motion $SNR_o \propto \sqrt{t}$. Equation (6.6) was used to calculate the SNR with orbital motion. After ~ 2 FWHM of movement, a maximum is reached and the observation can only be degraded by integrating further. Note that the fixed-aperture orbital motion case eventually goes down as $SNR \propto 1/\sqrt{t}$. For comparison we also show the results with an aperture expanding with the moving planet, which eventually reaches a limit of 0.75. In the bottom panel we show the fractional reduction in SNR due to orbital motion for the fixed radius photometric aperture.

In Figure 6.2 we plot the normalized SNR vs. time (measured in terms of FWHM of motion) with and without orbital motion and for both the fixed and expanding aperture cases. For the fixed aperture, after ~ 2 FWHM of orbital motion a maximum of 0.69 is reached, and from there noise is added faster than signal. This means that further integration only degrades the observation.

The expanding aperture SNR_{exp} exceeds the maximum of SNR_{fix} after about 8 FWHM of motion, and

$$\lim_{x \rightarrow \infty} 0.6 \frac{\sqrt{x}}{\sqrt{1 + (2/\pi)x}} = 0.6 \sqrt{\frac{\pi}{2}} \approx 0.75. \quad (6.9)$$

So if we integrate 4 times longer, adjusting the aperture size would allow us to gather a little more SNR , but only to a point. Given this large increase in telescope time for a relatively small improvement in SNR (only $\sim 9\%$ even if we integrate forever), and its better performance for smaller amounts of motion, the fixed-radius aperture will be our baseline for further analysis – keeping in mind that in some cases it may not be the true optimum.

The peak in SNR_{fix} (equation 6.6) sets the maximum nominal integration time before orbital motion will prevent us from achieving the science goal. That is $\Delta t_{max} = (SNR_{max}/0.6)^2$. If the observation of a *stationary planet* would require an integration time longer than Δt_{max} , then we can't achieve the desired SNR on an *orbiting planet*. This also sets the maximum orbital motion $\epsilon_{max} = v_{om} \Delta t_{max}$. From Figure 6.2 we find that $\epsilon_{max} = 1.3$ FWHM. If more than 1.3 FWHM of motion occurs during an observation, we will not achieve the required SNR .

We also show the fractional reduction in SNR in Figure 6.2. Almost no degradation occurs until after ~ 0.2 FWHM of motion has occurred. SNR is reduced by $\sim 1\%$ after 0.5 FWHM of motion, $\sim 5\%$ after 1.0 FWHM, and by $\sim 19\%$ after 2.0 FWHM of motion. We must now decide how much SNR loss we can accept in our observation.

The above analysis assumes a continuous integration. On a ground-based telescope one must consider that the maximum continuous integration time is $\lesssim 12$ hours, and in practice will likely be much shorter when performing high contrast

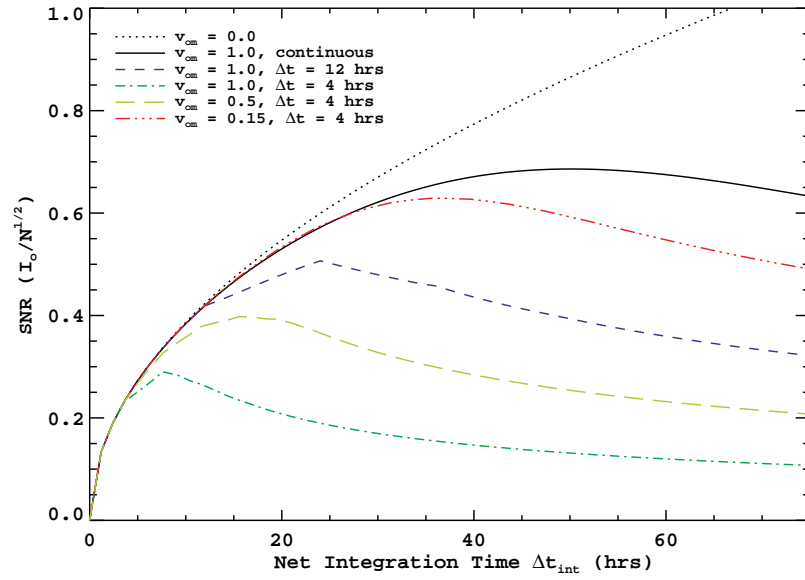


Figure 6.3 Here we show the impact of orbital motion when combined with finite nightly integration times. The SNR of a Gaussian PSF with and without orbital motion is plotted in arbitrary units vs Δt_{int} . The orbital speed v_{om} is given in FWHM day^{-1} .

AO corrected imaging. For instance, an exposure of 20 hours might have to be broken up over 4 or 5 or more nights, when considering the vagaries of seeing (required AO performance), airmass (either through transmission or r_0 requirements), rotation rate (for ADI), and weather. We can adapt the calculations for a ground based integration as follows

$$SNR_{gnd} = \frac{\int_0^{r_{ap}} \int_0^{2\pi} \left[\sum_{j=1}^{j=M} \int_{t_j}^{t_j + \Delta t_j} I_0 e^{(-4 \ln 2 ((r \cos \theta - v_{om} t - x_0)^2 + r^2 \sin^2 \theta))} dt \right] d\theta r dr}{\sqrt{N \pi r_{ap}^2 \Delta t_{int}}} \quad (6.10)$$

In this expression we have broken the observation up into M integration sets which start at times t_j and have lengths Δt_j . The total integration time is $\Delta t_{int} = \sum_{j=1}^{j=M} \Delta t_j$ and the total elapsed time of the observation is $\Delta t_{tot} = t_M + \Delta t_M - t_1$.

We plot the results for a few ground-based scenarios in Figure 6.3. As one can see, observations of planets with orbital motion will be significantly degraded from the ground. This problem, which has been negligible in the high contrast planet searches to date, only becomes worse as we consider larger telescopes and improvements in AO technology which allow searches at shorter wavelengths. We next analyze how this reduction in SNR will affect our ability to detect exoplanets by increasing the rate at which spurious detections occur.

6.3.3 Impact on Statistical Sensitivity

Now we turn to the problem of detecting a planet of a given brightness. A planet is considered detected if its flux is above some threshold SNR_t , which is chosen for statistical significance. The goal in choosing this threshold is to detect faint planets while minimizing the number of false alarms. For the purposes of this analysis we assume Gaussian statistics, in which case the false alarm probability (P_{FA}) per trial is

$$P_{FA} = \frac{1}{2} \text{erfc} \left(\frac{SNR}{\sqrt{2}} \right) \quad (6.11)$$

Typically, planet hunters use a threshold of $SNR = 5$, which gives $P_{FA} = 2.9 \times 10^{-7}$. The number of false alarms per star, the false alarm rate (FAR), is then

$$FAR = P_{FA} \times N_{trials}. \quad (6.12)$$

where N_{trials} is the number of statistical trials per star. Following Marois et al. (2008a), for a stationary planet N_{trials} is just the number of photometric apertures in the image. A typical Nyquist sampled detector of size 1024x1024 pixels has $N_{trials} \sim 8 \times 10^4$. Thus, an $SNR = 5$ threshold will result in $FAR \sim 0.02$ – about 1 false alarm for every 50 observations. In the speckle limited case with non-Gaussian statistics, FAR will be worse than this for the same SNR (Marois et al., 2008a). In any case, the FAR is the statistic which determines the efficiency of a search for exoplanets with direct imaging. A high FAR will cause us to waste telescope time following up spurious detections, while raising the SNR threshold to counter this limits the number of real planets we will detect.

The reduction of SNR caused by orbital motion confronts us with three options. Option I is to maintain the detection threshold constant and accept the loss of sensitivity. Option II is to lower the detection threshold to maintain sensitivity, accepting the increase in FAR. Option III is to correct for orbital motion, which as we will show also causes an increase in FAR.

Option I: Do Nothing

The default option is to do nothing, keeping our detection threshold set as if orbital motion is not significant. The drawback to this is that we will detect fewer planets. To quantify this we use the concept of completeness, that is the fraction of planets of a given brightness we detect. For Gaussian statistics and detection threshold $SNR_t = 5$, the search completeness is given by

$$C(\epsilon) = 1 - \frac{1}{2} \text{erfc} \left(\frac{SNR(\epsilon) - 5}{\sqrt{2}} \right). \quad (6.13)$$

where $\epsilon = v_{om} \Delta t$ is the amount of motion. In Figure 6.4 (top) we show the impact of orbital motion on search completeness. Maintaining the detection threshold lowers

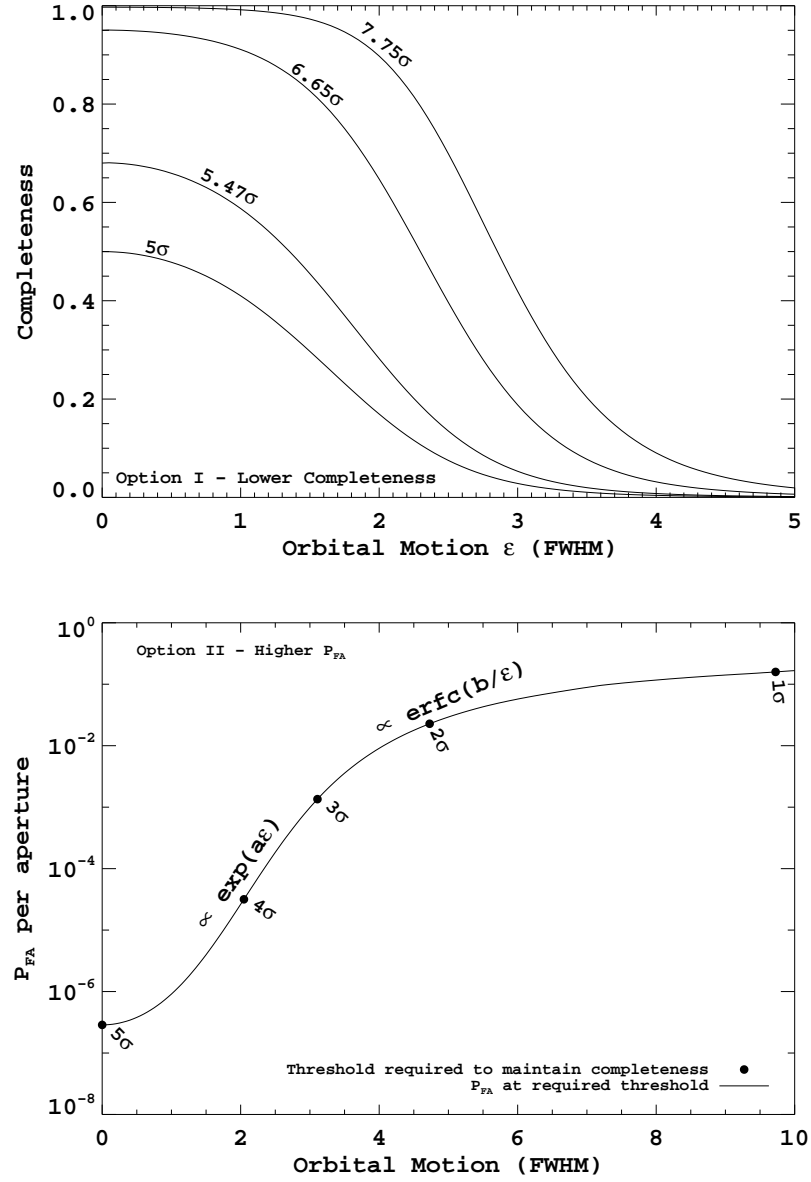


Figure 6.4 Top panel: completeness as a function of orbital motion if we maintain our detection threshold at 5σ . Planet brightness is expressed as the SNR at which we would be 50%(5σ), 68%(5.47σ), 95%(6.65σ), and 99.7% (7.75σ) complete with no orbital motion. Bottom panel: the increase in false alarm probability (P_{FA}) if we lower the detection threshold to maintain 50% completeness for an orbiting planet that would have a brightness of 5σ were it stationary. After ~ 1 FWHM of motion P_{FA} increases exponentially until ~ 4 FWHM where it becomes asymptotic to 0.5.

completeness. How much depends on the completeness level, with brighter planets being less affected. For planets bright enough to yield 95% completeness with no motion, significant reduction in the number of detections begins after ~ 1 FWHM of motion. For 99.7% completeness the impact becomes significant after ~ 1.5 FWHM.

Option II: Lower Threshold

Once orbital motion is recognized to be significant, a simple countermeasure would be to lower the detection SNR threshold in order to maintain completeness. The drawback to this option is that we have more false alarms, which must then be followed up using more telescope time. This results in a less efficient search. In Figure 6.4 (bottom) we show P_{FA} as a function of orbital motion, and denote the detection threshold we must use to maintain 50% completeness for a planet bright enough to give $SNR = 5$ were it stationary. Note that P_{FA} begins to increase exponentially after ~ 1 FWHM of motion. After ~ 4 FWHM P_{FA} begins approaching 0.5 asymptotically. Once $\epsilon \approx 2$ FWHM the number of false alarms per 1024x1024 image approaches 1.

Option III: De-orbit

Option III is to correct for orbital motion, hoping to maintain sensitivity while limiting the increase in P_{FA} . The essence of any such technique will be calculating the position of the planet during the observation, and de-orbiting in some way, say shift-and-add (SAA) on a sequence of images. The drawback of this approach is that it will produce more false alarms per observed star due to the increased number of trials, similar to lowering the detection threshold. If the orbit were precisely known, we could proceed with almost no impact on FAR . However, in the presence of uncertainties in orbital parameters or in a completely blind search we will have to consider many trial orbits. For now we can perform a “back-of-the-envelope” estimate of the number of possible orbits to understand how much FAR will increase. To do so, we begin by placing bounds on the problem.

We can first establish where on the detector we must consider orbital motion. At any separation r from the star, the slowest un-bound orbit will have the escape velocity. Since we know that physical separation is greater than or equal to projected separation, $r \geq \rho$, and that maximum projected speed will occur for inclination $i = 0$, we know that

$$v_{esc} = \sqrt{2}v_{FOC}(a \rightarrow \rho) \quad (6.14)$$

sets the upper limit on the projected focal plane speed of an object in a bound orbit. We can also set an upper limit on the amount of motion ϵ_{max} we can tolerate over the duration Δt_{tot} of the observation based on the SNR degradation it would cause. So we only need consider orbital motion when

$$\sqrt{2}v_{FOC}(\rho)\Delta t_{tot} > \epsilon_{max}. \quad (6.15)$$

From here we determine the upper limit on projected separation from the star for considering this problem:

$$\rho_{max} = 0.0136M_* \left(\frac{D}{\lambda d} \frac{\Delta t_{tot}}{\epsilon_{max}} \right)^2 \text{ AU}. \quad (6.16)$$

By the same logic, for any point closer than ρ_{max} the maximum possible change in position is

$$\Delta \rho_{max} \approx \sqrt{2}v_{FOC}(\rho)\Delta t_{tot} \text{ in FWHM}. \quad (6.17)$$

Then we must evaluate possible orbits ending anywhere in an area of $\pi(\Delta \rho_{max})^2$ FWHM² around an initial position.

These two limits set the statistical sensitivity of an attempt to de-orbit an observation. The number of different orbits, N_{orb} , will be determined by the area of the detector where orbital motion is non-negligible, and the size of the region around each point that we consider. That is

$$N_{orb} \propto \int_0^{\rho_{max}} \Delta \rho_{max}^2 \rho d\rho. \quad (6.18)$$

so

$$N_{orb} \propto \left(\frac{M_*}{\epsilon} \right)^2 \left(\frac{D}{\lambda d} \right)^4 \Delta t_{tot}^4. \quad (6.19)$$

In general $N_{\text{trials}} \propto N_{\text{orb}}$, so $FAR \propto P_{FA} \times N_{\text{orb}}$. Larger D , shorter λ , closer d , and smaller acceptable orbital motion ϵ will then all increase FAR^2 . Perhaps the most important feature of this result is that $N_{\text{orb}} \propto \Delta t_{\text{tot}}^4$ – **increasing integration time rapidly increases the FAR of a blind search**. Note that this is still less severe than the exponential increase in P_{FA} found for merely lowering the threshold. In the next section we will test these relationships after fully applying orbital mechanics, and see that they hold.

6.4 Blind Search: Recovering SNR after Orbital Motion

In this section we consider in detail a blind search, i.e. an observation of a star for which we have no prior knowledge of exoplanet orbits. We showed above that the problem is well constrained. Here we derive several ways to further limit the number of trial orbits we must consider. After that, we describe an algorithm for determining the orbital elements that must be considered and then discuss the results. Finally, we use this algorithm to de-orbit a sequence of simulated images and analyze the impact of correlations between trial orbits on FAR .

To provide numerical illustrations throughout this section we consider the problem of a 20 hour observation of α Cen A using the GMT at $10\mu\text{m}$. This scenario is loosely based on performance predictions made for the proposed TIGER instrument, a mid-IR diffraction limited imager for the GMT (Hinz et al., 2012). The details of these predictions are not important for our purposes, so we will only assert that this is a plausible case. There are other examples in the literature with similar integration times, such as the EPICS prediction we discussed earlier.

We assume that this 20 hr observation is broken up into five $\Delta t = 4$ hr exposures, spread over 7 nights or $\Delta t_{\text{tot}} = 6.2$ elapsed days from start to finish. The choice of Δt is essentially arbitrary, but we have good reasons to expect it to be shorter than an entire night. An important consideration is the planned use of ADI, and the

²Assuming background limited photometry with a diffraction limited PSF, we expect $\Delta t \propto 1/D^4$ (Hardy, 1998). All else being equal, larger telescopes are better when considering this problem

attendant need to obtain sufficient field rotation in a short enough time to provide good PSF calibration while avoiding self-subtraction (Marois et al., 2006). The effect of airmass on seeing through $r_0 \propto \cos(z)^{3/5}$, where z is the zenith angle, and hence on AO system performance, could also cause us to observe as near transit as possible. Efficiency will be affected by chopping and nodding, necessary for background subtraction at $10\mu\text{m}$. This will limit the net exposure time obtainable in one night..

Few ground-based astronomers would object to an assertion that we loose 2 nights out of 7 to weather. We could be observing in queue mode, such that these observations are only attempted when seeing is at least some minimal value, or precipitable water vapor is low. One can even imagine the opposite case at $10\mu\text{m}$, such that nights of the very best seeing are devoted to shorter wavelength programs. While this scenario may be somewhat contrived, we feel that it is both plausible and realistic. We now proceed to describe a technique that would mitigate the effects of orbital motion for our GMT example and should be applicable to other long exposure cases.

6.4.1 Limiting Trial Orbits

Here we derive limits on the semi-major axis and eccentricity of trial orbits to consider. These limits are based only on the amount of orbital motion tolerable for the science case, and do not represent physical limits on possible orbits around the star.

It is always true that $r \geq \rho$. This implies that, for any orbit, the separation of apocenter must obey $r_a \geq \rho$. This allows us to set a lower bound on a , a_{min} , given a choice of e through

$$\rho_{au} \leq a_{min}(1 + e) \quad (6.20)$$

which gives

$$a_{min} = 0.2063 \frac{\lambda d \rho}{D(1 + e)}. \quad (6.21)$$

The fastest speed in a bound planet's orbit will occur at pericenter, and using

the maximum tolerable motion ϵ_{max} during our observation of total elapsed time Δt_{tot} we can set an upper bound on a by noting that

$$v_{FOC}(a_{max}) \sqrt{\frac{1+e}{1-e}} \Delta t_{tot} \leq \epsilon_{max} \quad (6.22)$$

which leads to

$$a_{max} = \left(0.0834 \frac{D}{\lambda d}\right)^2 \frac{1+e}{1-e} M_* \left(\frac{\Delta t_{tot}}{\epsilon_{max}}\right)^2. \quad (6.23)$$

Using the GMT example: for $e = 0.0$, $a_{max} = 3.9$ AU; and for $e = 0.5$, $a_{max} = 11.8$ AU. Using Equation 6.16 we have a projected separation limit of $\rho_{max} = 7.7$ AU, so it is possible for these definitions to produce $a_{max} < a_{min}$ for certain choices of e at a given ρ . This condition tells us that at such a value of e no orbits can move fast enough to warrant consideration. Thus we can set a lower limit on e at projected separation ρ

$$e_{min} = \frac{1}{2} \sqrt{\xi^2 + 8\xi} - 1 - \frac{\xi}{2} \quad (6.24)$$

where we have simplified by pulling out

$$\xi = 29.66 \frac{\rho}{M_*} \left(\frac{\epsilon}{\Delta t}\right)^2 \left(\frac{\lambda d}{D}\right)^3. \quad (6.25)$$

In practice, we might consider eccentricity ranges with e_{max} less than 1, thus improving our sensitivity. Inputs to our choice of e_{max} could include some prior distribution of eccentricities, or dynamical stability considerations in binary star systems and systems with known outer companions.

6.4.2 Choosing Orbital Elements

Now we describe an algorithm for sampling the possible trial orbits over a set of M sequential images. For now, we assume no prior knowledge of orbital parameters. We will employ a simple grid search through the parameter space bounded as described above.

1. Determine the region around the star to consider using Eq. (6.16).

2. Identify regions of interest. In the best cases the orbital motion will be small enough that we will be able stack the images and search the result for regions with higher SNR (e.g. $SNR > 4$) and limit further analysis to those areas. In the worst cases orbital motion will be large enough that we will need to blindly apply this algorithm at each pixel within the bounding region identified in the previous step. In the present GMT- α Cen example we are in the former case.
3. For each region, choose a size, perhaps based on v_{esc} (as in Eq. 6.17).
4. Chose a starting point (x_1, y_1) , with $\rho_1 = \sqrt{x_1^2 + y_1^2}$. If we are proceeding pixel by pixel, then (x_1, y_1) describes the current pixel.
5. Choose $e \in e_{min}(\rho_1) \dots e_{max}$ using Equation (6.24) and assumptions about e_{max} .
6. Choose $a \in a_{min}(\rho_1, e) \dots a_{max}(e)$ using Equations (6.21) and (6.23).
7. Choose time of pericenter $\tau \in t_1 - P(M_*, a) \dots t_1$ where P is the orbital period and t_1 is the time of the first image. Now calculate the true anomaly $f(t_1; a, e, \tau, P)$ using *Kepler's* equation and physical separation using:

$$r = \frac{a(1 - e)}{1 + e \cos(f)} \quad (6.26)$$

8. if $e \neq 0$: Choose $\omega \in 0 \dots 2\pi$
if $e = 0$: set $\omega = 0$.
9. if $\sin(\omega + f) > 0$:

(a) Given e , a , τ , f , and ω , calculate

$$\cos i = \frac{\pm \sqrt{\frac{\rho^2}{r^2} - \cos^2(\omega + f)}}{\sin(\omega + f)} \quad (6.27)$$

$$\sin \Omega = \frac{y \cos(\omega + f) - x \sin(\omega + f) \cos i}{r(\cos^2(\omega + f) + \sin^2(\omega + f) \cos^2 i)} \quad (6.28)$$

$$\cos \Omega = \frac{y \sin(\omega + f) \cos i + x \cos(\omega + f)}{r(\cos^2(\omega + f) + \sin^2(\omega + f) \cos^2 i)} \quad (6.29)$$

where Ω should be determined in the correct quadrant.

- (b) We now have a complete set of elements, and so can SAA the sequence of images based on these orbits (one for each i). Doing so requires calculating the true anomaly f_j at the time of each image, and then calculating the projected orbital position of the prospective companion in each image.

10. if $\sin(\omega + f) = 0$, we do not have a unique solution for inclination. This is the special case where the planet is passing through the plane of the sky.

- (a) for $\omega + f = 0$ calculate Ω :

$$\sin \Omega = \frac{y}{r} \quad (6.30)$$

$$\cos \Omega = \frac{x}{r} \quad (6.31)$$

or for $\omega + f = \pi$ calculate Ω :

$$\sin \Omega = \frac{-y}{r} \quad (6.32)$$

$$\cos \Omega = \frac{-x}{r} \quad (6.33)$$

determining Ω in the correct quadrant.

- (b) Choose $i \in 0 \dots \pi$
- (c) We now have a complete set of elements, and so can SAA as in step 9b above.
- (d) Repeat steps 10b to 10c until all i chosen.
11. Repeat the above steps until the parameters ω , τ , a , and e are sufficiently sampled for each starting point.

6.4.3 De-orbiting: Unique Sequences of Whole-Pixel Shifts

The algorithm just described will produce a large number of trial orbits, many of which will be very similar. The information content of our image is set by the resolution of the telescope, so we can take advantage of this similarity to greatly reduce the number of statistical trials. This is done by grouping similar orbits into sequences of whole-pixel shift sequences, where the pixels are at least as small as $\text{FWHM}/2$. As we will see, we typically will want to oversample, to say $\text{FWHM}/3$, to ensure adequate SNR recovery.

We calculate the pixel-shift sequence for each orbit by determining which pixel the trial planet (or rather, the center of its PSF) lands on at each time step. Many orbits end up producing the same sequences of pixel-shifts, and we will keep only the unique ones for use in de-orbiting the observation. In Figure 6.5 we illustrate the outcome of the pixel-shift algorithm, showing two unique sequences and a few of the orbits that produced them.

To test the above algorithm and the pixel-shift technique, we used our GMT α Cen A example and determined the trial orbits for various separations and Δt s. We set $\epsilon_{max} = 0.5$ based on our earlier analysis of SNR . The results are summarized in Figure 6.6. The problem is generally well constrained in that we only have a finite search space for any initial point. The data used to construct Figure 6.6 are provided in Table 6.1. Comparing N_{orb} to N_{shifts} , note the large reduction in the number of trials ($\sim 10^8$ to $\sim 10^2$) due to combining similar orbits.

6.4.4 N_{orb} Scalings

In Figure 6.7 we plot the area of the detector which contains the possible trial orbits at $\rho_1 = 1.0$ AU vs. the total elapsed time Δt_{tot} . We conclude from this plot that the area around a given starting point is proportional to Δt_{tot}^2 . Also in Figure 6.7 we plot area vs separation from the star, and conclude that area is proportional to $1/\rho_1$. Taken together these results give confidence that the $N_{orb} \propto \Delta t_{tot}^4$ scaling derived earlier holds when we fully apply orbital mechanics rather than the escape

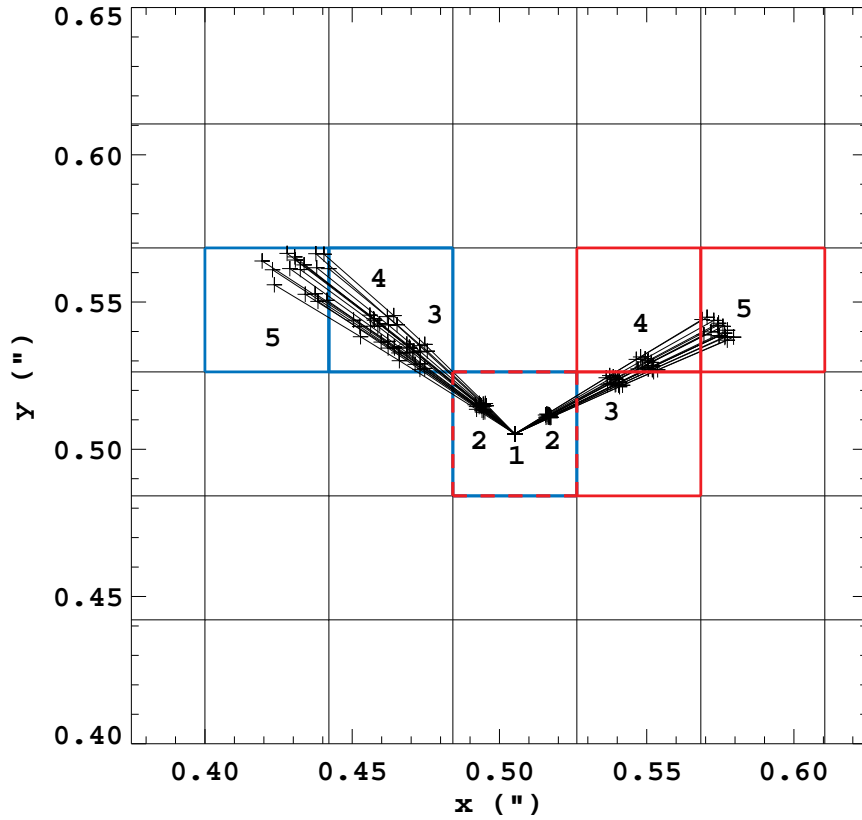


Figure 6.5 Two sequences of whole pixel-shifts, one in red and one in blue. We also show a few of the many orbits that produce these shift sequences. Once these shifts are determined, a set of 5 images can be de-orbited by shifting the images by the indicated sequence 5-4-3-2-1, that is the pixel containing the orbit in image 2 is shifted and added to the pixel containing the orbit in image 1, and likewise for images 3, 4, and 5. Of course, the entire image is shifted, not just single pixels.

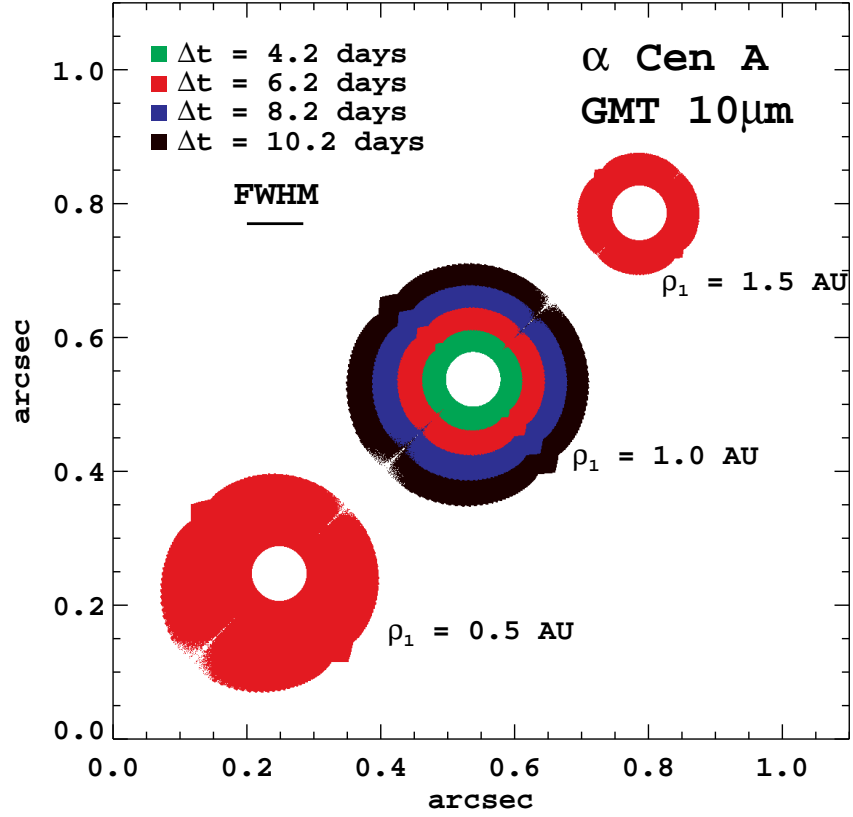


Figure 6.6 Example trial orbits for the GMT, working at $10\mu\text{m}$, observing $\alpha\text{Cen A}$. Plotted are the end points of orbits calculated using the algorithm given in Section 6.4.2 for the given initial projected planet separations ρ_1 and elapsed observation times Δt_{tot} . The red points show the effect of changing initial separation for a constant elapsed time. At 1 AU initial separation the colors correspond to different elapsed times as indicated in the legend. We further analyze these relationships in Table 6.1 and Figure 6.7. The results of the algorithm appear more complicated than the simple escape-velocity circle analysis in Section 6.3.3. The end-point clouds are not circularly symmetric about the starting point, and have some azimuthal structure. For instance there is a triangle extending azimuthally corresponding to face-on high- e orbits, and there are gaps along the radius from the star corresponding to i very near 90° . These structures are consequences of the chosen grid resolution.

Table 6.1 Results of applying the algorithm detailed in Section 6.4.2 for various separations and elapsed observation times. See also Figure 6.7. Note the dramatic reduction in the number of trials (N_{orb} vs. N_{shifts}) after combining similar orbits into whole-pixel shift sequences.

ρ_1 (AU)	Δt_{tot} (days)	No. Obs.	N_{orb}	N_{shifts}
0.5	6.0	5	2.7×10^8	285
1.0	2.0	5	4.1×10^8	14
1.0	4.0	5	4.1×10^8	76
1.0	6.0	5	4.1×10^8	134
1.0	8.0	5	4.1×10^8	253
1.5	6.0	5	5.2×10^8	90
1.0	2.0	3	4.1×10^8	10
1.0	4.0	5	4.1×10^8	78
1.0	6.0	7	4.1×10^8	292
1.0	8.0	9	4.1×10^8	815

velocity approximation.

Things are a bit more complicated when we consider the scaling of the number of whole-pixel shift sequences. We conducted two sets of trials at $\rho_1 = 1.0$ AU. In the first, the number of observations and their relative spacing was held constant regardless of Δt_{tot} . In the second set, the number of observations scaled with Δt_{tot} . As shown in Figure 6.7, when the number of observations is constant, the number of shifts scales as Δt_{tot}^2 , but when the number of observations grows with Δt_{tot} the number of shifts scales as roughly $\Delta t_{tot}^{3.6}$. Figure 6.7d shows that the number of shifts scales as $1/\rho_1$. Taken together, we see that for a constant number of observations the pixel-shift technique will follow the $N_{orb} \propto \Delta t_{tot}^4$ scaling. However, if the number of observations also scales with Δt_{tot} , then our results imply that $N_{orb} \propto \Delta t_{tot}^{5.6}$. The value of the exponent likely depends on the details of the observation sequence, but this has important implications for observation planning.

