

PLANETS AROUND SOLAR-TYPE STARS:  
METHODS FOR DETECTION AND CONSTRAINTS ON THEIR  
DISTRIBUTION  
FROM AN L' AND M BAND ADAPTIVE OPTICS IMAGING SURVEY

by  
Aren Nathaniel Heinze

---

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

2007

THE UNIVERSITY OF ARIZONA  
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Aren Nathaniel Heinze entitled "Planets Around Solar-Type Stars: Methods for Detection and Constraints on their Distribution from an L' and M Band Adaptive Optics Imaging Survey" and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

\_\_\_\_\_ Date: 2 August 2007  
Phil Hinz

\_\_\_\_\_ Date: 2 August 2007  
Michael Meyer

\_\_\_\_\_ Date: 2 August 2007  
Adam Burrows

\_\_\_\_\_ Date: 2 August 2007  
Laird Close

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.

\_\_\_\_\_ Date: 2 August 2007  
Dissertation Director: Phil Hinz

## 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: Aren Nathaniel Heinze

## ACKNOWLEDGMENTS

My first and ultimate thanks, now and forever, goes to my Creator and Redeemer, the triune God. He made the cosmos of which I am both a part and an observer for his glory, but also for my joy. He healed my heart and set me free to stand in awe of him and his works.

He has also blessed me with many opportunities which came through people – people who also deserve my heartfelt thanks.

My parents, who have always loved me and encouraged me. My Dad, who passed on to me his infinite curiosity, and who tried to answer every question I ever asked, even if it meant explaining calculus to an eleven year old. He also told me he thought I could make it through Caltech, when I didn't think so — and that started it all. My Mom, whose joy in my accomplishments and confidence in what I could do have always encouraged me. You're also one of the two bravest women I know.

My wife, who has encouraged me, believed I could do it, and sacrificed so that this dissertation and the degree that goes with it would become reality. You have not been dismayed when I have been, and your faith and good sense have blessed me and guarded me from more than I know or like to think.

My advisor, Phil Hinz, who is a great scientist and a good man. You have prodded me when I needed prodding, and been encouraging and forgiving the whole way.

Many other scientists and technicians with whom I feel very privileged to have been able to work: The designers and operators of the MMT's amazing AO system: Doug Miller, Matt Kenworthy, Phil Hinz (again), Vidhya Vaitheeswaran, Thomas Stalcup, Guido Brusa, and others with whom I worked less closely. My committee members: Michael Meyer, who jumped on every planet candidate with an eagerness to rival my own, and found information on them so quickly I still wonder how he did it. Adam Burrows, whose models of extrasolar planets have guided my research, and Laird Close, pioneer in AO science and planet detection. The other designers and builders of the Clio instrument: Phil Hinz (again), who made it happen; Melanie Freed, who worked on the early design and acquired and carefully documented much of the hardware; Suresh Sivanandam, who worked long hours with the hardware, saved the dichroic from an untimely death, and helped me brainstorm observing strategies and improve my image processing code; and Andy Breuninger, who met the formidable challenge of writing code to run the detector, and turned it into an efficient science instrument. Finally, the Telescope Operators of the MMT: John McAfee, Ale Milone, and Mike Alegria, surely members of one of the most exclusive skilled groups in the world, and the many other members of the MMT staff whose names, alas, I never thoroughly learned, but without whom that finicky, awesomely powerful telescope would never have done what I asked of it.

## DEDICATION

This dissertation is dedicated first of all to my Lord, the Creator of all the glory and wonder I have studied.

It is dedicated, secondly, to those I hope it will serve: the starwatchers who will come after me. They, like me, will spend lonely nights on mountains or with computers, using the technology and math developed by centuries of genius – standing on the shoulders of giants – trying to tease out the great secrets of the stars of God. They will carry on the great adventure in which I have had only a tiny part.

## TABLE OF CONTENTS

LIST OF FIGURES . . . . .	10
LIST OF TABLES . . . . .	17
ABSTRACT . . . . .	21
1 INTRODUCTION . . . . .	23
1.1 Poems from the Telescope . . . . .	23
1.2 Science Introduction . . . . .	29
2 DEEP $L'$ AND $M$ -BAND IMAGING FOR PLANETS AROUND VEGA AND $\epsilon$ ERIDANI . . . . .	37
2.1 Introduction . . . . .	37
2.2 Observations and Data Analysis . . . . .	41
2.3 Sensitivity Measurements and Source Detection Tests . . . . .	45
2.3.1 Calculating and Verifying our Sensitivity Results . . . . .	45
2.3.2 Final Sensitivity Results . . . . .	56
2.4 Vega: Comparison with Other Studies, and Upper Limits for Hy- pothetical Planets . . . . .	82
2.4.1 Comparing Our Sensitivity with Other Studies . . . . .	82
2.4.2 Upper Limits at the Locations of Hypothetical Planets . . . . .	87
2.5 $\epsilon$ Eri: Comparison with Other Studies, and Upper Limits for Hy- pothetical Planets . . . . .	90
2.5.1 Comparing Our Sensitivity with Other Studies . . . . .	90
2.5.2 Upper Limits at the Locations of Hypothetical Planets . . . . .	95
2.6 Conclusions . . . . .	100
3 THE SURVEY: DATA ACQUISITION, DATA QUALITY, CONFIRMED AND SUS- PECTED FAINT SOURCES . . . . .	104
3.1 Introduction . . . . .	104
3.2 Motivation for an $L'$ and $M$ Band Survey . . . . .	108
3.3 Our Survey Sample . . . . .	116
3.4 The Clio Instrument . . . . .	132
3.5 Observations . . . . .	136
3.6 Data Processing . . . . .	151
3.7 Sensitivity Analysis . . . . .	174
3.7.1 Three Sensitivity Estimators . . . . .	174
3.7.2 Our Choice of Estimator . . . . .	177
3.7.3 Testing the Three Estimators . . . . .	180
3.7.4 Blind Sensitivity Tests . . . . .	182
3.7.5 One Caveat to Our Sensitivity Estimates . . . . .	204

TABLE OF CONTENTS — <i>Continued</i>	
3.7.6	Sensitivity Estimators of Other Surveys . . . . . 210
3.8	Sensitivity Obtained Around Each Target . . . . . 217
3.9	Sensitivities Compared to Other Surveys . . . . . 278
3.10	Confirmed and Suspected Sources . . . . . 305
3.10.1	Our Source Detection Methods . . . . . 305
3.10.2	Confirmed and Suspected Sources . . . . . 310
3.11	Conclusion . . . . . 339
4	CONSTRAINTS ON PLANET ABUNDANCE FROM THE CLIO SURVEY OF 50 STARS . . . . . 344
4.1	Introduction . . . . . 344
4.2	The Survey Sample . . . . . 351
4.3	Statistical Distributions from RV Planets . . . . . 353
4.4	Theoretical Models of Giant Planets . . . . . 361
4.4.1	Two Sets of Models . . . . . 361
4.4.2	Absolute Magnitudes from the Burrows et al. (2003) Models 361
4.4.3	Interpolating Between Model Points . . . . . 363
4.4.4	Comparing the two Model Sets . . . . . 374
4.5	Introducing the Monte Carlo Simulations . . . . . 383
4.6	Monte-Carlo Simulations: A Detailed Look . . . . . 386
4.6.1	Constructing the Simulations . . . . . 386
4.6.2	Simulation Results . . . . . 390
4.7	Monte-Carlo Simulations: The Larger Picture . . . . . 415
4.7.1	Excluded Regions of Parameter Space . . . . . 415
4.7.2	Comparing Our Results to Other Surveys . . . . . 429
4.8	Conclusion . . . . . 436
5	WHERE FROM HERE? LESSONS LEARNED FROM THE CLIO SURVEY, AND THE FUTURE POTENTIAL OF THE $L'$ AND $M$ BANDS . . . . . 439
5.1	Introduction . . . . . 439
5.2	Clio Data Quality Under Various Conditions . . . . . 440
5.3	Comparing the Future Potential of Long and Short Wavelength Regimes . . . . . 449
5.4	Future Observing Plans . . . . . 461
A	THE ALGORITHMS USED BY THE LEGOLAS04 PROCESSING PACKAGE FOR CLIO IMAGES . . . . . 465
A.1	The Clio Camera, its Capabilities and Limitations . . . . . 466
A.2	Clio Data . . . . . 467
A.3	Overview of the Processing Steps for Clio Images . . . . . 474
A.4	Constructing the Nod-subtraction Pairings . . . . . 490

TABLE OF CONTENTS — <i>Continued</i>	
A.5	Dark Subtraction and Flatfielding . . . . . 496
A.6	Bad Pixel Fixing . . . . . 498
A.7	Nod Subtraction . . . . . 501
A.8	The Colfudge Algorithm . . . . . 502
A.8.1	History . . . . . 503
A.8.2	Colfudge . . . . . 504
A.8.3	Matt Kenworthy’s Algorithm . . . . . 507
A.8.4	Examples Images from Both Algorithms . . . . . 508
A.8.5	Dimming of Faint Sources By Colfudge . . . . . 518
A.9	Rotation and Shifting . . . . . 521
A.9.1	Calculation the Required Rotation . . . . . 521
A.9.2	Calculation the Required Shift . . . . . 523
A.9.3	Zero-Trim, Zero-Pad, Shift, and Rotate: A Single Operation 524
A.10	Image Stacking and the Creeping Mean . . . . . 526
A.11	Unsharp Masking . . . . . 530
A.12	Specialty Processing . . . . . 531
A.12.1	PSF Subtraction: The ‘b’ Reduction Method . . . . . 531
A.12.2	PSF Subtraction: The ‘c’ Reduction Method . . . . . 537
A.12.3	Pre-Stack Unsharp Masking: The ‘d’ Reduction Method . . 538
A.12.4	Combining Methods: The ‘e’ Reduction Method . . . . . 540
A.12.5	Removing Nod Artifacts: The ‘x’ Reduction Method . . . . . 542
A.12.6	Combining Methods: The ‘y’ Reduction Method . . . . . 544
A.12.7	Binary Star PSF Subtraction: the ‘f’ Reduction Method . . . 544
A.12.8	Combining Methods: the ‘g’ Reduction Method . . . . . 547
A.13	Conclusion . . . . . 548
<b>B</b>	<b>A MANUAL FOR THE LEGOLAS04 CLIO IMAGE PROCESSING PIPELINE . 550</b>
B.1	Disclaimer . . . . . 550
B.2	Taking the Data . . . . . 550
B.2.1	Planet Detection Data . . . . . 550
B.2.2	Monitoring Data Quality . . . . . 554
B.2.3	Darks . . . . . 555
B.2.4	Flats . . . . . 556
B.3	Introduction to Data Reduction . . . . . 556
B.4	Data Reduction: Darks . . . . . 557
B.5	Flats . . . . . 562
B.6	Overview of a legolas04 Run . . . . . 566
B.7	Making the Parameter Script . . . . . 572
B.8	Making the Image Script . . . . . 592
B.9	Running legolas04 (At Last!) . . . . . 604
B.10	Output of legolas04: Processed Images . . . . . 606

TABLE OF CONTENTS — <i>Continued</i>	
B.10.1 Code-Generated Files: Image Lists . . . . .	607
B.10.2 Post-legolas processing: cmcsets programs . . . . .	608
B.11 Post-legolas Processing: ‘f’ and ‘g’ Method Reductions . . . . .	617
B.12 Output of legolas04: Image Data Files . . . . .	626
B.13 Post-legolas processing: Source Detection . . . . .	629
B.14 Post-legolas processing: Sensitivity Estimation . . . . .	638
B.14.1 erradplot01 . . . . .	638
B.14.2 percentileplot01 . . . . .	654
B.14.3 Additional Auxiliary Programs . . . . .	661
REFERENCES . . . . .	666

## LIST OF FIGURES

2.1	Sensitivity contour map for the Vega $L'$ observations . . . . .	57
2.2	Sensitivity contour map for the Vega $M$ band observations . . . . .	58
2.3	Sensitivity contour map for our $\epsilon$ Eri $L'$ band observations . . . . .	59
2.4	Sensitivity contour map for our $\epsilon$ Eri $M$ band observations . . . . .	60
2.5	Sensitivity of the Vega $L'$ Observations in Magnitudes . . . . .	64
2.6	Sensitivity of the Vega $L'$ observations in MJ (using the Burrows et al. (2003) models) . . . . .	65
2.7	Sensitivity of the Vega $L'$ observations in MJ (using the Baraffe et al. (2003) models) . . . . .	66
2.8	Sensitivity of the Vega $L'$ observations in MJ (comparing the Burrows et al. (2003) and Baraffe et al. (2003) models) . . . . .	67
2.9	Sensitivity of the Vega $M$ band Observations in Magnitudes . . . . .	68
2.10	Sensitivity of the Vega $M$ band observations in MJ (using the Burrows et al. (2003) models) . . . . .	69
2.11	Sensitivity of the Vega $M$ band observations in MJ (using the Baraffe et al. (2003) models) . . . . .	70
2.12	Sensitivity of the Vega $M$ band observations in MJ (comparing the Burrows et al. (2003) and Baraffe et al. (2003) models) . . . . .	71
2.13	Comparison between the the Vega $L'$ and $M$ band observations, using the Burrows et al. (2003) models . . . . .	72
2.14	Comparison between the the Vega $L'$ and $M$ band observations, using the Baraffe et al. (2003) models . . . . .	73
2.15	Sensitivity of the $\epsilon$ Eri $L'$ Observations in Magnitudes . . . . .	74
2.16	Sensitivity of the $\epsilon$ Eri $L'$ observations in MJ (using the Burrows et al. (2003) models) . . . . .	75
2.17	Sensitivity of the $\epsilon$ Eri $L'$ observations in MJ (using the Baraffe et al. (2003) models) . . . . .	76
2.18	Sensitivity of the $\epsilon$ Eri $L'$ observations in MJ (comparing the Burrows et al. (2003) and Baraffe et al. (2003) models) . . . . .	77
2.19	Sensitivity of the $\epsilon$ Eri $M$ band Observations in Magnitudes . . . . .	78
2.20	Sensitivity of the $\epsilon$ Eri $M$ band observations in MJ (using the Burrows et al. (2003) models) . . . . .	79
2.21	Comparison between the the $\epsilon$ Eri $L'$ and $M$ band observations, using the Burrows et al. (2003) models . . . . .	80
2.22	Comparison between the the $\epsilon$ Eri $L'$ and $M$ band observations, using the Baraffe et al. (2003) models . . . . .	81
2.23	Comparison of sensitivities obtained around Vega with different techniques (using the Burrows et al. (2003) models) . . . . .	85