6.4.5 Recovering SNR

We next consider whether de-orbiting by whole-pixels adequately recovers SNR . To test this we “orbited” a Gaussian PSF on face-on orbits with various eccentricities, starting from pericenter. We then calculated shifts for detector samplings of 2, 3, and 4 pixels/FWHM, and then de-orbited by these shifts. The results are summarized in Table 6.2. On a critically sampled detector we only recover a 5σ planet to $\sim 4.9\sigma$, a 2% loss of SNR . At 3 pixels/FWHM we do much better, recovering SNR to 4.97 for low eccentricities, and 4.95 for higher eccentricities. Performance for 4 pixels/FHWM sampling is similar. A 2% loss of SNR nearly doubles P_{FA} , so it appears that we should oversample to at least 3 pixels/FWHM, either optically or by re-sampling images during data reduction. In our analysis we have assumed that the limiting noise source is background photons (PSF halo or sky), so we ignore the increased readout noise expected from oversampling.

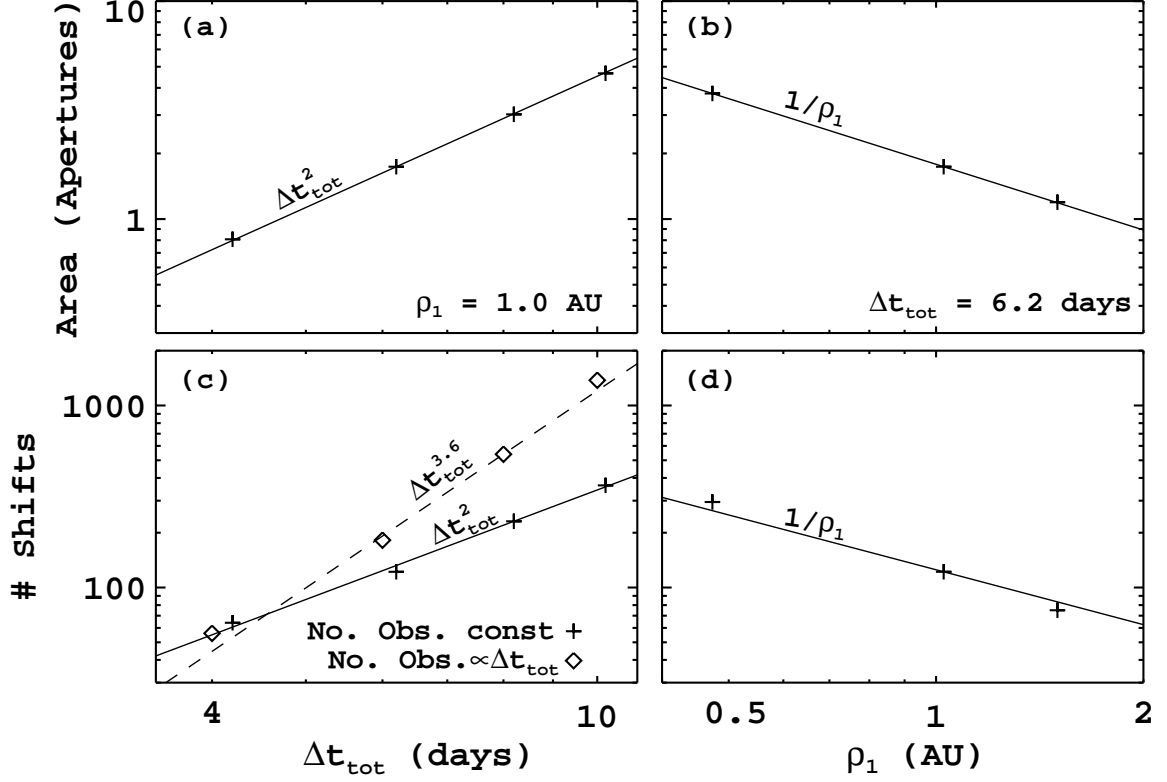


Figure 6.7 Scaling of the number of orbits and the number of resulting whole-pixel shifts with observation elapsed time and with distance from the star. These results demonstrate that the number of trial orbits $N_{\text{orb}} \propto \Delta t_{\text{tot}}^4$ scaling that we derived using the escape velocity holds when we rigorously apply orbital mechanics. Note though that the situation is more complicated with the number of shifts – if the number of observations increases with elapsed time then the number of shifts grows faster than Δt_{tot}^2 , implying that N_{orb} will increase faster than Δt_{tot}^4 . These scalings lead to one of our main, if seemingly obvious, conclusions: one must limit the elapsed time of an observation as much as possible when orbital motion is significant.

Table 6.2 SNR recovered after de-orbiting with whole-pixel shifts for various samplings.

Sampling (pix/FWHM)	<i>SNR</i> Recovered						
	e=0.0	e=0.1	e=0.2	e=0.3	e=0.5	e=0.7	e=0.9
2	4.89	4.89	4.89	4.88	4.86	4.86	4.86
3	4.97	4.96	4.95	4.94	4.94	4.95	4.95
4	4.97	4.97	4.97	4.97	4.94	4.92	4.92

6.4.6 Correlations And The True Impact On P_{FA}

As we have noted several times, the main impact of orbital motion is to reduce SNR , which in turn reduces our statistical sensitivity. If we attempt to de-orbit an observation in order to recover SNR , we do so at the cost of a large increase in the number of trials. Worst case, this results in a proportional increase in FAR since nominally $FAR = P_{FA} \times N_{orb}$. However, we expect significant correlation between trials of neighboring orbits and whole-pixel shifts. To investigate this, we performed a series of monte carlo experiments. A sequence of images with Gaussian noise was generated, and first stacked without shifting, hereafter called the naive-add. The same sequence was then shifted by each possible whole-pixel shift, assuming a 1AU initial separation around α Cen A. This experiment was conducted for observations with total elapsed times Δt_{tot} of 4.2, 6.2, 8.2, and 10.2 days, with samplings of 2, 3, and 4 pixels/FWHM.

We performed several tests on each sequence. The first was a simple threshold test on the naive-add, with the threshold set for the worst case orbital motion given by Equation 6.10 with $v_{om} = v_{esc}$. We performed simple aperture photometry, with a $r_{ap} = 1$ FWHM. As expected the resultant P_{FA1} is as predicted by Equation 6.11. The next test was to apply a 5σ threshold after de-orbiting by whole-pixel shifts and adding. If all shifts were completely uncorrelated, then we would expect $FAR = (2.9 \times 10^{-7}) \times N_{shifts}$, but as we predicted, shifts are correlated and P_{FA2} is lower than this.

The final test performed was to apply both thresholds in sequence, such that a detection is made only if the naive-add results in SNR greater than the threshold for worst case orbital motion, **and** the de-orbited SAA results in $SNR > 5$. This P_{FA3} is lower than either P_{FA1} or P_{FA2} , but still higher than if no orbital motion occurred.

The results of each trial are present in Table 6.3. Applying both threshold tests results in significant improvement over the naive-add in terms of FAR . Another interesting result is that sampling has only a minor impact on P_{FA3} . This makes

Table 6.3 False alarm probabilities after de-orbiting Gaussian noise images.

Δt_{tot} (days) ¹	SNR_t ²	N_{shifts} ³	P_{FA1} ⁴	P_{FA2} ⁵	P_{FA3} ⁶
2 pixels/FWHM					
4.2	4.635	64	1.74×10^{-6}	7.65×10^{-6}	8.06×10^{-7}
6.2	4.220	122	1.24×10^{-5}	1.52×10^{-5}	2.70×10^{-6}
8.2	3.330	231	4.40×10^{-4}	2.71×10^{-5}	9.93×10^{-6}
10.2	2.625	364	4.33×10^{-3}	4.03×10^{-5}	2.17×10^{-5}
3 pixels/FWHM					
4.2	4.635	108	2.11×10^{-6}	1.37×10^{-5}	4.80×10^{-7}
6.2	4.220	285	1.21×10^{-5}	3.39×10^{-5}	2.04×10^{-6}
8.2	3.330	496	4.31×10^{-4}	5.64×10^{-5}	9.96×10^{-6}
10.2	2.625	741	4.34×10^{-3}	8.15×10^{-5}	2.67×10^{-5}
4 pixels/FWHM					
4.2	4.635	217	1.78×10^{-6}	2.64×10^{-5}	4.44×10^{-7}
6.2	4.220	487	1.24×10^{-5}	5.61×10^{-5}	1.48×10^{-6}
8.2	3.330	844	4.35×10^{-4}	9.19×10^{-5}	1.14×10^{-5}
10.2	2.625	1315	4.32×10^{-3}	1.41×10^{-4}	3.15×10^{-5}

¹Elapsed time of the observation.² SNR threshold from Equation 6.10, using $v_{orb} = \sqrt{2}v_{FOC}$.³Number of unique whole-pixel shifts required to de-orbit.⁴False alarm probability for the naive-add, from MC experiment results. Expected values given by Equation 6.11.⁵False alarm probability after de-orbiting with whole-pixel shifts.⁶False alarm probability after testing both the naive-add and de-orbiting.

some sense as we expect the correlation of neighboring shifts to be set by the FWHM, not the sampling. So even though the accuracy of SNR recovery is improved, and quite a few more shifts are required, these shifts remain correlated across the same spatial scale resulting in little change in the overall FAR .

6.4.7 Impact on Completeness of the Double Test

There is *still* an impact on completeness, however, because we are now conducting two trials instead of one. This lowers the true positive probability (P_{TP}). Consider a

5σ planet on the worst case fastest possible orbit, for the 10.2 day elapsed time case. The threshold for the naive add is 2.625. We have a 50% probability of detecting this planet after the naive add. If it is detected on the first test, there is then some probability $P_{TP} < 1$ of detecting at $SNR \geq 5$ after de-orbiting. Worst case, this will be 50%, resulting in a net P_{TP} of 25%. In reality, it will be better than this as the two trials will be strongly correlated.

Even if this worst case of 25% were realized this is still significant improvement over Option I. A 2.6σ signal would only be detected 10% of the time with a 5σ threshold. Given the reduction in P_{FA} from 4.3×10^{-3} to 2.2×10^{-5} , likewise an improvement over Option II at 2.6σ , *it is clear that de-orbiting by whole-pixel shifts does improve our ability to detect an orbiting planet.* The situation will be even better for slower planets, and most of the area searched will not be subject to the worst case orbital speed. We leave a complete analysis of the impact on search completeness for future work. One can also imagine adjusting the thresholds to optimize completeness at the expense of worse P_{FA} .

6.4.8 Tractability of a Blind Search

We end this section by concluding that a blind search when orbital motion is significant is tractable. Orbital motion will make such a search less sensitive, both in terms of number of false alarms and in terms of completeness, but Keplerian mechanics gives us enough tools to bound the problem. As we have shown de-orbiting a sequence of observations can recover SNR to its nominal value, and we can do so while controlling the impact on statistical sensitivity. For the $\Delta t_{tot} = 6.2$ day observation, P_{FA3} was roughly a factor of 10 higher than if no orbital motion occurred. This increase only occurs over a bounded region around the star, so the net effect on FAR will be contained. Using this factor of 10 as the mean value over the $7.7 \text{ AU} = 69.1 \text{ FWHM radius}$ region around $\alpha \text{ Cen A}$ where orbital motion is significant, the FAR in this area will have gone from $\sim 1/1000$ to $\sim 1/100$ in our GMT/10 μm example. The key, though, appears to be to limit the elapsed time of the observation as the number of trials increases — decreasing sensitivity — proportionally to *at least*

Table 6.4 Orbital parameters for Gl 581d used in this analysis. We derived the values reported here from other parameters where necessary. Only the uncertainty in t_0 impacts our analysis. In both models the orbital period is 66.6 days.

Model	a (AU)	e	ω (deg)	σ_{t_0} (days)
Forveille et al. (2011)	0.218 ± 0.005	0.25 ± 0.09	356.0 ± 19.0	± 3.4
Vogt et al. (2012)	0.218 ± 0.005	0.0 ± 0.0	0.0 ± 0.0	± 7.45

Δt_{tot}^4 in a blind search.

The main caveat at this point in our analysis is that we have drawn the conclusion of tractability using Gaussian statistics. It is well known that speckle noise, which will often be the limiting noise source for high contrast imaging in the HZ, is not Gaussian and results in much higher P_{FA} for a given SNR (Marois et al., 2008a). Future work on this problem will need to take this into account.

Next we consider a more strongly bounded scenario, where we have significant prior information about the orbit of the planet from radial velocity surveys.

6.5 Cued Search: Using RV Priors

The situation is greatly improved if we have prior information, such as orbit parameters from RV or astrometry. Here we consider the case of Gliese 581d, and the previously discussed future observation of this planet by EPICS at the E-ELT (Kasper et al., 2010). There is some controversy surrounding the solution to the RV signal, and whether planet d even exists (Forveille et al., 2011; Vogt et al., 2012; Baluev, 2012). We show results for both the floating eccentricity Keplerian fits of Forveille et al. (2011)[hereafter F11], and the all circular interacting model of Vogt et al. (2012)[hereafter V12]. Doing so allows us to illustrate the impact of eccentricity on the analysis, and prevents us having to take a stand in a currently raging debate. The parameters used herein are listed in Table 6.4.

Instead of a grid search, we use a monte carlo (MC) method. The RV technique provides the parameters a , e , ω and t_0 or their equivalents. We can take the results of

fitting orbits to the RV signal, and the associated uncertainties, as prior distributions which we sample to form trial orbits. We will assume that all uncertainties are uncorrelated and are from Gaussian distributions.

We assume that the 20 hr integration is broken up over 6.2 nights based on the same logic discussed in Section 6.4. Kasper et al. (2010) actually assumed 20×1 hr observations based on the amount of rotation needed, but did not consider the effects of orbital motion over 20 days of a 67 day period (M. Kasper, personal communication (2012)).

6.5.1 Constraints

In order to minimize the number of trial orbits to consider, we can apply various constraints taking advantage of the information we have from the RV detection.

In the case of a multi-planet system dynamical analysis can place constraints on the inclination based on system stability. For Gl 581, Mayor et al. (2009) found the system was stable for $i > 30$. We can also make use of the geometric prior for inclination, where we expect $P_i = \sin(i)$ in a population of randomly oriented systems.

Since this is a reflected light observation, the orbital phase and its impact on the brightness of the planet must be considered. The planet's reflected flux is given by

$$F_p(\alpha) = F_* \left(\frac{R_p}{r} \right)^2 A_g(\lambda) \Phi(\alpha) \quad (6.34)$$

where F_* is the stellar flux, R_p is the planet's radius, r its separation, $A_g(\lambda)$ is the wavelength dependent geometric albedo, and Φ is the phase function at phase angle α . The phase angle is given by

$$\cos(\alpha) = \sin(f + \omega) \sin(i). \quad (6.35)$$

In general, determining the quantity $A_g(\lambda)\Phi(\alpha)$ requires atmospheric modeling (Cahoy et al., 2010). For now, we assume that Φ follows the Lambert phase function

$$\Phi(\alpha) = \frac{1}{\pi} [\sin(\alpha) + (\pi - \alpha) \cos(\alpha)] \quad (6.36)$$

We assume that the prediction of Kasper et al. (2010) was made for the planet at quadrature, $\alpha = \pi/2$, where $\Phi = 0.318$. We then require that the mean value of Φ during the observation be greater than this value - that is the planet is as bright or brighter than it is at quadrature.

6.5.2 Initial Detection

An important consideration in an RV-cued observation will be when to begin. As a first approximation, we assume that maximizing planet-star separation will maximize our sensitivity. This may not be true when working in reflected light due to the phase and separation dependent brightness of the planet in this regime. Proceeding with the approximation for now, we expect to plan this observation to be as close to apocenter as possible. In this case we will begin integrating 3.1 days before $t_0 + P/2$.

To understand the area where we will be searching for Gl 581d, we first conducted an MC experiment to calculate the possible positions of the planet at $t = t_0 + P/2 - 3.1$ days. To do so, we drew random values of a , e , w , and t_0 from Gaussian distributions with the parameters of Table 6.4. We drew a random value of i from the $\sin(i)$ distribution, and rejected any value of $i \leq 30$ based on the dynamical prior. Finally Ω was drawn from a uniform distribution in $0 \dots 2\pi$. This process was repeated 10^9 times, and the frequency at which starting points occur in the area around the star was recorded. The results are shown in Figure 6.8 for the V12 circular model and for the F11 eccentric model. The Figure shows the area which must be searched to obtain various completeness. For instance, if we desire 95% completeness in the V12 model, we must consider an area of 71 apertures. Since this $SNR = 5$ detection is broken up into 5 distinct integrations, our first attempt will have $SNR = 2.24$, giving a $FAR = 0.89$ for the first 4 hr integration. In other words, we should expect a false alarm in addition to a real detection.

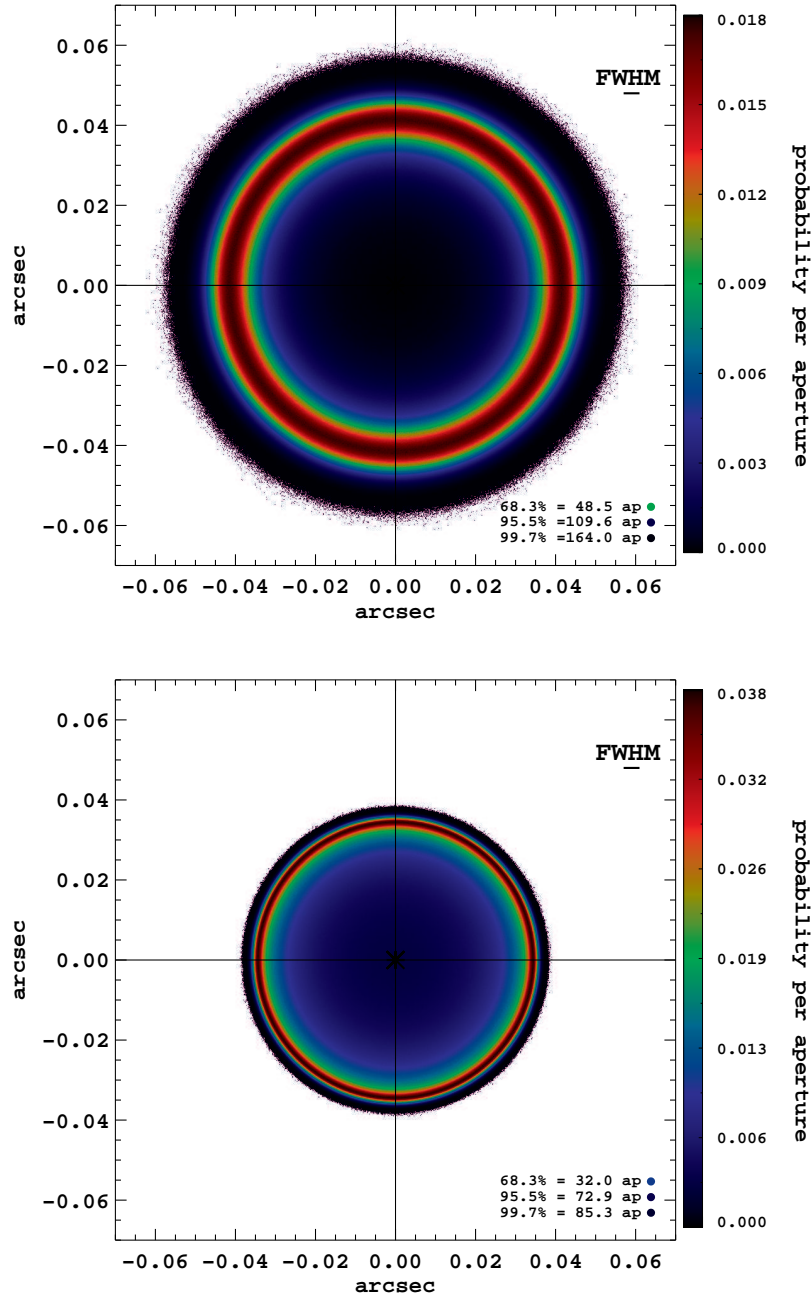


Figure 6.8 Possible starting points for Gl 581d, observed near apocenter. Top: using the parameters of Forveille et al. (2011)'s eccentric model. Bottom: assuming the parameters of Vogt et al. (2012)'s circular interacting model. The color shading is in units of probability per aperture (each aperture has area πFWHM^2). The legend indicates the color which encloses the given completeness intervals, and the enclosed area in apertures, which can be directly related to the false alarm rate as discussed in the text.

6.5.3 Calculating Orbits and Shifts

Now we assume that we have an initial detection at $SNR \sim 2.24$ within the highest probability regions³. In order to follow-up this detection over subsequent nights, we must determine the possible locations of the planet, constrained by the RV-derived orbital elements.

We proceed by choosing a , e , ω , and t_0 from Gaussian distributions as above. Now as long as $r > \rho$ we will have a unique solution for i and Ω given the randomly chosen parameters (see the blind search algorithm above). We take into account dynamical stability by rejecting any orbit which has $i \leq 30$. The orbit determined in this fashion was then projected 6.2 days into the future and the frequency of these final points was recorded. We show the result for the V11 model in Figure 6.9, top panel. Using the RV determined parameters and their uncertainties allows us to determine the probability density of orbit endpoints, and determine how much of the search space we must consider for a given completeness. The whole-pixel shifts were also calculated using a sampling of $FWHM/3$, and are shown in the legend. We also applied the blind search algorithm to this observation from the same starting point, and show the results for comparison in the bottom panel of Figure 6.9. As expected the RV priors significantly reduce the search space - we have 942 trial shift-sequences to consider instead of 12000.

Another important consideration here is that our initial 2.24σ detection will have a large position uncertainty, which we estimate by $\sigma_{\rho_0} = FWHM/SNR$. We added a random draw for the starting position, and repeated the MC experiment for F11 and also conducted a run for the V12 parameters. The results are shown in Figure 6.10. The number of shift sequences is much higher due to the uncertainty in the starting position caused by our low SNR initial detection, but we expect correlations to come to the rescue as in our α Cen example. To compare to Figure 6.9 keep in mind that the blind search would have to be applied to all 5500 pixels in the search space indicated by Figure 6.8.

³For the purposes of this analysis, we calculated initial separation ρ_1 using the mean parameters for each model and an inclination $i = 60$

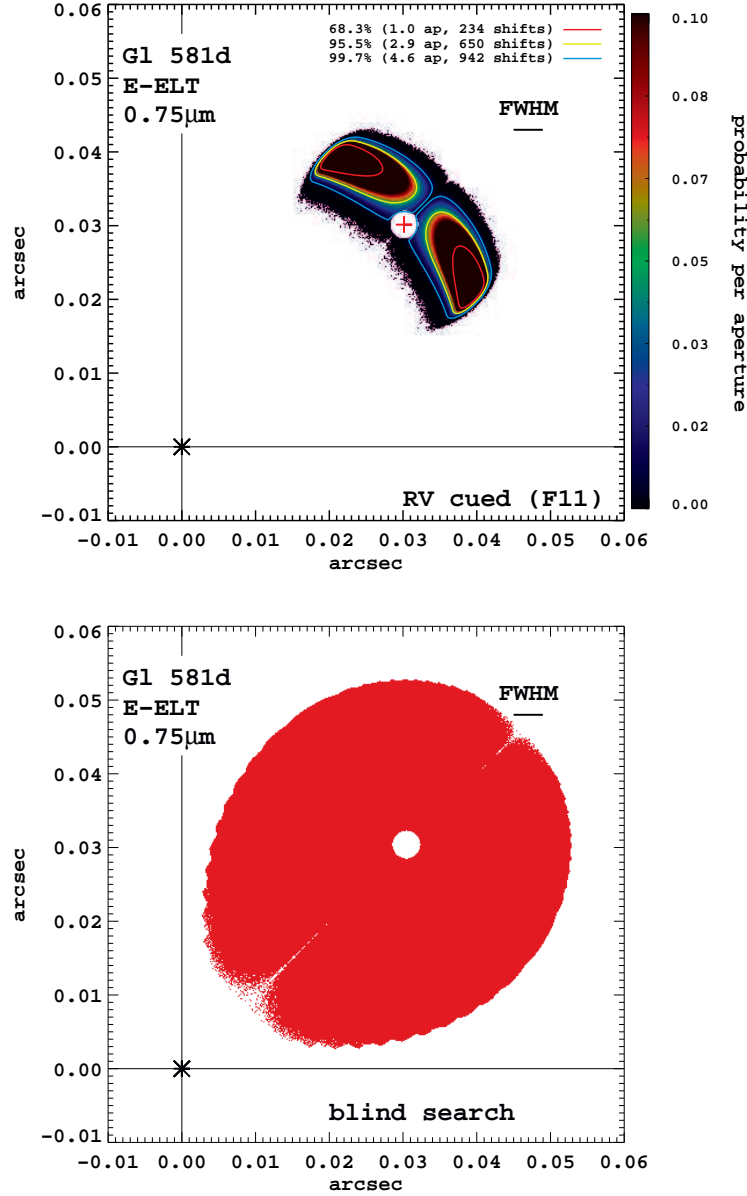


Figure 6.9 Trial orbits for Gl 581d, observed near maximum elongation. In the top panel we use the parameters of Forveille et al. (2011)’s eccentric model. The bottom panel shows the results for a blind search from the same starting point. The red cross shows the starting point, and the star is located at the origin. The top panel color shading is in units of probability per aperture (each aperture has area πFWHM^2). The legend indicates the color which encloses the given completeness intervals, the enclosed area in apertures, and the number of unique whole-pixel shift sequences which must be tried in order to de-orbit the observation. The number of shift sequences is directly related to the false alarm rate, and hence the sensitivity. For comparison, the blind search algorithm produced ~ 12000 shifts. RV cueing greatly improves our sensitivity in the presence of orbital motion.

As in the GMT/ α Cen example, we leave for future work a complete analysis of sensitivity and completeness. The large number of trial shifts calculated when we include uncertainty in the starting position motivates us to suggest that we will ultimately turn this analysis over to a much more robust optimization strategy, such as a Markov Chain Monte Carlo (MCMC) routine. Once an area of the image was identified with a high post-shift SNR , a MCMC analysis could determine the very best orbit and assign robust measures of significance to the result.

We also note that these results likely overestimate the number of trial orbits since we have assumed uncorrelated errors. In reality the RV best fit parameters are likely strongly correlated, which should act to reduce the number of orbits to consider.

6.6 Conclusions

In the coming campaigns to directly image planets in the HZs of nearby stars, orbital motion will be large enough to degrade our sensitivity. This effect has been ignorable in direct imaging campaigns to date, which have typically looked for wide separation planets. We have analyzed this issue in some detail, and shown that applying basic Keplerian orbital mechanics allows us to bound the problem sufficiently that we believe direct imaging in the HZ to be a tractable problem. Our main conclusions are:

- (1) When projected onto the focal plane, a planet in a face-on circular orbit moves with speed given by

$$v_{FOC} = 0.0834 \left(\frac{D}{\lambda d} \right) \sqrt{\frac{M_*}{a}} \text{ FWHM day}^{-1}. \quad (6.37)$$

In the HZ of nearby stars, especially when considering giant telescopes, speeds are high enough that planets will move significant fractions of a PSF FWHM during a single observation. This smears out the planet's flux resulting in a lower SNR .

- (2) In background limited photometry, an SNR maximum is reached after about ~ 2 FWHM of motion has occurred on the focal plane. From there, integrating

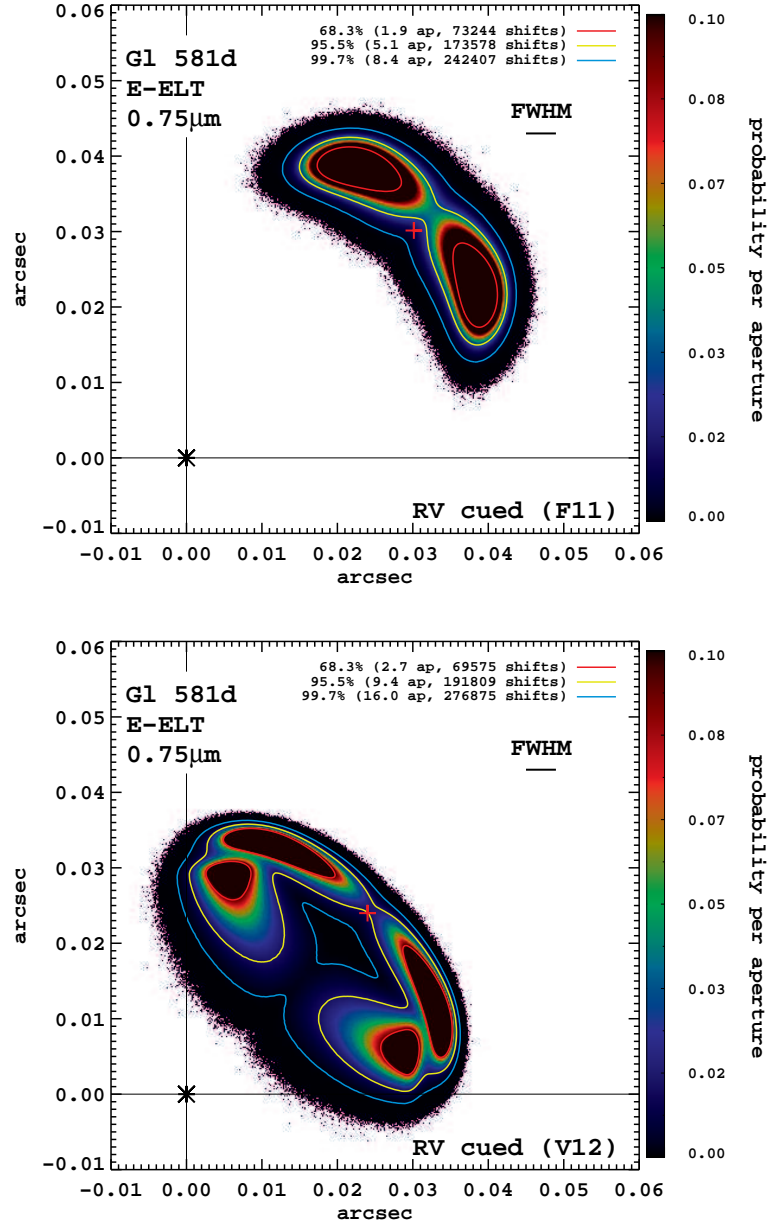


Figure 6.10 Trial orbits for Gl 581d, observed near maximum elongation, assuming the parameters of (top) Forveille et al. (2011)'s Keplerian eccentric model and (bottom) Vogt et al. (2012)'s circular interacting model. In this simulation we allowed the initial position to vary with standard deviation $\sigma_{x,y} = FWHM/SNR$. The red cross shows the starting point, and the star is located at the origin. The color shading is in units of probability per aperture (each aperture has area $\pi FWHM^2$). The legend indicates the color which encloses the given completeness intervals, the enclosed area in apertures, and the number of unique whole-pixel shift sequences which must be tried in order to de-orbit the observation. The number of shift sequences is directly related to the false alarm rate, and hence the sensitivity.

longer offers no improvement with a fixed-size aperture. Adapting the aperture could mitigate this to some extent, but at the cost of significantly longer exposure times.

(3) When SNR is reduced by orbital motion, we have three options. Option I is to do nothing, and accept the loss of completeness due to planets appearing fainter. Option II is to adjust our detection threshold at the cost of more false alarm detections. Option III is to de-orbit an observation, recovering SNR to its nominal value, but also at the cost of more false alarms.

(4) For exposure times of 10s of hours, we expect an observation to extend over several days under realistic assumptions about ground based observing. If we naively attempt to de-orbit such an observation, the false alarm rate per star will increase by at least $FAR \propto \Delta t_{tot}^4$, where Δt_{tot} is the total elapsed time of the observation.

(5) De-orbiting a sequence of shorter exposures is possible, and tractable. Taking advantage of strong correlations between trial orbits, we will realize increases in the FAR on the order of a factor of 10 in the region around a star where orbital motion matters. Since this will be a small, bounded region, this increase in FAR appears to be acceptable.

(6) Cueing from another detection method, such as RV, provides significant benefit. It allows us to initiate our search at the optimum time, and significantly reduces the size of the search space. Having prior distributions for some of the orbital elements will allow us to efficiently determine where and how to search to optimize completeness.

CHAPTER 7

CONCLUSION

In this dissertation we have explored the development and on-sky performance of the world’s newest AO system, MagAO, and its unique ability to contribute to exoplanet science. We started by introducing the MagAO system, and its visible wavelength science camera VisAO. VisAO is the first long-exposure diffraction-limited imager on a large telescope. MagAO is a near clone of the LBT AO systems, and was put through an extensive laboratory integration and testing period based on the LBT process. VisAO was used as the camera for all of these tests, and so was well characterized and a fully integrated part of the AO system by the end of this period. Once deployed at LCO, MagAO and VisAO proceeded to take the highest resolution filled-aperture image ever taken, splitting the 31 mas binary θ^1 Ori C in the Orion Trapezium Cluster for the first time.

Despite our optimism about the prospects for visible wavelength science, we were aware that AO systems frequently under-perform when they reach the telescope. Early in the development of VisAO, the LBT AO systems had not yet been proven on-sky. To prepare for possible disappointment, we developed our own version of Lucky imaging, called Real-Time Frame Selection (RTFS). Here we presented the theoretical justification of this technique, showing that any frame selection technique trades resolution for sensitivity. We then demonstrated the viability of RTFS using a fast shutter, both in simulations and in the laboratory. Luckily, MagAO and VisAO work at least as well as we had hoped, and so we ultimately do not use RTFS on sky.

VisAO’s detection of β Pictoris b, a $10 M_{Jup}$ exoplanet, demonstrates the high performance of this system. We achieved a contrast of 1.7×10^{-5} at only $0.47''$. This high contrast, close separation detection was made possible by the stable 40% optical Strehl ratio PSF in median seeing at LCO. This is the first detection of an

exoplanet from the ground with a CCD. Using our photometry, we can analyze new regions of the exoplanet’s atmosphere.