LIST OF FIGURES — *Continued*

2.24	Comparison of sensitivities obtained around Vega with different techniques (using the Baraffe et al. (2003) models) . . . . .	86
2.25	Comparison of sensitivities obtained around $\epsilon$ Eri with different techniques (using the Burrows et al. (2003) models) . . . . .	93
2.26	Comparison of sensitivities obtained around $\epsilon$ Eri with different techniques (using the Baraffe et al. (2003) models) . . . . .	94
3.1	Stellar Magnitude Distributions for Various Surveys . . . . .	133
3.2	Stellar Distance Distributions for Various Surveys . . . . .	134
3.3	Stellar Age Distributions for Various Surveys . . . . .	135
3.4	Nod Image, Position 1 . . . . .	140
3.5	Nod Image, Position 2 . . . . .	141
3.6	Difference of the 2 Nod Positions . . . . .	142
3.7	Comparison of the Creeping Mean with a Median/Sigma Clip. . . . .	158
3.8	$L'$ Image of $\epsilon$ Eri, Processed with Method 'a' . . . . .	165
3.9	$L'$ Image of $\epsilon$ Eri, Processed with Method 'd' . . . . .	166
3.10	$L'$ Image of $\epsilon$ Eri, Processed with Method 'x' . . . . .	167
3.11	$L'$ Image of $\epsilon$ Eri, Processed with Method 'y' . . . . .	168
3.12	$L'$ Image of $\epsilon$ Eri, Processed with Method 'a', Closer View . . . . .	169
3.13	$L'$ Image of $\epsilon$ Eri, Processed with Method 'b' . . . . .	170
3.14	$L'$ Image of $\epsilon$ Eri, Processed with Method 'e' . . . . .	171
3.15	$M$ Band Image of $\xi$ Boo AB, Processed with Method 'a' . . . . .	172
3.16	$M$ Band Image of $\xi$ Boo AB, Processed with Method 'f' . . . . .	173
3.17	Comparing Sensitivity Methods Using GJ 820A . . . . .	185
3.18	Comparing Sensitivity Methods Using GJ 625 . . . . .	186
3.19	Image from Blind Sensitivity Test on HD 29391 . . . . .	201
3.20	Sensitivity Contours for GJ 5 . . . . .	219
3.21	Sensitivity Contours for HD 1405 . . . . .	220
3.22	Sensitivity Contours for $\tau$ Ceti . . . . .	221
3.23	Sensitivity Contours for GJ 117 . . . . .	222
3.24	Sensitivity Contours for $\epsilon$ Eri . . . . .	223
3.25	Sensitivity Contours for $\epsilon$ Eri ( $M$ band) . . . . .	224
3.26	Sensitivity Contours for GJ 159 . . . . .	225
3.27	Sensitivity Contours for GJ 166BC . . . . .	226
3.28	Sensitivity Contours for HD 29391 . . . . .	227
3.29	Sensitivity Contours for GJ 211 . . . . .	228
3.30	Sensitivity Contours for GJ 216A . . . . .	229
3.31	Sensitivity Contours for BD+20 1790 . . . . .	230
3.32	Sensitivity Contours for GJ 278C . . . . .	231
3.33	Sensitivity Contours for GJ282A . . . . .	232
3.34	Sensitivity Contours for GJ 311 . . . . .	233

LIST OF FIGURES — *Continued*

3.35	Sensitivity Contours for HD 77407AB . . . . .	234
3.36	Sensitivity Contours for HD 78141 . . . . .	235
3.37	Sensitivity Contours for GJ 349 . . . . .	236
3.38	Sensitivity Contours for GJ 355 . . . . .	237
3.39	Sensitivity Contours for GJ 354.1A . . . . .	238
3.40	Sensitivity Contours for GJ 380 . . . . .	239
3.41	Sensitivity Contours for GJ 410 . . . . .	240
3.42	Sensitivity Contours for HD 96064AB . . . . .	241
3.43	Sensitivity Contours for GJ 450 . . . . .	242
3.44	Sensitivity Contours for BD+60 1417 . . . . .	243
3.45	Sensitivity Contours for HD 113449 . . . . .	244
3.46	Sensitivity Contours for GJ 505AB . . . . .	245
3.47	Sensitivity Contours for GJ 519 . . . . .	246
3.48	Sensitivity Contours for GJ 3860 . . . . .	247
3.49	Sensitivity Contours for GJ 564 . . . . .	248
3.50	Sensitivity Contours for HD 133002 . . . . .	249
3.51	Sensitivity Contours for $\xi$ Boo AB . . . . .	250
3.52	Sensitivity Contours for $\xi$ Boo AB ( <i>M</i> band) . . . . .	251
3.53	Sensitivity Contours for HD 139813 . . . . .	252
3.54	Sensitivity Contours for GJ 625 . . . . .	253
3.55	Sensitivity Contours for GJ 659A . . . . .	254
3.56	Sensitivity Contours for GJ 659B . . . . .	255
3.57	Sensitivity Contours for GJ 702AB . . . . .	256
3.58	Sensitivity Contours for 61 Cyg A . . . . .	257
3.59	Sensitivity Contours for 61 Cyg A ( <i>M</i> band) . . . . .	258
3.60	Sensitivity Contours for 61 Cyg B . . . . .	259
3.61	Sensitivity Contours for BD+48 3686 . . . . .	260
3.62	Sensitivity Contours for GJ 879 . . . . .	261
3.63	Sensitivity Contours for HD 220140AB . . . . .	262
3.64	Sensitivity Contours for GJ 896AB . . . . .	263
3.65	Min. detectable planet mass vs separation in AU for survey targets.	264
3.66	Min. detectable planet mass vs separation in AU for survey targets.	265
3.67	Min. detectable planet mass vs separation in AU for survey targets.	266
3.68	Min. detectable planet mass vs separation in AU for survey targets.	267
3.69	Min. detectable planet mass vs separation in AU for survey targets.	268
3.70	Min. detectable planet mass vs separation in AU for survey targets.	269
3.71	Min. detectable planet mass vs separation in AU for survey targets.	270
3.72	Min. detectable planet mass vs separation in AU for survey targets.	271
3.73	Min. detectable planet mass vs separation in AU for survey targets.	272
3.74	Min. detectable planet mass vs separation in AU for survey targets.	273
3.75	Min. detectable planet mass vs separation in AU for survey targets.	274

LIST OF FIGURES — *Continued*

3.76	Min. detectable planet mass vs separation in AU for survey targets.	275
3.77	Min. detectable planet mass vs separation in AU for survey targets.	276
3.78	Min. detectable planet mass vs separation in AU for survey targets.	277
3.79	$\Delta$ -Mag vs Separation Plot for HD 96064 in Various Surveys . . . . .	282
3.80	$\Delta$ -Mag vs Separation Plot for $\epsilon$ Eri in Various Surveys . . . . .	283
3.81	Min. Detectable Planet Mass vs Separation Plot for HD 96064 in Various Surveys (using the Burrows et al. (2003) models) . . . . .	284
3.82	Min. Detectable Planet Mass vs Separation Plot for HD 96064 in Various Surveys (using Baraffe et al. (2003) models) . . . . .	285
3.83	Min. Detectable Planet Mass vs Separation Plot for $\epsilon$ Eri in Various Surveys (using the Burrows et al. (2003) models) . . . . .	288
3.84	Min. Detectable Planet Mass vs Separation Plot for $\epsilon$ Eri in Various Surveys (using the Baraffe et al. (2003) models) . . . . .	289
3.85	Min. Detectable Planet Mass vs Separation Plot for GJ 117 in Various Surveys (using the Burrows et al. (2003) models) . . . . .	290
3.86	Min. Detectable Planet Mass vs Separation Plot for GJ 117 in Various Surveys (using the Baraffe et al. (2003) models) . . . . .	292
3.87	Min. Detectable Planet Mass vs Separation Plot for 61 Cyg A in Various Surveys (using the Burrows et al. (2003) models) . . . . .	293
3.88	Min. Detectable Planet Mass vs Separation Plot for 61 Cyg A in Various Surveys (using the Baraffe et al. (2003) models) . . . . .	294
3.89	Min. Detectable Planet Mass vs Separation Plot for 61 Cyg A in Various Surveys with our $\epsilon$ Eri mag. sens. (using the Burrows et al. (2003) models) . . . . .	296
3.90	Min. Detectable Planet Mass vs Separation Plot for 61 Cyg A in Various Surveys with our $\epsilon$ Eri mag. sens. (using the Baraffe et al. (2003) models) . . . . .	297
3.91	Min. Detectable Planet Mass vs Separation Plot for GJ 450 in Various Surveys (using the Burrows et al. (2003) models) . . . . .	299
3.92	Min. Detectable Planet Mass vs Separation Plot for GJ 450 in Various Surveys (using the Baraffe et al. (2003) models) . . . . .	300
3.93	Image of BD+20 1790 showing the detected companions . . . . .	313
3.94	Image of GJ 354.1A showing the detected companion . . . . .	314
3.95	Image of HD 96064 showing the detected companion . . . . .	315
3.96	Image of BD+60 1417 showing the detected companion . . . . .	316
3.97	Image of GJ 3860 showing the detected companion . . . . .	317
3.98	Image of GJ 564 showing the detected companion . . . . .	318
3.99	Image of HD133002 showing the detected companion . . . . .	319
3.100	Image of 61 Cyg A showing the detected companions . . . . .	320
3.101	Image of 61 Cyg B showing the detected companion . . . . .	321
3.102	POSS image of 61 Cyg system . . . . .	322

LIST OF FIGURES — *Continued*

4.1	Histogram of RV Planet Masses with Power Law Overlaid . . . . .	358
4.2	Histogram of RV Planet Semimajor Axes with Power Law Overlaid . . . . .	359
4.3	Histogram of RV Planet Eccentricities with Model Overlaid . . . . .	360
4.4	Clio $L'$ and M-band Filter Transmission Curves . . . . .	364
4.5	$L'$ Magnitudes vs Age from Burrows et al. (2003) . . . . .	375
4.6	$M$ band Magnitudes vs Age from Burrows et al. (2003) . . . . .	376
4.7	$L'$ Magnitudes vs Age from Baraffe et al. (2003) . . . . .	377
4.8	$M$ band Magnitudes vs Age from Baraffe et al. (2003) . . . . .	378
4.9	$L'$ Magnitudes vs Mass from Burrows et al. (2003) . . . . .	379
4.10	$M$ Band Magnitudes vs Mass from Burrows et al. (2003) . . . . .	380
4.11	$L'$ Magnitudes vs Mass from Baraffe et al. (2003) . . . . .	381
4.12	$M$ Band Magnitudes vs Mass from Baraffe et al. (2003) . . . . .	382
4.13	Comparing the Burrows and Baraffe Models at $L'$ . . . . .	384
4.14	Comparing the Burrows and Baraffe Models at $M$ . . . . .	385
4.15	Example Sensitivity Plot (GJ 117 at $L'$ ) . . . . .	391
4.16	Planet Mass Histogram for Simulated Planets . . . . .	396
4.17	Planet Semimajor Histogram for Simulated Planets . . . . .	397
4.18	Orbital Eccentricity Histogram for Simulated Planets . . . . .	398
4.19	Orbital Inclination Histogram for Simulated Planets . . . . .	399
4.20	Apparent Magnitude Histogram for Simulated Planets . . . . .	400
4.21	Separation Histogram for Simulated Planets . . . . .	401
4.22	Detection Significance Histogram for Simulated Planets . . . . .	402
4.23	Orbital Phase Histogram for Simulated Planets . . . . .	403
4.24	Orbits of Detected Planets from Simulation #1. . . . .	412
4.25	Orbits of Detected Planets from Simulation #2. . . . .	413
4.26	Orbits of Detected Planets from Simulation #3. . . . .	414
4.27	Orbits of Detected Planets from Simulation #4. . . . .	416
4.28	Orbits of Detected Planets from Simulation #5. . . . .	417
4.29	Orbits of Detected Planets from Simulation #6. . . . .	418
4.30	Orbits of Detected Planets from Simulation #7. . . . .	419
4.31	Orbits of Detected Planets from Simulation #8. . . . .	420
4.32	Orbits of Detected Planets from Simulation #9. . . . .	421
4.33	Orbits of Detected Planets from Simulation #10. . . . .	422
4.34	Orbits of Detected Planets from Simulation #11. . . . .	423
4.35	Orbits of Detected Planets from Simulation #12. . . . .	424
4.36	Orbits of Detected Planets from Simulation #13. . . . .	425
4.37	Orbits of Detected Planets from Simulation #14. . . . .	426
4.38	Orbits of Detected Planets from Simulation #15. . . . .	427
4.39	Orbits of Detected Planets from Simulation #16. . . . .	428
4.40	Planetary Distributions Ruled Out by the Null Result (using Burrows et al. (2003) planet models) . . . . .	430

LIST OF FIGURES — *Continued*

4.41	Planetary Distributions Ruled Out by the Null Result (using Baraffe et al. (2003) planet models) . . . . .	431
4.42	Planetary Distributions Ruled Out by the Null Result (using Burrows et al. (2003) planet models) #2 . . . . .	432
4.43	Planetary Distributions Ruled Out by the Null Result (using Burrows et al. (2003) planet models) #3 . . . . .	433
5.1	Sensitivity as a function of exposure time . . . . .	444
5.2	Sensitivity as a function of parallactic rotation . . . . .	445
5.3	Peak background-limited sensitivity as a function of sky brightness . . . . .	446
5.4	$\Delta$ -mag sensitivity as a function of PSF FWHM . . . . .	448
A.1	Nod Image, Position 1 . . . . .	471
A.2	Nod Image, Position 2 . . . . .	472
A.3	Difference of the 2 Nod Positions . . . . .	473
A.4	Processing Example Image: GJ 450, Raw. . . . .	475
A.5	Processing Example Image: GJ 450, Dark Subtracted and Flatfielded . . . . .	476
A.6	Processing Example Image: GJ 450, Pixel Fixing Round # 1 . . . . .	477
A.7	Processing Example Image: GJ 450, Nod-Subtracted . . . . .	478
A.8	Processing Example Image: GJ 450, Pixel Fixing Round # 2. . . . .	479
A.9	Processing Example Image: GJ 450, Colfudge. . . . .	480
A.10	Processing Example Image: GJ 450, Final Shift and Rotation. . . . .	481
A.11	Processing Example Image: $\epsilon$ Eri, Raw. . . . .	482
A.12	Processing Example Image: $\epsilon$ Eri, Dark Subtracted and Flatfielded . . . . .	483
A.13	Processing Example Image: $\epsilon$ Eri, Pixel Fixing Round # 1 . . . . .	484
A.14	Processing Example Image: $\epsilon$ Eri, Nod-Subtracted . . . . .	485
A.15	Processing Example Image: $\epsilon$ Eri, Pixel Fixing Round # 2. . . . .	486
A.16	Processing Example Image: $\epsilon$ Eri, Colfudge. . . . .	487
A.17	Processing Example Image: $\epsilon$ Eri, Final Shift and Rotation. . . . .	488
A.18	Column Algorithm Example: GJ 450 Before Column Noise Removal . . . . .	509
A.19	Column Algorithm Example: GJ 450 After Colfudge . . . . .	510
A.20	Column Algorithm Example: GJ 450 After Matt Kenworthy's Algorithm . . . . .	511
A.21	Column Algorithm Example: $\epsilon$ Eri Before Column Noise Removal . . . . .	513
A.22	Column Algorithm Example: $\epsilon$ Eri After Colfudge . . . . .	514
A.23	Column Algorithm Example: $\epsilon$ Eri After Matt Kenworthy's Algorithm . . . . .	515
A.24	Column Algorithm Example: Final Stack of $\epsilon$ Eri Images Processed with Colfudge . . . . .	516
A.25	Column Algorithm Example: Final Stack of $\epsilon$ Eri Images Processed with the 4-column Pattern Algorithm . . . . .	517

LIST OF FIGURES — *Continued*

A.26 Comparison of the Creeping Mean with a Median/Sigma Clip. . .	529
A.27 Example of Unsharp Masking. . . . .	532
A.28 Comparing the 'a' and 'b' Processing Methods. . . . .	535
A.29 Comparing the 'a' and 'c' Processing Methods. . . . .	539
A.30 Comparing the 'a' and 'd' Processing Methods. . . . .	541
A.31 Comparing the 'a' and 'x' Processing Methods. . . . .	543
A.32 Comparing the 'x' and 'y' Processing Methods. . . . .	545
B.1 Example Postage Stamp Image . . . . .	602

## LIST OF TABLES

2.1	Observations of Science Targets: Basic Parameters . . . . .	43
2.2	Observations of Science Targets: Data Acquired . . . . .	44
2.3	Vega L'-band fake planet experiment. . . . .	49
2.3	Vega L'-band fake planet experiment. . . . .	50
2.4	Vega M-band fake planet experiment. . . . .	51
2.4	Vega M-band fake planet experiment. . . . .	52
2.5	$\epsilon$ Eridani L'-band fake planet experiment. . . . .	53
2.5	$\epsilon$ Eridani L'-band fake planet experiment. . . . .	54
3.1	Astronomical Infrared Filter Bands . . . . .	105
3.2	Ages and Age References for Observed Survey Targets . . . . .	119
3.2	Ages and Age References for Observed Survey Targets . . . . .	120
3.2	Ages and Age References for Observed Survey Targets . . . . .	121
3.3	Distances, Ages, and Spectral Types of Observed Survey Targets . . . . .	123
3.3	Distances, Ages, and Spectral Types of Observed Survey Targets . . . . .	124
3.3	Distances, Ages, and Spectral Types of Observed Survey Targets . . . . .	125
3.4	Positions and Magnitudes of Observed Survey Targets . . . . .	127
3.4	Positions and Magnitudes of Observed Survey Targets . . . . .	128
3.4	Positions and Magnitudes of Observed Survey Targets . . . . .	129
3.5	Observations of Science Targets: Basic Parameters . . . . .	144
3.5	Observations of Science Targets: Basic Parameters . . . . .	145
3.5	Observations of Science Targets: Basic Parameters . . . . .	146
3.6	Observations of Science Targets: Data Acquired . . . . .	147
3.6	Observations of Science Targets: Data Acquired . . . . .	148
3.6	Observations of Science Targets: Data Acquired . . . . .	149
3.7	Photometric Calibrations for Clio Observing Runs . . . . .	150
3.7	Photometric Calibrations for Clio Observing Runs . . . . .	151
3.8	Different Methods of Sensitivity Estimation Tested on GJ 820A . . . . .	183
3.9	Different Methods of Sensitivity Estimation Tested on GJ 625 . . . . .	184
3.10	GJ 450 L'-band fake planet experiment. . . . .	190
3.10	GJ 450 L'-band fake planet experiment. . . . .	191
3.11	Vega L'-band fake planet experiment. . . . .	192
3.11	Vega L'-band fake planet experiment. . . . .	193
3.12	Vega M-band fake planet experiment. . . . .	194
3.12	Vega M-band fake planet experiment. . . . .	195
3.13	$\epsilon$ Eridani L'-band fake planet experiment. . . . .	196
3.13	$\epsilon$ Eridani L'-band fake planet experiment. . . . .	197
3.14	HD 29391 L'-band fake planet experiment. . . . .	198
3.14	HD 29391 L'-band fake planet experiment. . . . .	200

LIST OF TABLES — *Continued*

3.15	Flux Loss in Processing from Simulated GJ 5 Experiment . . . . .	207
3.16	Background-limited Planet Mass Sensitivities in Different Bands . .	301
3.16	Background-limited Planet Mass Sensitivities in Different Bands . .	303
3.17	Confirmed Sources in Our Survey . . . . .	312
3.18	List of Suspected Sources . . . . .	325
3.18	List of Suspected Sources . . . . .	326
3.18	List of Suspected Sources . . . . .	327
3.18	List of Suspected Sources . . . . .	328
3.18	List of Suspected Sources . . . . .	329
3.18	List of Suspected Sources . . . . .	330
3.18	List of Suspected Sources . . . . .	331
3.18	List of Suspected Sources . . . . .	332
3.18	List of Suspected Sources . . . . .	333
3.18	List of Suspected Sources . . . . .	334
3.18	List of Suspected Sources . . . . .	335
3.18	List of Suspected Sources . . . . .	336
3.18	List of Suspected Sources . . . . .	337
3.18	List of Suspected Sources . . . . .	338
4.1	Astronomical Infrared Filter Bands . . . . .	352
4.2	Power Law Fits to the Masses and Semimajor Axes of RV Planets .	357
4.3	$L'$ Band Absolute Mags from Burrows et al. (2003) . . . . .	365
4.4	$M$ Band Absolute Mags from Burrows et al. (2003) . . . . .	366
4.5	$L'$ Absolute Mags from Baraffe et al. (2003) . . . . .	367
4.5	$L'$ Absolute Mags from Baraffe et al. (2003) . . . . .	368
4.6	$M$ Band Absolute Mags from Baraffe et al. (2003) . . . . .	369
4.6	$M$ Band Absolute Mags from Baraffe et al. (2003) . . . . .	371
4.7	Probabilities of Detecting $N$ Planets . . . . .	392
4.8	Statistics of Detected Planets . . . . .	392
4.9	Percentage of Detected Planets Found Around Each Star . . . . .	407
4.9	Percentage of Detected Planets Found Around Each Star . . . . .	408
4.9	Percentage of Detected Planets Found Around Each Star . . . . .	409
4.10	Stars Binned by Fraction of Detected Planets . . . . .	410
4.11	Importance of the $M$ Band Data . . . . .	411
5.1	Performance of Bigger Telescopes . . . . .	453
5.1	Performance of Bigger Telescopes . . . . .	454
5.2	Performance of $L'$ Surveys vs Lafrenière et al. (2007b) . . . . .	456
5.3	Distances at which various wavelength bands become superior . .	458
5.3	Distances at which various wavelength bands become superior . .	460
5.4	Distances, Ages, and Spectral Types of Survey Targets . . . . .	462

LIST OF TABLES — *Continued*

A.1	Nod Subtraction Pairing Option 1 . . . . .	491
A.2	Nod Subtraction Pairing Option 2 . . . . .	492
A.3	Actual Nod Subtraction Algorithm, Example # 1 . . . . .	494
A.4	Actual Nod Subtraction Algorithm, Example # 2 . . . . .	495
B.1	Explanation of zerocheckm01 script . . . . .	559
B.1	Explanation of zerocheckm01 script . . . . .	560
B.2	Explanation of flatcheckm01 script . . . . .	564
B.2	Explanation of flatcheckm01 script . . . . .	565
B.3	Explanation of legoparswrite04 script . . . . .	575
B.3	Explanation of legoparswrite04 script . . . . .	576
B.3	Explanation of legoparswrite04 script . . . . .	577
B.4	Explanation of legoparswrite03 script . . . . .	587
B.4	Explanation of legoparswrite03 script . . . . .	588
B.4	Explanation of legoparswrite03 script . . . . .	589
B.5	Explanation of scriptwritelas02 script . . . . .	594
B.5	Explanation of scriptwritelas02 script . . . . .	595
B.6	Explanation of starcenters01 script . . . . .	599
B.6	Explanation of starcenters01 script . . . . .	600
B.7	Explanation of cmcunsh01 script . . . . .	610
B.7	Explanation of cmcunsh01 script . . . . .	611
B.8	Explanation of cmcsets01 script . . . . .	614
B.9	Explanation of psfsub02 script . . . . .	620
B.9	Explanation of psfsub02 script . . . . .	621
B.9	Explanation of psfsub02 script . . . . .	622
B.10	Explanation of psfsub02 script . . . . .	634
B.10	Explanation of psfsub02 script . . . . .	635
B.10	Explanation of psfsub02 script . . . . .	636
B.11	Explanation of errradplot01 script . . . . .	642
B.11	Explanation of errradplot01 script . . . . .	643
B.11	Explanation of errradplot01 script . . . . .	644
B.11	Explanation of errradplot01 script . . . . .	645
B.11	Explanation of errradplot01 script . . . . .	646
B.11	Explanation of errradplot01 script . . . . .	647
B.11	Explanation of errradplot01 script . . . . .	648
B.11	Explanation of errradplot01 script . . . . .	649
B.11	Explanation of errradplot01 script . . . . .	650
B.12	Explanation of percentileplot01 script . . . . .	657
B.12	Explanation of percentileplot01 script . . . . .	658
B.12	Explanation of percentileplot01 script . . . . .	659
B.12	Explanation of percentileplot01 script . . . . .	660

LIST OF TABLES — *Continued*

B.13 Explanation of convpercen01 script . . . . .	663
B.13 Explanation of convpercen01 script . . . . .	664

## ABSTRACT

We have attempted adaptive optics (AO) imaging of planets around nearby stars in the  $L'$  and  $M$  bands, using the Clio instrument on the MMT. The MMT AO system, with its deformable secondary mirror, offers uniquely low background AO-corrected images in these bands. This allowed us to explore a wavelength regime that has not been well utilized in searches for extrasolar planets, but offers some advantages over the more commonly used shorter-wavelength  $H$  band regime.

We have taken deep  $L'$  and  $M$  band images of the interesting debris disk stars Vega and  $\epsilon$  Eri. Our observations of  $\epsilon$  Eri attain better sensitivity to low mass planets within 3 arcseconds of the star than any other AO observations to date. At 1.7 arcsec, the maximum separation of the known planet  $\epsilon$  Eri b, our  $M$  band sensitivity corresponds to objects only 9-16 times brighter than the predicted brightness of this planet.  $M$  is by far the most promising band for directly imaging this planet for the first time, though Clio would require a multi-night integration.

We have carried out a survey of 50 nearby stars, using mostly the  $L'$  band. The survey objective was to determine whether power law fits to the statistics of planet mass  $m$  and orbital semimajor axis  $a$  from radial velocity (RV) surveys apply when extrapolated to orbital radii beyond the outer limits of RV sensitivity. Given  $dN/dm \propto m^{-1.44}$ , our survey null result rules out  $dN/da \propto a^{-0.2}$  extending beyond 155 AU, or  $dN/da$  constant extending beyond 70 AU, at the 95% confidence level. We have not placed as tight constraints on the planet distributions as the best  $H$  band surveys. However, we have probed older planet populations and by using a different wavelength regime have helped diversify results against model uncertainties.