MagAO’s contributions to exoplanet science have only just begun. With its unique O/IR wavelength coverage and stable high-contrast PSF, we will be able to use MagAO to probe the Habitable Zones of select nearby stars for the first time. In the thermal-IR, we are sensitive to gas-giant exoplanets with very low masses — as small as $0.03 M_{Jup}$. With VisAO, we will take the first high-resolution, high-contrast images of the close surroundings of these stars in visible light. While searching for planets, these observations will teach us many lessons about the challenges of long exposures in this regime. We are “blazing the trail” for the next generation of giant telescopes.

We closed with an analysis of one problem that the giant telescope planet-hunters will have to deal with: orbital motion. As we begin to probe the HZ, we will find that planets move enough over the course of a single observation that they smear, resulting in reduced S/N. This problem is correctable, but not without some work. We showed that Keplerian mechanics can be used to bound the problem, and recover most of the lost sensitivity.

When I started graduate school in the Fall of 2008, the number of known exoplanets was counted in the tens, and the imaging of the HR8799, Fomalhaut, and β Pic planets had not yet been announced. In the last 5 years the number of exoplanets has grown to over 3000, mainly thanks to *Kepler*, and we have had some success with direct imaging. It is an exciting time in exoplanet science, and as we continue to improve our instruments it will continue to be so.

APPENDIX A

POINT SPREAD FUNCTION RADIOMETRY AND PHOTOMETRY

Here I present a collection of derivations useful for analyzing AO corrected point spread functions (PSFs). These are generally straightforward and can be found in one form or another in various textbooks. They are given here to provide a common notation and common units.

A.1 PSF Modeling

The obscured Airy pattern is

$$I(x) = \frac{I_o}{(1 - \epsilon^2)^2} \left(\frac{2J_1(\pi x)}{\pi x} - \frac{2\epsilon J_1(\epsilon\pi x)}{\pi x} \right)^2. \quad (\text{A.1})$$

where ϵ is the obscuration fraction and x is the radial coordinate in units of λ/D . The quantity represented by I is the irradiance, and has units of energy per time per area, or power per area. For the purpose of modeling the performance of photon-detecting arrays of pixels at the focus of a telescope, the natural unit for power is photons/sec/ $(\lambda/D)^2$. Here we are using λ/D as the unit of length in the focal plane.

Expanding Equation (A.1) gives us

$$I(x) = \frac{4I_o}{\pi^2(1 - \epsilon^2)^2} \left[\left(\frac{J_1(\pi x)}{x} \right)^2 - 2\epsilon \frac{J_1(\pi x)J_1(\pi\epsilon x)}{x^2} + \epsilon^2 \left(\frac{J_1(\pi\epsilon x)}{x} \right)^2 \right].$$

which we can integrate to find the power enclosed by the Airy pattern at radius x'

$$P_{enc}(x') = \frac{8I_o}{\pi(1 - \epsilon^2)^2} \left[\int_0^{x'} \left(\frac{J_1(\pi x)}{x} \right)^2 x dx + \epsilon^2 \int_0^{x'} \left(\frac{J_1(\pi\epsilon x)}{x} \right)^2 x dx - 2\epsilon \int_0^{x'} \frac{J_1(\pi x)J_1(\pi\epsilon x)}{x} dx \right]$$

which simplifies to

$$P_{enc}(x') = \frac{4I_o}{\pi(1 - \epsilon^2)^2} \left[P_c(\pi x) + \epsilon^2 P_c(\pi\epsilon x) - 4\epsilon \int_0^{x'} \frac{J_1(\pi x)J_1(\pi\epsilon x)}{x} dx \right] \quad (\text{A.2})$$

where we have followed Mahajan (1986) in defining

$$P_c(y) = 1 - J_0^2(y) - J_1^2(y).$$

Now if we let $x' \rightarrow \infty$ this becomes

$$P_{tot} = \frac{4I_o}{\pi(1 - \epsilon^2)^2} \left[1 + \epsilon^2 - 4\epsilon \int_0^\infty \frac{J_1(x)J_1(\epsilon x)}{x} dx \right]$$

so

$$P_{tot} = \frac{4I_o}{\pi(1 - \epsilon^2)}$$

Now we have the peak of the Airy pattern in terms of P_{tot} .

$$I_o = \frac{\pi P_{tot}(1 - \epsilon^2)}{4} \quad (\text{A.3})$$

P_{tot} is conveniently expressed as the photon flux from a star in a given bandpass:

$$P_{tot} = F_0 10^{-0.4m} \frac{\pi}{4} D^2 (1 - \epsilon^2)$$

where F_0 is the flux, in photons $m^{-2} \text{ sec}^{-1}$, of a 0 magnitude star taking into account filter transmission, QE, etc.

Another useful quantity when applying the above expressions is the plate-scale of the detector. Each pixel has a fixed projected angular size, measured in arcseconds, which we denote here with ps . The area of a pixel, in λ/D units, is just

$$A_{pix} = \left(\frac{psD}{0.2063\lambda} \right)^2$$

where λ is in μm and D is in m. For VisAO $ps \approx 0.00791$ so

$$A_{pix}^{VisAO} = \frac{0.249}{\lambda^2}.$$

A.2 Characterizing the PSF

Several quantities are used to characterize the PSF, including the full-width at half maximum ($FWHM$), the radius of the first minimum, and the fraction of the total

Quantity	$\epsilon = 0$	$\epsilon = 0.11$	$\epsilon = 0.29$
$FWHM$	$1.03 \lambda/D$	$1.02 \lambda/D$	$0.98 \lambda/D$
1^{st} Min.	$1.22 \lambda/D$	$1.20 \lambda/D$	$1.12 \lambda/D$
1^{st} Max.	$1.63 \lambda/D$	$1.63 \lambda/D$	$1.62 \lambda/D$
$I(1^{st} \text{ Max.})/I_o$	0.0175	0.0212	0.0455
$\mathcal{E}_P(FWHM/2)$	0.475	0.467	0.432
$\mathcal{E}_P(1^{st} \text{ Min.})$	0.838	0.824	0.747

Table A.1 Quantities typically used to characterize a PSF. For comparison, the values for an unobscured aperture, and for $\epsilon = 0.11$ (LBT & MMT) and $\epsilon = 0.29$ (MagAO) are given.

power contained within various radii ($\mathcal{E}_P(r)$). These quantities depend on the central obscuration. Using Equations (A.1) and (A.2) yields the values in Table A.1.

It is also convenient to introduce the notation

$$FWHM = \kappa \frac{\lambda}{D}$$

where κ is the coefficient in the first row of Table A.1. While typically negligible, on the MagAO system where the central obscuration $\epsilon = 0.29$, there is actually a 5% change in the $FWHM$ of the Airy pattern - a detectable difference of, e.g., 0.2 pixels at 1 micron on VisAO. In Section A.3 below we will see other ways that κ should be taken into account.

A.3 Tip & Tilt Errors, FWHM, and Strehl Ratio

If we have residual tip & tilt (TT) errors in the focal plane after AO correction, the result will be a smeared out PSF. We can model TT errors (jitter) as a displacement in the position of the PSF center, described by a 2D Gaussian probability distribution

$$p_{TT}(r) = \frac{1}{2\pi\sigma_{TT}^2} e^{-r^2/(2\sigma_{TT}^2)}$$

where σ_{TT} is the RMS displacement of the PSF center. Now we can assume that the PSF is itself a Gaussian, with $FWHM = \kappa\lambda/D$ (κ defined above), which means

a Gaussian width parameter (a.k.a. standard deviation) of

$$\sigma_{PSF} = \frac{FWHM}{2\sqrt{2\ln(2)}}.$$

The resulting smeared out image will just be the convolution of the TT probability distribution and the PSF. As such it will have a width parameter given by $\sigma^2 = \sigma_{PSF}^2 + \sigma_{TT}^2$. This yields the following estimate for the FWHM of a PSF with TT errors

$$FWHM_{TT} = \sqrt{\left(\kappa \frac{\lambda}{D}\right)^2 + 8\ln(2)\sigma_{TT}^2}. \quad (\text{A.4})$$

TT errors will also be evident as a lower peak height, i.e. a lower Strehl ratio. The degraded peak height can be estimated by calculating

$$I_{o,TT} = \frac{1}{2\pi\sigma_{TT}^2} \int_0^\infty I_o e^{-r^2/(2\sigma_{PSF}^2)} e^{-r^2/(2\sigma_{TT}^2)} r dr d\theta.$$

Now $S_{TT} = I_{o,TT}/I_o$ so

$$S_{TT} = \frac{1}{1 + \frac{8\ln(2)}{\kappa^2} \left(\frac{\sigma_{TT}}{\lambda/D}\right)^2}. \quad (\text{A.5})$$

It is worth noting that these two expressions are slightly different than others that have appeared in the literature. Some of the difference can be explained by other authors (e.g. Tyson (2011), Hardy (1998)) using $1.22\lambda/D$ as the FWHM - this is actually the radius of the first Airy minimum for an unobscured aperture. (Hardy, 1998) gives 5.17 instead of $8\ln(2)/1.03^2$, which appears to just be a rounding difference. Another subtle difference is the constant $8\ln(2)$ vs $\pi^2/2$ given by Sandler et al. (1994). These are numerically somewhat similar, 5.5 compared to 4.9. I have so far been unable to find an explanation for this discrepancy.

Again we note the importance of κ for MagAO. In Equation (A.5) it enters as a square, and so has over a 10% impact - MagAO is more sensitive to tip-tilt errors than an unobscured aperture.

A.4 Gaussian PSF Photometry

We use the ‘‘CCD equation’’ to write the signal to noise ratio

$$SNR = \frac{I_{tot}}{\sqrt{N_{tot}}}$$

Where I_{tot} is the total signal collected, and N_{tot} is the variance of the measurement. For a Gaussian PSF, the SNR , in an aperture of radius x , is

$$SNR = \frac{2\pi I_o}{\sqrt{\pi x^2 N}} \int_0^x \exp\left(-\frac{r^2}{2\sigma^2}\right) r dr$$

where I_o is the peak of the Gaussian, N is the variance per unit area, and we have already integrated over the polar angle θ . The width parameter σ (a.k.a. standard deviation) is related to the full-width-at-half-maximum as

$$\sigma = \frac{FWHM}{2\sqrt{2\ln(2)}}$$

Integrating gives

$$SNR = \frac{\sqrt{\pi}}{4\ln(2)} \frac{FWHM^2}{x\sqrt{N}} \left(1 - \exp\left(\frac{-4\ln(2)x^2}{FWHM^2}\right)\right) I_o$$

Differentiating with respect to aperture radius x leads to the transcendental equation

$$\exp\left(\frac{-4\ln 2}{FWHM^2}x^2\right) = \frac{1}{1 + \frac{8\ln(2)}{FWHM^2}x^2}$$

which has a minimum at $x = 0$ and a maximum at $x \approx 0.7FWHM$.

If we set $x = FWHM$ we have

$$SNR = 0.599 \frac{FWHM}{\sqrt{N}} I_o.$$

It is convenient to note that in the peak pixel $SNR_{pk} = I_o/\sqrt{N}$, so

$$SNR_{ap} = 0.599FWHM \times SNR_{pk}$$

is the correction from peak pixel to the full aperture SNR_{ap} .

A.5 Propagation of Errors

A.5.1 Absolute Magnitude

The absolute magnitude is given by

$$M_* = m_* + 5\log(d_{pc}) - 5$$

where d_{pc} is the distance of the star in parsecs. Propagation of errors for this quantity yields

$$\sigma_{M_*}^2 = \sigma_{m_*}^2 + \left(\frac{5}{\ln(10)} \right)^2 \frac{\sigma_{d_{pc}}^2}{d_{pc}^2}$$

A.5.2 Physical Photometry

The flux-density of a star in a given bandpass is

$$F_* = F_0 10^{-m_*/2.5}$$

where m_* is the apparent magnitude of the star in the bandpass and F_0 is the flux of an $m_* = 0$ mag star in the bandpass (the zero-point). Propagation of errors for this quantity gives

$$\sigma_{F_*}^2 = \left(\frac{\partial F_*}{\partial F_0} \right)^2 \sigma_{F_0}^2 + \left(\frac{\partial F_*}{\partial m_*} \right)^2 \sigma_{m_*}^2.$$

The partial derivatives are

$$\begin{aligned} \frac{\partial F_*}{\partial F_0} &= 10^{-m_*/2.5} \\ \frac{\partial F_*}{\partial m_*} &= \frac{-\ln(10)}{2.5} 10^{-m_*/2.5} F_0 \end{aligned}$$

so

$$\sigma_{F_*}^2 = \left[\sigma_{F_0}^2 + \left(\frac{\ln(10)}{2.5} \right)^2 F_0^2 \sigma_{m_*}^2 \right] 10^{-2m_*/2.5}$$

or

$$\frac{\sigma_{F_*}}{F_*} = \left[\frac{\sigma_{F_0}^2}{F_0^2} + \left(\frac{\ln(10)}{2.5} \right)^2 \sigma_{m_*}^2 \right]^{1/2}. \quad (\text{A.6})$$

Similarly, if F_* is known the magnitude of a star is

$$m_* = -2.5 \log_{10} \left(\frac{F_*}{F_0} \right).$$

Propagation of errors gives

$$\sigma_{m_*}^2 = \left(\frac{\partial m_*}{\partial F_*} \right)^2 \sigma_{F_*}^2 + \left(\frac{\partial m_*}{\partial F_0} \right)^2 \sigma_{F_0}^2$$

and the partial derivatives are

$$\begin{aligned}\frac{\partial m_*}{\partial F_*} &= \frac{-2.5}{F_0 \ln(10)} \\ \frac{\partial m_*}{\partial F_0} &= \frac{2.5}{F_* \ln(10)}\end{aligned}$$

so we have

$$\sigma_{m_*}^2 = \left(\frac{2.5}{\ln(10)} \right)^2 \left[\left(\frac{\sigma_{F_*}}{F_*} \right)^2 + \left(\frac{\sigma_{F_0}}{F_0} \right)^2 \right]. \quad (\text{A.7})$$

APPENDIX B

SYNTHETIC PHOTOMETRY AND CONVERSIONS

In this appendix we provide details of our synthetic photometry. The primary purpose of this analysis is to verify the methodology used for our analysis of β Pic b, but we also determine transformations between various filter systems used in brown dwarf and exoplanet imaging which may be useful to others.

B.1 Filters

Here we describe the process by which we assembled synthetic filter profiles for this analysis. In general, we obtained a transmission profile, and determined an atmospheric transmission profile appropriate for the site. Table B.1 summarizes the atmosphere assumptions and models used. We finally converted to photon-weighted “relative spectral response” (RSR) curves, using the following equation (Bessell, 2000)

$$T(\lambda) = \frac{1}{hc} \lambda T_0(\lambda) \quad (\text{B.1})$$

where T_0 is the raw energy-weighted profile.

In the following subsections we describe details particular to the different photometric systems and passbands. Comparisons of the filters in each band pass are shown in Figures B.1, B.2, B.3 and B.4.

B.1.1 The Y Band

The Y band was first defined in Hillenbrand et al. (2002). We follow Liu et al. (2012) and assume that the UKIDSS Y filter defines the MKO system passband, as the largest number of published observations in this passband are from there (cf. Burningham et al. (2013)). This is a slightly narrow version of the filter. The

Table B.1 Atmospheres.

System	Model	Airmass	PWV (mm)	Notes
VisAO	ATRAN, Cerro Pachon	1.0	2.3	1,2
2MASS	PLEXUS	1.0	5.0	3
MKO	ATRAN, Mauna Kea	1.0	1.6	1
UKIDSS		1.3	1.0	4
NACO	Paranal-like	1.0	2.3	5
NICI	ATRAN, Cerro Pachon	1.0	2.3	1

Notes:

[1] Lord (1992)

[2] C. Manqui elevation is 2380 m, C. Pachon is 2700 m.

[3] Cohen et al. (2003)

[4] Hewett et al. (2006)

[5] 0.4-6.0 μm atmospheric transmission for Paranal from ESO.

UKIDSS Y RSR curve is provided in Hewett et al. (2006), which is already photo-normalized and includes an atmosphere appropriate for Mauna Kea.

We also consider the the unfortunately named Z filter used at Subaru/IRCS and Keck/NIRC2, which is actually in the Y window rather than in the traditionally optical Z/z band. To add to the confusion the filter has been labeled with a lower-case z , which is how it is referred to in Currie et al. (2011a), but the scanned filter curves and Alan Tokunaga’s website¹ indicate that it was meant to be capital Z . In any case, it is a narrow version of the Y filter of Hillenbrand et al. (2002). Here we follow Liu et al. (2012) and refer to it as $z_{1.1}$ to emphasize its location in the Y window. We used the same atmospheric assumptions as for the MKO system (see below).

The VisAO Y_S (Y -short) passband is defined by a Melles-Griot long-pass dichroic filter at $\sim 950\mu\text{m}$ (LPF-950) on the blue side, and limited by CCD QE ($\sim 1.1\mu\text{m}$) on the red side. This places the Y_S filter on the blue edge of the Y atmospheric window. We convolved the transmission curve from the manufacturer catalog with the quantum efficiency (QE) for our EEV CCD47-20 with near-IR coating, and

¹http://www.ifa.hawaii.edu/~tokunaga/MKO-NIR_filter_set.html

included the effects of 3 Al reflections. We used the ATRAN model atmosphere (Lord, 1992), provided by Gemini Observatory ², for Cerro Pachon with 2.3 mm precipitable water vapor (PWV). Cerro Pachon, $\sim 2700\text{m}$, is slightly higher than the Magellan site at Cerro Manqui, $\sim 2380\text{m}$ (D. Osip, private communication), so this will somewhat underestimate atmospheric absorption.

B.1.2 The 2MASS System

The 2MASS J , H , and K_S transmission and RSR profiles were collected from the 2MASS website³. The RSR profiles are from Cohen et al. (2003). They used an atmosphere based on the PLEXUS model for 1.0 airmass. This model does not use a parameterization corresponding directly to PWV, but according to the website it is equivalent to 5.0 mm PWV.

B.1.3 The MKO System

We used the Mauna Kea filter profiles provided by the NSFCam website⁴ for the MKO J , H , K_S , and K passbands. According to Alan Tokunaga’s website these correspond to the 1998 production run of these filters. We again used the ATRAN model atmosphere from Gemini, now for Mauna Kea with 1.6 mm precipitable water vapor (PWV) at 1.0 airmass.

B.1.4 The NACO System

We obtained transmission profiles for NACO from the instrument website⁵. We used the “Paranal-like” atmosphere provided by ESO⁶, which is for 1.0 airmass and 2.3 mm PWV. The NACO filters are close to the 2MASS system. As shown in Figures B.2 through B.4 there are subtle differences, which are somewhat more pronounced once the atmosphere appropriate for each site is included.

²<http://www.gemini.edu/?q=node/10789>

³http://www.ipac.caltech.edu/2mass/releases/allsky/doc/sec6_4a.html

⁴<http://irtfweb.ifa.hawaii.edu/~nsfcam/filters.html>

⁵<http://www.eso.org/sci/facilities/paranal/instruments/naco/inst/filters.html>

⁶http://www.eso.org/sci/facilities/eelt/science/drm/tech_data/data/atm_abs/

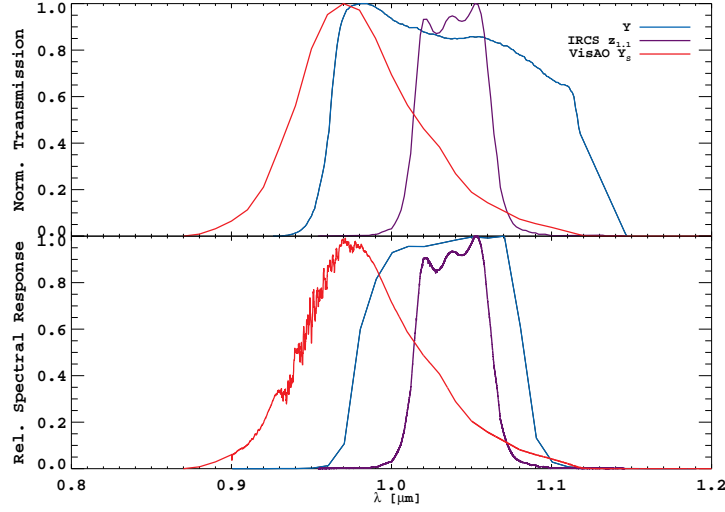


Figure B.1 Comparison of filters in the Y atmospheric window. The top panel shows the curves prior to applying the atmosphere and photo-normalizing. The Y curve in the top panel is from Hillenbrand et al. (2002). The bottom panel shows the curves after being multiplied by atmospheric transmission and converted to RSR, and in the case of Y itself this is the UKIDSS curve from Hewett et al. (2006).

B.1.5 The NICI System

Profiles for the NICI filters were obtained from the instrument websites for NICI⁷ and NIRI⁸. The Cerro Pachon ATRAN atmosphere was used, with 1.0 airmass and 2.3 mm PWV. The NICI J , H , and K_S bandpasses are intended to be in the MKO system, but due to differences in the altitude and transmission between Cerro Pachon and Mauna Kea, as well as subtle differences between the filter profiles we used, there are noticeable differences in the filters. As we will see these can be appreciable for cool brown dwarfs.

⁷<http://www.gemini.edu/sciops/instruments/nici/>

⁸<http://www.gemini.edu/sciops/instruments/niri/>

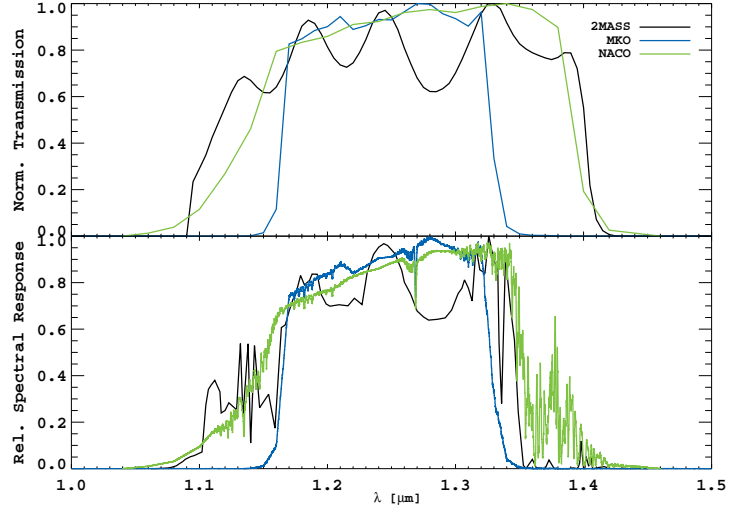


Figure B.2 Comparison of filters in the J atmospheric window. The top panel shows raw transmission profiles, and the bottom shows our final RSR curves.

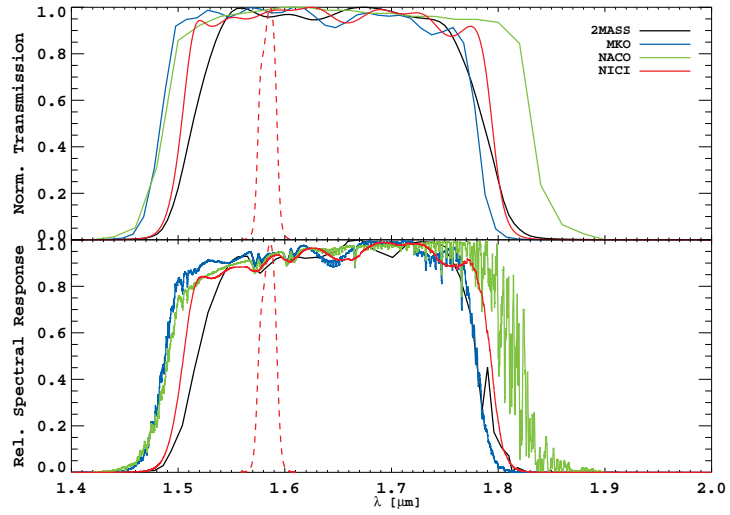


Figure B.3 Comparison of filters in the H atmospheric window. The top panel shows raw transmission profiles, and the bottom shows our final RSR curves.

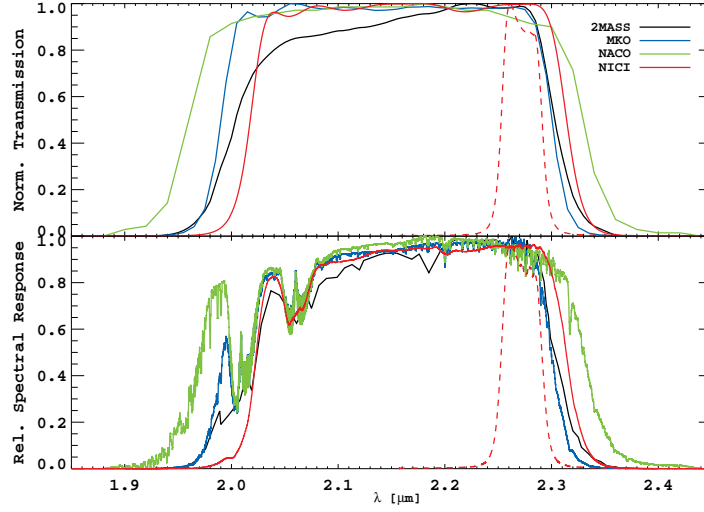


Figure B.4 Comparison of filters in the K atmospheric window. The top panel shows raw transmission profiles, and the bottom shows our final RSR curves.

B.2 Synthetic Photometry

Having collected or synthesized RSR profiles for each bandpass, we next proceed to calculate various quantities of interest, and ultimately fluxes, in each. The effective wavelength λ_0 of each filter is calculated as

$$\lambda_0 = \frac{\int_0^\infty \lambda R(\lambda) d\lambda}{\int_0^\infty R(\lambda) d\lambda}$$

where $T(\lambda)$ is filter RSR. The effective width $\Delta\lambda$, such that

$$F_\lambda(\lambda_0) \Delta\lambda = \int_0^\infty F_\lambda(\lambda) R(\lambda) d\lambda$$

was also calculated, where F_λ is the flux in the bandpass. Finally we integrated the filter profiles with the HST CALSPEC spectrum of Vega from Bohlin (2007) to determine the flux densities of a 0 mag star in each filter. These calculations are summarized in Table B.2.

Now to calculate the magnitude of some object with a spectrum given by $F_{\lambda, obj}$

Table B.2 Synthetic Photometric System Characteristics

Filter	λ_0 (μm)	$\Delta\lambda$ (μm)	0 mag F_λ (10^{-6} ergs/s/ $cm^2/\mu\text{m}$)
<i>Y</i> Band			
VisAO Y_S	0.986	0.087	6.91
MKO Y	1.032	0.101	5.99
IRCS $z_{1.1}$	1.039	0.049	5.89
<i>J</i> Band			
2MASS J	1.241	0.163	3.23
MKO J	1.249	0.145	3.11
NACO J	1.256	0.192	3.1
<i>H</i> Band			
NICI $CH_{4,1\%S}$	1.584	0.017	1.31
MKO H	1.634	0.277	1.22
2MASS H	1.651	0.251	1.18
NACO H	1.656	0.308	1.18
NICI H	1.658	0.270	1.18
<i>K</i> Band			
MKO K_S	2.156	0.272	0.45
NACO K_S	2.160	0.323	0.45
2MASS K_S	2.166	0.262	0.443
NICI K_S	2.176	0.268	0.436
MKO K	2.206	0.293	0.414
NICI K_{cont}	2.272	0.038	0.366

we calculate

$$m = -2.5 \log \left[\frac{\int_0^\infty R(\lambda) F_{\lambda, obj} d\lambda}{\int_0^\infty R(\lambda) F_{\lambda, vega} d\lambda} \right] \quad (\text{B.2})$$

using the Vega spectrum of Bohlin (2007).

B.3 Photometric Conversions

To quantify the differences between these systems, and to accurately compare results for objects with measured in the different systems, we used the library of brown dwarf spectra we compiled from various sources (described in Chapter 4). We calculated the magnitudes in each of the various filters and then fit a 4th or 5th order polynomial to the results. Our notation is

$$m_1 - m_2 = c_0 + c_1 SpT + c_2 SpT^2 + c_3 SpT^3 + c_4 SpT^4 + c_5 SpT^5. \quad (\text{B.3})$$

where SpT is the spectral type given by

$$SpT = 0...9, \text{ for } M0...M9$$

$$SpT = 10...19, \text{ for } L0...L9$$

$$SpT = 20...29, \text{ for } T0...T9$$

We provide the coefficients determined in this manner for a variety of transformations in Table B.3. Next we discuss some of the results.

B.3.1 Converting 2MASS to MKO

To begin with, we consider the conversions from 2MASS to MKO. There are many objects with measurements in both systems, which allows us to directly compare our synthetic photometry to actual measurements. We here use the compilation of Dupuy and Liu (2012). This also allows a comparison to the previous work of Stephens and Leggett (2004), who employed similar methodology to ours but with

fewer objects. The results are shown graphically in Figure B.5. In all three bands our synthetic photometry and fit appear to be a good match to the measurements. In J our results appear to be an improvement over Stephens and Leggett (2004), and in H and K either fit appears to be reasonable. These results give confidence that our synthetic photometry reproduces the variations in these systems reasonably well.

B.3.2 Y Band Conversions

We analyze the variation with spectral type of between the Y_{MKO} and the VisAO Y_S and IRCS/NIRC2 $z_{1.1}$ filters. In our analysis we did not use these conversions, but provide them here to further characterize the bandpasses.

B.3.3 J Band Conversions

We present the relationship between the J_{MKO} and the J_{NACO} bandpasses in Figure B.7. Note that the differences are quite significant for later brown dwarf spectral types.

B.3.4 H Band Conversions

The shifts in the various H bandpasses are presented in Figure B.8. These shifts are generally small, and within the errors of typical exoplanet photometry. Note, however, the large shift for later spectral types between the narrow band $CH4_{1\%S}$ filter of NICI.

B.3.5 K Band Conversions

We show the differences in the various K_S bandpasses in Figure B.9. These conversions result in small shifts for the M and L dwarfs, but there are more noticeable differences for T dwarfs.

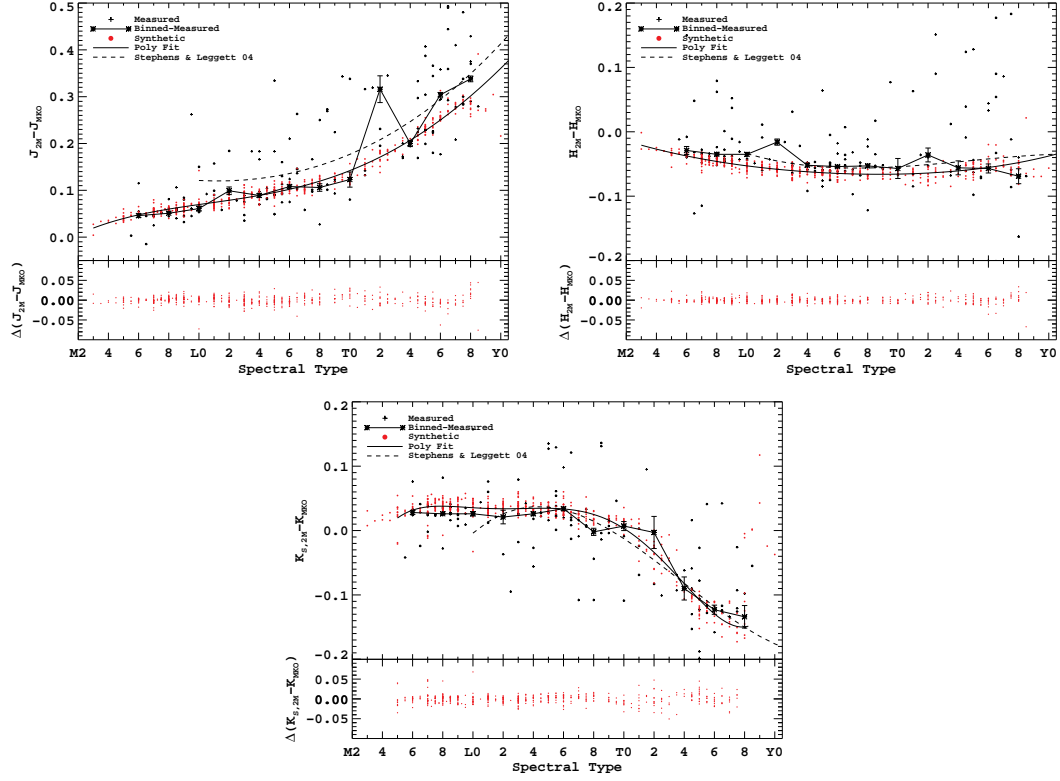


Figure B.5 2MASS to MKO Conversions. Here we show our synthetic photometry (red points) in the 2MASS and MKO systems and compare to the measurements made in the 2 systems (crosses, from Dupuy and Liu (2012)). We also plot the binned median of the measurements. Our polynomial fit is shown as the solid black line. For comparison we show the fit determined by Stephens and Leggett (2004), who also used synthetic photometry, albeit with fewer objects. Our fit to the synthetic photometry appears to be a better match to the actual measurements in J. In the H and K bands both fits appear to be acceptable, with Stephens and Leggett (2004) being somewhat better for M and L dwarfs in H. The coefficients of our fits are given in Table B.3. These results give confidence in our synthetic photometry.