We have developed careful and well-tested observing, image processing, sensitivity analysis, and source detection methods, and helped advance  $L'$  and  $M$  band AO astronomy. These wavelengths will become increasingly important with the advent of new giant telescopes sensitive to interesting, low-temperature planets with red  $H - L'$  and  $H - M$  colors.

## CHAPTER 1

## INTRODUCTION

## 1.1 Poems from the Telescope

This dissertation is in part the record of a long and challenging project: the Clio 50-star survey at the MMT. As such it would not be complete without the inclusion of two poems I wrote along the way – poems descriptive of my joy and hope in God, my deep respect for my fellow astronomers and the staff who support us, and the challenge and adventure of the astronomical endeavor.

*The First Poem*

*Written at the end of the unsuccessful Clio run of December 2005.*

Mist Blows.

Wind wails.

Small cold rain that won't quite turn to snow flies with that wind.

The cold has seeped into my bones.

And now it's over.

Even if the sky cleared now, at 2:30,

We could never recover before dawn.

And so I look back:

Three nights – no stars.

The code I wrote but could not use.

The instrument failure we still can't understand.

The hours trying,

Fearing the sky would clear.  
Partially fixed at last, we were ready,  
For the stars that did not come.

Of this no word is breathed,  
In books of popular astronomy,  
Or science shows on television.

This effort to be ready.  
This fearing our desire.  
Until at last, worn out but ready we wait,  
For the stars that do not come.

Their pages and segments fill with star trails:  
The dreams and real discoveries,  
When from the wild tops of the mountains,  
Our sight leaps to the depths of space.

But not of this they tell:  
This halo round the cloud-blurred moon.  
The T. O. playing video games.  
The stubborn waiting or the frantic work,  
To make all ready,  
For the stars that do not come.

No dreams came true this run.  
And yet – and yet, to be here – is that not a dream?  
To speak as we have spoken of Your heavens,  
Mysteries we may yet seek another time.

Though I and one other alone may own them Yours,  
Yet not the less Your star-watchers we are.

And in between the nights,  
I have seen the ranges with names I do not know,  
Rugged beneath the flying clouds that cast,  
A light that shows their wildness as the Sun,  
Or Moon of starry nights could never show.  
One morning over one distant, nameless peak,  
The dawn glowed with a red that touched my heart,  
And stirred my blood with what I cannot name.  
Another, nearer peaks,  
And dizzy slopes were clothed,  
In flying swaths of wind-torn cloud.  
And as they flew the clouds told,  
Wonder of airy gulfs and strong ramparts of stone.

So Lord, You Who denied my dream,  
Have You denied it?  
Or from Whom came the dawn-glow and the wild clouds?  
Or strength to ready for what did not come?  
The chance to stand beneath this steel and glass,  
And speak with other watchers of the stars?  
Frail if we be, and frail all our work,  
Yet here, and a few places else,  
Least frail our tools and least (oh how!) unworthy,  
To search the depths of heaven where You wrought,

Wonders You hid that we might now disclose,  
 If with our trembling art of glass and wire,  
 Steel and charge and air and mind and force,  
 The thought of many, and their genius,  
 Brought to bear upon one common task –  
 If with this feeble, mighty best that we can do,  
 You combine skies that let Your secrets pass,  
 Another time.

*The author would be distressed if the line referring to the T.O. were considered a criticism. It is only intended to illustrate how hopeless the weather was: there was nothing else for him to do.*

***The Second Poem***

***Written during the successful run of September 2006, on the occasion of seeing a provocative source in a (incorrectly rotated, as it turned out) data set on GJ 5.***

I write this before I know,  
 Really and truly in the middle of science:  
 About to discover,  
 Or again be disappointed.

God, who made the stars, gave me this chance,  
 The telescope is big as a house.  
 Dozens of brilliant people built it and still make it run.  
 There is none like it in the world.

My code uses math from centuries of genius,  
 Assembled by my thought and others'

To guide a thunderous torrent of numbers  
Into an image of a distant star.

Humbled, I know I am only a small part  
Of the human endeavor.

And even it: this telescope, astronomy,  
The brilliance of countless scholars and engineers, through centuries,  
Reaching a focus (one of many) on this mountain,  
Even all this is frail – smaller and frailer than dust,  
Beneath the mighty stars of God.

He is all – and all our effort at its truest,  
Is only that we may more deeply stand in awe.

And now, in the image of star GJ 5,  
There was a thing that made my heart pound.  
Discovery or noise burst?  
I have made the images that will tell.  
And I dare not yet look.

For here, after 9 years study of astronomy,  
I still feel like the 15 year-old boy,  
Who would go with beating heart up to the desk at Sam's Club,  
To get his pictures.

Did they come out this time – or would he only learn  
A new obscure thing that can go wrong in photography?

My Lord, by whose mercy alone I dare say those words,  
I am Yours – Yours are the stars.

Your will be done...

Please...

And the answer is no yet again.

Blessed be your name, oh Lord.

For all discoveries are rubbish,

Compared to the surpassing greatness

Of knowing You.

## 1.2 Science Introduction

Extrasolar planets have captured the popular and scientific imagination for decades. Reasons for this are myriad. The most significant one is that as more was learned about the planets of our own solar system and it became clear that none of its planets except Earth would be hospitable to advanced life forms, hopes and fears of intelligent extraterrestrials came to be centered around worlds orbiting other stars. Whether or not other civilizations orbit other suns, extrasolar planets might be places humans could eventually colonize — the only habitable oases in a universe that is vast and dangerous, though hauntingly beautiful. But, I suspect, beyond these considerations, fascination with extrasolar planets is simply a part of the intense human curiosity that drives all of science. Planets, extrasolar or not, are just interesting. Our solar system includes sun-baked, atmosphere-less Mercury with its oversized iron core, the crushing, acid-laced atmosphere of Venus, temperate and lovely Earth with its oddly large moon, the fantastic calderas and canyons of low pressure desert-planet Mars, the hyperactive volcanoes of Io, the endless storms of Jupiter, Saturn's impossibly thin rings, tilted Uranus with its odd collection of satellites, Neptune with its large, retrograde, cryovolcanic moon Triton, and then a bunch of deep-frozen bodies at vast heliocentric distances, which occasionally fall inward and claim their day in the sun as splendid but ephemeral comets. With all this orbiting one 'ordinary' star, we can't help but be curious what a whole galaxy of planetary systems may have to offer.

The pulsar planets (Wolszczan & Frail, 1992) and the massive planet/brown dwarf orbiting HD 114762 (Latham et al., 1989) make the exact date of the first 'real' extrasolar planet discovery difficult to assign, but certainly the discovery of 51 Pegasi b by Michel Mayor and Didier Queloz (Mayor & Queloz, 1995) be-

gan the current explosion of extrasolar planet science. For those who hoped for planets bearing alien life, or offering a good place for future human starfarers to land, the discovery of the first extrasolar planet was not exactly made to order. For those who hoped new planetary systems would be fascinating and bizarre, it was!

The planet 51 Peg B orbits its host star in about 4 days. Mercury, the planet closest to our own sun, takes 90 days to orbit: 51 Peg B is about eight times closer to its star than Mercury is to the sun. And unlike Mercury, which is a rocky body only 50% larger than Earth's moon, the mass of 51 Peg B is about half that of Jupiter.

New discoveries followed rapidly on the heels of the Mayor & Queloz (1995) result, until as of this writing almost 200 planetary systems are known, a number of them containing multiple planets (see <http://exoplanet.eu/catalog.php>). The vast majority of these have been detected by the radial velocity (RV) method used in the Mayor & Queloz (1995) result. In this method, the radial velocity of the host star is measured with great precision using a simultaneous fit to thousands of spectral lines. If there is an orbiting planet the reflex motion of the star in response to the planet's gravity will cause low amplitude, periodic variations in the stellar radial velocity, which are detected by the minute doppler shifts they induce in the spectral lines.

With such a large number of known radial velocity planets, detailed statistical analyses of the population become possible (Fischer & Valenti, 2005). Among the results are that planets are far more common around stars of high metallicity, which probably indicates they must accrete around rocky cores. The very short period 'hot Jupiters' of which 51 Peg B is the prototype form a distinct population, but there is also a more numerous population of planets extending out to larger

orbital radii. The orbits of the hot jupiters are apparently circularized by tidal interactions with the host star, but RV planets further from their stars can have very eccentric orbits (see for example Juric & Tremaine (2007)). This is surprising, because the planets in our own solar system all have nearly circular orbits. Our own system may thus represent an unlikely outcome in planet formation — an unusually well behaved planetary system. The norm may be more chaotic, with giant planets on highly eccentric orbits.

Unfortunately, the completeness of the RV surveys drops to zero for orbital radii comparable to or greater than that of our own Jupiter (5 AU). This is because the orbital periods of these distant planets are so long that the enabling technology for RV planet surveys has not been around long enough for a good, full-period detection — and also because the amplitude of the RV variations falls off with increasing planetary orbital radius. Thus we find ourselves in the situation of having excellent population statistics only for planets in orbits very unlike those of the giant planets in our own solar system. We are still unable to answer the question of whether our solar system is unusual or not. It may be rare to have giant planets only in large, circular orbits. On the other hand, a large population of planets in Jupiter- and Saturn-like orbits may await discovery just beyond the outer sensitivity limits of the RV surveys.

This question can be addressed by planet imaging surveys. While RV surveys are not sensitive to planets in distant, long-period orbits, these are the very planets that appear sufficiently well separated from the stellar glare to be directly resolved. The methods are complementary in principle.

Theoretical models (Burrows et al. (2003), Baraffe et al. (2003)) indicate that giant planets can be imaged orbiting young, nearby stars using infrared (IR) instruments on the current generation of large telescopes equipped with adaptive

optics (AO). They would be imaged not by reflected light from the host star, as the human eye sees the planets of our own solar system, but by the IR glow from their own interiors. The interiors of giant planets are warm from the gravitational potential energy converted to internal heat in their initial accretion and subsequent slow contraction. They emit a slowly-fading glow at IR wavelengths for several billion years (Gyr) after their formation. Giant planets can most easily be imaged in young star systems, where they are still relatively bright.

Planet-imaging observations of even the most promising star systems are extremely technically challenging, and obtain sensitivity only to planets considerably more massive than the lowest-mass ones detected in tight orbits by RV surveys. As of this writing there has been no unambiguous imaging detection of a planet orbiting an ordinary, solar-type star, though objects straddling the planet/brown dwarf boundary (Neuhäuser et al., 2005; Chauvin et al., 2005; Biller et al., 2006b) or orbiting very low mass primaries (Song et al., 2006) have been imaged.

Many planet-imaging surveys have been carried out in the near-IR *H* band (1.6  $\mu\text{m}$ ) and related wavelengths (Neuhäuser et al., 2000; Masciadri et al., 2005; Biller et al., 2006a; Geißler et al., 2007; Biller et al., 2007; Lafrenière et al., 2007b). Models indicate this is a favorable wavelength regime because molecular absorption in the planet atmospheres at other wavelengths causes a disproportionate fraction of the total IR emission to come out in the *H* band. However, the *H* band is far from the blackbody peak of objects the expected temperatures of giant extrasolar planets. The much longer-wavelength *L'* (3.8  $\mu\text{m}$ ) and *M* (4.8  $\mu\text{m}$ ) bands are located closer to the peak, and offer much more favorable planet/star flux ratios, which aids in the detection of planets against the bright glare of their host stars.

Observations at the  $L'$  and  $M$  bands are also desirable to diversify the investment of planet imaging effort, in case planets are fainter than expected at the  $H$  band. Unexpected chemistry or cloud structure in planetary atmospheres could easily result in lower than expected  $H$  band fluxes. This is not to say there are no model uncertainties at longer wavelengths; no mature extrasolar planets have yet been discovered and the models are thus without observational constraint at all wavelengths. Brown dwarfs with spectral types ranging from T2 to T8 ( $T_{\text{eff}}$  from 1300 to 700K) are found observationally to be 0.2-0.7 magnitudes fainter at  $M$  band wavelengths (4.5-5.0  $\mu\text{m}$ ) than expected from theory, probably due to above-equilibrium CO concentrations (Leggett et al., 2007; Reid & Cruz, 2002). The CO enhancement probably vanishes below some  $T_{\text{eff}}$  cutoff less than 700 K. Only the youngest and most massive extrasolar planets are expected to have  $T_{\text{eff}}$  hotter than 700 K (Burrows et al., 2003; Baraffe et al., 2003), so the CO absorption in the  $M$  band may not be a feature of most planetary spectra. Increased CO would enhance the  $L'$  flux or leave it unaltered. The  $L'$  and  $M$  bands, especially considered together may be more robust since they are closer to the peak of blackbody emission for objects at the expected  $T_{\text{eff}}$  values of detectable giant extrasolar planets. In any case observations at a range of wavelengths increase the likelihood of detecting planets, and the confidence with which constraints may be placed based on null results.

The MMT AO system, because of its adaptive secondary mirror, delivers light to the science camera with fewer warm-mirror reflections than any other AO system. This allows it to deliver a uniquely low thermal background at the  $L'$  and  $M$  band wavelengths, which are strongly affected by thermal glow from uncooled optics. The Clio camera (Freed et al. (2004), Sivanandam et al. (2006)) was developed as an  $L'$  and  $M$  band AO imager expressly to take advantage of the low

thermal background the MMT AO system provides.

We have used Clio to take deep  $L'$  and  $M$  band images of two very bright stars of particular interest,  $\epsilon$  Eri and Vega. The long wavelength bands excel for observations of bright stars, since it is for the brightest primaries that the favorable planet/star flux ratio they deliver is most helpful. The  $L'$  and  $M$  bands are also at their best for very nearby stars such as  $\epsilon$  Eri, because for such stars the sensitivity extends down to intrinsically faint, cool objects with very red  $H - L'$  and  $H - M$  colors.

We have also carried out a planet imaging survey of 50 nearby, moderate-age stars using Clio mostly at the  $L'$  band but with some  $M$  band observations. The goal of this survey was to determine whether the power law distributions of masses and orbital semimajor axes that have been fit to the population of RV-detected planets still apply when extrapolated out to larger orbital radii. We sought to define a sample of stars around which our observations would be sensitive enough that a null result would rule out the best-fit extrapolations from the known RV planets. Since we wanted to compare our results directly to those from RV surveys, we tried to choose stars in a similar range of spectral types: that is, F, G, and K stars. The ability of the  $L'$  and  $M$  bands to detect very low-temperature objects around bright stars led us to choose a more nearby, older set of stars than all other planet imaging surveys to date. We thus probed more mature systems, while other surveys with younger samples have probed systems that are still undergoing considerable dynamical evolution.

In the course of the Vega and  $\epsilon$  Eri observations and the larger survey (which includes  $\epsilon$  Eri in its sample), we honed observing and data analysis strategies for  $L'$  and  $M$  band observations, which are in many ways more technically challenging than is work in the  $H$  band regime. We have also carefully analyzed ways

of estimating the point-source sensitivity of planet search images, compared the methods used in previous surveys with those we have developed, and discussed which methods are most reliable.

In Chapter 2 we briefly introduce our instrument and observing methods, and describe our observations of Vega and  $\epsilon$  Eridani. We discuss the possibility of detecting the known RV/astrometric planet of  $\epsilon$  Eridani (Benedict et al., 2006) with future deep  $M$  band images.

In Chapter 3 we discuss our larger survey. We explain how the sample was chosen, and give detailed explanations of our observing, processing, and sensitivity estimation techniques. Since we did in fact obtain a null result, our sensitivity estimation methods were critical in obtaining and validating the upper limits on which the Chapter 4 results are based. Our discussion of them is intended both to guide future workers and to aid others in interpreting our non-detections. We also present the specific sensitivity acquired around each star, as well as lists of possible sources that were detected with too high a false-alarm probability to warrant followup but may still correspond to real sources that will be detected in future, more sensitive surveys.

Chapter 4 presents the key scientific result of our survey. We discuss how our null result, combined with the statistics of radial velocity surveys, constrains the distribution of extrasolar giant planets from small to large orbital radii.

In Appendix A we give a detailed, nuts-and-bolts description of our image processing algorithms. These are quite involved and some fairly useful ideas have not, so far as we know, been implemented or suggested elsewhere. They may help future observers obtain the best possible point-source sensitivity from the typically ugly images that come from  $L'$  and  $M$  band detectors.

Appendix B consists of a detailed manual for running the image processing software mostly written by A. H. to process the survey data.

## CHAPTER 2

DEEP L' AND M-BAND IMAGING FOR PLANETS AROUND VEGA AND  $\epsilon$  ERIDANI

Using the recently commissioned Clio camera on the MMT, we have obtained long integrations in the L' and M-bands on the interesting debris-disk stars Vega and  $\epsilon$  Eridani to search for orbiting planets or brown dwarfs. Vega and  $\epsilon$  Eri are attractive targets for such a search for three reasons. First, they have dusty debris disks (Aumann et al., 1984; Gillett et al., 1984; Gillett et al., 1985; Aumann, 1985), which imply the presence of asteroids or comets — i.e. possible leftovers of planet formation (Backman & Paresce, 1993; Holland et al., 1998; Deller & Maddison, 2005). Second, in both cases asymmetries in the dust disks have led to hypotheses of giant planets orbiting at large radii, gravitationally sculpting the dust (Wilner, 2004; Wyatt, 2003; Deller & Maddison, 2005; Marsh et al., 2006; Ozernoy et al., 2000; Quillen & Thorndike, 2002). Third, a planet orbiting  $\epsilon$  Eri has already been detected by radial velocity (RV) (Hatzes et al., 2000) and astrometry (Benedict et al., 2006). We briefly describe our methods of reducing and analyzing Clio data and report the sensitivity we have obtained, including careful 'blind' tests we have performed with simulated planets to quantify our ability to make secure detections of faint objects. Finally, we discuss how the upper limits we have set to any planets or brown dwarfs orbiting these two stars compare to upper limits set by other observations, to the planets that have been hypothesized to account for the dust distributions, and to the known planet orbiting  $\epsilon$  Eri.

## 2.1 Introduction

Early space based observations with the IRAS satellite identified four bright, nearby stars with strong IR excesses:  $\beta$  Pic, Vega, Fomalhaut, and  $\epsilon$  Eri (Aumann

et al., 1984; Gillett et al., 1984; Gillett et al., 1985; Aumann, 1985). The only reasonable explanation for these excesses is that the systems contain substantial dust, which is warmed by starlight to temperatures that cause it to have a blackbody peak in the infrared. It then radiates brightly in the IR because of the vast total surface area of the countless small grains (see for example Deller & Maddison (2005); Backman (1996); Li & Lunine (2003)).

For each of these stars the timescale for dust to be removed from the system by interaction with the starlight is much shorter than the system age (Deller & Maddison, 2005; Backman, 1996; Li & Lunine, 2003). The dust cannot be primordial, but must be continually generated by the grinding down of larger bodies such as asteroids. The systems are therefore said to have 'debris disks'. They are very interesting because the existence of the debris disk indicates the star has at least an asteroid belt, and probably a more extensive planetary system, because it is unlikely that an asteroid belt would form without planets also forming, or that it would continue to grind down without being stirred by ongoing gravitational interactions with planets.

Theoretical models (e.g. Burrows et al. (2003) and Baraffe et al. (2003)) predict that it should be possible to make direct images of giant planets orbiting nearby, young stars, using the current generation of large ground-based telescopes with adaptive optics (AO). These observations are only possible at near infrared wavelengths from about 1-5  $\mu\text{m}$ , where giant planets are self-luminous due to the gravitational energy converted to internal heat in their formation and subsequent slow contraction. Because giant planets radiate this energy away over time, they become cooler and fainter as they age. It is natural, therefore, to target the youngest nearby stars in direct imaging surveys for self-luminous giant planets.

Vega and  $\epsilon$  Eridani are two nearby, relatively young stars that are interesting for additional reasons. Both have dust disks whose survival time against Poynting-Robertson drag and radiation pressure is much shorter than the age of the systems, indicating that they must have been produced fairly recently (or are continuously being produced) by collisions in asteroidal or cometary belts (Deller & Maddison, 2005). These are presumably the leftovers from planet-formation, so there should be planets in both systems. Additionally, asymmetries in the dust distributions in both systems have led to hypotheses that the dust is being gravitationally sculpted by giant planets orbiting at large distances (Wyatt, 2003; Wilner, 2004; Deller & Maddison, 2005; Marsh et al., 2006; Ozernoy et al., 2000; Quillen & Thorndike, 2002). In the case of  $\epsilon$  Eridani there is the further very interesting fact that the star is orbited by the one of the first extrasolar planets for which good radial velocity and astrometric signatures have been detected (Benedict et al., 2006). A full orbit may thus be computed, including ephemerides for separation and position angle:  $\epsilon$  Eridani b is the most promising case yet in which direct imaging surveys know precisely where to look, although because of its low mass (1.5 MJ) and small maximum separation from the star (1.7 arcsec), directly imaging the planet is a challenging goal that has not yet been achieved (Janson et al., 2007). The star is also a very promising target for direct imaging searches because of the possible existence of other planets in more distant orbits (Deller & Maddison, 2005; Benedict et al., 2006).

Most imaging searches for extrasolar planets to date have used either the  $H$  band (1.5 - 1.8  $\mu\text{m}$ ) or other filters in the same wavelength regime (see for example Neuhäuser et al. (2000); Masciadri et al. (2005); Biller et al. (2006a); Geißler et al. (2007); Biller et al. (2007); Lafrenière et al. (2007b)). The magnitude vs mass tables of Baraffe et al. (2003) and the theoretical spectra of Burrows et al. (2003) show

clearly why the  $H$  band is usually chosen: giant planets are predicted to be very bright at these wavelengths, much brighter than black bodies at their effective temperatures. Detector formats are large, technology well developed, and sky backgrounds faint at the  $H$  band relative to longer wavelengths.

However, theoretical models indicate that the planet/star flux ratio is much more favorable at the longer wavelength  $L'$  and  $M$  bands (3.4-4.1  $\mu\text{m}$  and 4.5-5.0  $\mu\text{m}$ , respectively). For planets at sufficiently large separations, or planets orbiting faint stars, the planet/star flux ratio is not relevant. Rather, it is the planet's brightness relative to the sky background and/or detector read noise that matters. In this regime the very high sky background in the  $L'$  and  $M$  bands prevents them from being as sensitive as the  $H$  band and nearby wavelengths. However, close to very bright stars the *background* becomes irrelevant and only the planet/star flux ratio matters. Under these circumstances using the longer wavelengths makes sense.

Vega is a magnitude 0.0 standard star and is among the brightest stars in the sky at almost any wavelength.  $\epsilon$  Eri, while not impressive at visible wavelengths, is a very bright magnitude 1.9 at  $H$  band. The stars are therefore excellent targets for Clio, an  $L'$  and  $M$  band optimized AO camera that had its first light on the MMT in June 2005 (Hinz et al., 2006). We have made deep  $\sim 1$  hour integrations in both the  $L'$  and  $M$  bands on both stars. Our  $M$ -band observations are the deepest ground-based images ever made in this band.

In Section 2.2 we present our observations and data analysis strategy. In Section 2.3, we describe of our methods of analyzing our sensitivity, and present the results of blind sensitivity tests that showed we obtained near 92% completeness for sources at  $10\sigma$  significance, 78% completeness for  $7\sigma$  sources, and very low completeness for  $5\sigma$  sources. We note that no other planet-imaging papers

to date present such careful blind tests in their sensitivity analyses, and that the fact that our tests resulted in somewhat lower completeness values at each significance level than might have been expected suggests such tests should always be attempted and may result in a need to revise some sensitivity estimates to more conservative values.

In Section 2.4 we compare the sensitivity we have obtained around Vega to the expected brightness of planets that have been hypothesized to explain the dust distribution. We also compare our sensitivity to that attained in other deep observations of Vega. In Section 2.5, we present the same comparisons for  $\epsilon$  Eri, and in Section 2.6 we present the conclusions of our study.

## 2.2 Observations and Data Analysis

The Clio instrument we used for our observations has been well described elsewhere (Freed et al. (2004), Sivanandam et al. (2006), and Hinz et al. (2006)). Our methods for observing and data analysis are detailed in Chapter 3. We will present only a brief overview here.

The MMT AO system delivers a lower thermal background than other AO systems because it uses the world's first deformable secondary mirror, thereby avoiding the multiple warm-mirror reflections (each adding to the thermal background) that are needed in AO systems where the deformable mirror is not the secondary. This unique property makes the MMT ideal for AO observations in wavelengths such as the  $L'$  and  $M$  bands that are strongly affected by thermal glow. Clio was developed to take advantage of this to make sensitive planet-search observations in these bands.

**Observing Strategy:** We carry out  $L'$  and  $M$  band imaging with Clio using the technique of nod-subtraction, in which we take images of each star in two differ-

ent telescope positions offset typically by about 5.5 arcsec, and then subtract the images taken in one position from those taken in the other to remove sky artifacts. Since the star is present on images taken in both positions, both provide useful science data. Nod-subtraction does result in a dark negative image of the star reducing the sensitivity in part of each image, but the area affected is fractionally small, far (5.5 arcsec) from the star, where planets are less likely to be found, and can be placed away from objects of potential interest by a good choice of the nod direction. We also have alternative ways of processing noded data that do away with the dark images entirely.

We typically nod the telescope every 2-5 minutes, which appears to be fast enough that any variations in the sky background are sampled and satisfactorily removed. We take 5 or 10 images in each nod position, each of which typically represents about 20 seconds worth of data. A full data set consists of 100-500 such images.

We choose the exposure for most of the images so that the sky background level is about 70% of the detector full-well. At such exposure times the cores of bright stars such as Vega and  $\epsilon$  Eri are saturated, but optimal sensitivity is obtained to faint point sources beyond the saturation radii. When possible, we interleave a few nod cycles of shorter exposures yielding unsaturated star images into the sequence of longer exposure images. This is very useful because the short exposures provide a good measurement of the PSF under the exact conditions of the particular observing sequence. This was possible with  $\epsilon$  Eri, but Vega proved too bright for us reasonably to obtain unsaturated images. We used other stars observed close in time to our Vega observations to provide a reference PSF for the Vega data.

Tables 2.1 and 2.2 give details of our Vega observations. The June 2006 *M*

Table 2.1. Observations of Science Targets: Basic Parameters

Star	Date Obs	Band	Clio int(msec)	Coadds	# Images
Vega	April 12, 2006	<i>L'</i>	2000	10	160
Vega	April 13, 2006	<i>M</i>	200	90	110
Vega	June 10, 2006	<i>M</i>	100	50	558
Vega	June 11, 2006	<i>M</i>	120	100	180
$\epsilon$ Eri	September 09, 2006	<i>M</i>	130	100	180
$\epsilon$ Eri	September 11, 2006	<i>L'</i>	1500	15	184

Note. — Clio int(msec) refers to the nominal single-frame exposure time in Clio. The integrate-while-reading mode used in high efficiency science imaging causes the true single-frame exposure time to be about 59.6 msec longer than the nominal exposures listed here. Coadds is the number of images internally coadded by Clio to produce a single 2-D FITS image.

band Vega observations, for reasons that are not entirely clear, had far higher sky noise than the April 2006 data, and therefore were not used in calculating the final sensitivity.