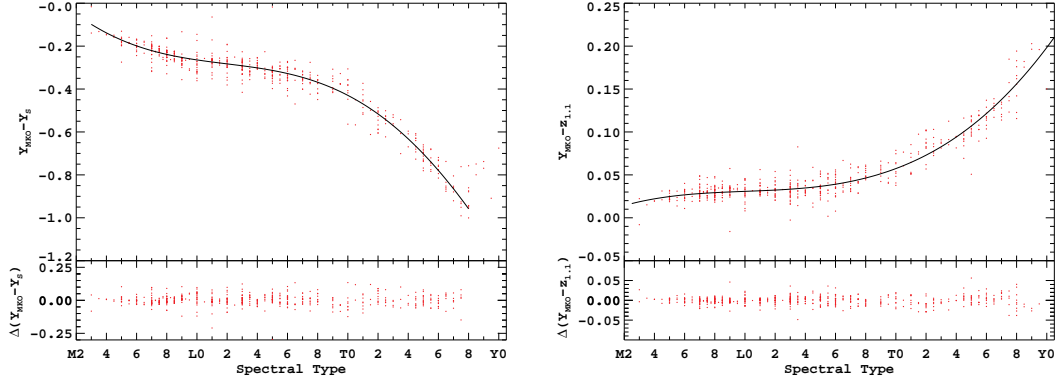


Figure B.6 Here we show the relationships between the Y_{MKO} and the VisAO Y_S and IRCS/NIRC2 $z_{1.1}$ filters. The coefficients of our fits are given in Table B.3.

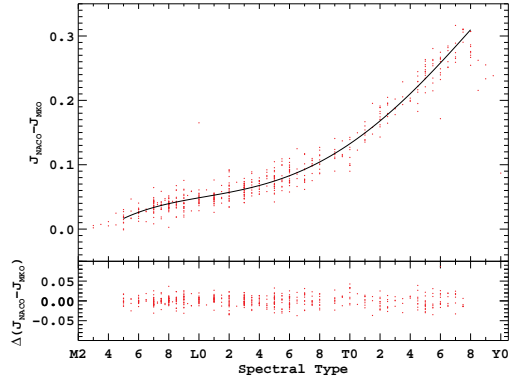


Figure B.7 Here we show the relationship between the J_{MKO} and the J_{NACO} bandpasses. The coefficients of our fit are given in Table B.3.

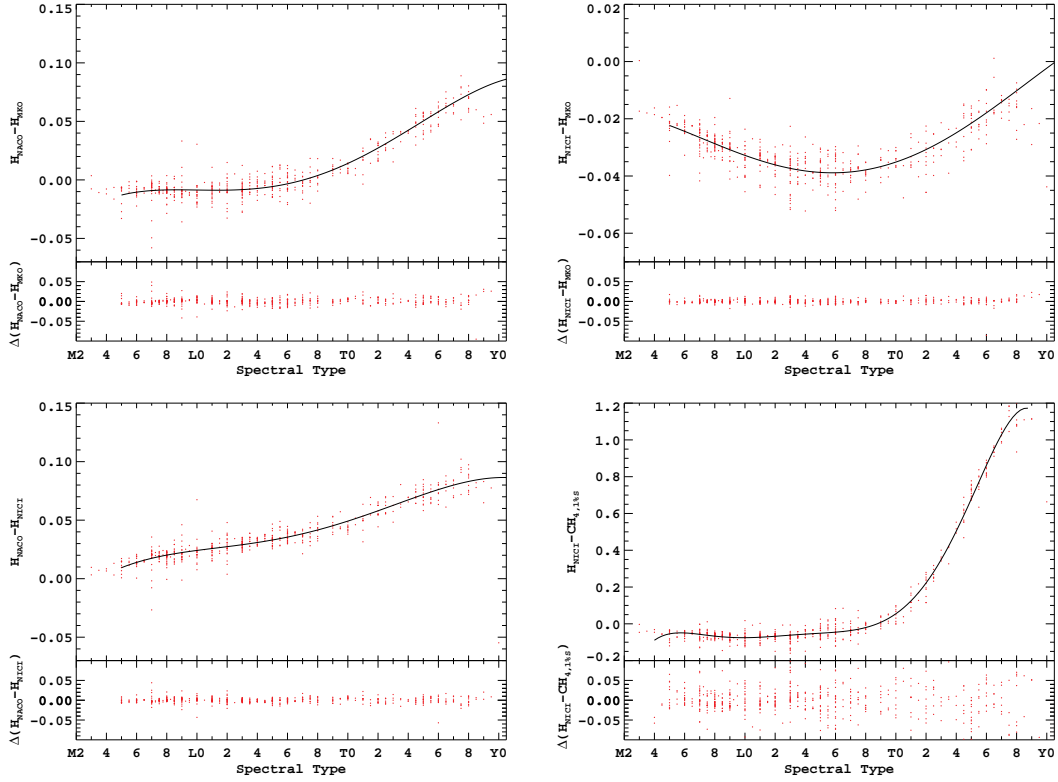


Figure B.8 Here we show the relationship between the H_{MKO} , H_{NACO} , and H_{NICI} bandpasses. These conversions are generally small, and within the errors of typical exoplanet photometry. We also compare the H_{NICI} filter and NICI's $CH_{4,1.6}$ narrow band filter, which can have dramatic shifts for the T dwarfs. The coefficients of our fits are given in Table B.3.

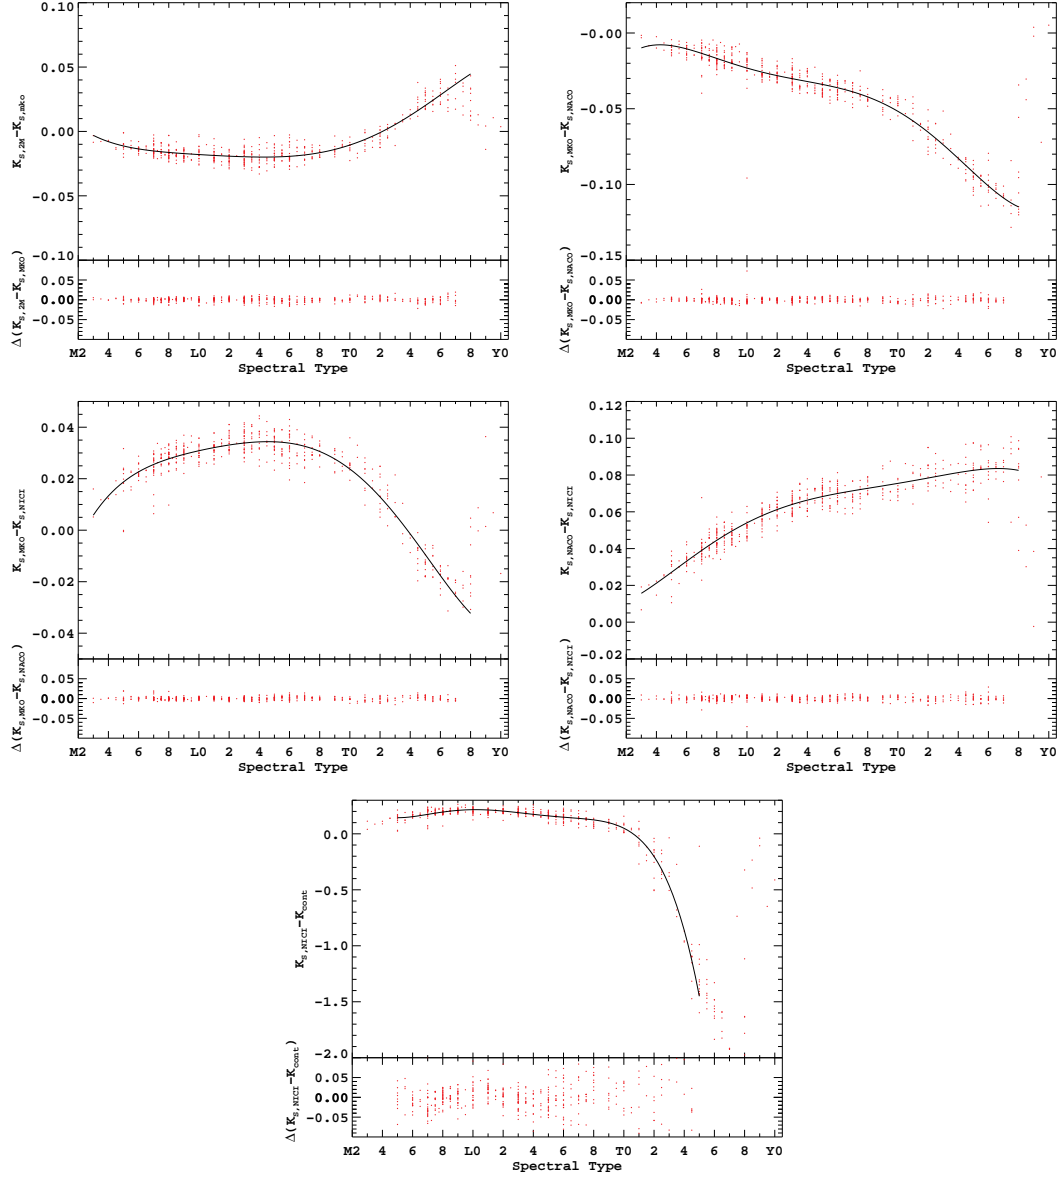


Figure B.9 Here we show the relationship between the $K_{S,2MASS}$, $K_{S,MKO}$, $K_{S,NACO}$, and $K_{S,NICI}$ bandpasses. These corrections are generally small, but for later T dwarfs begin to be significant. We also compare the $K_{S,NICI}$ filter and NICI's K_{cont} narrow band filter. Note, the large shift for later spectral types in this filter is in the opposite direction (brighter) than for the $CH4_{1\%S}$ filter (fainter) of NICI. The coefficients of our fits are given in Table B.3.

Table B.3 Photometric conversion coefficients.

Filters	c_0	c_1	c_2	c_3	c_4	c_5	σ
$Y_{MKO} - Y_S$	0.0664	-0.0671	0.00423	-4.19×10^{-5}	-4.88×10^{-6}	8.48×10^{-8}	0.048
$Y - z_{1.1}$	0.00226	0.00727	-0.000671	2.37×10^{-5}	-7.09×10^{-8}	0.0	0.012
$J_{2M} - J_{MKO}$	-0.0281	0.02	-0.00159	6.16×10^{-5}	-5.51×10^{-7}	0.0	0.012
$J_{NACO} - J_{MKO}$	-0.0866	0.0334	-0.0032	0.00014	-1.78×10^{-6}	0.0	0.015
$H_{2M} - H_{MKO}$	0.00165	-0.00858	0.000398	-1.02×10^{-5}	1.68×10^{-7}	0.0	0.008
$H_{NACO} - H_{MKO}$	-0.0553	0.0157	-0.00184	8.61×10^{-5}	-1.24×10^{-6}	0.0	0.012
$H_{NICI} - H_{MKO}$	-0.0172	0.000395	-0.000389	2.25×10^{-5}	-3.14×10^{-7}	0.0	0.008
$H_{NACO} - H_{NICI}$	-0.0381	0.0153	-0.00145	6.36×10^{-5}	-9.23×10^{-7}	0.0	0.007
$H_{NICI} - CH_{4,1\%S}$	-1.23	0.671	-0.147	0.0159	-0.000908	0.0	0.037
$K_{S,2M} - K_{MKO}$	-0.313	0.145	-0.0227	0.00167	-5.79×10^{-5}	7.41×10^{-7}	0.013
$K_{S,2M} - K_{S,MKO}$	0.027	-0.0154	0.00226	-0.00017	6.08×10^{-6}	-7.75×10^{-8}	0.006
$K_{S,MKO} - K_{S,NACO}$	-0.0433	0.0211	-0.00424	0.00034	-1.23×10^{-5}	1.59×10^{-7}	0.007
$K_{S,MKO} - K_{S,NICI}$	-0.0371	0.0214	-0.00294	0.000214	-7.65×10^{-6}	9.92×10^{-8}	0.005
$K_{S,NACO} - K_{S,NICI}$	0.00619	0.00025	0.00131	-0.000125	4.6×10^{-6}	-6×10^{-8}	0.007
$K_{S,NICI} - K_{cont}$	1.05	-0.517	0.107	-0.00988	0.000424	-6.91×10^{-6}	0.057

APPENDIX C

FOUR DECADES OF IRC +10216: EVOLUTION OF A CARBON RICH DUST
SHELL RESOLVED AT $10\mu\text{M}$ WITH MMT ADAPTIVE OPTICS AND
MIRAC4

C.1 Introduction

This appendix has previously been published in the *Astrophysical Journal* (Males et al., 2012b). This research formed part of my 2nd year project. It is included here mainly to demonstrate that AO works at thermal wavelengths, and that we have experience working with challenging thermal backgrounds.

C.1.1 The carbon star IRC +10216

When stars of low to intermediate mass are in the last stages of nuclear burning on the asymptotic giant branch (AGB) of the Hertzsprung-Russel (HR) diagram, they are typically characterized by high luminosity, which varies with long periods (1 – 2yrs), and mass loss. The high mass loss rates, up to $\dot{M} \sim 10^{-4} M_{\odot} \text{ yr}^{-1}$, ultimately lead to the termination of nuclear burning and produce dusty, often optically thick, circumstellar envelopes (CSEs) which provide one of the key observational features of AGB stars. Of particular concern here, the CSE is often the dominant source of light in the near and mid-infrared (IR) and contains much information about the evolution and mass loss history of the enshrouded star. For thorough treatments of AGB stars and their evolution see Habing and Olofsson (2003) and Herwig (2005), and references therein.

The carbon star IRC +10216 (CW Leo) is perhaps the best studied example of an AGB star. Since its discovery by the $2.2 \mu\text{m}$ survey in 1969 (Becklin et al., 1969), IRC +10216 has been recognized as a star enshrouded by a thick CSE. It exhibits large ($> 2\text{X}$) changes in luminosity over its 649 day cycle (Le Bertre, 1992) and is

extremely bright in the mid-IR ($> 10^4$ Jy). IRC +10216 is classified as a carbon star (Herbig and Zappala, 1970), implying that the ratio of carbon to oxygen in its photosphere is greater than 1. It is believed to be in the final transitional stage between the thermal pulse (TP) AGB and the post-AGB/planetary nebula stage (Skinner et al., 1998; Osterbart et al., 2000). More recently Melnick et al. (2001) reported the detection of warm water vapor in the CSE of IRC +10216, and Decin et al. (2010) have reported the detection of many water lines in the CSE by the Herschel satellite (Pilbratt et al., 2010).

IRC +10216 is clearly a fascinating and well studied object, and it is impossible to fully review the extensive literature on it here. As such we focus mainly on the N band atmospheric window, which is bounded by water vapor at $\lambda \lesssim 8\mu\text{m}$ and CO_2 at $\lambda \gtrsim 14\mu\text{m}$, and on high-spatial resolution imaging and interferometry of the CSE. In carbon stars, the N-band spectrum usually shows the emission feature of SiC.

C.1.2 SiC dust

Around 40 years ago the production of SiC (Friedemann, 1969; Gilman, 1969) and the presence of its emission feature near $11\mu\text{m}$ (Gilra, 1971) were predicted. This feature was then discovered in the N band spectra of carbon stars (Hackwell, 1972), including in IRC +10216 by Treffers and Cohen (1974, hereafter TC74). In this paper we present observational evidence that either the spectroscopic SiC feature in IRC +10216, or the underlying continuum, has undergone a significant change in the last 15 to 20 years, so we will briefly discuss some of the previous work attempting to connect the properties of this feature to the evolutionary state of the underlying AGB stars. The Infrared Astronomical Satellite Low-Resolution Spectrometer (IRAS/LRS) provided a wealth of data in the mid-IR spectral region, including a catalog of AGB star spectra. These data have been used extensively to study the $\sim 11\mu\text{m}$ SiC feature¹ of carbon stars, generally finding a positive

¹We adopt the nomenclature “ $\sim 11\mu\text{m}$ ” of Speck, Thompson, and Hofmeister (2005) to indicate the varied peak wavelengths of this feature.

correlation between dust continuum temperature and the strength of the emission peak (Baron et al., 1987; Chan and Kwok, 1990; Sloan et al., 1998).

A common feature of these efforts has been an attempt to relate the SiC feature and other characteristics of the mid-IR dust spectra to the long term evolution of the host AGB star. Thompson et al. (2006) provide a useful review of this work, and use the more recent Infrared Space Observatory (ISO, Kessler et al., 1996) Short Wavelength Spectrometer (SWS, ISO Handbook, Vol V) data set to further investigate correlations between the SiC peak strength, peak wavelength, and dust continuum temperature. They ultimately conclude that there are no useful correlations, and blame poor continuum fitting for the previous results.

C.1.3 IRC +10216 in the spatial domain

IRC +10216 has also provided many fascinating results in the spatial domain. Given its extreme brightness in the mid-IR, it was an early target for interferometry, and has more recently been subject to intense study in the near-IR. At wider spatial scales, visible wavelength imaging has shown an extended dusty envelope composed of multiple shells. We will now briefly review some of these results, with particular interest in their implications for the process and variability of mass loss from IRC +10216.

Deep optical observations have shown that IRC +10216 is surrounded by multiple dusty shells, which can be seen scattering ambient galactic light out to separations of $\sim 200''$. Maunon and Huggins (1999) analyzed these shells in B and V band images from the Canada France Hawaii Telescope (CFHT) on Mauna Kea and concluded that some process modulates the mass loss on a timescale of 200-800 years, and later found evidence for timescales as short as 40 years (Maunon and Huggins, 2000) using Hubble Space Telescope (HST) imaging. Leão et al. (2006) used deep Very Large Telescope (VLT) V band images to show that the shells can be resolved into even smaller structures. These shells appear to be only approximately spherical and are azimuthally incomplete, indicating that the mass loss is not isotropic. HST imaging of the inner $\sim 10''$ shows a nearly bipolar structure, reminiscent of the typical but

poorly understood structure of planetary nebulae (Skinner et al., 1998).

High spatial resolution observations in the near-IR have produced a fascinating picture of the inner portions of the dusty envelope around IRC +10216 (which are invisible in the optical). Using speckle-masking interferometry in the K' band Weigelt et al. (1998) found the inner $1/2''$ to be composed of at least 5 distinct clumps, indicating an inhomogeneous recent mass loss history. Haniff and Buscher (1998) then presented diffraction limited imaging data which showed that between 1989 and 1997 these clumps had undergone significant evolution, exhibiting relative motion and some either appearing or becoming brighter. Tuthill et al. (2000) showed significant relative motion of various components of the dust, with possible acceleration, based on 7 epochs of sparse aperture mask interferometric imaging in K band using the Keck I telescope². Interestingly these authors found no evidence for new dust production during these observations.

Relative motion within the inner regions of the dust shell were also found by Osterbart et al. (2000), who argued that this evolution was not related to the ~ 2 year luminosity cycle in any simple way. In a related effort, extensive radiative transfer modeling was conducted by Men'shchikov et al. (2001) taking into account much of the archival multi-wavelength data set (including spectral and spatial information). Men'shchikov et al. (2002) used their model to explain the time evolution reported by Osterbart et al. (2000). A key conclusion from this study is that since its discovery IRC +10216 has been undergoing an intense period of mass loss, probably starting ~ 50 years earlier. They also concluded that the mass loss rate had recently increased.

IRC +10216 has been repeatedly studied by interferometers in the mid-IR. McCarthy et al. (1980, hereafter MHL80) measured visibilities at 2.2, 3.5, 5.0, 8.4, 10.2, 11.1, 12.5 and $\sim 20\mu\text{m}$, at several epochs and position angles (PAs). They found evidence for asymmetry at the short wavelengths, indicating an elongation along PA $\sim 25^\circ$, which matched the early optical images and has been confirmed repeatedly by later observations (Skinner et al., 1998; Leão et al., 2006). No evidence of this

²See the movie: <http://www.physics.usyd.edu.au/~gekko/irc10216.html>

elongation was found at $11.1\mu\text{m}$ however. It was also noted that apparent size, but not morphology, changes with photometric phase.

This object has also been observed at $\sim 11\mu\text{m}$ using the UC Berkeley Infrared Spatial Interferometer (ISI). Danchi et al. (1990) generally confirmed the large change in visibilities with photometric phase found by MHL80, and argued that dust was being formed much closer to the star than previous studies had found. Using data from the ISI taken ~ 10 years later, Monnier et al. (2000) found that the inner radius of the dust had moved away from the star. This result was based on model fits to the visibilities, and led them to conclude that no new dust was being formed for most of the 1990's. This appears to contradict the radiative transfer based mass loss predictions of Men'shchikov et al. (2002). Most recently the ISI detected some asymmetry at $11.15\mu\text{m}$ using baselines of up to 12m (Chandler et al., 2007).

C.1.4 New results from the MMT

Here we present new spatially resolved mid-IR photometry and spectroscopy of IRC +10216 with high resolution, AO corrected, spatial information, obtained at the MMT on Mt. Hopkins, AZ, in 2009 and 2010. We first describe our observations and data reduction, paying particular attention to the correction needed when observing an extended object with a spectroscopic slit and a diffraction limited beam. We also review nearly four decades of measurements of the spectrum of IRC +10216. We then discuss our new results in context with the previous work on IRC +10216.

C.2 Observations and data reduction

We observed IRC +10216 at two epochs separated by approximately 1 year, using the 4th generation Mid-Infrared Array Camera (MIRAC4), fed by the MMT Adaptive Optics (MMTAO) thermally efficient adaptive secondary mirror (Wildi et al., 2003). Infrared light first passes through the Bracewell Infrared Nulling Cryostat (BLINC,

Hinz et al. (2000))³, with visible light being reflected to the visible wavelength wavefront sensor of the AO system. This system routinely achieves $\sim 98\%$ Strehl ratios at $10\mu\text{m}$ (Close et al., 2003), and can super-resolve structure smaller than its diffraction limit (Biller et al., 2006; Skemer et al., 2008). In addition to imaging, MIRAC4 has a grism spectroscopy mode described in Skemer et al. (2009).

C.2.1 2009 bandpass photometry

We observed IRC +10216 on 13 Jan 2009 UT with the imaging mode of BLINC/MIRAC4 using its fine plate scale (0.055 arcsec/pixel). Conditions were photometric, with excellent seeing, estimated to be better than FWHM=0.5" at V from the AO acquisition camera. To avoid saturation from the extremely bright source ($\sim 40,000\text{Jy}$) we read out MIRAC4's array with a 0.008s frame time. IRC +10216 is optically faint (R mag > 15), and we were unable to close the MMTAO loop. As a result these observations were taken with the adaptive secondary in its static position, which uses a pre-determined set of actuator commands to hold the mirror shape. We took data in the typical fashion for MIRAC4: chopping using the BLINC internal chopper and telescope nods in the perpendicular direction. In the case of IRC +10216 we set the nod amplitude to be large enough that only one pair of chops was on the detector since the object was expected to be significantly extended. Observations of the standard star μ UMa were taken immediately after IRC +10216 in identical fashion, but for μ UMa the nod amplitude was set so that all four positions were on the detector to increase observing efficiency. Table C.1 lists the filters and airmasses for these observations.

The data were reduced by first applying a custom artifact removal script developed for the MIRAC4 detector (Skemer et al., 2008), which also performs the background subtraction of the chop-nod sets. Each frame was then inspected to look for bad chops (caused by the chopper sticking) and excessive pattern noise from the detector. Frames with these problems were discarded. Photometry was

³Further information on MIRAC4 and BLINC can be found at <http://zero.as.arizona.edu/miracblinc>

conducted on the individual images, rather than registering and combining, to allow an empirical estimation of the uncertainties from the artifact reduction and background subtraction processes. We used the DAOPHOT package in IRAF, and selected the best photometric aperture for the standard and object based on the mean “curve of growth” for counts vs. aperture radius. Since IRC +10216 is extended, the aperture where the curve flattened was always much wider than for the PSF.

Following Skemer et al. (2010) we applied a telluric correction to the photometry using transmission curves provided by Gemini Observatory⁴ calculated with the ATRAN model atmosphere code (Lord, 1992), and an estimate of 3mm precipitable water vapor (PWV) and the airmass of the observations. The PWV assumption is supported by contemporaneous PWV measurements taken on Kitt Peak (74 km west-northwest), and as noted in Skemer et al. (2010) the correction at these wavelengths is generally insensitive to PWV. In the $9.79\mu\text{m}$ filter the correction was +2.5% (due to telluric ozone), and in all others it was $< 1\%$. Finally, we normalized the photometry by the Cohen et al. (1996) flux for μ UMa. The results are presented in Table C.2.

The uncertainty in our photometry was calculated in similar fashion to that used for grism spectroscopy in Skemer et al. (2010) and below (Section C.2.2), with the exception that for bandpass photometry we did not assume a correlated global uncertainty. As discussed above we performed photometry on individual frames, which provides an empirical measurement uncertainty for both the object and standard. Thus our measurement uncertainty includes the random effects of detector artifacts and our removal procedure. We then add in quadrature the mean uncertainty in the μ UMa standard flux across the filter bandpass from Cohen et al. (1996), and use the telluric calibration uncertainties from Skemer et al. (2010), which were measured the following night. The values used and the final total 1σ uncertainty are included in Table C.2.

When compared to our (normalized) grism spectrum from a year later, the pho-

⁴<http://www.gemini.edu/sciops/telescopes-and-sites/observing-condition-constraints/>

tometry from $9.8\mu\text{m}$ to $12.5\mu\text{m}$ matches very well. In the $8.7\mu\text{m}$ filter, however, there is a $\sim 30\%$ discrepancy between the bandpass photometry and the grism data, as well as with archival data. We are suspicious of this data point since it represents a high counts regime of the detector not well understood, but we do not yet have any specific reason to discard it. We discuss this further in Appendix A.

C.2.2 2010 grism spectroscopy

We observed IRC +10216 nearly one year later on 1 Jan 2010, UT in the grism spectroscopy mode of MIRAC4, using a $1''$ slit. With this configuration MIRAC4 has a spectral resolution of $R \sim 125$ and a spatial resolution of $\lambda/D \sim 0.32''$ at $10\mu\text{m}$. The detector wavelength scale was calibrated at the telescope using a well characterized polystyrene sample and fitting a quadratic function to the measured centroids of features in the spectrum. The coarse platescale used for grism work was measured using the binary α Gem on 2 Jan 2010, UT and elements from the USNO Sixth Catalog of Orbits of Visual Binary Stars⁵ (Hartkopf et al., 2001). We found a value of $0.107''/\text{pixel}$.

Conditions on 1 Jan 2010 UT were photometric. Through a combination of excellent seeing conditions and IRC +10216 being near its brightness maximum, we were able to lock the MMTAO system on IRC +10216 with a loop speed of 25Hz. We set the frame time to 0.008s, and to ensure that we could take advantage of the diffraction limited information being delivered by the MMTAO system we read out each 0.008s frame. Due to its extreme brightness at $10\mu\text{m}$ only a few of these short frames were needed to provide sufficient S/N, and this data taking mode allows us to reject frames with bad slit alignment due to residual tip/tilt errors and frames with excessive artifacts.

Observations of the standard μ UMa were challenging for nearly opposite reasons. Ordinarily one tries to operate the AO system with identical parameters between PSF and science object, but in the optical μ UMa saturated the wavefront sensor (WFS) at speeds slower than 100Hz. At $10\mu\text{m}$ μ UMa is a factor of ~ 500 fainter

⁵<http://ad.usno.navy.mil/wds/orb6.html>

(even though it is one of the brightest $10\mu\text{m}$ standards), so longer integrations are required to efficiently build S/N. Table C.1 lists the details of these observations.

Determining Strehl ratio for these observations is problematic. We did not take data in the imaging mode (other than for slit alignment) because of the limited time (~ 1 hr) that conditions were good enough to lock the MMTAO system on this faint star. Without two dimensional imaging data it is difficult to directly measure Strehl ratio from our PSF observations. In addition, since we necessarily operated the AO system with different parameters, any such measurement would not apply to IRC +10216. At $10\mu\text{m}$ the dominant wavefront error term will be from loop delay (servo error), even on an optically faint target such as IRC +10216. Based on the very high Strehl ratios routinely achieved by MMTAO and MIRAC4 (98%), the WFS integration times (40ms), and our use of short exposures, we estimate the Strehl ratio of our IRC +10216 observations to be $\sim 80\%$.

The Moon moved closer to IRC +10216 on the following night (2 Jan) and seeing was somewhat worse, so we were unable to lock MMTAO on IRC +10216. Though we took seeing limited data, we find that the good spatial information provided by AO is necessary to adequately correct for differential slit loss between the point source standard and a resolved IRC +10216. We did, however, take AO-on spectra of the standards μ UMa and β Gem, which we use to calibrate our slit loss correction procedure. Details of these observations are also included in Table C.1.

Reduction of grism data is similar to the imaging procedure described above, except that our data were taken with nods only, as the chopper is unnecessary for very bright sources and can sometimes cause slit-misalignment. We used the same artifact removal script, and the images at each position angle were registered and median combined. Our fully reduced images of IRC+10216 and μ UMa are presented in Figure C.1.

Our first step in analyzing the data was to fit the spatial profiles of the PSF and IRC +10216 at each detector row. We found that a Lorentzian is a good fit for IRC +10216 out to wide separations from the peak, generally achieving $\chi^2_\nu < 2$ across the entire wavelength range. As expected a Gaussian was good for the core of the

Airy pattern of our PSF standard. We used profile plots (an example of which is shown in Figure C.2) to assess the quality of this analysis. The chosen functions describe the core of the objects well, and we find that the Lorentzian full-width at half-maximum (FWHM) is a meaningful proxy for the size of IRC +10216 relative to the PSF. Based on this conclusion we show FWHM vs. wavelength in Figure C.3, where we see that the PSF was essentially diffraction limited. The comparison is not perfect due to the difference in AO system parameters between the two objects, but it is clear that IRC +10216 is extended. It is also apparent that the dependence of size on wavelength is much more complicated than mere λ/D scaling due to diffraction.

Regarding data reduction, an important conclusion to draw from Figure C.3 is that one cannot simply divide by a point source standard to calibrate this extended object when using a slit, as a different amount of light is lost due to the slit, and this effect depends on wavelength in a non-analytic way. To quantify the effect of the slit we constructed a surface of revolution for IRC +10216 at each wavelength using the 1-D spatial profile. We then calculated the fraction of flux enclosed by the 1" slit, taking into account the width of the aperture used to extract flux at each wavelength. The results of these calculations are shown in Figure C.4, for each position angle.

For the PSF, in addition to the photometric standard taken on 1 Jan, we used the AO-on standard observations from the following night in order to improve S/N. Each standard was analyzed independently, then we took the median of the results at each detector row (i.e. wavelength). We compare the outcome of this procedure to that expected based on the theoretical Airy pattern for the MMT, which we processed in similar fashion, in Figure C.4.

To correct for the differential slit-loss, we calculate the slit-loss correction factor (SLCF) as the ratio of the enclosed flux of the PSF to that of IRC +10216. We use the theoretically calculated curve for the actual aperture due to the relatively noisier empirical results for the PSF. We use the average of the IRC +10216 results to suppress noise, ignoring the small possible source asymmetry highlighted by our

FWHM curves since it will cause only a small difference in the results. The resultant SLCF curve is shown in Figure C.4.

We performed aperture photometry on IRC +10216 and the standard. As with the bandpass photometry we corrected for airmass and used the Cohen et al. (1996) standard spectrum of μ UMa to calibrate the results. In the PA=107.0 spectrum we found a $-1.7\%/ \mu\text{m}$ slope compared to the other three, likely due to a slight offset in the slit or possibly a period of worse AO correction. This slope was removed. We show the raw spectrum at each position angle prior to applying the SLCF in Figure C.5. Most of the noise in the spectrum appears to be correlated noise, i.e. it is identical in all four position angles, implying it comes from the standard. A simple Poisson noise calculation indicates that we should have achieved slightly greater S/N in the μ UMa observation. The higher noise in the standard is most likely due to its lower flux relative to the background and the interaction of this with the MIRAC4 detector artifacts.

The SLCF was applied to each PA, and then we re-binned by 7 pixels using the median as in Skemer et al. (2010). Though this sacrifices some spectral resolution, it has the benefit of increasing the S/N in each bin and allowing a robust empirical estimate of the uncertainties. We used the same prescription for calculating local measurement error as Skemer et al. (2010), estimating the Gaussian 1σ error from the 2nd and 6th ordered values in each bin. The global bias (correlated error) reported by Cohen et al. (1996) for μ UMa is negligible, so we add their total uncertainty in quadrature to the measurement error to calculate the total local uncertainty.

Global telluric error was estimated from the four spectra, which were taken at different airmasses. We find results similar to Skemer et al. (2010): 2.7% outside the ozone feature and 10% inside. We do not include a separate local telluric error as this will be included in our measurement error. Finally we adopt a 5% global systematic error term from the SLCF procedure which is based on the scatter in the IRC +10216 enclosed fraction results.

The fully calibrated and slit-loss corrected results are listed in Table C.3 and

shown in Figure C.5, along with the local and total uncertainties. The effect of the SLCF can be seen, in addition to the overall increase in flux the SLCF has effects on the shape of the spectrum compared to the uncorrected curves. It highlights an apparent “bump” at $\sim 9\mu\text{m}$. The SLCF also reveals a steeper negative slope longer than $\sim 11\mu\text{m}$. This slope matches the bandpass photometry from 2009 very well, giving us confidence in our slit loss correction procedure. We discuss these features in more detail in section C.4.

C.3 Archival data

Since its discovery, IRC +10216 has been observed many times, at nearly every wavelength available to astronomers. In this paper we concentrate mainly on the region of N-band accessible through the Earth’s atmosphere. Some of the earliest observations at these wavelengths were of IRC +10216, and it has been observed regularly over the last four decades, though, ironically, increases in telescope size and detector sensitivity may be curtailing this somewhat due to the dynamic range required to avoid saturation. It would be impossible to account for all of the work done on this object. Here we use a sample of N-band spectra, and several datasets of bandpass photometry.