**Data Analysis:** Our Clio image processing pipeline is described in Chapter 3 and in more detail in Appendix A. Here we briefly state that the baseline processing involves dark subtraction, flat fielding, nod subtraction, several iterations of different types of deviant ('hot') pixel removal, a pattern noise correction, shifting, rotation, and zeropadding in a single bicubic spline operation, stacking using

Table 2.2. Observations of Science Targets: Data Acquired

Star	Date	Band	Exposure(sec)	Mean Airmass	Rotation
Vega	06/04/12	L'	3295.4	1.018	80.63°
Vega	06/04/13	M	2570.0	1.026	36.39°
Vega	06/06/10	M	4452.8	1.034	72.36°
Vega	06/06/11	M	3232.8	1.054	25.53°
$\epsilon$ Eri	06/09/09	M	3412.8	1.334	23.41°
$\epsilon$ Eri	06/09/11	L'	4304.5	1.342	36.92°

Note. — The observations in June were plagued with high sky noise, which has yet to be explained. Adding them to the April M-band data on Vega did not significantly increase the sensitivity to objects far from the star, though in the speckle dominated regime near the star, the sensitivity did increase by about 40 %.

a creeping mean algorithm with standard 20% rejection (see Chapter 3 or Appendix A), and finally unsharp masking using a gaussian kernel 3-4 times wider than the PSF.

We typically use six variations on this baseline processing for each data set; all six were used for Vega and  $\epsilon$  Eri. Two involve PSF subtraction somewhat analogous to the angular differential imaging described in Marois et al. (2006). Three involve pre-stack unsharp masking of each individual image as well as unsharp masking of the final combined image. Two involve the creation of a special master sky frame that eliminates the problem of the negative stellar images from the nod subtraction, at the cost of slightly higher sky noise.

The flux loss from a faint PSF in the most extreme of our six processing methods appears to be 22-30%, depending on the sharpness of the input PSF. Sharp PSFs lose less flux. Very unfortunately placed PSFs may lose more than 30% but less than 50% of their flux in rare cases. The baseline processing typically loses 16-18% of the flux. This issue is discussed in detail in Section 3.7.5 and Appendix A.

The end result of our processing is six master images. Each is optimally sensitive to sources with a slightly different set of characteristics; together, they also provide confirmation of sources that appear in multiple differently processed images. See Chapter 3 for details.

## 2.3 Sensitivity Measurements and Source Detection Tests

### 2.3.1 Calculating and Verifying our Sensitivity Results

From each of the six master images we create a separate sensitivity map. We do this by fitting a PSF from the unsaturated data to the noise on the master image, centering it in turn on every pixel of the master image, to create a best-fit PSF

amplitude map of the master image. This map contains the amplitude of the best-fit PSF centered at every point of the fully processed science image. The noise in the amplitude map accurately reflects the PSF-scale noise in the original image that would confuse detection of real point sources. We calculate the sensitivity at every point in the original image by computing the RMS in an 8 pixel radius aperture about that point on the amplitude map image (for regions too close to the star, where a disk would not produce accurate results, we use a 45-pixel arc at constant radius from the star instead). Calculating the RMS on the PSF amplitude image rather than the original processed master image takes into account spatial correlations between pixels (that is, the fact that the noise in adjacent pixels is not independent). This is a large effect in the case of speckle noise. We note that most planet-imaging papers to date have not used a sensitivity estimator that takes into account correlated noise. This can cause a significant overestimate of the sensitivity close to the star. The only paper we are aware of prior to this work that used an estimator that was properly able to account for correlated noise was Lafrenière et al. (2007b).

We quote  $10\sigma$  limits based on our sensitivity estimator. As described above, these limits are determined by measuring the noise on our final images, with careful attention to the correlated noise that affects our ability to detect point sources with the true PSF expected for the particular data set. Our processing would remove some flux from real sources in the images. In our sensitivity estimation we do not apply an explicit correction for these flux losses. Our limits may thus correspond to a true formal significance of 7-8  $\sigma$  rather than the nominal  $10\sigma$ .

After a sensitivity map has been constructed from each of our six master images for a given data set, we combine the six maps to a single master sensitivity image. To do this we first smooth the maps slightly and then perform a 'mini-

mum combine' in which the best sensitivity obtained at a given location with any method goes into the master sensitivity image.

To test the accuracy of our sensitivity estimator, we conducted blind tests in which fake planets were inserted into the raw data. The altered images were then run through our six standard processing methods, and the 'planets' were detected using both automatic and manual methods by an experimenter who knew neither their positions nor their number. These planets were inserted at fixed nominal significance levels of  $10\sigma$ ,  $7\sigma$ , and  $5\sigma$  based on the master sensitivity image. We conducted such tests for the  $L'$  and  $M$  band Vega data, and the  $L' \epsilon$  Eri data, but not for the  $M$  band  $\epsilon$  Eri images. The results of the tests are given in Tables 2.3, 2.4, and 2.5. Fake planets in these tables are classified as 'Confirmed', 'Noticed', or 'Unnoticed'. 'Confirmed' means the source was confidently detected and would certainly be worthy of long-exposure followup observations at the MMT. There is no significant doubt that confirmed sources are real. 'Noticed' means the source was flagged by our automatic detection algorithm, or noticed as a possible real object during the purely manual phase of planet-searching. 'Unnoticed' means it was not flagged nor noticed manually. The category of noticed sources is subdivided into 'Noticed; not confirmable' and 'Noticed; denied'. 'Not confirmable' means that on careful manual investigation it was not clear whether the source was real or was due to a noise fluctuation or an artifact. These sources would be considered too likely to be unreal to be worth long-exposure followup. Most unaltered data sets have a few sources in this category; it is certain the vast majority are spurious. 'Denied' means that on careful manual investigation it seemed that the entire appearance of the source could be explained by an artifact, such as a ghost or a superspeckle. The denial of one  $10\sigma$  source in Table 2.3 was apparently an operator mistake: the source seemed incon-

sistent in its appearance through the data set, and was for that reason thought to be spurious. In fact it was a 'real' inserted planet.

For the  $L' \epsilon$  Eri sensitivity test an error (in the form of a mirror flip of the sensitivity image with respect to the data) caused the planets not to be inserted at exactly the fixed 10, 7, and  $5\sigma$  values intended. The deviations were not great, however, and the problem of how to evaluate the results consistently with the other tests was solved by considering all planets fainter than  $5.92\sigma$  as  $5\sigma$  planets, all planets between  $5.92$  and  $8.36\sigma$  as  $7\sigma$  planets, and all planets brighter than  $8.36\sigma$  as  $10\sigma$  planets.  $5.92$  and  $8.36$  are the logarithmic averages of 5 and 7, and 7 and 10, respectively.

Table 2.3. Vega L'-band fake planet experiment.

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
1.27	9.87	>20 MJ	39.42 MJ	10.00 $\sigma$	Confirmed
2.19	12.53	18.62 MJ	11.22 MJ	10.00 $\sigma$	Confirmed
2.31	12.28	>20 MJ	11.62 MJ	10.00 $\sigma$	Confirmed
2.79	12.64	17.74 MJ	11.04 MJ	10.00 $\sigma$	Confirmed
3.00	13.16	13.85 MJ	10.20 MJ	10.00 $\sigma$	Confirmed
3.01	12.92	15.60 MJ	10.59 MJ	10.00 $\sigma$	Confirmed
4.44	14.52	7.63 MJ	6.35 MJ	10.00 $\sigma$	Confirmed
5.09	14.73	7.08 MJ	5.84 MJ	10.00 $\sigma$	Confirmed
5.15	15.05	6.37 MJ	4.98 MJ	10.00 $\sigma$	Confirmed
6.16	15.15	6.15 MJ	4.79 MJ	10.00 $\sigma$	Noticed and denied
8.44	15.27	5.91 MJ	4.57 MJ	10.00 $\sigma$	Confirmed
8.46	14.82	6.87 MJ	5.59 MJ	10.00 $\sigma$	Confirmed
9.14	14.81	6.89 MJ	5.62 MJ	10.00 $\sigma$	Confirmed
10.40	15.44	5.60 MJ	4.26 MJ	10.00 $\sigma$	Confirmed
1.89	12.25	>20 MJ	11.67 MJ	7.00 $\sigma$	Confirmed
2.60	13.00	15.01 MJ	10.46 MJ	7.00 $\sigma$	Confirmed
3.65	13.75	10.30 MJ	8.49 MJ	7.00 $\sigma$	Confirmed
4.71	14.08	8.95 MJ	7.49 MJ	7.00 $\sigma$	Confirmed
5.25	15.17	6.11 MJ	4.76 MJ	7.00 $\sigma$	Noticed; not confirmable

Table 2.3—Continued

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
6.27	15.53	5.43 MJ	4.09 MJ	$7.00\sigma$	Noticed and denied
6.46	15.23	5.98 MJ	4.64 MJ	$7.00\sigma$	Confirmed
1.23	10.79	>20 MJ	27.23 MJ	$5.00\sigma$	Unnoticed
1.44	10.92	>20 MJ	25.96 MJ	$5.00\sigma$	Confirmed
4.29	14.45	7.81 MJ	6.51 MJ	$5.00\sigma$	Noticed; not confirmable
4.50	15.32	5.82 MJ	4.48 MJ	$5.00\sigma$	Unnoticed
4.70	15.73	5.06 MJ	3.73 MJ	$5.00\sigma$	Unnoticed
7.27	15.02	6.44 MJ	5.05 MJ	$5.00\sigma$	Confirmed
8.14	15.90	4.78 MJ	3.44 MJ	$5.00\sigma$	Unnoticed

Note. — Planets confirmed: 13/14 at  $10\sigma$ ; 5/7 at  $7\sigma$ ; 2/7 at  $5\sigma$ . Planets noticed: 14/14 at  $10\sigma$ ; 7/7 at  $7\sigma$ ; 3/7 at  $5\sigma$ . Planet magnitude to mass conversion carried out for the two mass columns using the models of Burrows et al. (2003) and Baraffe et al. (2003), respectively.

Table 2.4. Vega M-band fake planet experiment.

Sep (asec)	<i>M</i> Band Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
0.46	6.90	> 20 MJ	>100 MJ	10.00 $\sigma$	Confirmed
1.13	9.11	> 20 MJ	74.59 MJ	10.00 $\sigma$	Confirmed
1.19	9.05	> 20 MJ	76.72 MJ	10.00 $\sigma$	Confirmed
2.39	12.43	10.66 MJ	9.80 MJ	10.00 $\sigma$	Confirmed
3.60	12.95	7.13 MJ	7.31 MJ	10.00 $\sigma$	Confirmed
3.77	13.09	6.44 MJ	6.78 MJ	10.00 $\sigma$	Confirmed
6.31	13.35	5.31 MJ	5.98 MJ	10.00 $\sigma$	Noticed; not confirmable
6.49	13.15	6.16 MJ	6.60 MJ	10.00 $\sigma$	Confirmed
7.96	13.08	6.49 MJ	6.81 MJ	10.00 $\sigma$	Confirmed
10.61	12.41	10.83 MJ	9.90 MJ	10.00 $\sigma$	Confirmed
10.84	11.83	17.23 MJ	11.35 MJ	10.00 $\sigma$	Confirmed
11.89	12.84	7.74 MJ	7.80 MJ	10.00 $\sigma$	Confirmed
0.74	7.65	> 20 MJ	>100 MJ	7.00 $\sigma$	Confirmed
0.84	8.10	> 20 MJ	>100 MJ	7.00 $\sigma$	Confirmed
2.31	12.53	9.82 MJ	9.27 MJ	7.00 $\sigma$	Confirmed
2.45	12.88	7.51 MJ	7.62 MJ	7.00 $\sigma$	Noticed; not confirmable
2.65	12.81	7.90 MJ	7.94 MJ	7.00 $\sigma$	Noticed; not confirmable
3.69	12.93	7.24 MJ	7.40 MJ	7.00 $\sigma$	Confirmed
3.99	13.22	5.85 MJ	6.38 MJ	7.00 $\sigma$	Confirmed

Table 2.4—Continued

Sep (asec)	<i>M</i> Band Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
9.48	13.07	6.54 MJ	6.84 MJ	7.00 $\sigma$	Confirmed
2.15	12.52	9.90 MJ	9.32 MJ	5.00 $\sigma$	Unnoticed
2.19	12.74	8.34 MJ	8.26 MJ	5.00 $\sigma$	Unnoticed
3.16	13.42	5.02 MJ	5.75 MJ	5.00 $\sigma$	Confirmed
3.23	12.91	7.35 MJ	7.49 MJ	5.00 $\sigma$	Noticed; denied
5.62	14.30	2.61 MJ	3.18 MJ	5.00 $\sigma$	Noticed; not confirmable

Note. — Planets confirmed: 11/12 at 10 $\sigma$ ; 6/8 at 7 $\sigma$ ; 1/5 at 5 $\sigma$ . Planets noticed: 12/12 at 10 $\sigma$ ; 8/8 at 7 $\sigma$ ; 3/5 at 5 $\sigma$ . Planet magnitude to mass conversion carried out for the two mass columns using the models of Burrows et al. (2003) and Baraffe et al. (2003), respectively.

Table 2.5.  $\epsilon$  Eridani L'-band fake planet experiment.