We present these data in the context of the light curve parameters of Le Bertre (1992), where the period $P = 649$ days, and phase $\phi = 0$ at JD 2447483 (where ϕ varies from 0 to 1). Based on the spread reported in the various filters used, and other determinations (e.g. 638 days in Dyck et al. (1991)), the period is uncertain by ~ 10 days. This should be kept in mind when comparing widely separated measurements, i.e. nearly 21 cycles have occurred between the TC74 data and our 2010 measurement so the relative phase between them could be off by 30% or more. This is less of a concern for more closely spaced data and we are not attempting a light-curve analysis here, rather we claim that the ~ 2 year Mira variability isn’t the source of the changes we discuss. For our observations the star’s luminosity was at $\phi = 0.27$ on 13 Jan., 2009 and $\phi = 0.89$ on 1 Jan., 2010, assuming Le Bertre

(1992)’s parameters.

C.3.1 Introduction to the spectral datasets

Here we collect a sample of N-band spectra, choosing some of the earliest measurements, two space-based observations (IRAS/LRS and ISO/SWS), and a set of observations taken on the same instrument (CGS3 at UKIRT) repeatedly over a short period of time. We briefly describe these datasets here and any processing we did. It is worth noting that none of these observations are affected by slit loss, as they either used no slit or had large apertures.

The spectrum of Treffers and Cohen

The spectrum of TC74 was taken with a scanning Michelson interferometer (i.e. a Fourier transform spectrometer (FTS)) on a 2.2m telescope on Mauna Kea on 15 and 16 Feb., 1973 ($\phi = 0.13$). They used the Moon as a telluric standard, and reported in arbitrary flux per unit wavenumber (F_ν) with a resolution of 2cm^{-1} ($\sim 0.02\mu\text{m}$). The gap in the spectral fragments was in the original data, and though not commented on by TC74 is almost certainly due to the telluric ozone feature. Below $\sim 8\mu\text{m}$ the spectrum appears to be unreliable due to telluric water vapor.

The Spectrophotometry of Merrill and Stein

Merrill and Stein (1976, hereafter MS76) used a circular variable filter (CVF) photometer to observe IRC +10216 from Mt. Lemmon, AZ. A date is not given for this observation, but from the publication date and other dates given in the paper we can infer that IRC +10216 was observed no later than 1975. Based on information provided by the anonymous referee we believe the most likely date for this observation was early 1973. We extracted the data from their Figure 2, and converted from λF_λ units to arbitrary F_ν .

The IRAS LRS Spectrum

The IRAS LRS spectrum for IRC +10216 was extracted from a database maintained by Kevin Volk⁶. We spliced the blue and red fragments together and applied a correction for the spectral shape of the IRAS standard α Tau using the procedure of Cohen et al. (1992). It is not possible to assign a single epoch to IRAS observations, so we adopt the range 1 Feb. to 1 Nov. 1983 ($\phi = 0.74 - 0.16$).

The ISO/SWS Spectrum

We retrieved the reduced ISO/SWS observation of IRC +10216 from the ISO archive⁷, taken on 2 June 1996 ($\phi = 0.24$). The ISO data presented in this paper are from the Highly Processed Data Product (HPDP) set called ‘High resolution processed and defringed SWS01s’, available for public use in the ISO Data Archive (Frieswijk et al., 2007). The data we are interested in span detector bands 2C (7.0-12.0 μ m) and 3A (12.0-16.5 μ m) (ISO Handbook, Vol V). Though the pipeline attempts some defringing in band 3, we applied a 0.01 μ m binning (averaging) to the data to reduce fringing, which is especially prominent in band 2C. This spectrum has been published previously by Cernicharo et al. (1999) who used an earlier reduction pipeline and did not discuss the 8 to 13 μ m region.

Spectra from UKIRT

Monnier, Geballe, and Danchi (1998) (hereafter MGD98) obtained spectra of IRC +10216 at 4 epochs from 1994 to 1996 as part of a survey of variability in late type stars, using the Cooled Grating Spectrometer 3 (CGS3) at UKIRT on Mauna Kea. Several stars were used as photometric standards at each epoch. We considered these measurements separately, and also averaged the 4 spectra for comparison with our data, first normalizing each to 10.55 μ m and using 0.1 μ m bins to compensate for the slight changes in wavelength scale between observations. The luminosity phases

⁶http://www.iras.ucalgary.ca/~volk/getlrs_plot.html

⁷<http://iso.esac.esa.int/>

of these spectra were $\phi = 0.38, 0.47, 0.96$, and 0.27 . MGD98 noted some small fluctuations in the spectral slope (across the $10\mu\text{m}$ window) with phase and though they speculated that the changes were due to a rapidly changing dust condensation zone, they could not rule out poor calibrations as the cause. We note that their data from 22 June 1996 matches the 2 June 1996 ISO/SWS data fairly well.

C.3.2 Bandpass photometry archives

Bandpass photometry can provide a useful check on spectra, which can be plagued by such things as uncertain slopes or slit effects. For these purposes we require photometry taken in several filters at the same epoch (to within a few days) so that any apparent changes with wavelength are not caused by the variation in overall brightness. There are several datasets in the literature which contain measurements of IRC +10216 across the $10\mu\text{m}$ window. We use these primarily to confirm the normalized shapes of the spectra discussed above, as the large error bars and uncertainties in normalization make epoch to epoch comparisons of the photometry difficult. In all cases, we used the Vega spectrum of Cohen et al. (1995) to reduce magnitudes reported in the literature. We describe our normalization method in detail below.

Strecker and Ney

Strecker and Ney (1974, hereafter SN74) observed IRC +10216 at 5 epochs in 1973 at 8.6 , 10.7 , and $12.2\mu\text{m}$ (as well as other points outside the N band) from the O'Brien observatory in Minnesota, USA. Their measurements in Jan, Mar and Apr 1973 provide a nearly contemporaneous check on the spectral shape found by TC74, and provide a useful comparison to the MS76 spectrum. Estimated errors were reported as $\pm 20\%$, which make individual points nearly useless for comparing to spectra. To overcome this we average the three points from early 1973, and the two points from late 1973, after applying the normalization procedure described below. We could average all 5, however this method allows for the possibility of short term

(< 9 months) variability in the spectral shape of IRC +10216.

McCarthy, Howell, and Low

MHL80 reported measurements of IRC +10216's brightness at many epochs in the late 1970s, taken on Kitt Peak, Arizona, using 4, 2.3, 1.5, and 1 m telescopes. These data were taken in support of their interferometric size measurements. At only 2 of these epochs (17 Dec., 1977 and 18 Nov., 1978) were measurements made at enough points across the 10 μm window to be useful for shape comparisons with our spectra. We use the MHL80 photometry, with estimated errors of $\pm 10\%$, for comparison with the TC74 spectrum.

The Photometry of Le Bertre

Le Bertre (1997, hereafter LB97) obtained bandpass photometry using the European Southern Observatory 1m telescope at La Silla Observatory, Chile, spanning 1985-1988. These are the same data used in part in the production of our adopted light curve parameters in Le Bertre (1992). LB97 used filters with central wavelength of 8.38, 9.69, 10.36, and 12.89 μm and reported errors of 10%, 10%, 10%, and 15% respectively. These data are used here to compare to the IRAS/LRS spectrum.

TIRCAM

IRC +10216 was observed in January 1993 by Busso et al. (1996) using TIRCAM, a mid-IR camera equipped with a 10x64 array, on the 1.5m Telescopio Italiano Infrarosso at Gornergrat (TIRGO), Switzerland. The filters used had central wavelengths of 8.8, 9.8, 11.7, and 12.5 μm , with errors of 7%, 7%, 15%, and 15% reported for IRC +10216. We compare the photometry of Busso et al. (1996) to the UKIRT spectra of MGD98 and the ISO/SWS spectrum.

C.3.3 Comparison of Archival Data

Our main reason for including this archival data is comparison with our new results. First, though, we can compare the various measurements to each other. Have the changes our measurements reveal been observed before? Are these changes part of the regular variability of this object? The data sets we have extracted from the literature were taken at various points in IRC +10216's two year brightness variations. To account for this we first normalize the spectra at $\lambda = 10.55\mu\text{m}$, averaging across the MIRAC4 $10.55\mu\text{m}$ filter bandpass. This area appears to have had a very stable spectral slope throughout the nearly forty years of observations we consider here.

Normalizing the photometry is a bit more challenging. Each instrument used had a different photometric system, and authors did not always report results in all filters at each epoch. Since we are most interested in analyzing the shape of the spectrum for $\lambda > 11\mu\text{m}$, we proceed by first fitting a line to the spectra from $8 - 11\mu\text{m}$ after they were normalized to the MIRAC4 $10.55\mu\text{m}$ bandpass. The spectra all appear to be roughly linear and similar in slope across this region, though with noticeable variation at $\lambda < 9\mu\text{m}$. We then normalize the photometry to this line (which has $F\nu(10.55\mu\text{m}) = 1$), using the best-fit normalization factor for each epoch. We also propagate errors from the fitting procedure to the new normalized photometric points.

In Figure C.6 we show the bandpass photometry of SN74, the FTS spectrum of TC74, the CVF spectrophotometry of MS76, and the photometry of MHL80. We also show the spectrum of IRC +10216 in 2010 as measured by MIRAC4 in this work. For $\lambda \lesssim 11\mu\text{m}$ the photometry appears to agree nicely with the spectra, but from $12 - 13\mu\text{m}$ it is noticeably brighter in both SN74 and MHL80 - though consistent with the spectra at the $\sim 2\sigma$ level. This could be explained by a slope offset in the TC74 FTS spectrum, however the MS76 CVF spectrum would not likely have such an artifact. Given the variability of IRC +10216, variations in spectral shape hinted at by this plot might be associated with the 649 day cycle of IRC +10216's

luminosity. We note, however, that the SN74 and the first three TC74 points were taken at nearly identical times, and our best guess at the epoch of MS76 indicates it was taken very close to these data sets as well. This would require rapid short term variability in the spectral shape over time scales much shorter than the 649 day brightness variation.

A comparison of data from the 1980s is provided in Figure C.7. The photometry of LB97 and the IRAS/LRS spectrum agree well within the 1σ error, with the exception of the June 1985 point which appears to have been strongly affected by atmospheric O_3 . As in Figure C.6 we see that the photometry is generally consistent with the 2010 shape (and that measured by MS76) within the 2σ uncertainty.

We continue our decade by decade comparisons with Figure C.8, which shows data from the 1990s. In this case the data agree quite well across the entire $10\mu\text{m}$ window, and we note especially the agreement between the UKIRT and ISO/SWS spectra taken twenty days apart in June 1996. These measurements span three and a half years, and two full 649 day periods, so unlike the previous decades we can say with some confidence that there is no variation in shape, large enough to explain our 2009/2010 results at $\lambda \gtrsim 11\mu\text{m}$, occurring as part of the regular 649 day variability of IRC +10216 during this time period.

C.4 Discussion

C.4.1 Changes in the $10\mu\text{m}$ spectrum of IRC +10216

Figure C.9 shows a comparison of nearly four decades of N-band spectra of IRC +10216 with the data normalized to $F_\nu(10.55\mu\text{m}) = 1$. We also include our 2009 photometry, which matches our 2010 grism spectrum very well. It is clear that a significant change at wavelengths longer than $\sim 11\mu\text{m}$ was recorded in our 2009 and 2010 data when compared to the mid-1990s, and that a similar shape was observed in the early 1970s. The negative slope of the spectrum has become steeper and the continuum is lower at wavelengths redder than $13\mu\text{m}$. The close match between our bandpass photometry and grism spectrum gives confidence that this is not merely

a calibration error, and since they are taken 1 year apart at different luminosity phases the shape appears to be stable at the time of our observations.

In Figure C.10 we plot the flux ratio $F_\nu(12.5\mu\text{m})/F_\nu(10.55\mu\text{m})$ vs. time, where flux at $12.5\mu\text{m}$ was calculated as the mean between 12 and $13\mu\text{m}$ for each of the spectra. This figure illustrates the change in the shape of the spectrum over time. The change from 1996 to 2009 does not occur as part of the 649 day Mira variability, as evidenced by the mid 1990s data. We discuss two possible interpretations of this plot further below.

Monnier et al. (1999) found long term changes in three carbon stars, including IRC +10216 which had the smallest change. They used the MS76, IRAS/LRS, and MGD98 data, and reported a change in spectral slope across the $10\mu\text{m}$ window ($8 - 13\mu\text{m}$) from the early 1970s to 1996. We now see a change in the opposite direction from 1996 to 2009 in IRC +10216's spectrum.

Perhaps the simplest interpretation of these results is that we have recorded an episode of irregular variability (i.e. recurring but not periodic changes) in the spectrum, rather than a trend. Due to the sparse sampling and relatively long time periods between measurements (e.g. the gap from 1978 to 1983, or from 1996 to 2009) we can make no statements about how often this irregular variability occurs. In the case of the 2009-2010 MIRAC4 data this new shape lasts for at least one year, or half the period. Given the clumpy structure of the CSE (Weigelt et al., 1998), the anisotropic nature of the mass loss history (Leão et al., 2006), and the rapid variations seen in the inner regions of the CSE (Haniff and Buscher, 1998; Tuthill et al., 2000), relatively rapid and irregular variability in IRC +10216's spectrum might be expected.

We also consider the possibility that these changes are occurring over a longer term. Though sparsely sampled in time, figure C.10 has the appearance of a smooth change over ~ 40 years. Other circumstantial evidence for the longer term variability is provided in the findings of previous studies we discussed in section C.1.3. Maunron and Huggins (2000) found evidence for ~ 40 year modulation in the expanding dust shells around IRC +10216 from deep V band observations. Men'shchikov et al.

(2002) claimed that the current episode of dust production began roughly in the 1950s, and that the mass loss rate had recently increased. Monnier et al. (2000) claimed that dust production had stopped by the end of the 1990s in contradiction to Men'shchikov et al. (2002). We won't attempt to resolve these contradictory modeling results, but rather take them as evidence that something had changed in the mass loss rate at the end of the 1990s. This idea, coupled with the correspondence between the periods evident in the V band observations, the hypothesized start date of the current mass loss episode, and the timing of the changes evident in the spectrum by 2010, supports the possibility of a longer term change in IRC +10216.

In any case, we can place the shape of IRC +10216's N band spectrum in context with other carbon stars using the "Carbon-Rich Dust Sequence" of Sloan, Little-Marenin, and Price (1998) (hereafter SLMP98). Their system is based on 96 carbon-rich variable stars observed by IRAS/LRS, from which they subtract a 2400K blackbody to remove the stellar continuum. The blackbody is fit to the wavelength range $7.67\text{--}8.05\mu\text{m}$. After subtraction the spectra were normalized, and then were grouped by inspection according to the shape of the $\sim 11\mu\text{m}$ SiC feature and the presence and strength of the $9\mu\text{m}$ feature (which we acknowledge is likely not real). In Figure C.11 we show their sequence, formed from averaging and smoothing each spectrum in the class, as the solid black curves. As we discussed in Section C.1.2, in light of the results of Thompson et al. (2006) we do not treat this sequence as reflecting the evolutionary state of these carbon stars. Nevertheless, we have found the system of SLMP98 to be a useful atlas of the SiC spectra in carbon stars and as an aid to interpreting our results.

In Figure C.11, we show the IRAS/LRS spectrum of IRC +10216 in red, which was classified by SLMP98 as "Red", and cited as the prototype of that class. We also show our 2010 spectrum in blue, continuum subtracted and normalized according to the above prescription. In this framework, it appears that the center of the SiC peak has shifted blue-ward, and the $9\mu\text{m}$ region is enhanced, which appears to be true regardless of whether the minimum at $\sim 9.5\mu\text{m}$ is caused by ozone (it almost

certainly is). In 2010, IRC +10216 is a better match for the “Broad 2” (Br2) class in the SLMP98 system. Even though SLMP98 does not represent an astrophysical sequence for carbon stars, we see that the shape we observed does occur in other carbon stars.

Regardless of whether this change occurs irregularly on timescales of a few years, or represents multi-decade variability, such variability could easily be overlooked in similar Carbon stars. The SLMP98 system is especially useful here, as it provides a catalog of spectra and targets for follow-up observations to test this possibility. We find that the IRAS/LRS spectra of sources RV Cen (Br1) and CR Gem (Br2) are good qualitative matches for our 2010 MIRAC4 spectrum. We also checked the SiC feature of LP And, the one other source in the Red classification that was also observed by ISO/SWS, and found that it does not appear to have changed between the two observations. All of the “Red” and “Broad” sources deserve future observations in this wavelength range to check for any irregular variability or long-term changes.

In addition to follow up observations of IRC +10216 and similar carbon stars, fully understanding the changes reported here will require detailed modeling of the dusty CSE of IRC +10216. Given the wealth of data on this object and its complex asymmetric structure, such modeling is beyond the scope of the current paper (see Men’shchikov et al. (2001) for an example of such a comprehensive effort).

C.4.2 The spatial signature of SiC emission

In Figure C.12 we place our 2010 spectrum and corresponding FWHM measurement on a common wavelength axis. Here we have “deconvolved” IRC +10216 by subtracting the FWHM of the PSF in quadrature, in order to estimate its intrinsic size. IRC +10216 exhibits an apparent increase in size, quite separate from the effect of diffraction alone, in the wavelength range of the SiC emission feature. Fully understanding the spatial signature of the spectral emission feature will require additional observations at different luminosity phases, including 2D imaging, as well as detailed modeling.

A likely interpretation is that we have observed the effect of radiative transfer through the CSE. SiC has higher opacity at the wavelengths of the feature, hence we observe photons produced further from the star at those wavelengths, which causes an apparent increase in size in the feature's part of the spectrum. The change in apparent size can be thought of as mapping the optical depth of the CSE as a function of wavelength.

We can use these results to establish the plausibility of the changes in the spectrum being caused by the outflow of material from the star. If we treat the estimated intrinsic size at $13\mu\text{m}$, $\sim 70\text{AU}$, as the diameter of the CSE, and take the time for the changes in the CSE to occur as 12.5 years (1996.5 to 2009), we find an estimated outflow velocity of 13 km sec^{-1} . Estimates in the literature for the expansion velocity of IRC +10216's CSE range from 12 to 17 km sec^{-1} (Men'shchikov et al., 2001, and references therein). Using their model, Men'shchikov et al. (2001) calculated the deprojected radial velocities of the clumps observed by Osterbart et al. (2000) as $\sim 15\text{ km sec}^{-1}$. Hence, we find that the changes in the spectrum from 1996 to 2009 can plausibly be explained by evolution of the CSE, and the resultant estimated outflow velocity is in good agreement with previous estimates.

A comment on the possible asymmetry evident in Figure C.3 should be made. The FWHM is smaller in the East-West direction than in the North-South direction, roughly indicating an elongation towards the N-NE, exactly as expected from imaging studies at shorter wavelengths. We are cautious with this result, however. An important consequence of the profiles of IRC +10216 following a Lorentzian is that small offsets in the slit will cause changes in the apparent shape of the object. Whereas when a 2D Gaussian is sliced somewhere off the peak the same parameters (i.e. FWHM) describe the resulting curve, the same is not true for a Lorentzian. Since we were not able to repeat the observations at each position angle due to time limitations, we have no way to estimate the uncertainties in the individual FWHM curves at different position angles due to slit alignment.

C.5 Conclusion

We have presented new photometric and spectroscopic measurements of the well-studied carbon star IRC +10216 in the N band atmospheric window. When compared to nearly 4 decades of prior observation we find that a significant change (decrease in brightness) appears to have occurred in the $11 - 13.5\mu\text{m}$ region of the spectrum, which includes the SiC emission feature, between 1996 and 2009. Measurements taken in early 1970s appear to match the 2009/2010 shape, but data from the 1980s and 1990s does not. We discussed two possible explanations for these changes. We may have observed an episode of irregular variability distinct from IRC +10216's regular ~ 2 year Mira variability. We also consider it a possibility that we have observed a long term change occurring over several decades.

Critical to our reduction of the grism spectrum was the stable, high Strehl, diffraction limited information provided by the MMTAO system, which was needed to correct for the differential slit loss between the extended source IRC +10216 and the point source standard. This spatial information, which allows us to analyze size vs. wavelength, shows that the SiC emission feature has a clear spatial signature in the dust surrounding IRC +10216. The CSE exhibits an increase in apparent size of $\sim 30\%$ between 10.2 and $11.6\mu\text{m}$ compared to the continuum on either side of the SiC feature. This is likely tracing the higher optical depth due to SiC in the ~ 70 AU CSE. We used this estimate of the object's intrinsic size to establish that the observed spectrum change over 12.5 years can plausibly be associated with the evolution of the dusty CSE of IRC +10216 given $12\text{-}17\text{ km s}^{-1}$ outflow velocities.

We thank the anonymous referee for a thorough review and for providing insight into 1970s spectrophotometry, which resulted in a much improved manuscript. We thank John Monnier for providing data from UKIRT in tabular form. We thank John Bieging for his comments on a draft of this manuscript. JRM is grateful for the generous support of the Phoenix ARCS Foundation. AJS acknowledges the generous support of the NASA GSRP program. This work and LMC were supported by the NASA Origins program, and the NSF AAG and TSIP programs.

C.6 Appendix: The possibly erroneous $8.7\mu\text{m}$ photometry from 2009

Here we further discuss the MIRAC4 2009 $8.7\mu\text{m}$ photometry data point, which appears significantly over-luminous in Figure C.9. We see no evidence that a change in weather or seeing affected the $8.7\mu\text{m}$ PSF measurement without affecting the others, and there has so far been no evidence that this filter has a leak during other BLINC/MIRAC4 observations. Nevertheless, since this data point was not confirmed by our follow-up grism data and is in a high-flux regime of the detector not well understood, nor tested by any of our other data, we remain suspicious of the $8.7\mu\text{m}$ photometry.

The $8.7\mu\text{m}$ filter had $\sim 85\%$ higher peak counts than the next brightest $10.55\mu\text{m}$ filter (both are above the background), due in part to its width ($\sim 40\%$ wider than $10.55\mu\text{m}$), as well as differences in detector quantum efficiency. This led us to suspect that the most likely culprit for the discrepancy would be non-linearity of the MIRAC4 detector, which does exhibit an increase in slope at higher fluxes. As part of the normal preparation for observing a linearity measurement was performed in a laboratory at Steward Observatory one week prior to these observations. Unfortunately the bias level appears to have changed in the intervening period, which prevents us from directly applying the curve to our data. Ordinarily this is of no consequence when using chop and nod background subtraction, so it was not noticed until long after the observations were complete. We can still perform a worst case analysis though, and decide what effect, if any, non-linearity has on the $8.7\mu\text{m}$ measurement.

We start by assuming that the peak counts value in the $10.55\mu\text{m}$ image is the last linear value and that all pixels in the $8.7\mu\text{m}$ data above this value have a different slope. With this definition the fraction of flux in non-linear pixels in the $8.7\mu\text{m}$ image is $F_{NL} = 32\%$. We can then estimate the change in slope ΔL required to produce the change in total flux: $\Delta L = \frac{\Delta F}{F_{NL}}$. Since we are trying to explain a discrepancy of $\Delta F = 30\%$ we need a slope change of nearly 100% . Figure C.13 shows the laboratory linearity measurement, along with a fit to the lower portion

of the curve. The data have been bias subtracted using the fit. We also show a line with 12% higher slope, which represents the worst case prior to saturation. The arrows on the plot indicate the peak counts per read in the 8.7 filter and in the 10.55 filter, where the 8.7 point is from raw counts prior to background subtraction (and so includes the unknown bias level) and the 10.55 point is background subtracted. Figure C.13 demonstrates that even in the worst case scenario where every non-linear pixel has a 12% higher slope, non-linearity can explain at most $\sim 4\%$ of the excess flux in the $8.7\mu\text{m}$ filter. We see that the non-linearity in fact likely causes a less than 1% error. Given this result, we have no reason to reject the $8.7\mu\text{m}$ data point out of hand due to non-linearity.

Finally, we note that the archival photometry presented in Section C.3 has several examples of apparent excesses at $\sim 8.7\mu\text{m}$, but such a feature never appears in the spectra. This points to unquantified systematics in broadband photometry in this region, which is bounded closely by variable water vapor and ozone. We have tried applying a correction for differences in spectral shape between object and standard, and assuming large changes in PWV between object and standard using the ATRAN model, and so far have not found an explanation for these excesses.

At this point in time, the evidence is inconclusive and we remain suspicious of the 2009 MIRAC4 $8.7\mu\text{m}$ filter photometry. The change from 2009 to 2010 would require a significant decoupling at this wavelength from the rest of the spectrum with regards to the regular 649 day variability. Though we have ruled out non-linearity as a cause, the per-pixel flux achieved is the highest ever observed with MIRAC4 and we cannot yet rule out changes in the read-out artifacts (e.g. cross-talk) at higher flux. Further observations, with both bandpass photometry and grism data taken at the same epoch, are required to fully understand our 2009 data point.

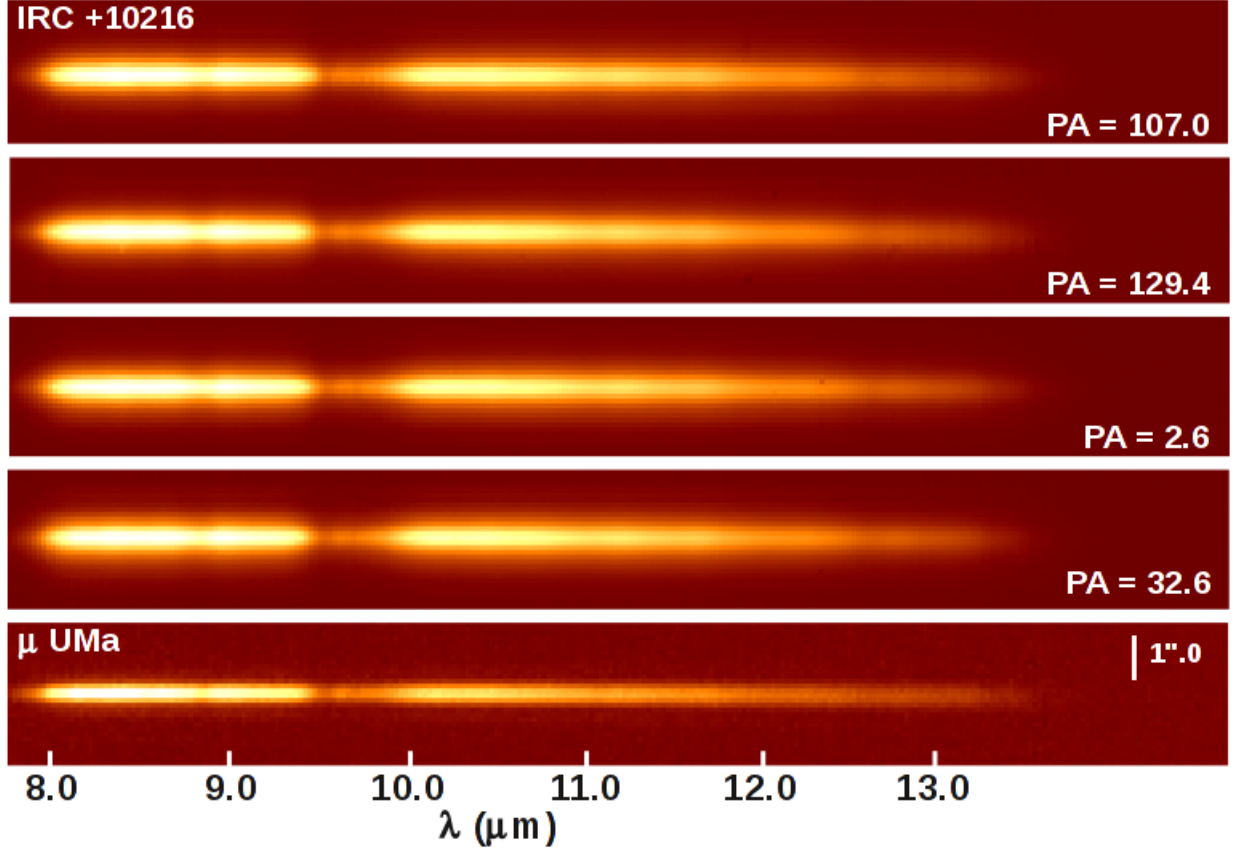


Figure C.1 MIRAC4 grism observations of IRC +10216 and μ UMa with the MM-TAO loop closed. Here we present $\sim 0.3''$ diffraction limited spatial information in the vertical direction and $R \sim 125$ spectral information in the horizontal. See Table C.1 and the text for the details of the observations, especially AO system parameters which were necessarily different due to the relative optical brightness of the two sources. Compared to the PSF, IRC +10216 is clearly resolved. See Figures C.2 and C.3 for the results of extracting profiles in the spatial direction and Figure C.5 for the fully reduced spectrum of IRC +10216. Note the impact of telluric ozone absorption between 9 and $10\mu\text{m}$, and the decreasing sensitivity starting at $13\mu\text{m}$.

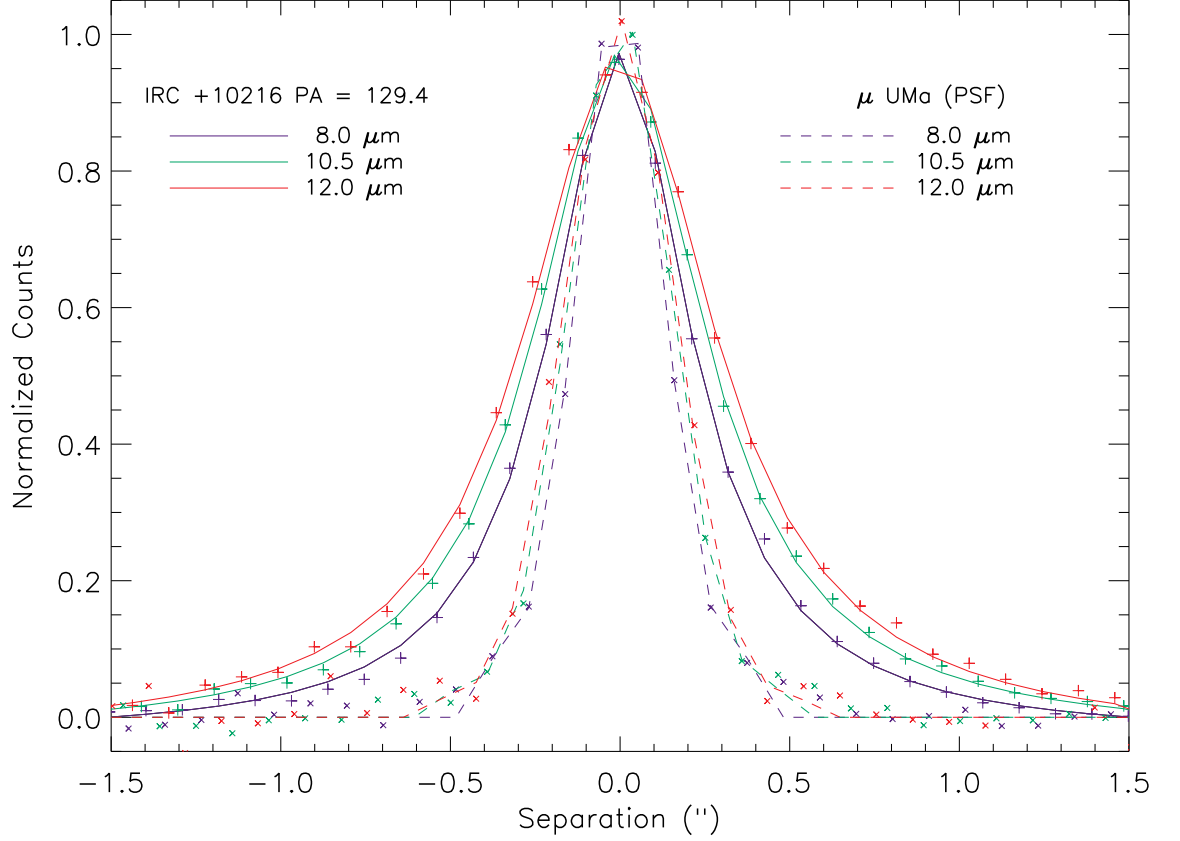


Figure C.2 Normalized spatial profiles of IRC +10216 and the PSF standard μ UMa, at three discrete wavelengths (i.e. single detector rows) for a single position angle (129.4). The data are denoted by x's for μ UMa and +'s for IRC +10216. The PSF core is well fit by a Gaussian (dashed lines), as expected for a well corrected Airy disk (we don't fit past the first airy minimum, which can be seen along with the first airy ring at $\sim 0.5''$). IRC +10216 is well described by a Lorentzian profile (solid lines), though there are apparent correlated discrepancies at wider separations. This result, and similar results for the other position angles, gives us confidence that IRC +10216 is resolved and the FWHM determined by fitting a Lorentzian gives a meaningful proxy for object size vs. wavelength.