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
0.64	9.66	>20 MJ	27.79 MJ	10.48 $\sigma$	Confirmed
0.72	9.92	>20 MJ	25.39 MJ	10.08 $\sigma$	Noticed; not confirmable
0.84	10.21	>20 MJ	22.70 MJ	8.64 $\sigma$	Confirmed
1.37	12.44	11.55 MJ	9.64 MJ	8.53 $\sigma$	Confirmed
2.1	13.36	8.37 MJ	6.73 MJ	11.16 $\sigma$	Confirmed
3.37	15.18	4.69 MJ	3.15 MJ	8.40 $\sigma$	Confirmed
3.69	15.24	4.62 MJ	3.07 MJ	10.19 $\sigma$	Confirmed
4.76	15.31	4.53 MJ	2.98 MJ	12.10 $\sigma$	Confirmed
6.13	15.09	4.80 MJ	3.28 MJ	11.01 $\sigma$	Confirmed
9.42	15.52	4.28 MJ	2.81 MJ	8.57 $\sigma$	Confirmed
1.27	12.36	11.86 MJ	9.91 MJ	7.83 $\sigma$	Confirmed
1.67	12.94	9.72 MJ	7.95 MJ	7.39 $\sigma$	Confirmed
1.83	13.07	9.29 MJ	7.56 MJ	6.33 $\sigma$	Confirmed
2.83	14.92	5.01 MJ	3.52 MJ	6.52 $\sigma$	Confirmed
2.97	14.64	5.48 MJ	3.92 MJ	6.30 $\sigma$	Confirmed
3.27	14.98	4.94 MJ	3.44 MJ	7.58 $\sigma$	Confirmed
4.27	15.41	4.41 MJ	2.90 MJ	6.85 $\sigma$	Noticed; not confirmable
4.29	15.82	3.94 MJ	2.55 MJ	7.90 $\sigma$	Confirmed
7.52	15.88	3.89 MJ	2.50 MJ	6.32 $\sigma$	Confirmed

Table 2.5—Continued

Sep	Mass	Mass	Det.		
(asec)	$L'$ Mag	(Bur)	(Bar)	Sig.	Status
10.04	15.62	4.15 MJ	2.72 MJ	$8.25\sigma$	Confirmed
10.3	15.76	3.99 MJ	2.60 MJ	$6.06\sigma$	Unnoticed
8.98	16.13	3.68 MJ	2.29 MJ	$3.90\sigma$	Unnoticed

Note. — Planets confirmed: 9/10 at  $10\sigma$ ; 9/11 at  $7\sigma$ ; and 0/1 at  $5\sigma$ . Planets noticed: 10/10 at  $10\sigma$ ; 10/11 at  $7\sigma$ ; and 0/1 at  $5\sigma$ . Planet magnitude to mass conversion carried out for the two mass columns using the models of Burrows et al. (2003) and Baraffe et al. (2003), respectively.

The masses in Tables 2.3 through 2.5 are calculated by converting our sensitivity values to  $L'$  and  $M$  band magnitudes and then converting the magnitudes to planet masses in Jupiter Masses (MJ) by interpolating the planet models of Burrows et al. (2003) and Baraffe et al. (2003) for the appropriate mass, age, and distance of the two stars. We have adopted an age of 0.56 Gyr (Fischer, 1998) and a distance of 3.27 pc (Perryman et al., 1997) for  $\epsilon$  Eri; and an age of 0.3 Gyr (Song et al., 2001) and a distance of 7.756 pc (Perryman et al., 1997) for Vega.

The end result of the three blind sensitivity tests was that at  $10\sigma$ , 33 of 36 total inserted planets were recovered, and all 36 were noticed (and confirmed, found unconfirmable, or denied). At  $7\sigma$ , 20 of 26 total inserted sources were confirmed and 25 of the 26 were noticed. At  $5\sigma$ , 3 of 13 total inserted planets were confirmed,

and 6 of the 13 were noticed. These statistics give us a 92% completeness at  $10\sigma$ , a 77% completeness at  $7\sigma$ , and a 23% completeness at  $5\sigma$ . We note that if we had quoted  $5\sigma$  sensitivities without conducting a blind sensitivity test we would have severely overestimated our true sensitivity. Most papers in the field of planet-imaging surveys do in fact quote  $5\sigma$  limits, and do not verify their validity by a blind test.

In our sensitivity experiment low-significance planets were far more likely to be suspected than confirmed, but many spurious sources were suspected as well. Only one object was tentatively confirmed that did not turn out to be an inserted planet. It was detected at very low significance against the intense speckle noise just outside the saturation radius of the April  $M$  band image of Vega. The brightness an object must have to be detected this close to Vega corresponds to a low-mass star rather than a planet or even a brown dwarf. It seems very unlikely that such an object could be present in an orbit of a few-year period around Vega without having been noticed astrometrically. Our June 2006  $M$  band images of Vega, though considerably less sensitive than the April ones at large separations from the star, should be sensitive enough at least to suggest the source. It does not seem that they do. It appears, therefore, that the detection is spurious. At some level, this is a false positive in our sensitivity tests.

This is unfortunate, as in general we want the false positive rate to be very low so that large amounts of telescope time are not wasted following up spurious sources. We note, however, that the June observations were made simply to get more data on Vega, not to follow up the source, which was not noticed in the April data until much later. Also, this one apparent false positive was noted as more tentative than the vast majority of the other ‘real’ detections of input planets.

It thus appears that our detection strategy does indeed have a low false alarm

probability, and delivers high completeness at  $10\sigma$  and moderate completeness at  $7\sigma$ .

### 2.3.2 Final Sensitivity Results

We have converted the master sensitivity maps described above into magnitude contour images. We quote sensitivities in apparent magnitudes based on photometric standard star observations, rather than  $\Delta$ -magnitudes relative to the primary. We present our  $L'$  and  $M$  band Vega results in Figures 2.1 and 2.2, with the approximate position of the hypothetical planet from Wilner (2004) marked with a white 'X'. Figures 2.3 and 2.4 present the analogous results for  $\epsilon$  Eri. Our Vega  $M$  band observation is the deepest ground-based  $M$  band observation ever made.

We have further translated our master sensitivity map for each data set into sensitivity curves, and plotted them in various ways in Figures 2.5 through 2.22. In making the curves, we wished to take into account the fact that our sensitivity varies azimuthally as well as radially, due to the negative nod subtraction images, ghosts, and the different distances to the edge of the valid data region in different directions. Accordingly, at each radius we computed the 0th, 50th, and 90th percentile sensitivity, and have plotted all three curves in many of our figures. The 0th percentile is the worst sensitivity attained at a given radius; that is, the sensitivity in the noisiest part of a circle of that radius centered on the primary star. The 50th percentile is the median sensitivity at that radius, and the 90th percentile is the sensitivity that is exceeded only in the cleanest 10% of the image at that radius. It is always easy to identify the 0th, 50th, and 90th percentile curves in our plots, since they nest by definition.

In all cases we have plotted nominal  $10\sigma$  limits. The sensitivity in the plot increases with separation from the star but then decreases again as the edge of the

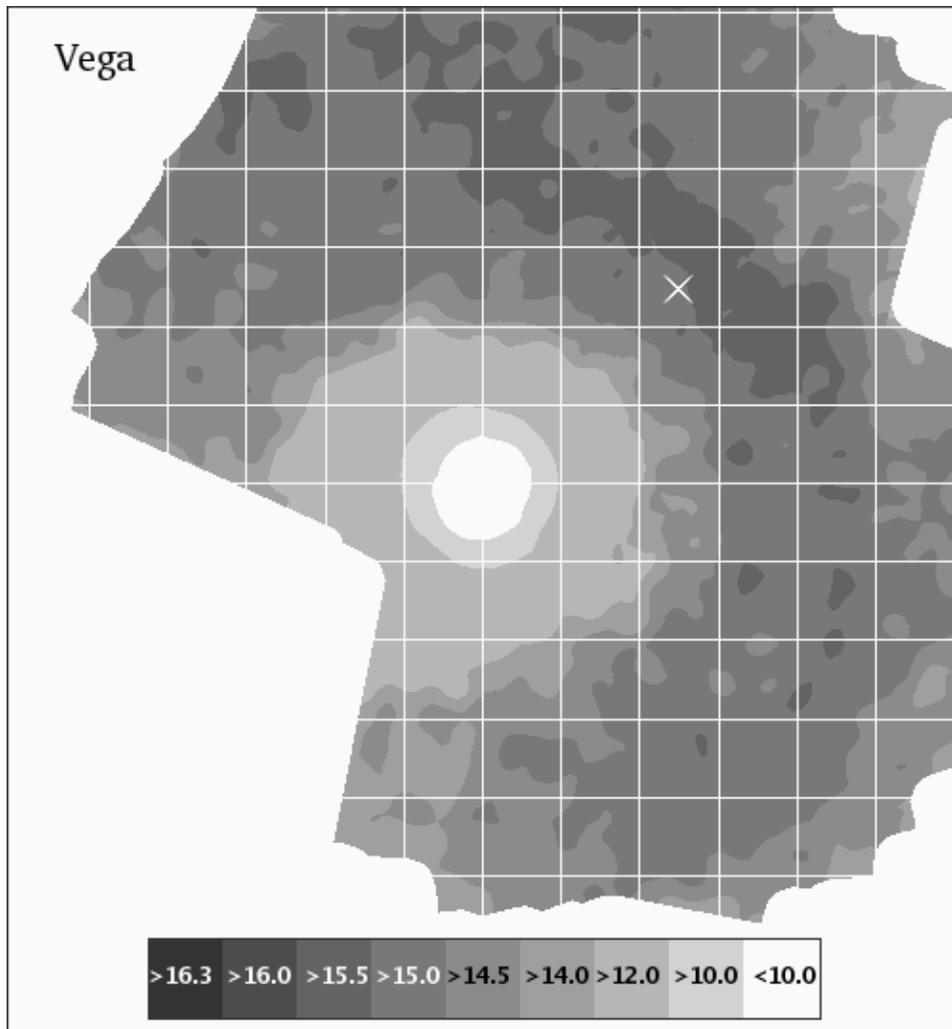


Figure 2.1  $10\sigma$  sensitivity contour map for our Vega  $L'$  observations in magnitudes. The grid squares are  $2 \times 2$  arcsec. The approximate location of the hypothetical planet from Wilner (2004) is marked with a white 'X'. The best areas in this image give sensitivity to objects fainter than  $L' = 15.5$ . In planet masses using the Burrows et al. (2003) models,  $L' = 15.5$  is 5.48 MJ, 15.0 is 6.48 MJ, 14.5 is 7.68 MJ, 14.0 is 9.28 MJ, and 12.0 and 10.0 are more than 20 MJ. The Baraffe et al. (2003) models give the masses as 4.14, 5.10, 6.39, 7.74, 15.23, and 37.48 MJ, respectively.

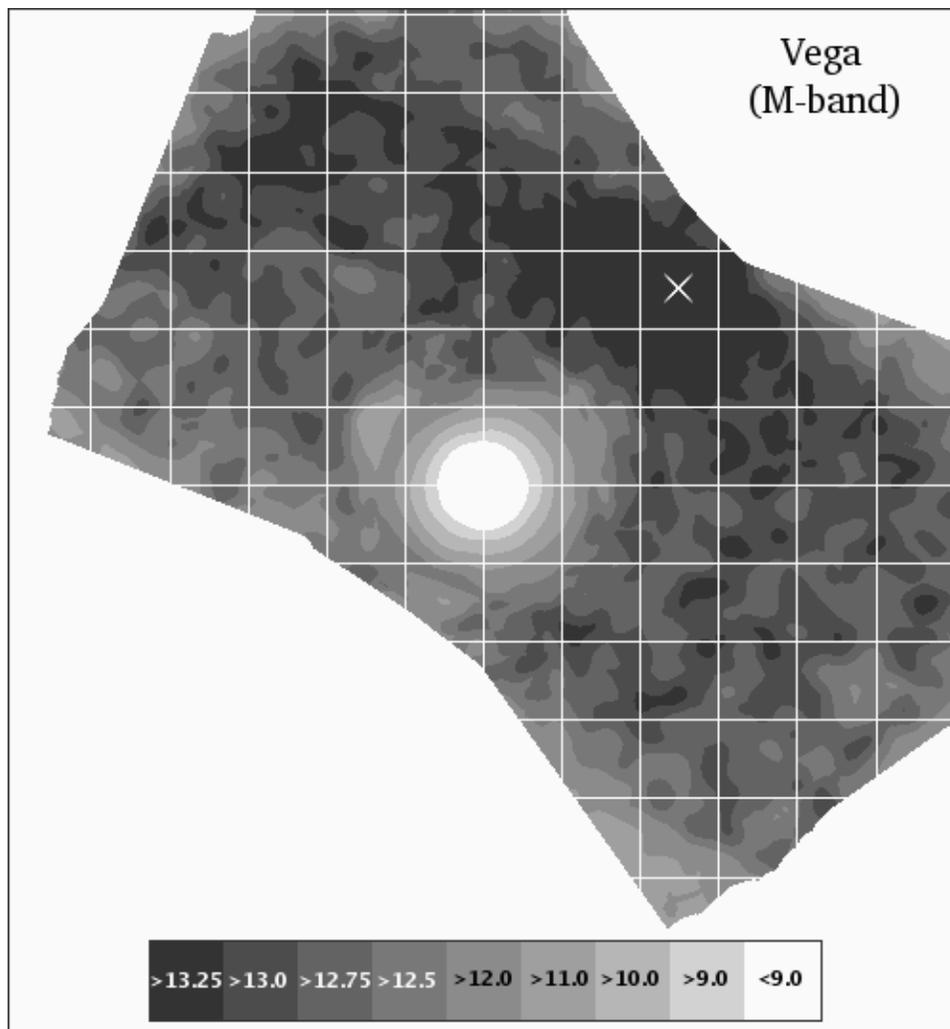


Figure 2.2  $10\sigma$  sensitivity contour map for our Vega  $M$  observations in magnitudes. The grid squares are  $2 \times 2$  arcsec. The approximate location of the hypothetical planet from Wilner (2004) is marked with a white 'X'. Using the Burrows et al. (2003) models,  $M = 13.25$  corresponds to a planet mass of 5.70 MJ, 13.0 to 6.85 MJ, 12.75 to 8.25 MJ, 12.0 to 15.11 MJ, and 11.0, 10.0, and 9.0 to masses above 20 MJ. Using the Baraffe et al. (2003) models, the masses become 6.23, 7.06, 8.20, 9.41, 10.93, 25.11, 49.15, and 78.36 MJ, respectively.