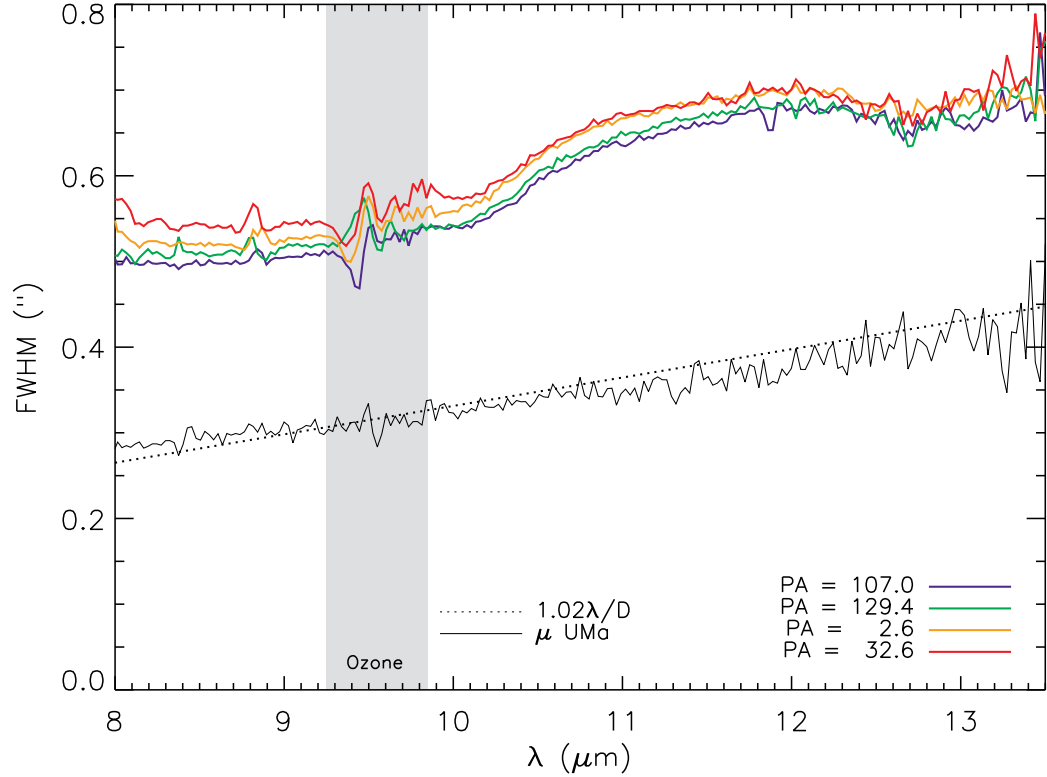


Figure C.3 The results of fitting profiles of the images presented in Figure C.1 as a function of wavelength, plotted as FWHM. The PSF core was fit with a Gaussian, which is expected to match a well corrected Airy pattern inside the first Airy minimum. For comparison we plot the predicted result for a circular pupil 6.35m in diameter with an 11% central obscuration (i.e. the MMT with the adaptive secondary, dotted line). Though the slope of the line does not match perfectly (likely due to a stop reducing the effective diameter or changing the central obscuration) it shows that the MMTAO system reached the diffraction limit for these observations. IRC +10216 was fit with a Lorentzian, which, though chosen for no astrophysical reason, matches our data well. The fits show clear evidence of a size change with wavelength, distinct from the effect of diffraction, between 10.2 and 12.6 μm , which matches the SiC emission feature. To avoid confusion we have indicated the spectral region typically impacted by telluric ozone. Also note the small feature at $\sim 8.8\mu\text{m}$ which can be attributed to a sharp feature in the detector QE.

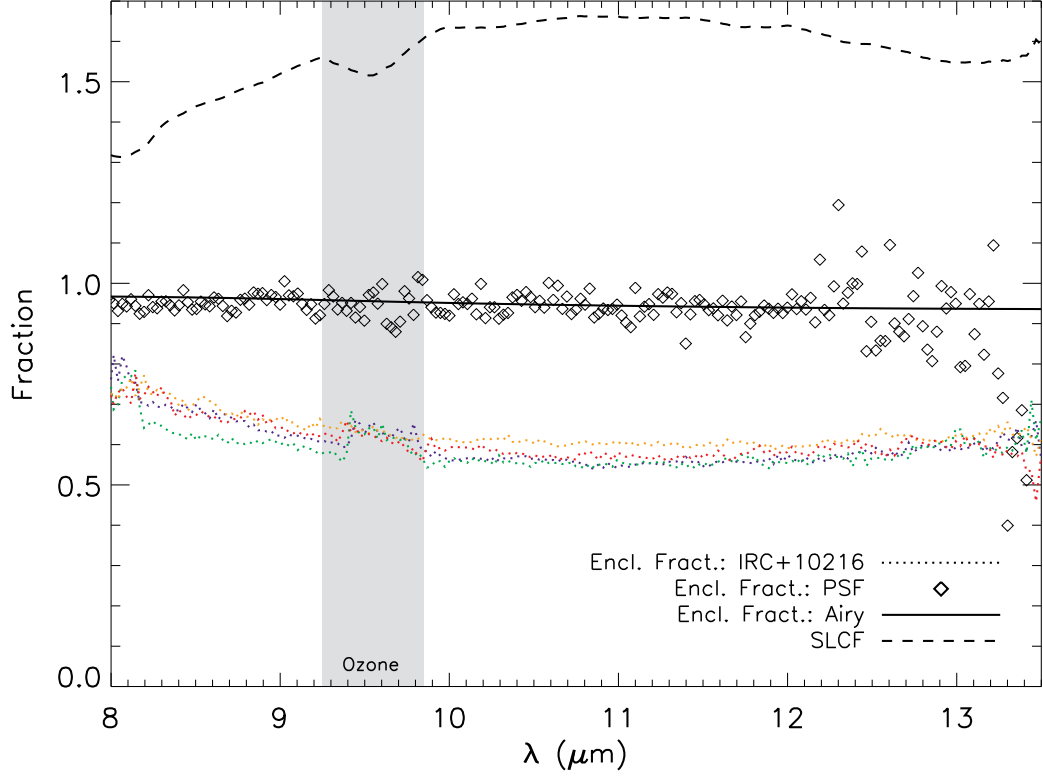


Figure C.4 Slit-loss correction calculations. We show the empirically calculated flux enclosed by the slit and photometric aperture for IRC +10216 (colored dots, using the same colors as Figure C.3 to denote PA), and the median of the 5 AO-on PSF standards obtained (diamonds), four of which are from the night after the IRC +10216 data were taken. Also plotted are the expected results for a centrally obscured Airy pattern, which we use for our final correction factor calculation to avoid introducing noise in our spectrum. Finally, we show the resultant slit-loss correction factor (SLCF), which we multiply with the spectrum of IRC +10216.

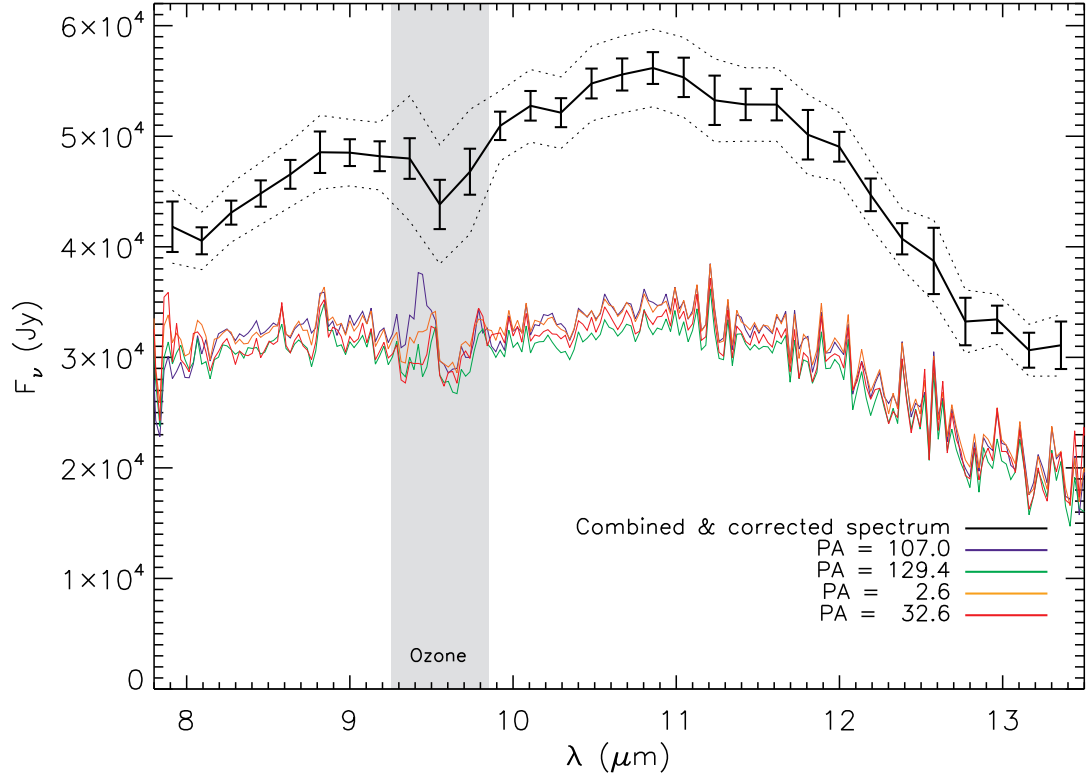


Figure C.5 Calibrated flux before and after correction for differential slit loss. The lower curves show the raw calibrated flux, before applying the SLCF, for each of the four slit position angles. The top curve is our fully corrected median combined spectrum, which takes into account the differential slit loss of the extended object compared to the PSF standard. See Figure C.4 and the text for further discussion of the SLCF. The error bars denote the local error, and the dashed lines denote our total uncertainty, which in addition to the local error includes the global (correlated) uncertainties. The average flux from $8 - 13\mu\text{m}$ is 47611 Jy, which is very similar to the value of 47627 Jy obtained by Monnier et al. (1998) at similar phase (near maximum brightness).

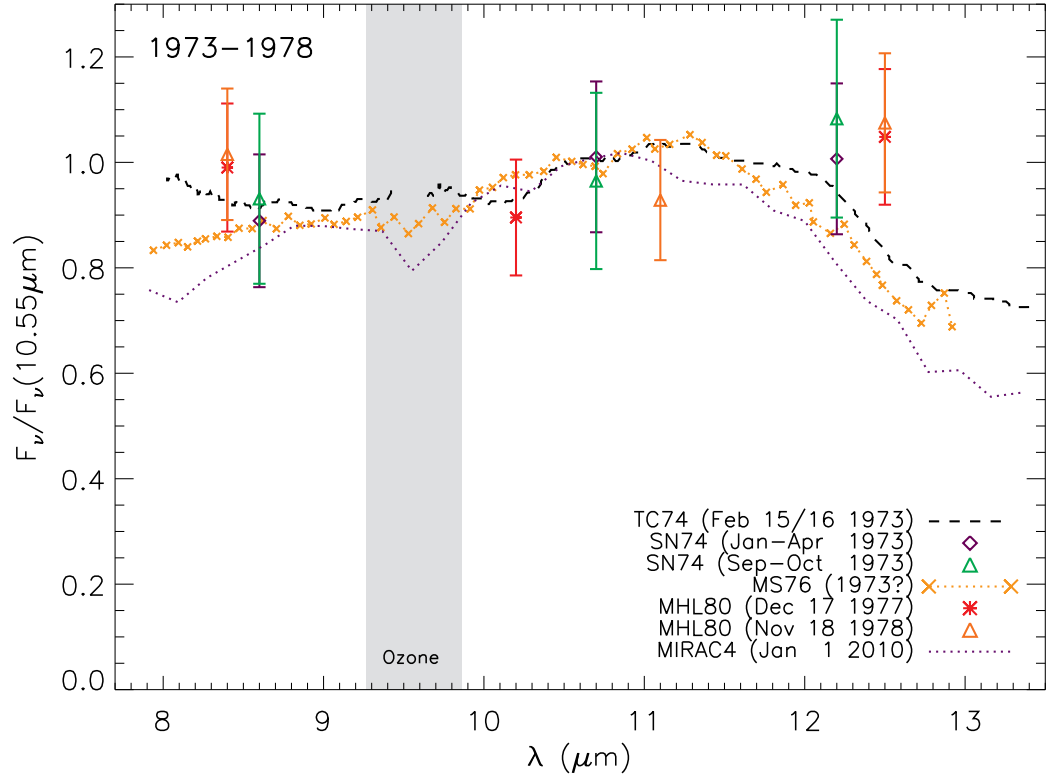


Figure C.6 IRC +10216 $10\mu\text{m}$ spectra and photometry from 1973 to 1978 normalized to $F_\nu(10.55\mu\text{m}) = 1$. Here we compare the bandpass photometry of Strecker and Ney (1974, SN74), the FTS spectrum of Treffers and Cohen (1974, TC74), and the CVF photometry of Merrill and Stein (1976, MS76), and the bandpass photometry of McCarthy et al. (1980, MHL80). We also show the shape of the spectrum in 2010 as reported in this work for reference. The five SN74 photometry epochs have been averaged (3 points from Jan-Apr, and 2 points in Sep-Oct) to reduce the 20% uncertainty in the individual points. Note that the Jan-Apr SN74 data and the TC74 data are essentially contemporaneous. The photometry generally supports the spectrum shape obtained by TC74, and is also consistent with the MS76 spectrum within 2σ uncertainty.

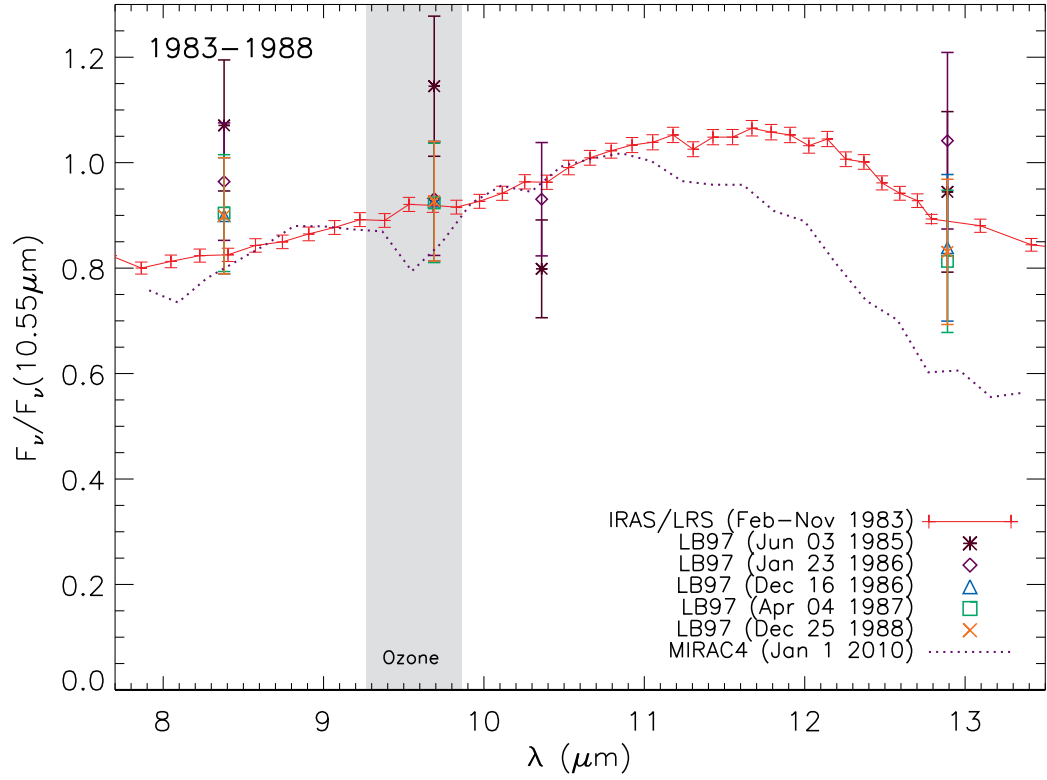


Figure C.7 IRC +10216 $10\mu\text{m}$ spectra from 1983 to 1988. Here we compare the space-based IRAS/LRS spectrum, and the bandpass photometry of Le Bertre (1997, LB97), normalized to $F_\nu(10.55\mu\text{m}) = 1$. The photometry appears to match the IRAS/LRS spectrum well, though as in Fig C.6 it is consistent with our 2010 data at the 2σ level. The LB97 points are taken far enough apart in time that we do not average in case there is short term variation in the shape.

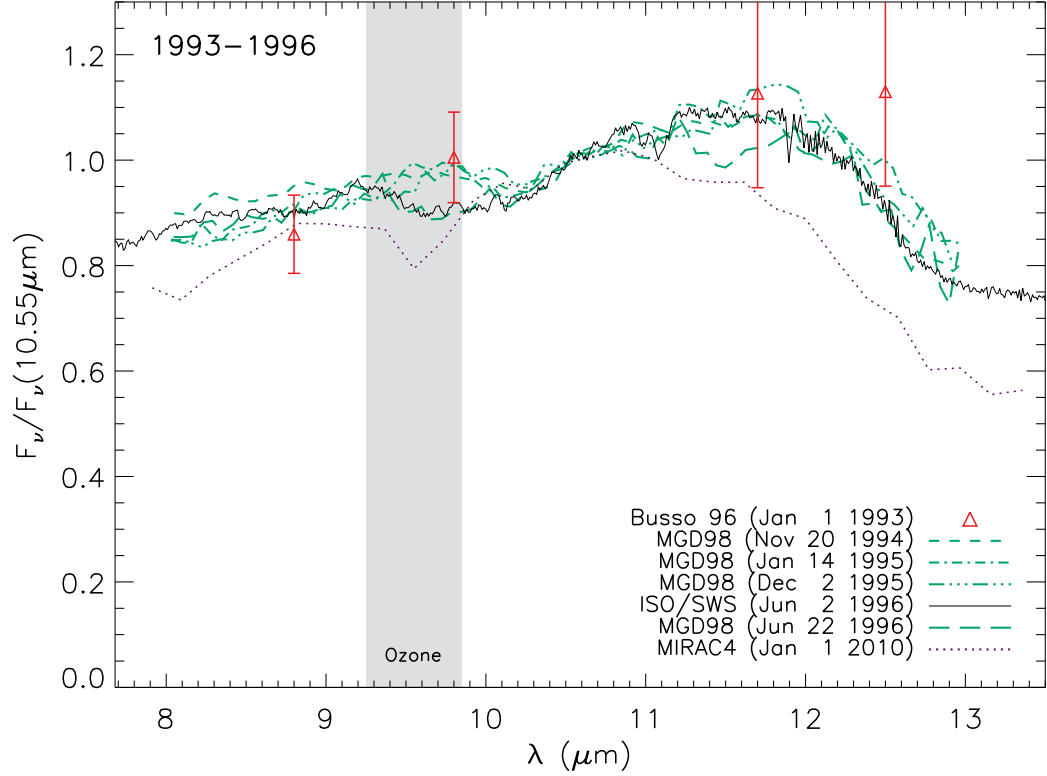


Figure C.8 IRC +10216 $10\mu\text{m}$ spectra from 1993 to 1996. Here we compare the bandpass photometry of Busso et al. (1996), the UKIRT spectra of Monnier et al. (1998, MGD98), and the space-based spectrum obtained by ISO/SWS. We normalized the data to $F_\nu(10.55\mu\text{m}) = 1$. All three data sets are in good agreement during this period, which spans three and a half years and well samples nearly two luminosity periods. The comparison with our 2010 data clearly shows that the change in the spectrum at $\lambda > 11\mu\text{m}$ is not simply associated with the regular 649 day variation in brightness of IRC +10216.

0₃

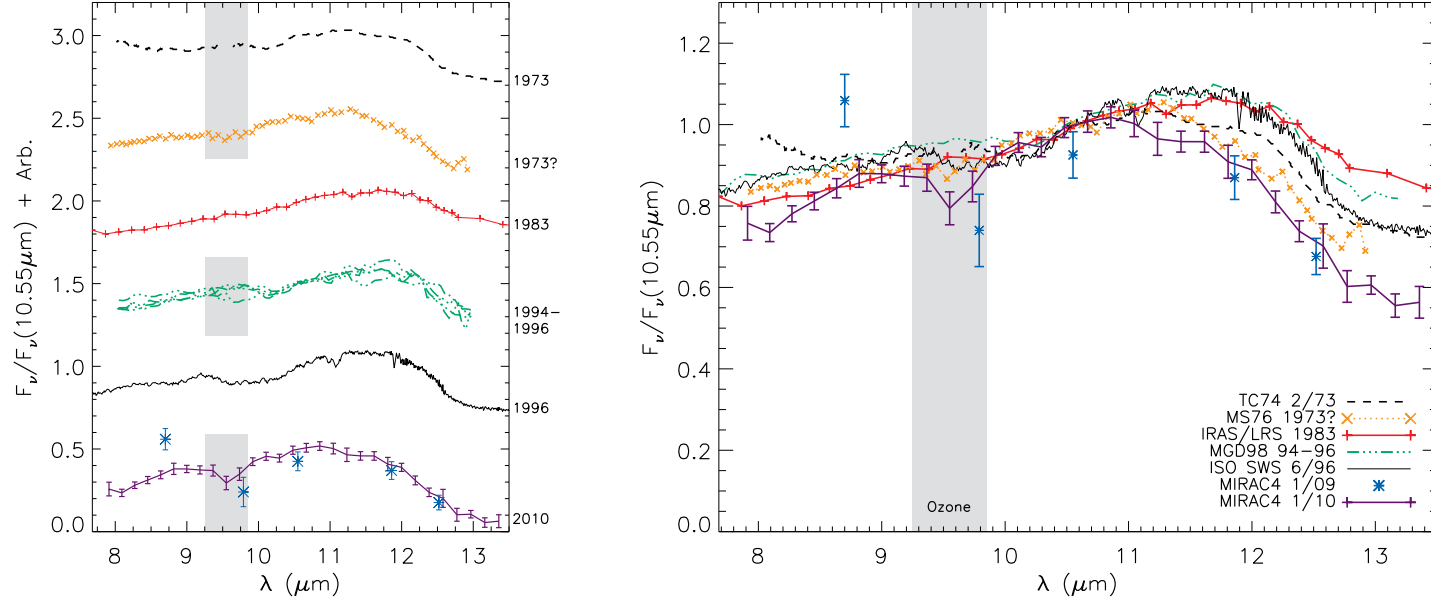


Figure C.9 Our new N band spectrum and photometry compared to previous observations of IRC +10216 spanning nearly 4 decades. As in Figures C.6-C.8 the data have been normalized at $10.55\mu\text{m}$. At left we have added an arbitrary constant to offset each epoch. The MIRAC4 photometry and grism spectrum, taken a year apart, match very well from $9.8\mu\text{m}$ to $12.5\mu\text{m}$. We present the same data at right without the offset. Of all the data from prior epochs, the spectrophotometry of MS76 is most similar to the 2009/2010 MIRAC4 data red-ward of $11\mu\text{m}$. The archival data rule out these changes being simply related to the regular 649 day Mira variability exhibited by IRC +10216.

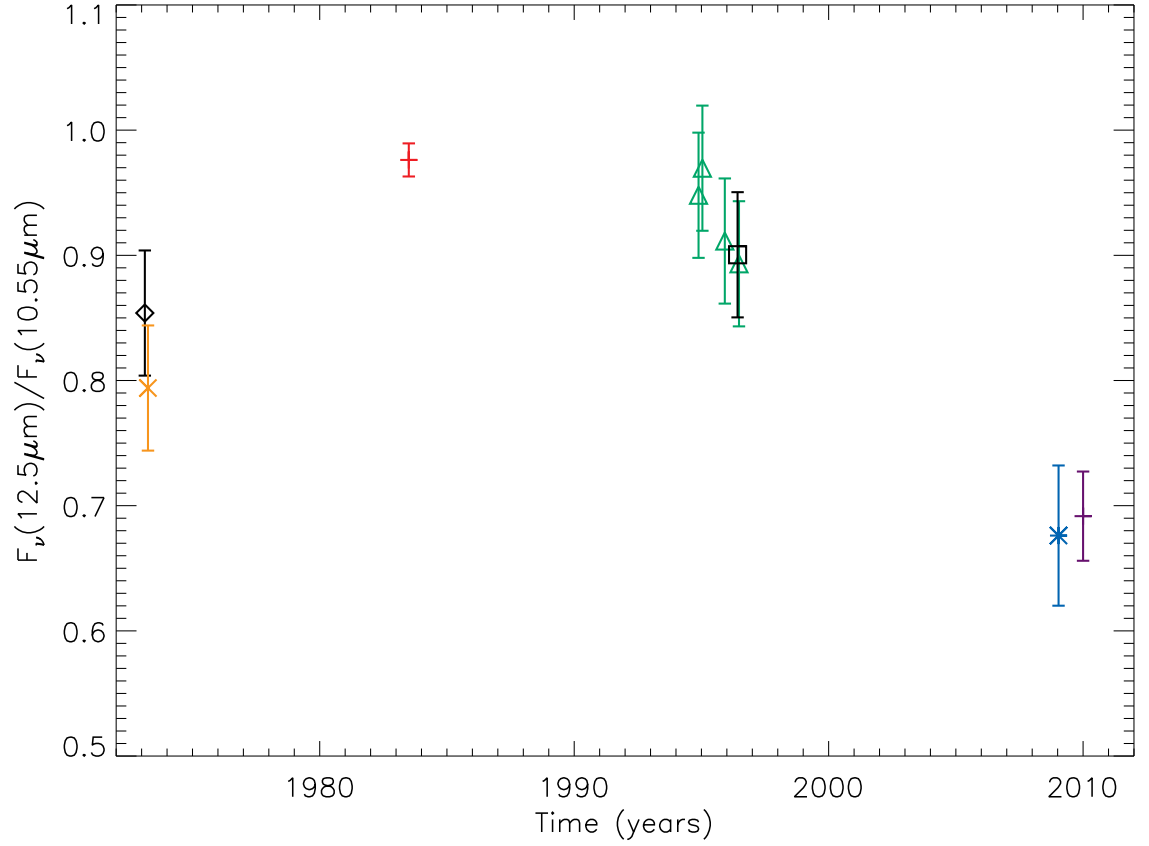


Figure C.10 The flux ratio $F_\nu(12.5\mu\text{m})/F_\nu(10.55\mu\text{m})$ vs. time. Colors are the same as in Figure C.9. Flux at $12.5\mu\text{m}$ was calculated as the mean between 12 and $13\mu\text{m}$. MIRAC4 and IRAS/LRS errors are as given. We adopt local or relative error of $\pm 5\%$ for the other data sets where such errors were not given. This plot illustrates the change in the shape of the spectrum over time, highlighting variability not associated with the regular 649 day Mira luminosity variations of IRC +10216. Whether this is a recurring spectrum shape which occurs at irregular intervals, or a longer term ($\gtrsim 40$ year) periodicity cannot be determined from the available data.

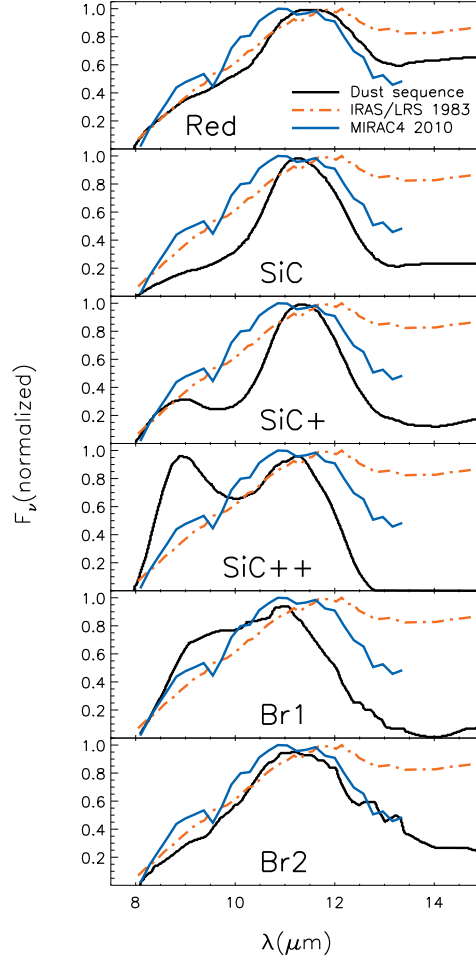


Figure C.11 The current state of IRC +10216 plotted on the “Carbon-Rich Dust Sequence” classification system of SLMP98. This system was based on the IRAS/LRS spectra of 96 carbon-rich AGB stars, and involves subtracting a 2400K blackbody (an approximation for the stellar continuum), normalizing, and visually inspecting the resulting curves. The heavy black curves are the summed and smooth spectra used to illustrate the sequence, and we show the 1983 IRAS/LRS spectrum (dotted red) and our 2010 MIRAC4 spectrum (blue). SLMP98 cited IRC +10216 as the prototype of the Red class, but it now (2010) appears to be a better match to the Broad 2 (Br2) spectra for $\lambda > 11\mu\text{m}$. Though the SLMP98 system does not represent an astrophysical sequence for C stars, it is useful in this case to show that our measurement of IRC +10216’s spectrum matches other C stars.

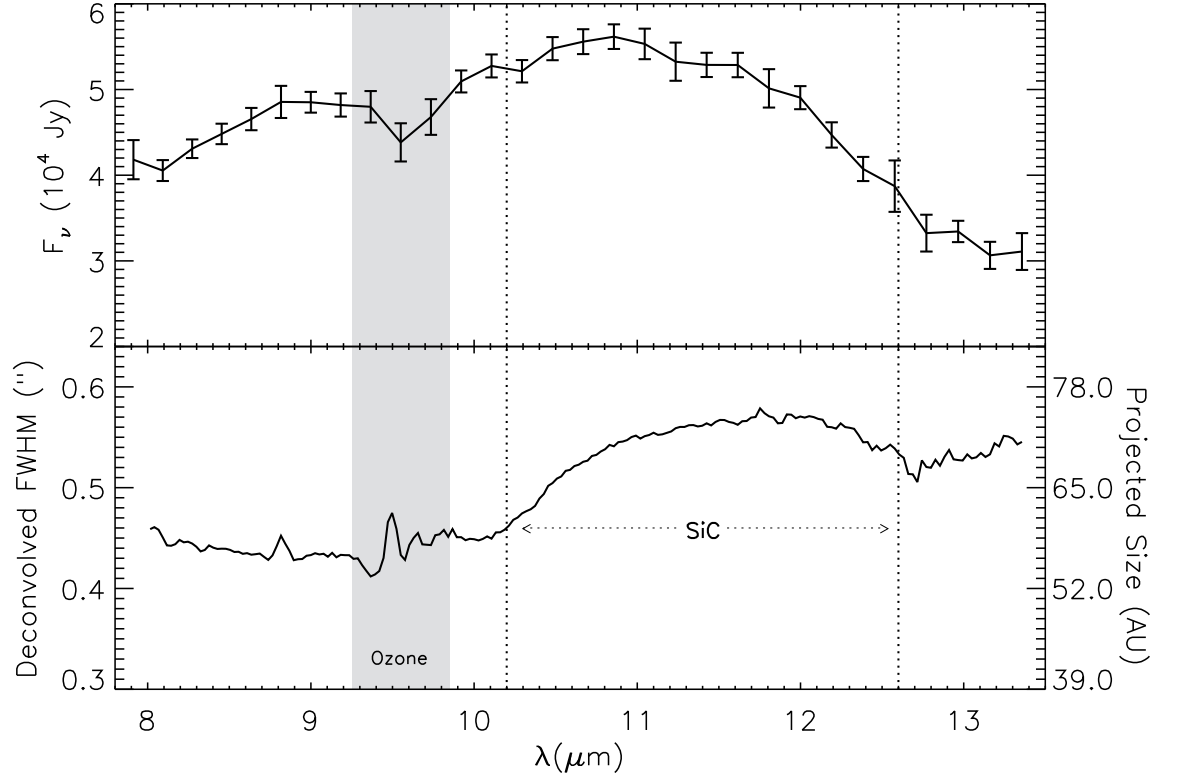


Figure C.12 The spectral and spatial signatures of SiC dust around IRC +10216. The well known SiC spectral feature can be seen in the MIRAC4 grism spectrum from 2010 in the top panel. In the bottom panel we have deconvolved the FWHM of IRC +10216 by subtracting the PSF FWHM in quadrature in order to estimate its intrinsic size, after averaging the four PAs. To provide a physical scale, we follow Men'shchikov et al. (2001) and adopt 130pc for the distance to IRC +10216 and calculate the projected size corresponding to the FWHM. An increase in the size of IRC +10216, clearly corresponding to the SiC feature, is evident (we have used the typical bounds for this feature as found by Clément et al. (2003)).

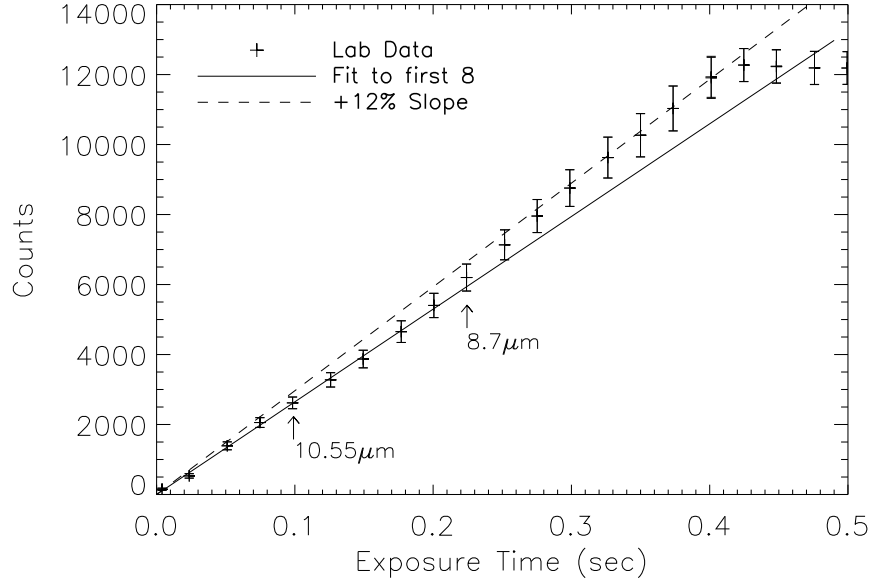


Figure C.13 MIRAC4 detector linearity measurement. The data were taken prior to the 2009 observations in a laboratory, and show that the detector becomes non-linear at higher fluxes, exhibiting an increase in slope which could be an explanation for the high flux detected at $8.7\mu\text{m}$ (see Figure C.9). We do not apply the curve directly to the data due to an unnoticed change in the detector bias that occurred between this measurement and the observations. The solid line is the fit to the first 8 data points. This fit was used to bias subtract the data. The dashed line has a 12% higher slope, chosen to illustrate a worst case scenario where every pixel in the $8.7\mu\text{m}$ that is brighter than the peak in the $10.55\mu\text{m}$ data has that slope. We also note the actual peak counts per read in the $8.7\mu\text{m}$ filter prior to background subtraction, and the peak counts per read in the $10.55\mu\text{m}$ after background subtraction (where the arrows intersect the solid line). The peak pixel may have become slightly non-linear, but the integrated non-linearity effect was likely $< 1\%$, causing us to rule it out as an explanation for the high flux at $8.7\mu\text{m}$.

Table C.1 Observations of IRC +10216 and standards.

Object	Filter (μm) ¹	Airmass	Pos. Angle ²	AO Speed (Hz)	No. Frames ³	Total Exp. Time (sec) ³
2009 13 Jan UT						
IRC +10216	8.7	1.35	N/A	Off	8	1.76
IRC +10216	9.79	1.38	N/A	Off	16	3.52
IRC +10216	10.55	1.32	N/A	Off	12	2.64
IRC +10216	11.86	1.40	N/A	Off	4	0.88
IRC +10216	12.52	1.31	N/A	Off	12	2.64
μ UMa	8.7	1.23	N/A	Off	24	5.28
μ UMa	9.79	1.22	N/A	Off	20	4.40
μ UMa	10.55	1.18	N/A	Off	20	4.40
μ UMa	11.86	1.20	N/A	Off	24	5.28
μ UMa	12.52	1.19	N/A	Off	16	3.52
2010 1 Jan UT						
IRC +10216	Grism	1.06	107.0	25	5	0.040
IRC +10216	Grism	1.10	129.4	25	2	0.016
IRC +10216	Grism	1.12	2.6	25	8	0.064
IRC +10216	Grism	1.25	32.6	25	6	0.032
μ UMa	Grism	1.02	73.1	100	6	10.000
2010 2 Jan UT ⁴						
β Gem	Grism	1.01	144.3	150	20	200.0
β Gem	Grism	1.07	164.0	550	8	80.0
μ UMa	Grism	1.03	133.7	550	8	80.0
μ UMa	Grism	1.06	92.1	550	8	80.0

¹Filter widths are given in Table C.2.²Position angle of the slit.³After rejecting bad chops, frames with excessive pattern noise, and bad slit alignment.⁴Data from this night were only used to check our slit loss correction procedure.

Table C.2 Bandpass photometry of IRC +10216 from 13 Jan 2009 UT.

Filter (μm)	Width ¹ (μm)	F_ν (Jy)	Obj ² σ (%)	PSF ³ σ (%)	Std. ⁴ σ (%)	Atm. ⁵ σ (%)	Total σ (%)
8.7	(8.08-9.32)	45099	1.63	0.78	2.4	4	5.0
9.79	(9.33 - 10.25)	31514	2.03	1.50	2.4	11	11.5
10.55	(10.06 - 11.04)	39408	1.50	1.52	2.4	4	5.1
11.86	(11.29 - 12.43)	37035	1.65	1.37	2.4	4	5.1
12.52	(11.94 - 13.11)	28790	1.50	2.10	2.9	4	5.6

¹Half power points of the manufacturer provided curves.

²The measurement uncertainty in the IRC +10216 photometry, estimated empirically.

³The measurement uncertainty in the μ UMa photometry, estimated empirically.

⁴Mean value of the total uncertainty given by Cohen et al. (1996) between the half power points.

⁵Based on the global and local telluric uncertainties of Skemer et al. (2010) from the following night.

Table C.3 Grism photometry from 1 Jan 2010 UT.

λ (μm)	F_ν (Jy)	Meas. ¹ σ (%)	Local ² σ (%)	Global ³ σ (%)	Total ⁴ σ (%)
7.913	41810.8	4.94	5.46	2.7	8.24
8.092	40541.0	1.79	3.01	2.7	6.84
8.273	43087.7	0.71	2.52	2.7	6.64
8.453	44816.2	1.12	2.66	2.7	6.69
8.635	46551.8	1.16	2.80	2.7	6.70
8.817	48551.0	3.03	3.86	2.7	7.26
8.999	48513.1	0.80	2.50	2.7	6.65
9.183	48190.2	1.44	2.80	2.7	6.75
9.366	47983.4	2.90	3.83	10.0	12.03
9.551	43828.8	4.51	5.09	10.0	12.51
9.735	46788.7	3.72	4.45	10.0	12.25
9.921	50939.4	0.94	2.52	2.7	6.66
10.107	52751.5	1.02	2.54	2.7	6.68
10.294	52128.7	0.82	2.51	2.7	6.65
10.481	54768.2	0.64	2.47	2.7	6.63
10.668	55583.9	0.95	2.61	2.7	6.67
10.857	56168.0	1.00	2.56	2.7	6.67
11.046	55322.7	2.06	3.22	2.7	6.91
11.235	53247.8	3.50	4.19	2.7	7.47
11.425	52871.8	1.28	2.68	2.7	6.72
11.616	52857.9	1.35	2.70	2.7	6.73
11.807	50134.5	3.80	4.47	2.7	7.62
11.999	49050.2	1.41	2.74	2.7	6.75
12.191	44697.1	2.19	3.30	2.7	6.95
12.384	40722.9	2.58	3.47	2.7	7.09
12.577	38720.9	7.14	7.75	2.7	9.72
12.771	33237.0	5.53	6.46	2.7	8.61
12.966	33431.4	1.61	3.72	2.7	6.79
13.161	30643.0	3.92	5.16	2.7	7.67
13.357	31091.5	6.04	6.91	2.7	8.95

¹Measurement uncertainty, including IRC +10216 and the standard μ UMa.²Total local uncertainty, including 2.31-3.35% uncertainty from the μ UMa spectrum of Cohen et al. (1996).³The estimated global telluric calibration uncertainty.⁴Includes 5% systematic uncertainty from the SLCF.

REFERENCES

- Allers, K. N. and M. C. Liu (2013). A Near-infrared Spectroscopic Study of Young Field Ultracool Dwarfs. *ApJ*, **772**, 79. doi:10.1088/0004-637X/772/2/79.
- Anderson, J. and I. R. King (2006). PSFs, Photometry, and Astronomy for the ACS/WFC. Technical report.
- Augereau, J. C., R. P. Nelson, A. M. Lagrange, J. C. B. Papaloizou, and D. Mouillet (2001). Dynamical modeling of large scale asymmetries in the beta Pictoris dust disk. *A&A*, **370**, pp. 447–455. doi:10.1051/0004-6361:20010199.
- Baluev, R. V. (2012). The impact of red noise in radial velocity planet searches: Only three planets orbiting GJ581? *ArXiv e-prints*.
- Baraffe, I., G. Chabrier, T. S. Barman, F. Allard, and P. H. Hauschildt (2003). Evolutionary models for cool brown dwarfs and extrasolar giant planets. The case of HD 209458. *A&A*, **402**, pp. 701–712. doi:10.1051/0004-6361:20030252.
- Baranec, C., R. Riddle, A. N. Ramaprakash, N. Law, S. Tendulkar, S. Kulkarni, R. Dekany, K. Bui, J. Davis, M. Burse, H. Das, S. Hildebrandt, S. Punnadi, and R. Smith (2012). Robo-AO: autonomous and replicable laser-adaptive-optics and science system. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8447 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. doi:10.1117/12.924867.
- Barman, T. S., B. Macintosh, Q. M. Konopacky, and C. Marois (2011). The Young Planet-mass Object 2M1207b: A Cool, Cloudy, and Methane-poor Atmosphere. *ApJ*, **735**, L39. doi:10.1088/2041-8205/735/2/L39.
- Baron, Y., R. Papoular, M. Jourdain de Muizon, and B. Pegourie (1987). An analysis of the emission features of the IRAS low-resolution spectra of carbon stars. *A&A*, **186**, pp. 271–279.
- Batygin, K. and D. J. Stevenson (2013). Mass-Radius Relationships for Very Low Mass Gaseous Planets. *ArXiv e-prints*.
- Becklin, E. E., J. A. Frogel, A. R. Hyland, J. Kristian, and G. Neugebauer (1969). The Unusual Infrared Object IRC+10216. *ApJ*, **158**, pp. L133+. doi:10.1086/180450.

- Benedict, G. F., B. E. McArthur, G. Gatewood, E. Nelan, W. D. Cochran, A. Hatzes, M. Endl, R. Wittenmyer, S. L. Baliunas, G. A. H. Walker, S. Yang, M. Kürster, S. Els, and D. B. Paulson (2006). The Extrasolar Planet ϵ Eridani b: Orbit and Mass. *AJ*, **132**, pp. 2206–2218. doi:10.1086/508323.
- Bernstein, G. M., D. E. Trilling, R. L. Allen, M. E. Brown, M. Holman, and R. Malhotra (2004). The Size Distribution of Trans-Neptunian Bodies. *AJ*, **128**, pp. 1364–1390. doi:10.1086/422919.
- Bessel, M. S. (1990). BVRI photometry of the Gliese Catalogue stars. *A&AS*, **83**, pp. 357–378.
- Bessell, M. S. (1990). UBVRI passbands. *PASP*, **102**, pp. 1181–1199. doi:10.1086/132749.
- Bessell, M. S. (2000). The Hipparcos and Tycho Photometric System Passbands. *PASP*, **112**, pp. 961–965. doi:10.1086/316598.
- Beust, H., A. Vidal-Madjar, and R. Ferlet (1991). The Beta Pictoris protoplanetary system. XII - Planetary perturbations in the disk and star-grazing bodies. *A&A*, **247**, pp. 505–515.
- Beust, H., A. Vidal-Madjar, R. Ferlet, and A. M. Lagrange-Henri (1990). The Beta Pictoris circumstellar disk. X - Numerical simulations of infalling evaporating bodies. *A&A*, **236**, pp. 202–216.
- Biller, B. A., L. M. Close, A. Li, M. Marengo, J. H. Bieging, P. M. Hinz, W. F. Hoffmann, G. Brusa, and D. Miller (2006). Resolving the Dusty Circumstellar Structure of the Enigmatic Symbiotic Star CH Cygni with the MMT Adaptive Optics System. *ApJ*, **647**, pp. 464–470. doi:10.1086/505194.
- Biller, B. A., L. M. Close, E. Masciadri, E. Nielsen, R. Lenzen, W. Brandner, D. McCarthy, M. Hartung, S. Kellner, E. Mamajek, T. Henning, D. Miller, M. Kenworthy, and C. Kulesa (2007). An Imaging Survey for Extrasolar Planets around 45 Close, Young Stars with the Simultaneous Differential Imager at the Very Large Telescope and MMT. *ApJS*, **173**, pp. 143–165. doi:10.1086/519925.
- Blandford, et al. (2010). *Astro2010*.
- Boccaletti, A., A.-M. Lagrange, M. Bonnefoy, R. Galicher, and G. Chauvin (2013). Independent confirmation of β Pictoris b imaging with NICI. *A&A*, **551**, L14. doi:10.1051/0004-6361/201321255.
- Bohlin, R. C. (2007). HST Stellar Standards with 1% Accuracy in Absolute Flux. In Sterken, C. (ed.) *The Future of Photometric, Spectrophotometric and Polarimetric*

- Standardization*, volume 364 of *Astronomical Society of the Pacific Conference Series*, p. 315.
- Bonnefoy, M., A. Boccaletti, A.-M. Lagrange, F. Allard, C. Mordasini, H. Beust, G. Chauvin, J. H. V. Girard, D. Homeier, D. Apai, S. Lacour, and D. Rouan (2013). The near-infrared spectral energy distribution of β Pictoris b. *ArXiv e-prints*.
- Bonnefoy, M., A.-M. Lagrange, A. Boccaletti, G. Chauvin, D. Apai, F. Allard, D. Ehrenreich, J. H. V. Girard, D. Mouillet, D. Rouan, D. Gratadour, and M. Kasper (2011). High angular resolution detection of β Pictoris b at 2.18 μ m. *A&A*, **528**, L15. doi:10.1051/0004-6361/201016224.
- Born, M. and E. Wolf (1999). *Principles of Optics*.
- Bovaird, T. and C. H. Lineweaver (2013). Exoplanet Predictions Based on the Generalised Titius-Bode Relation. *ArXiv e-prints*.
- Brown, R. A. (2004). Obscurational Completeness. *ApJ*, **607**, pp. 1003–1013. doi:10.1086/383586.
- Brown, R. A. (2005). Single-Visit Photometric and Obscurational Completeness. *ApJ*, **624**, pp. 1010–1024. doi:10.1086/429124.
- Brown, R. A. and R. Soummer (2010). New Completeness Methods for Estimating Exoplanet Discoveries by Direct Detection. *ApJ*, **715**, pp. 122–131. doi:10.1088/0004-637X/715/1/122.
- Bruntt, H., T. R. Bedding, P.-O. Quirion, G. Lo Curto, F. Carrier, B. Smalley, T. H. Dall, T. Arentoft, M. Bazot, and R. P. Butler (2010). Accurate fundamental parameters for 23 bright solar-type stars. *MNRAS*, **405**, pp. 1907–1923. doi:10.1111/j.1365-2966.2010.16575.x.
- Burningham, B., C. V. Cardoso, L. Smith, S. K. Leggett, R. L. Smart, A. W. Mann, S. Dhital, P. W. Lucas, C. G. Tinney, D. J. Pinfield, Z. Zhang, C. Morley, D. Saumon, K. Aller, S. P. Littlefair, D. Homeier, N. Lodieu, N. Deacon, M. S. Marley, L. van Spaandonk, D. Baker, F. Allard, A. H. Andrei, J. Canty, J. Clarke, A. C. Day-Jones, T. Dupuy, J. J. Fortney, J. Gomes, M. Ishii, H. R. A. Jones, M. Liu, A. Magazzú, F. Marocco, D. N. Murray, B. Rojas-Ayala, and M. Tamura (2013). 76 T dwarfs from the UKIDSS LAS: benchmarks, kinematics and an updated space density. *MNRAS*, **433**, pp. 457–497. doi:10.1093/mnras/stt740.
- Burrows, C. J., J. E. Krist, K. R. Stapelfeldt, and WFPC2 Investigation Definition Team (1995). HST Observations of the Beta Pictoris Circumstellar Disk. In *American Astronomical Society Meeting Abstracts*, volume 27 of *Bulletin of the American Astronomical Society*, p. 1329.

- Busso, M., L. Origlia, M. Marengo, P. Persi, M. Ferrari-Toniolo, G. Silvestro, L. Corcione, M. Tapia, and J. Bohigas (1996). Mid-infrared imaging of AGB star envelopes. I. Colors, classification and mass-loss estimates. *A&A*, **311**, pp. 253–263.
- Cahoy, K. L., M. S. Marley, and J. J. Fortney (2010). Exoplanet Albedo Spectra and Colors as a Function of Planet Phase, Separation, and Metallicity. *ApJ*, **724**, pp. 189–214. doi:10.1088/0004-637X/724/1/189.
- Carillet, M., C. Vérinaud, S. Esposito, A. Riccardi, A. Puglisi, B. Femenía, and L. Fini (2003). Performance of the first-light adaptive optics system of LBT by means of CAOS simulations. volume 4839 of *Proc. SPIE*, pp. 131–139. doi:10.1117/12.479721.
- Carillet, M., C. Vérinaud, B. Femenía, A. Riccardi, and L. Fini (2005). Modelling astronomical adaptive optics - I. The software package CAOS. *MNRAS*, **356**, pp. 1263–1275. doi:10.1111/j.1365-2966.2004.08524.x.
- Cavarroc, C., A. Boccaletti, P. Baudoz, T. Fusco, and D. Rouan (2006). Fundamental limitations on Earth-like planet detection with extremely large telescopes. *A&A*, **447**, pp. 397–403. doi:10.1051/0004-6361:20053916.
- Cernicharo, J., I. Yamamura, E. González-Alfonso, T. de Jong, A. Heras, R. Escribano, and J. Ortigoso (1999). The ISO/SWS Spectrum of IRC +10216: The Vibrational Bands of C_2H_2 and HCN. *ApJ*, **526**, pp. L41–L44. doi:10.1086/312360.
- Chan, S. J. and S. Kwok (1990). Evolution of infrared carbon stars. *A&A*, **237**, pp. 354–368.
- Chandler, A. A., K. Tatebe, D. D. S. Hale, and C. H. Townes (2007). The Radiative Pattern and Asymmetry of IRC +10216 at 11 μ m Measured with Interferometry and Closure Phase. *ApJ*, **657**, pp. 1042–1045. doi:10.1086/510796.
- Chauvin, G., A.-M. Lagrange, H. Beust, M. Bonnefoy, A. Boccaletti, D. Apai, F. Allard, D. Ehrenreich, J. H. V. Girard, D. Mouillet, and D. Rouan (2012a). Orbital characterization of the β Pictoris b giant planet. *A&A*, **542**, A41. doi:10.1051/0004-6361/201118346.
- Chauvin, G., A.-M. Lagrange, H. Beust, M. Bonnefoy, A. Boccaletti, D. Apai, F. Allard, D. Ehrenreich, J. H. V. Girard, D. Mouillet, and D. Rouan (2012b). Orbital characterization of the β Pictoris b giant planet. *A&A*, **542**, A41. doi:10.1051/0004-6361/201118346.
- Chiang, E. I. and M. E. Brown (1999). Keck Pencil-Beam Survey for Faint Kuiper Belt Objects. *AJ*, **118**, pp. 1411–1422. doi:10.1086/301005.

- Clément, D., H. Mutschke, R. Klein, and T. Henning (2003). New Laboratory Spectra of Isolated β -SiC Nanoparticles: Comparison with Spectra Taken by the Infrared Space Observatory. *ApJ*, **594**, pp. 642–650. doi:10.1086/376864.
- Close, L. M., B. Biller, W. F. Hoffmann, P. M. Hinz, J. H. Biegging, F. Wildi, M. Lloyd-Hart, G. Brusa, D. Fisher, D. Miller, and R. Angel (2003). Mid-Infrared Imaging of the Post-Asymptotic Giant Branch Star AC Herculis with the Multiple Mirror Telescope Adaptive Optics System. *ApJ*, **598**, pp. L35–L38. doi:10.1086/380429.
- Close, L. M., J. R. Males, D. Kopon, V. Gasho, K. B. Follette, P. Hinz, K. Morzinski, A. Uomoto, T. Hare, A. Riccardi, S. Esposito, A. Puglisi, E. Pinna, L. Busoni, C. Arcidiacono, M. Xompero, R. Briguglio, F. Quiros-Pacheco, and J. Argomedo (2012a). First closed-loop visible AO test results for the advanced adaptive secondary AO system for the Magellan Telescope: MagAO’s performance and status. volume 8447 of *Proc. SPIE*.
- Close, L. M., J. R. Males, D. A. Kopon, V. Gasho, K. B. Follette, P. Hinz, K. Morzinski, A. Uomoto, T. Hare, A. Riccardi, S. Esposito, A. Puglisi, E. Pinna, L. Busoni, C. Arcidiacono, M. Xompero, R. Briguglio, F. Quiros-Pacheco, and J. Argomedo (2012b). First closed-loop visible AO test results for the advanced adaptive secondary AO system for the Magellan Telescope: MagAO’s performance and status. *Proc. SPIE*, **8447**, 84470X. doi:10.1117/12.926545.
- Close, L. M., A. Puglisi, J. R. Males, C. Arcidiacono, A. Skemer, J. C. Guerra, L. Busoni, G. Brusa, E. Pinna, D. L. Miller, A. Riccardi, D. W. McCarthy, M. Xompero, C. Kulesa, F. Quiros-Pacheco, J. Argomedo, J. Brynnel, S. Esposito, F. Mannucci, K. Boutsia, L. Fini, D. J. Thompson, J. M. Hill, C. E. Woodward, R. Briguglio, T. J. Rodigas, R. Briguglio, P. Stefanini, G. Agapito, P. Hinz, K. Follette, and R. Green (2012c). High-resolution Images of Orbital Motion in the Orion Trapezium Cluster with the LBT AO System. *ApJ*, **749**, 180. doi:10.1088/0004-637X/749/2/180.
- Cohen, M., R. G. Walker, and F. C. Witteborn (1992). Spectral irradiance calibration in the infrared. II - Alpha Tau and the recalibration of the IRAS low resolution spectrometer. *AJ*, **104**, pp. 2030–2044. doi:10.1086/116379.
- Cohen, M., W. A. Wheaton, and S. T. Megeath (2003). Spectral Irradiance Calibration in the Infrared. XIV. The Absolute Calibration of 2MASS. *AJ*, **126**, pp. 1090–1096. doi:10.1086/376474.
- Cohen, M., F. C. Witteborn, D. F. Carbon, J. K. Davies, D. H. Wooden, and J. D. Bregman (1996). Spectral Irradiance Calibration in the Infrared. VII. New Composite Spectra, Comparison with Model Atmospheres, and Far-Infrared Extrapolations. *AJ*, **112**, pp. 2274–+. doi:10.1086/118180.

- Cohen, M., F. C. Witteborn, R. G. Walker, J. D. Bregman, and D. H. Wooden (1995). Spectral Irradiance Calibration in the Infrared.IV. 1.2-35 micron spectra of six standard stars. *AJ*, **110**, pp. 275–+. doi:10.1086/117517.
- Currie, T., A. Burrows, Y. Itoh, S. Matsumura, M. Fukagawa, D. Apai, N. Madhusudhan, P. M. Hinz, T. J. Rodigas, M. Kasper, T.-S. Pyo, and S. Ogino (2011a). A Combined Subaru/VLT/MMT 1-5 μ m Study of Planets Orbiting HR 8799: Implications for Atmospheric Properties, Masses, and Formation. *ApJ*, **729**, 128. doi:10.1088/0004-637X/729/2/128.
- Currie, T., A. Burrows, N. Madhusudhan, M. Fukagawa, J. H. Girard, R. Dawson, R. Murray-Clay, S. Kenyon, M. Kuchner, S. Matsumura, R. Jayawardhana, J. Chambers, and B. Bromley (2013). A Combined VLT and Gemini Study of the Atmosphere of the Directly-Imaged Planet, beta Pictoris b. *ArXiv e-prints*.
- Currie, T., C. Thalmann, S. Matsumura, N. Madhusudhan, A. Burrows, and M. Kuchner (2011b). A 5 μ m Image of β Pictoris b at a Sub-Jupiter Projected Separation: Evidence for a Misalignment Between the Planet and the Inner, Warped Disk. *ApJ*, **736**, L33. doi:10.1088/2041-8205/736/2/L33.
- Daigle, O., C. Carignan, J. Gach, C. Guillaume, S. Lessard, C. Fortin, and S. Blais-Ouellette (2009). Extreme Faint Flux Imaging with an EMCCD. *PASP*, **121**, pp. 866–884. doi:10.1086/605449.
- Danchi, W. C., M. Bester, C. G. Degiacomi, P. R. McCullough, and C. H. Townes (1990). Location and phase of dust formation in IRC + 10216 indicated by 11 micron spatial interferometry. *ApJ*, **359**, pp. L59–L63. doi:10.1086/185795.
- Dawson, R. I., R. A. Murray-Clay, and D. C. Fabrycky (2011). On the Misalignment of the Directly Imaged Planet β Pictoris b with the System’s Warped Inner Disk. *ApJ*, **743**, L17. doi:10.1088/2041-8205/743/1/L17.
- Decin, L., M. Agúndez, M. J. Barlow, F. Daniel, J. Cernicharo, R. Lombaert, E. de Beck, P. Royer, B. Vandenbussche, R. Wesson, E. T. Polehampton, J. A. D. L. Blommaert, W. De Meester, K. Exter, H. Feuchtgruber, W. K. Gear, H. L. Gomez, M. A. T. Groenewegen, M. Guélin, P. C. Hargrave, R. Huygen, P. Imhof, R. J. Ivison, C. Jean, C. Kahane, F. Kerschbaum, S. J. Leeks, T. Lim, M. Matsuura, G. Olofsson, T. Posch, S. Regibo, G. Savini, B. Sibthorpe, B. M. Swinyard, J. A. Yates, and C. Waelkens (2010). Warm water vapour in the sooty outflow from a luminous carbon star. *Nature*, **467**, pp. 64–67. doi:10.1038/nature09344.
- Diolaiti, E., O. Bendinelli, D. Bonaccini, L. Close, D. Currie, and G. Parmeggiani (2000). Analysis of isoplanatic high resolution stellar fields by the StarFinder code. *A&AS*, **147**, pp. 335–346. doi:10.1051/aas:2000305.

- Ducati, J. R., C. M. Bevilacqua, S. B. Rembold, and D. Ribeiro (2001). Intrinsic Colors of Stars in the Near-Infrared. *ApJ*, **558**, pp. 309–322. doi:10.1086/322439.
- Dumusque, X., F. Pepe, C. Lovis, D. Ségransan, J. Sahlmann, W. Benz, F. Bouchy, M. Mayor, D. Queloz, N. Santos, and S. Udry (2012). An Earth-mass planet orbiting α Centauri B. *Nature*, **491**, pp. 207–211. doi:10.1038/nature11572.
- Dupuy, T. J. and M. C. Liu (2012). The Hawaii Infrared Parallax Program. I. Ultracool Binaries and the L/T Transition. *ApJS*, **201**, 19. doi:10.1088/0067-0049/201/2/19.
- Dyck, H. M., J. A. Benson, R. R. Howell, R. R. Joyce, and C. Leinert (1991). The light curve and changes in the circumstellar envelope around IRC + 10216. *AJ*, **102**, pp. 200–207. doi:10.1086/115866.
- Eggl, S., E. Pilat-Lohinger, B. Funk, N. Georgakarakos, and N. Haghighipour (2013). Circumstellar habitable zones of binary-star systems in the solar neighbourhood. *MNRAS*, **428**, pp. 3104–3113. doi:10.1093/mnras/sts257.
- Esposito, S., A. Riccardi, F. Quirós-Pacheco, E. Pinna, A. Puglisi, M. Xompero, R. Briguglio, L. Busoni, L. Fini, P. Stefanini, G. Brusa, A. Tozzi, P. Ranfagni, F. Pieralli, J. C. Guerra, C. Arcidiacono, and P. Salinari (2010). Laboratory characterization and performance of the high-order adaptive optics system for the Large Binocular Telescope. *AO*, **49**, pp. G174+. doi:10.1364/AO.49.00G174.
- Faherty, J. K., E. L. Rice, K. L. Cruz, E. E. Mamajek, and A. Núñez (2013). 2MASS J035523.37+113343.7: A Young, Dusty, Nearby, Isolated Brown Dwarf Resembling a Giant Exoplanet. *AJ*, **145**, 2. doi:10.1088/0004-6256/145/1/2.
- Floyd, D. J. E., J. Thomas-Osip, and G. Prieto (2010). Seeing, Wind, and Outer Scale Effects on Image Quality at the Magellan Telescopes. *PASP*, **122**, pp. 731–742. doi:10.1086/653740.
- Fortney, J. J., I. Baraffe, and B. Militzer (2011). *Giant Planet Interior Structure and Thermal Evolution*, pp. 397–418.
- Forveille, T., X. Bonfils, X. Delfosse, R. Alonso, S. Udry, F. Bouchy, M. Gillon, C. Lovis, V. Neves, M. Mayor, F. Pepe, D. Queloz, N. C. Santos, D. Ségransan, J. M. Almenara, H. Deeg, and M. Rabus (2011). The HARPS search for southern extra-solar planets XXXII. Only 4 planets in the Gl 581 system. *ArXiv e-prints*.
- Fried, D. L. (1978). Probability of getting a lucky short-exposure image through turbulence. *JOSA*, **68**, pp. 1651–1657.
- Friedemann, C. (1969). Evolution of silicon carbide particles in the atmospheres of carbon stars. *Physica*, **41**, pp. 139–143. doi:10.1016/0031-8914(69)90247-X.

- Frieswijk, W. F., R. F. Shipman, F. Lahuis, and F. Hormuth (2007). SWS AOT-1 High Resolution Processing: Documentation. Technical report, European Space Agency. http://ida.esac.esa.int:8080/hpdp/technical_reports/technote52.pdf.
- Fukugita, M., T. Ichikawa, J. E. Gunn, M. Doi, K. Shimasaku, and D. P. Schneider (1996). The Sloan Digital Sky Survey Photometric System. *AJ*, **111**, p. 1748. doi:10.1086/117915.
- Garrel, V., O. Guyon, P. Baudoz, F. Martinache, F. Vogt, Y. Takashi, Y. Kaito, and F. Cantalloube (2010). The Subaru coronagraphic extreme AO (SCExAO) system: visible imaging mode. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7736 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. doi:10.1117/12.857874.
- Gilman, R. C. (1969). On the Composition of Circumstellar Grains. *ApJ*, **155**, pp. L185+. doi:10.1086/180332.
- Gilra, D. P. (1971). Composition of Interstellar Grains. *Nature*, **229**, pp. 237–241. doi:10.1038/229237a0.
- Gladysz, S., J. Christou, N. Law, R. Dekany, M. Redfern, and C. Mackay (2008a). Lucky imaging and speckle discrimination for the detection of faint companions with adaptive optics. volume 7015, of *Proc. SPIE*. doi:10.1117/12.788442.
- Gladysz, S., J. C. Christou, L. W. Bradford, and L. C. Roberts (2008b). Temporal Variability and Statistics of the Strehl Ratio in Adaptive-Optics Images. *PASP*, **120**, pp. 1132–1143. doi:10.1086/592787.
- Gladysz, S., J. C. Christou, and M. Redfern (2006). Characterization of the Lick adaptive optics point spread function. volume 6272, of *Proc. SPIE*. doi:10.1117/12.671605.
- Golimowski, D. A., D. R. Ardila, J. E. Krist, M. Clampin, H. C. Ford, G. D. Illingworth, F. Bartko, N. Benítez, J. P. Blakeslee, R. J. Bouwens, L. D. Bradley, T. J. Broadhurst, R. A. Brown, C. J. Burrows, E. S. Cheng, N. J. G. Cross, R. Demarco, P. D. Feldman, M. Franx, T. Goto, C. Gronwall, G. F. Hartig, B. P. Holden, N. L. Homeier, L. Infante, M. J. Jee, R. A. Kimble, M. P. Lesser, A. R. Martel, S. Mei, F. Menanteau, G. R. Meurer, G. K. Miley, V. Motta, M. Postman, P. Rosati, M. Sirianni, W. B. Sparks, H. D. Tran, Z. I. Tsvetanov, R. L. White, W. Zheng, and A. W. Zirm (2006). Hubble Space Telescope ACS Multiband Coronagraphic Imaging of the Debris Disk around β Pictoris. *AJ*, **131**, pp. 3109–3130. doi:10.1086/503801.

- Golimowski, D. A., S. T. Durrance, and M. Clampin (1993). Coronagraphic imaging of the Beta Pictoris circumstellar disk - Evidence of changing disk structure within 100 AU. *ApJ*, **411**, pp. L41–L44. doi:10.1086/186907.
- Guyon, O., F. Martinache, E. J. Cady, R. Belikov, K. Balasubramanian, D. Wilson, C. S. Clergeon, and M. Mateen (2012). How ELTs will acquire the first spectra of rocky habitable planets. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8447 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. doi:10.1117/12.927181.
- Habing, H. J. and H. Olofsson (eds.) (2003). *Asymptotic giant branch stars*. Astronomy and Astrophysics. Springer.
- Hackwell, J. A. (1972). Long wavelength spectrometry and photometry of M, S and C-stars. *A&A*, **21**, pp. 239–248.
- Haniff, C. A. and D. F. Buscher (1998). Variable sub-arcsecond structure in the circumstellar envelope of IRC+10216. *A&A*, **334**, pp. L5–L8.
- Hardy, J. W. (1998). *Adaptive Optics for Astronomical Telescopes*.
- Hartkopf, W. I., B. D. Mason, and C. E. Worley (2001). The 2001 US Naval Observatory Double Star CD-ROM. II. The Fifth Catalog of Orbits of Visual Binary Stars. *AJ*, **122**, pp. 3472–3479. doi:10.1086/323921.
- Heap, S. R., D. J. Lindler, T. M. Lanz, R. H. Cornett, I. Hubeny, S. P. Maran, and B. Woodgate (2000). Space Telescope Imaging Spectrograph Coronagraphic Observations of β Pictoris. *ApJ*, **539**, pp. 435–444. doi:10.1086/309188.
- Herbig, G. H. and R. R. Zappala (1970). Near-Infrared Spectra of NML Cygnt and IRC+10216. *ApJ*, **162**, pp. L15+. doi:10.1086/180613.
- Herwig, F. (2005). Evolution of Asymptotic Giant Branch Stars. *ARA&A*, **43**, pp. 435–479. doi:10.1146/annurev.astro.43.072103.150600.
- Hewett, P. C., S. J. Warren, S. K. Leggett, and S. T. Hodgkin (2006). The UKIRT Infrared Deep Sky Survey ZY JHK photometric system: passbands and synthetic colours. *MNRAS*, **367**, pp. 454–468. doi:10.1111/j.1365-2966.2005.09969.x.
- Hillenbrand, L. A., J. B. Foster, S. E. Persson, and K. Matthews (2002). The Y Band at 1.035 Microns: Photometric Calibration and the DwarfStellar/Substellar Color Sequence. *PASP*, **114**, pp. 708–720. doi:10.1086/341699.
- Hinz, P., J. Codona, O. Guyon, W. Hoffmann, A. Skemer, J. Hora, V. Tolls, A. Boss, A. Weinberger, P. Arbo, T. Connors, O. Durney, T. McMahon, M. Montoya, and V. Vaitheeswaran (2012). TIGER: a high contrast infrared imager for the Giant Magellan Telescope. In *SPIE*, volume 8446 of *SPIE*. doi:10.1117/12.926751.

- Hinz, P. M., J. R. P. Angel, N. J. Woolf, W. F. Hoffmann, and D. W. McCarthy (2000). BLINC: a testbed for nulling interferometry in the thermal infrared. volume 4006 of *Proc. SPIE*, pp. 349–353.
- ISO Handbook, Vol V (2003). *SWS - The Short Wavelength Spectrometer*. European Space Agency., 2 edition.
- Johns, M., P. McCarthy, K. Raybould, A. Bouchez, A. Farahani, J. Filgueira, G. Jacoby, S. Sackett, and M. Sheehan (2012). Giant Magellan Telescope: overview. *Proc. SPIE*, **8444**, 84441H. doi:10.1117/12.926716.
- Johnson, H. L. and R. I. Mitchell (1975). Thirteen-color photometry of 1380 bright stars. *RMxAA*, **1**, pp. 299–324.
- Kalas, P., J. R. Graham, E. Chiang, M. P. Fitzgerald, M. Clampin, E. S. Kite, K. Stapelfeldt, C. Marois, and J. Krist (2008). Optical Images of an Exosolar Planet 25 Light-Years from Earth. *Science*, **322**, pp. 1345–. doi:10.1126/science.1166609.
- Kalas, P. and D. Jewitt (1995). Asymmetries in the Beta Pictoris Dust Disk. *AJ*, **110**, p. 794. doi:10.1086/117565.
- Kasper, M., J.-L. Beuzit, C. Verinaud, R. G. Gratton, F. Kerber, N. Yaitskova, A. Boccaletti, N. Thatte, H. M. Schmid, C. Keller, P. Baudoz, L. Abe, E. Allart, J. Antichi, M. Bonavita, K. Dohlen, E. Fedrigo, H. Hanenburg, N. Hubin, R. Jager, V. Korkiakoski, P. Martinez, D. Mesa, O. Preis, P. Rabou, R. Roelfsema, G. Salter, M. Tecza, and L. Venema (2010). EPICS: direct imaging of exoplanets with the E-ELT. *Proc. SPIE*, **7735**, p. 7735E. doi:10.1117/12.856850.
- Kasting, J. F., D. P. Whitmire, and R. T. Reynolds (1993). Habitable Zones around Main Sequence Stars. *ICARUS*, **101**, pp. 108–128. doi:10.1006/icar.1993.1010.
- Kepler, J. (1609). *Astronomia nova*.
- Kessler, M. F., J. A. Steinz, M. E. Anderegg, J. Clavel, G. Drechsel, P. Estaria, J. Faelker, J. R. Riedinger, A. Robson, B. G. Taylor, and S. Ximénez de Ferrán (1996). The Infrared Space Observatory (ISO) mission. *A&A*, **315**, pp. L27–L31.
- Kirkpatrick, J. D., M. C. Cushing, C. R. Gelino, R. L. Griffith, M. F. Skrutskie, K. A. Marsh, E. L. Wright, A. Mainzer, P. R. Eisenhardt, I. S. McLean, M. A. Thompson, J. M. Bauer, D. J. Benford, C. R. Bridge, S. E. Lake, S. M. Petty, S. A. Stanford, C.-W. Tsai, V. Bailey, C. A. Beichman, J. S. Bloom, J. J. Bochanski, A. J. Burgasser, P. L. Capak, K. L. Cruz, P. M. Hinz, J. S. Kartaltepe, R. P. Knox, S. Manohar, D. Masters, M. Morales-Calderón, L. A. Prato, T. J. Rodigas, M. Salvato, S. D. Schurr, N. Z. Scoville, R. A. Simcoe, K. R. Stapelfeldt,

- D. Stern, N. D. Stock, and W. D. Vacca (2011). The First Hundred Brown Dwarfs Discovered by the Wide-field Infrared Survey Explorer (WISE). *ApJS*, **197**, 19. doi:10.1088/0067-0049/197/2/19.
- Kopparapu, R. K., R. Ramirez, J. F. Kasting, V. Eymet, T. D. Robinson, S. Mahadevan, R. C. Terrien, S. Domagal-Goldman, V. Meadows, and R. Deshpande (2013). Habitable Zones around Main-sequence Stars: New Estimates. *ApJ*, **765**, 131. doi:10.1088/0004-637X/765/2/131.
- Krist, J. (2003). ACS WFC & HRC fielddependent PSF variations due to optical and charge diffusion effects. Technical report.
- Lafrenière, D., R. Doyon, C. Marois, D. Nadeau, B. R. Oppenheimer, P. F. Roche, F. Rigaut, J. R. Graham, R. Jayawardhana, D. Johnstone, P. G. Kalas, B. Macintosh, and R. Racine (2007a). The Gemini Deep Planet Survey. *ApJ*, **670**, pp. 1367–1390. doi:10.1086/522826.
- Lafrenière, D., C. Marois, R. Doyon, D. Nadeau, and É. Artigau (2007b). A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging. *ApJ*, **660**, pp. 770–780. doi:10.1086/513180.
- Lagrange, A.-M., A. Boccaletti, J. Milli, G. Chauvin, M. Bonnefoy, D. Mouillet, J. C. Augereau, J. H. Girard, S. Lacour, and D. Apai (2012). The position of β Pictoris b position relative to the debris disk. *A&A*, **542**, A40. doi:10.1051/0004-6361/201118274.
- Lagrange, A.-M., M. Bonnefoy, G. Chauvin, D. Apai, D. Ehrenreich, A. Boccaletti, D. Gratadour, D. Rouan, D. Mouillet, S. Lacour, and M. Kasper (2010). A Giant Planet Imaged in the Disk of the Young Star β Pictoris. *Science*, **329**, pp. 57–. doi:10.1126/science.1187187.
- Lagrange, A.-M., D. Gratadour, G. Chauvin, T. Fusco, D. Ehrenreich, D. Mouillet, G. Rousset, D. Rouan, F. Allard, É. Gendron, J. Charton, L. Mugnier, P. Rabou, J. Montri, and F. Lacombe (2009). A probable giant planet imaged in the β Pictoris disk. VLT/NaCo deep L'-band imaging. *A&A*, **493**, pp. L21–L25. doi:10.1051/0004-6361:200811325.
- Lagrange-Henri, A. M., A. Vidal-Madjar, and R. Ferlet (1988). The Beta Pictoris circumstellar disk. VI - Evidence for material falling on to the star. *A&A*, **190**, pp. 275–282.
- Law, N. M., R. G. Dekany, C. D. Mackay, A. M. Moore, M. C. Britton, and V. Velur (2008). Getting lucky with adaptive optics: diffraction-limited resolution in the

- visible with current AO systems on large and small telescopes. volume 7015, of *Proc. SPIE*. doi:10.1117/12.788621.
- Law, N. M., C. D. Mackay, and J. E. Baldwin (2006). Lucky imaging: high angular resolution imaging in the visible from the ground. *A&A*, **446**, pp. 739–745. doi:10.1051/0004-6361:20053695.
- Law, N. M., C. D. Mackay, R. G. Dekany, M. Ireland, J. P. Lloyd, A. M. Moore, J. G. Robertson, P. Tuthill, and H. C. Woodruff (2009). Getting Lucky with Adaptive Optics: Fast Adaptive Optics Image Selection in the Visible with a Large Telescope. *APJ*, **692**, pp. 924–930. doi:10.1088/0004-637X/692/1/924.
- Leão, I. C., P. de Laverny, D. Mékarnia, J. R. de Medeiros, and B. Vandame (2006). The circumstellar envelope of IRC+10216 from milli-arcsecond to arcmin scales. *A&A*, **455**, pp. 187–194. doi:10.1051/0004-6361:20054577.
- Le Bertre, T. (1992). Carbon-star lightcurves in the 1-20 micron range. *A&AS*, **94**, pp. 377–398.
- Le Bertre, T. (1997). Optical and infrared observations of 23 carbon-rich stars. Modelling of the circumstellar dust shells. *A&A*, **324**, pp. 1059–1070.
- Lecavelier Des Etangs, A., M. Deleuil, A. Vidal-Madjar, R. Ferlet, C. Nitschelm, B. Nicolet, and A. M. Lagrange-Henri (1995). β Pictoris: evidence of light variations. *A&A*, **299**, p. 557.
- Lecavelier Des Etangs, A., G. Perrin, R. Ferlet, A. Vidal Madjar, F. Colas, C. Buil, F. Sevre, J. E. Arlot, H. Beust, A. M. Lagrange Henri, J. Lecacheux, M. Deleuil, and C. Gry (1993). Observation of the Central Part of the Beta-Pictoris Disk with an Anti-Blooming CCD. *A&A*, **274**, p. 877.
- Lecavelier Des Etangs, A. and A. Vidal-Madjar (2009). Is β Pictoris b the transiting planet of November 1981? *A&A*, **497**, pp. 557–562. doi:10.1051/0004-6361/200811528.
- Leconte, J., R. Soummer, S. Hinkley, B. R. Oppenheimer, A. Sivaramakrishnan, D. Brenner, J. Kuhn, J. P. Lloyd, M. D. Perrin, R. Makidon, L. C. Roberts, Jr., J. R. Graham, M. Simon, R. A. Brown, N. Zimmerman, G. Chabrier, and I. Baraffe (2010). The Lyot Project Direct Imaging Survey of Substellar Companions: Statistical Analysis and Information from Nondetections. *ApJ*, **716**, 1551. doi:10.1088/0004-637X/716/2/1551.
- Lissauer, J. J., D. Ragozzine, D. C. Fabrycky, J. H. Steffen, E. B. Ford, J. M. Jenkins, A. Shporer, M. J. Holman, J. F. Rowe, E. V. Quintana, N. M. Batalha, W. J. Borucki, S. T. Bryson, D. A. Caldwell, J. A. Carter, D. Ciardi, E. W.

- Dunham, J. J. Fortney, T. N. Gautier, III, S. B. Howell, D. G. Koch, D. W. Latham, G. W. Marcy, R. C. Morehead, and D. Sasselov (2011). Architecture and Dynamics of Kepler’s Candidate Multiple Transiting Planet Systems. *ApJS*, **197**, 8. doi:10.1088/0067-0049/197/1/8.
- Liu, M. C., T. J. Dupuy, B. P. Bowler, S. K. Leggett, and W. M. J. Best (2012). Two Extraordinary Substellar Binaries at the T/Y Transition and the Y-band Fluxes of the Coolest Brown Dwarfs. *ApJ*, **758**, 57. doi:10.1088/0004-637X/758/1/57.
- Liu, M. C., Z. Wahhaj, B. A. Biller, E. L. Nielsen, M. Chun, L. M. Close, C. Ftaclas, M. Hartung, T. L. Hayward, F. Clarke, I. N. Reid, E. L. Shkolnik, M. Tecza, N. Thatte, S. Alencar, P. Artymowicz, A. Boss, A. Burrows, E. de Gouveia Dal Pino, J. Gregorio-Hetem, S. Ida, M. J. Kuchner, D. Lin, and D. Toomey (2010). The Gemini NICI Planet-Finding Campaign. *Proc. SPIE*, **7736**, p. 7736E. doi:10.1117/12.858358.
- Lloyd-Hart, M. (2000). Thermal Performance Enhancement of Adaptive Optics by Use of a Deformable Secondary Mirror. *PASP*, **112**, pp. 264–272. doi:10.1086/316503.
- Lord, S. D. (1992). Technical Memorandum 103957. Technical report, NASA.
- Lowrance, P. J., E. E. Becklin, G. Schneider, J. D. Kirkpatrick, A. J. Weinberger, B. Zuckerman, C. Dumas, J.-L. Beuzit, P. Plait, E. Malumuth, S. Heap, R. J. Terile, and D. C. Hines (2005). An Infrared Coronagraphic Survey for Substellar Companions. *AJ*, **130**, pp. 1845–1861. doi:10.1086/432839.
- Macintosh, B. A., A. Anthony, J. Atwood, N. Barriga, B. Bauman, K. Caputa, J. Chilcote, D. Dillon, R. Doyon, J. Dunn, D. T. Gavel, R. Galvez, S. J. Goodsell, J. R. Graham, M. Hartung, J. Isaacs, D. Kerley, Q. Konopacky, K. Labrie, J. E. Larkin, J. Maire, C. Marois, M. Millar-Blanchaer, A. Nunez, B. R. Oppenheimer, D. W. Palmer, J. Pazder, M. Perrin, L. A. Poyneer, C. Quirez, F. Rantakyro, V. Reshtov, L. Saddlemyer, N. Sadakuni, D. Savransky, A. Sivaramakrishnan, M. Smith, R. Soummer, S. Thomas, J. K. Wallace, J. Weiss, and S. Wiktorowicz (2012). The Gemini Planet Imager: integration and status. *Proc. SPIE*, **8446**, 84461U. doi:10.1117/12.926721.
- Mackay, C., A. Basden, and M. Bridgeland (2004). Astronomical imaging with L3CCDs: detector performance and high-speed controller design. volume 5499 of *Proc. SPIE*, pp. 203–209. doi:10.1117/12.550448.
- Madhusudhan, N., A. Burrows, and T. Currie (2011). Model Atmospheres for Massive Gas Giants with Thick Clouds: Application to the HR 8799 Planets and Predictions for Future Detections. *ApJ*, **737**, 34. doi:10.1088/0004-637X/737/1/34.

- Mahajan, V. N. (1986). Uniform versus Gaussian beams: a comparison of the effects of diffraction, obscuration, and aberrations. *JOSAA*, **3**(4), pp. 470–485.
- Males, J. R., L. M. Close, D. Kopon, V. Gasho, and K. Follette (2010). Frame selection techniques for the Magellan adaptive optics VisAO camera. volume 7736 of *Proc. SPIE*. doi:10.1117/12.858055.
- Males, J. R., L. M. Close, D. Kopon, F. Quiros-Pacheco, A. Riccardi, M. Xompero, A. Puglisi, V. Gasho, K. M. Morzinski, and K. B. Follette (2012a). Laboratory demonstration of real time frame selection with Magellan AO. **8447**, 844742. doi:10.1117/12.926698.
- Males, J. R., L. M. Close, A. J. Skemer, P. M. Hinz, W. F. Hoffmann, and M. Marengo (2012b). Four Decades of IRC +10216: Evolution of a Carbon-rich Dust Shell Resolved at 10 μ m with MMT Adaptive Optics and MIRAC4. *ApJ*, **744**, 133. doi:10.1088/0004-637X/744/2/133.
- Males, J. R., A. J. Skemer, and L. M. Close (2013). Direct Imaging in the Habitable Zone and the Problem of Orbital Motion. *ApJ*, **771**, 10. doi:10.1088/0004-637X/771/1/10.
- Marois, C., D. Lafrenière, R. Doyon, B. Macintosh, and D. Nadeau (2006). Angular Differential Imaging: A Powerful High-Contrast Imaging Technique. *ApJ*, **641**, pp. 556–564. doi:10.1086/500401.
- Marois, C., D. Lafrenière, B. Macintosh, and R. Doyon (2008a). Confidence Level and Sensitivity Limits in High-Contrast Imaging. *ApJ*, **673**, pp. 647–656. doi:10.1086/523839.
- Marois, C., B. Macintosh, T. Barman, B. Zuckerman, I. Song, J. Patience, D. Lafrenière, and R. Doyon (2008b). Direct Imaging of Multiple Planets Orbiting the Star HR 8799. *Science*, **322**, pp. 1348–. doi:10.1126/science.1166585.
- Marois, C., B. Zuckerman, Q. M. Konopacky, B. Macintosh, and T. Barman (2010). Images of a fourth planet orbiting HR 8799. *Nature*, **468**, pp. 1080–1083. doi:10.1038/nature09684.
- Mauron, N. and P. J. Huggins (1999). Multiple shells in the circumstellar envelope of IRC+10216. *A&A*, **349**, pp. 203–208.
- Mauron, N. and P. J. Huggins (2000). Multiple shells in IRC+10216: shell properties. *A&A*, **359**, pp. 707–715.
- Mayor, M., X. Bonfils, T. Forveille, X. Delfosse, S. Udry, J.-L. Bertaux, H. Beust, F. Bouchy, C. Lovis, F. Pepe, C. Perrier, D. Queloz, and N. C. Santos (2009). The

- HARPS search for southern extra-solar planets. XVIII. An Earth-mass planet in the GJ 581 planetary system. *A&A*, **507**, pp. 487–494. doi:10.1051/0004-6361/200912172.
- Mayor, M., M. Marmier, C. Lovis, S. Udry, D. Ségransan, F. Pepe, W. Benz, J. . Bertaux, F. Bouchy, X. Dumusque, G. Lo Curto, C. Mordasini, D. Queloz, and N. C. Santos (2011). The HARPS search for southern extra-solar planets XXXIV. Occurrence, mass distribution and orbital properties of super-Earths and Neptune-mass planets. *ArXiv e-prints*.
- McCarthy, D. W., R. Howell, and F. J. Low (1980). Spatial spectra of IRC +10216 from 2.2 to 20 microns deviations from spherical symmetry. *ApJ*, **235**, pp. L27–L31. doi:10.1086/183151.
- Melnick, G. J., D. A. Neufeld, K. E. S. Ford, D. J. Hollenbach, and M. L. N. Ashby (2001). Discovery of water vapour around IRC+10216 as evidence for comets orbiting another star. *Nature*, **412**, pp. 160–163.
- Men'shchikov, A. B., Y. Balega, T. Blöcker, R. Osterbart, and G. Weigelt (2001). Structure and physical properties of the rapidly evolving dusty envelope of jAS-*TROBJ*_{IRC +10 216}/ASTROBJ_{IRC +10 216} reconstructed by detailed two-dimensional radiative transfer modeling. *A&A*, **368**, pp. 497–526. doi:10.1051/0004-6361:20000554.
- Men'shchikov, A. B., K. Hofmann, and G. Weigelt (2002). jASTROBJ_{IRC +10 216}/ASTROBJ_{IRC +10 216} in action: Present episode of intense mass-loss reconstructed by two-dimensional radiative transfer modeling. *A&A*, **392**, pp. 921–929. doi:10.1051/0004-6361:20020954.
- Merrill, K. M. and W. A. Stein (1976). 2 - 14 μ m stellar spectrophotometry II. Stars from the 2 μ m infrared sky survey. *PASP*, **88**, pp. 294–307. doi:10.1086/129946.
- Mitchell, R. I. and H. L. Johnson (1969). Thirteen-color narrow-band photometry of one thousand bright stars. *Communications of the Lunar and Planetary Laboratory*, **8**, pp. 1–49.
- Mohanty, S., R. Jayawardhana, N. Huélamo, and E. Mamajek (2007). The Planetary Mass Companion 2MASS 1207-3932B: Temperature, Mass, and Evidence for an Edge-on Disk. *ApJ*, **657**, pp. 1064–1091. doi:10.1086/510877.
- Monnier, J. D., W. C. Danchi, D. S. Hale, E. A. Lipman, P. G. Tuthill, and C. H. Townes (2000). Mid-Infrared Interferometry on Spectral Lines. II. Continuum (Dust) Emission Around IRC +10216 and VY Canis Majoris. *ApJ*, **543**, pp. 861–867. doi:10.1086/317126.

- Monnier, J. D., T. R. Geballe, and W. C. Danchi (1998). Temporal Variations of Midinfrared Spectra in Late-Type Stars. *ApJ*, **502**, pp. 833–+. doi:10.1086/305945.
- Monnier, J. D., T. R. Geballe, and W. C. Danchi (1999). Mid-Infrared Spectra of Late-Type Stars: Long-Term Evolution. *ApJ*, **521**, pp. 261–270. doi:10.1086/307530.
- Morzinski, K., L. C. Johnson, D. T. Gavel, B. Grigsby, D. Dillon, M. Reinig, and B. A. Macintosh (2010). Performance of MEMS-based visible-light adaptive optics at Lick Observatory: closed- and open-loop control. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7736 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. doi:10.1117/12.857444.
- Mouillet, D., J. D. Larwood, J. C. B. Papaloizou, and A. M. Lagrange (1997). A planet on an inclined orbit as an explanation of the warp in the Beta Pictoris disc. *MNRAS*, **292**, p. 896.
- Murray, C. D. and A. C. M. Correia (2010). *Keplerian Orbits and Dynamics of Exoplanets*, pp. 15–23.
- Oppenheimer, B. R., C. Baranec, C. Beichman, D. Brenner, R. Burruss, E. Cady, J. R. Crepp, R. Dekany, R. Fergus, D. Hale, L. Hillenbrand, S. Hinkley, D. W. Hogg, D. King, E. R. Ligon, T. Lockhart, R. Nilsson, I. R. Parry, L. Pueyo, E. Rice, J. E. Roberts, L. C. Roberts, Jr., M. Shao, A. Sivaramakrishnan, R. Soummer, T. Truong, G. Vasisht, A. Veicht, F. Vescelus, J. K. Wallace, C. Zhai, and N. Zimmerman (2013). Reconnaissance of the HR 8799 Exosolar System. I. Near-infrared Spectroscopy. *ApJ*, **768**, 24. doi:10.1088/0004-637X/768/1/24.
- Oppenheimer, B. R. and S. Hinkley (2009). High-Contrast Observations in Optical and Infrared Astronomy. *ARAA*, **47**, pp. 253–289. doi:10.1146/annurev-astro-082708-101717.
- Osterbart, R., Y. Y. Balega, T. Blöcker, A. B. Men’shchikov, and G. Weigelt (2000). The dynamical evolution of the fragmented, bipolar dust shell around the carbon star IRC +10 216 . Rapid changes of a PPN-like structure? *A&A*, **357**, pp. 169–176.
- Parker, A. H. and J. J. Kavelaars (2010). Pencil-Beam Surveys for Trans-Neptunian Objects: Novel Methods for Optimization and Characterization. *PASP*, **122**, pp. 549–559. doi:10.1086/652424.
- Pickles, A. J. (1998). A Stellar Spectral Flux Library: 1150-25000 Å. *PASP*, **110**, pp. 863–878. doi:10.1086/316197.

- Pilbratt, G. L., J. R. Riedinger, T. Passvogel, G. Crone, D. Doyle, U. Gageur, A. M. Heras, C. Jewell, L. Metcalfe, S. Ott, and M. Schmidt (2010). Herschel Space Observatory. An ESA facility for far-infrared and submillimetre astronomy. *A&A*, **518**, pp. L1+. doi:10.1051/0004-6361/201014759.
- Quanz, S. P., M. R. Meyer, M. A. Kenworthy, J. H. V. Girard, M. Kasper, A.-M. Lagrange, D. Apai, A. Boccaletti, M. Bonnefoy, G. Chauvin, P. M. Hinz, and R. Lenzen (2010). First Results from Very Large Telescope NACO Apodizing Phase Plate: 4 μ m Images of The Exoplanet β Pictoris b. *ApJ*, **722**, pp. L49–L53. doi:10.1088/2041-8205/722/1/L49.
- Racine, R. (1996). The Telescope Point Spread Function. *PASP*, **108**, pp. 699–+. doi:10.1086/133788.
- Racine, R., G. A. H. Walker, D. Nadeau, R. Doyon, and C. Marois (1999). Speckle Noise and the Detection of Faint Companions. *PASP*, **111**, pp. 587–594. doi:10.1086/316367.
- Roberts, L. C., Jr. and C. R. Neyman (2002). Characterization of the AEOS Adaptive Optics System. *PASP*, **114**, pp. 1260–1266. doi:10.1086/343221.
- Roddier, F. (1999). Adaptive Optics in Astronomy. chapter 2: Imaging through the atmosphere, p. 9. Cambridge University Press.
- Rodgers, C. T., R. Canterna, J. A. Smith, M. J. Pierce, and D. L. Tucker (2006). Improved u'g'r'i'z' to UBVR_CI_C Transformation Equations for Main-Sequence Stars. *AJ*, **132**, pp. 989–993. doi:10.1086/505864.
- Roelfsema, R., H. M. Schmid, J. Pragt, D. Gisler, R. Waters, A. Bazzon, A. Baruffolo, J.-L. Beuzit, A. Boccaletti, J. Charton, C. Cumani, K. Dohlen, M. Downing, E. Elswijk, M. Feldt, C. Groothuis, M. de Haan, H. Hanenburg, N. Hubin, F. Joos, M. Kasper, C. Keller, J. Kragt, J.-L. Lizon, D. Mouillet, A. Pavlov, F. Rigal, S. Rochat, B. Salasnich, P. Steiner, C. Thalmann, L. Venema, and F. Wildi (2010). The ZIMPOL high-contrast imaging polarimeter for SPHERE: design, manufacturing, and testing. *Proc. SPIE*, **7735**. doi:10.1117/12.857045.
- Sandler, D. G., S. Stahl, J. R. P. Angel, M. Lloyd-Hart, and D. McCarthy (1994). Adaptive optics for diffraction-limited infrared imaging with 8-m telescopes. *JOSAA*, **11**, pp. 925–945.
- Skemer, A. J. and L. M. Close (2011). Sirius B Imaged in the Mid-infrared: No Evidence for a Remnant Planetary System. *ApJ*, **730**, 53. doi:10.1088/0004-637X/730/1/53.

- Skemer, A. J., L. M. Close, P. M. Hinz, W. F. Hoffmann, T. P. Greene, J. R. Males, and T. L. Beck (2010). ISM Dust Grains and N-band Spectral Variability in the Spatially Resolved Subarcsecond Binary UY Aur. *APJ*, **711**, pp. 1280–1290. doi:10.1088/0004-637X/711/2/1280.
- Skemer, A. J., L. M. Close, P. M. Hinz, W. F. Hoffmann, M. A. Kenworthy, and D. L. Miller (2008). Evidence for Misaligned Disks in the T Tauri Triple System: 10 μ m Superresolution with MMTAO and Markov Chains. *APJ*, **676**, pp. 1082–1087. doi:10.1086/527555.
- Skemer, A. J., P. M. Hinz, W. F. Hoffmann, L. M. Close, S. Kendrew, R. J. Mathar, R. Stuik, T. P. Greene, C. E. Woodward, and M. S. Kelley (2009). A Direct Measurement of Atmospheric Dispersion in N-band Spectra: Implications for Mid-IR Systems on ELTs. *PASP*, **121**, pp. 897–904. doi:10.1086/605312.
- Skinner, C. J., M. Meixner, and M. Bobrowsky (1998). The birth of a planetary nebula around the carbon star IRC+10216. *MNRAS*, **300**, pp. L29–L33. doi:10.1046/j.1365-8711.1998.300004129.x.
- Sloan, G. C., I. R. Little-Marenin, and S. D. Price (1998). The carbon-rich dust sequence - Infrared spectral classification of carbon stars. *AJ*, **115**, pp. 809–+. doi:10.1086/300205.
- Smith, B. A. and R. J. Terrile (1984). A circumstellar disk around Beta Pictoris. *Science*, **226**, pp. 1421–1424. doi:10.1126/science.226.4681.1421.
- Smith, J. A., D. L. Tucker, S. Kent, M. W. Richmond, M. Fukugita, T. Ichikawa, S.-i. Ichikawa, A. M. Jorgensen, A. Uomoto, J. E. Gunn, M. Hamabe, M. Watanabe, A. Tolea, A. Henden, J. Annis, J. R. Pier, T. A. McKay, J. Brinkmann, B. Chen, J. Holtzman, K. Shimasaku, and D. G. York (2002). The u'g'r'i'z' Standard-Star System. *AJ*, **123**, pp. 2121–2144. doi:10.1086/339311.
- Song, I., G. Schneider, B. Zuckerman, J. Farihi, E. E. Becklin, M. S. Bessell, P. Lowrance, and B. A. Macintosh (2006). HST NICMOS Imaging of the Planetary-mass Companion to the Young Brown Dwarf 2MASSW J1207334-393254. *ApJ*, **652**, pp. 724–729. doi:10.1086/507831.
- Soummer, R., L. Pueyo, and J. Larkin (2012). Detection and Characterization of Exoplanets and Disks Using Projections on Karhunen-Loève Eigenimages. *ApJ*, **755**, L28. doi:10.1088/2041-8205/755/2/L28.
- Speck, A. K., G. D. Thompson, and A. M. Hofmeister (2005). The Effect of Stellar Evolution on SiC Dust Grain Sizes. *ApJ*, **634**, pp. 426–435. doi:10.1086/496955.

- Stephens, D. C. and S. K. Leggett (2004). JHK Magnitudes for L and T Dwarfs and Infrared Photometric Systems. *PASP*, **116**, pp. 9–21. doi:10.1086/381135.
- Stephens, D. C., S. K. Leggett, M. C. Cushing, M. S. Marley, D. Saumon, T. R. Geballe, D. A. Golimowski, X. Fan, and K. S. Noll (2009). The 0.8–14.5 μm Spectra of Mid-L to Mid-T Dwarfs: Diagnostics of Effective Temperature, Grain Sedimentation, Gas Transport, and Surface Gravity. *ApJ*, **702**, pp. 154–170. doi:10.1088/0004-637X/702/1/154.
- Strecker, D. W. and E. P. Ney (1974). 0.9–18-micron photometry of the 14 CIT objects. *AJ*, **79**, pp. 1410–1415. doi:10.1086/111693.
- Thomas-Osip, J. E., G. Prieto, M. Johns, and M. M. Phillips (2008). Giant Magellan Telescope site evaluation and characterization at Las Campanas Observatory. volume 7012, of *Proc. SPIE*. doi:10.1117/12.789863.
- Thompson, G. D., A. B. Corman, A. K. Speck, and C. Dijkstra (2006). Challenging the Carbon Star Dust Condensation Sequence: Anarchist C Stars. *ApJ*, **652**, pp. 1654–1673. doi:10.1086/508417.
- Tokovinin, A. (2002). From Differential Image Motion to Seeing. *PASP*, **114**, pp. 1156–1166. doi:10.1086/342683.
- Traub, W. A. (2012). Terrestrial, Habitable-zone Exoplanet Frequency from Kepler. *ApJ*, **745**, 20. doi:10.1088/0004-637X/745/1/20.
- Traub, W. A. and B. R. Oppenheimer (2011). *Direct Imaging of Exoplanets*, pp. 111–156.
- Treffers, R. and M. Cohen (1974). High-resolution spectra of cool stars in the 10- and 20-micron regions. *ApJ*, **188**, pp. 545–552. doi:10.1086/152746.
- Tuthill, P. G., J. D. Monnier, W. C. Danchi, and B. Lopez (2000). Smoke Signals from IRC +10216. I. Milliarcsecond Proper Motions of the Dust. *ApJ*, **543**, pp. 284–290. doi:10.1086/317082.
- Tyson, R. K. (2011). *Principles of adaptive optics*.
- Vigan, A., J. Patience, C. Marois, M. Bonavita, R. J. De Rosa, B. Macintosh, I. Song, R. Doyon, B. Zuckerman, D. Lafrenière, and T. Barman (2012). The International Deep Planet Survey. I. The frequency of wide-orbit massive planets around A-stars. *A&A*, **544**, A9. doi:10.1051/0004-6361/201218991.
- Vogt, S. S., R. P. Butler, and N. Haghighipour (2012). GJ 581 update: Additional Evidence for a Super-Earth in the Habitable Zone. *ArXiv e-prints*.

- von Braun, K., T. S. Boyajian, S. R. Kane, G. T. van Belle, D. R. Ciardi, M. López-Morales, H. A. McAlister, T. J. Henry, W.-C. Jao, A. R. Riedel, J. P. Subasavage, G. Schaefer, T. A. ten Brummelaar, S. Ridgway, L. Sturmann, J. Sturmann, J. Mazingue, N. H. Turner, C. Farrington, P. J. Goldfinger, and A. F. Boden (2011). Astrophysical Parameters and Habitable Zone of the Exoplanet Hosting Star GJ 581. *ApJ*, **729**, L26. doi:10.1088/2041-8205/729/2/L26.
- Weigelt, G., Y. Balega, T. Bloeker, A. J. Fleischer, R. Osterbart, and J. M. Winters (1998). 76mas speckle-masking interferometry of IRC+10216 with the SAO 6m telescope: Evidence for a clumpy shell structure. *A&A*, **333**, pp. L51–L54.
- Whaley, R. C. and A. Petitet (2005). Minimizing development and maintenance costs in supporting persistently optimized BLAS. *Software: Practice and Experience*, **35**(2), pp. 101–121.
- Wildi, F. P., G. Brusa, M. Lloyd-Hart, L. M. Close, and A. Riccardi (2003). First light of the 6.5-m MMT adaptive optics system. volume 5169 of *Proc. SPIE*, pp. 17–25. doi:10.1117/12.507687.
- Wittenmyer, R. A., C. G. Tinney, R. P. Butler, S. J. O’Toole, H. R. A. Jones, B. D. Carter, J. Bailey, and J. Horner (2011a). The Frequency of Low-mass Exoplanets. III. Toward η at Short Periods. *ApJ*, **738**, 81. doi:10.1088/0004-637X/738/1/81.
- Wittenmyer, R. A., C. G. Tinney, S. J. O’Toole, H. R. A. Jones, R. P. Butler, B. D. Carter, and J. Bailey (2011b). On the Frequency of Jupiter Analogs. *ApJ*, **727**, 102. doi:10.1088/0004-637X/727/2/102.
- Yamamoto, N., D. Kinoshita, T. Fuse, J.-I. Watanabe, and K. Kawabata (2008). A Deep Sky Survey of Edgeworth Kuiper Belt Objects with an Improved Shift-and-Add Method. *PASJ*, **60**, pp. 285–.
- Zsom, A., S. Seager, and J. de Wit (2013). Towards the Minimum Inner Edge Distance of the Habitable Zone. *ArXiv e-prints*.