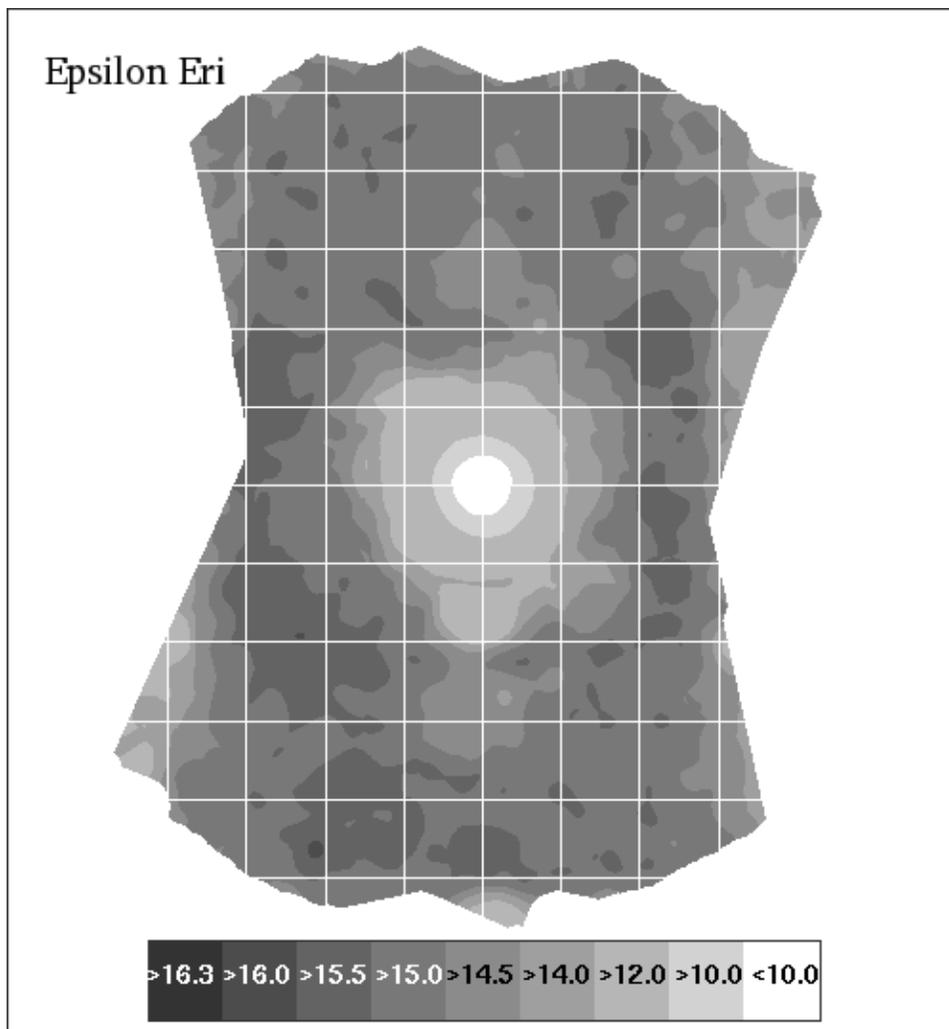


Figure 2.3  $10\sigma$  sensitivity contour map for our  $\epsilon$  Eri  $L'$  observations in magnitudes. The grid squares are  $2 \times 2$  arcsec. Tiny areas within this image attain sensitivity better than  $L' = 16.0$ ; in wider areas the limit is between 15.5 and 16.0. Converted to planet masses using the Burrows et al. (2003) models,  $L' = 16.0$  is 3.79 MJ, 15.5 is 4.30 MJ, 15.0 is 4.92 MJ, 14.5 is 5.72 MJ, 14.0 is 6.71 MJ, 12.0 is 13.49 MJ, and 10.0 is more than 20 MJ. With the Baraffe et al. (2003) models the masses become 2.40, 2.82, 3.41, 4.16, 5.14, 10.65, and 24.67 MJ, respectively.

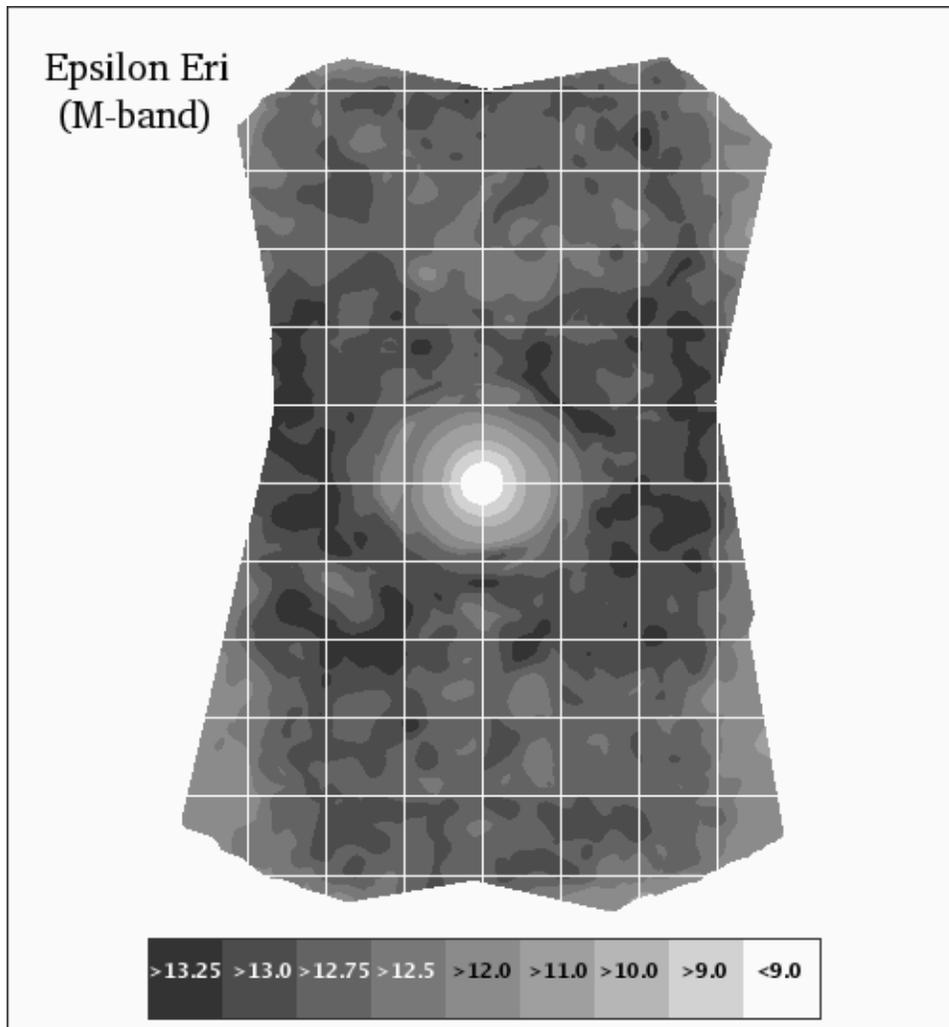


Figure 2.4  $10\sigma$  sensitivity contour map for our  $\epsilon$  Eri  $M$  observations in magnitudes. The grid squares are  $2 \times 2$  arcsec. Converted to planet masses using the models of Burrows et al. (2003), an  $M$  band magnitude of 13.25 corresponds to a planet mass of 2.76 MJ, 13.0 to 3.07 MJ, 12.75 to 3.54 MJ, 12.5 to 4.00 MJ, 12.0 to 5.37 MJ, 11.0 to 11.02 MJ, and 10.0 and 9.0 to masses above 20 MJ. Using the Baraffe et al. (2003) models the masses become 3.03, 3.49, 3.94, 4.58, 6.12, 10.32, 18.51, and 39.53 MJ, respectively.

good data region (ie, useful field on the master stacked images) is reached. The noise goes up at the edge of the useful field because due to the shifts and rotations required to register the images, the coverage (number of images supplying data to a given pixel) goes down near the edge of the field.

In Figure 2.5, we plot vs the 0th, 50th, and 90th percentile sensitivity we obtained in our  $L'$  observations of Vega in units of magnitudes, vs the separation from Vega in arcseconds. Note that since Vega has approximately magnitude 0.0 at every band, these are also  $\Delta$ -magnitude values. The simulated planets from the blind sensitivity test are shown in addition the sensitivity curves, with a legend indicating their input sensitivity and whether they were confirmed or suspected.

In Figures 2.6 and 2.7 we plot the same sensitivity curves and simulated planets as in Figure 2.5, but this time the sensitivities are in MJ based on the Burrows et al. (2003) and Baraffe et al. (2003) models for the two figures, respectively, and they are plotted vs projected separation in AU rather than arcseconds. Note that the Baraffe et al. (2003) models indicate a somewhat better sensitivity for our data than the Burrows et al. (2003) models. This is shown more clearly in Figure 2.8, where we plot the median sensitivities converted to planet masses according to the Burrows et al. (2003) models (solid line) and the Baraffe et al. (2003) models (dashed line).

Figure 2.9 shows the magnitude sensitivity vs separation in arcseconds for our April  $M$  band observations of Vega, with the planets from the sensitivity test and a legend for identifying them shown as well. Figures 2.10 and 2.11 present the same information, but the sensitivity is plotted in terms of planet mass against projected separation in AU, with the magnitudes converted to masses using the models of Burrows et al. (2003) and Baraffe et al. (2003) for the two figures respectively. The differences between the two model sets are much less at  $M$  band

than at  $L'$ , as is made clear in Figure 2.12, showing the median sensitivity vs. projected separation converted to mass using the Burrows et al. (2003) (solid line) and Baraffe et al. (2003) (dashed line) planet models.

Finally, Figures 2.13 and 2.14 compare the median mass sensitivities of the  $L'$  and  $M$  band Vega observations with magnitudes converted to planet masses using the models of Burrows et al. (2003) and Baraffe et al. (2003), respectively, for the two figures. It is interesting to note that according to the Burrows et al. (2003) models better sensitivity is in general obtained in the  $M$  band, while the Baraffe et al. (2003) models favor the  $L'$  band instead. The disagreement is not slight. An ambiguity is introduced by the fact that we integrated the Burrows et al. (2003) model spectra using the MKO filters used in Clio, but the model magnitudes from Baraffe et al. (2003) are based on the Johnson-Glass  $L'$  and the Johnson  $M$  band filters. From Bessell & Brett (1988) it appears that the Johnson  $M$  band filter is indistinguishable from the MKO, and Reid & Cruz (2002) indicate that only subtle differences exist between the various standard  $L'$  filters. However, the Burrows et al. (2003) models indicate that the flux from extrasolar planets changes rapidly across the  $L'$  bandpass, so we cannot rule out subtle filter differences being partly responsible for the disagreement between the models. We note, however, that the Burrows et al. (2003) and Baraffe et al. (2003) models also differ in the  $T_{\text{eff}}$  values assigned to planets at a given mass and age, and that the  $L'$  band should be more sensitive to these variations as well. It is unlikely that differences in the filter sets alone are responsible for the differing predictions of the models.

Figure 2.15 shows the  $L'$  sensitivity of our  $\epsilon$  Eri observations in magnitudes vs separation in arcseconds. The simulated planets from the sensitivity test are also shown, with a legend as in the previous figures. Note that the magnitudes in this plot may be converted to  $\Delta$ -magnitudes by subtracting the  $L'$  magnitude of  $\epsilon$  Eri,

which is about 1.72. Figures 2.16 and 2.17 present the same information as 2.15, with the magnitude limits converted to planet masses limits using the models of Burrows et al. (2003) and Baraffe et al. (2003), respectively, for the two figures. The sensitivities are plotted against separation in projected AU rather than arcsec. Note that the Baraffe et al. (2003) models indicate a much better sensitivity for our observations than the Burrows et al. (2003) models. This is made even more clear in Figure 2.18, in which the median mass sensitivity from the Burrows et al. (2003) models (solid line) is compared with that from the Baraffe et al. (2003) models (dashed line).

Since there was no blind sensitivity test for the  $\epsilon$  Eri  $M$  band data, and since there is very little difference between the predictions of the Burrows et al. (2003) and Baraffe et al. (2003) models for  $M$  band sensitivity, we show only two figures for this data set. Figure 2.19 shows the sensitivity in  $M$  band magnitudes plotted against the separation in arcsec, and Figure 2.20 shows the same data with the sensitivities converted to mass limits using the Burrows et al. (2003) models, and plotted against projected separation in AU.

Finally, we present Figures 2.21 and 2.22, comparing the mass sensitivities we attained around  $\epsilon$  Eri at  $L'$  (solid line) and the  $M$  band (dashed line), according to the models of Burrows et al. (2003) and Baraffe et al. (2003), respectively, for the two figures. Note that  $M$  is at least comparable with  $L'$  under the Baraffe et al. (2003) model set and is much better than  $L'$  according to the Burrows et al. (2003) models. The  $L'$  band is not a clear winner according either model set. This is in contrast to the Vega results, where the Baraffe et al. (2003) models gave a clear advantage to the  $L'$  band and the Burrows et al. (2003) models gave the same to the  $M$  band.

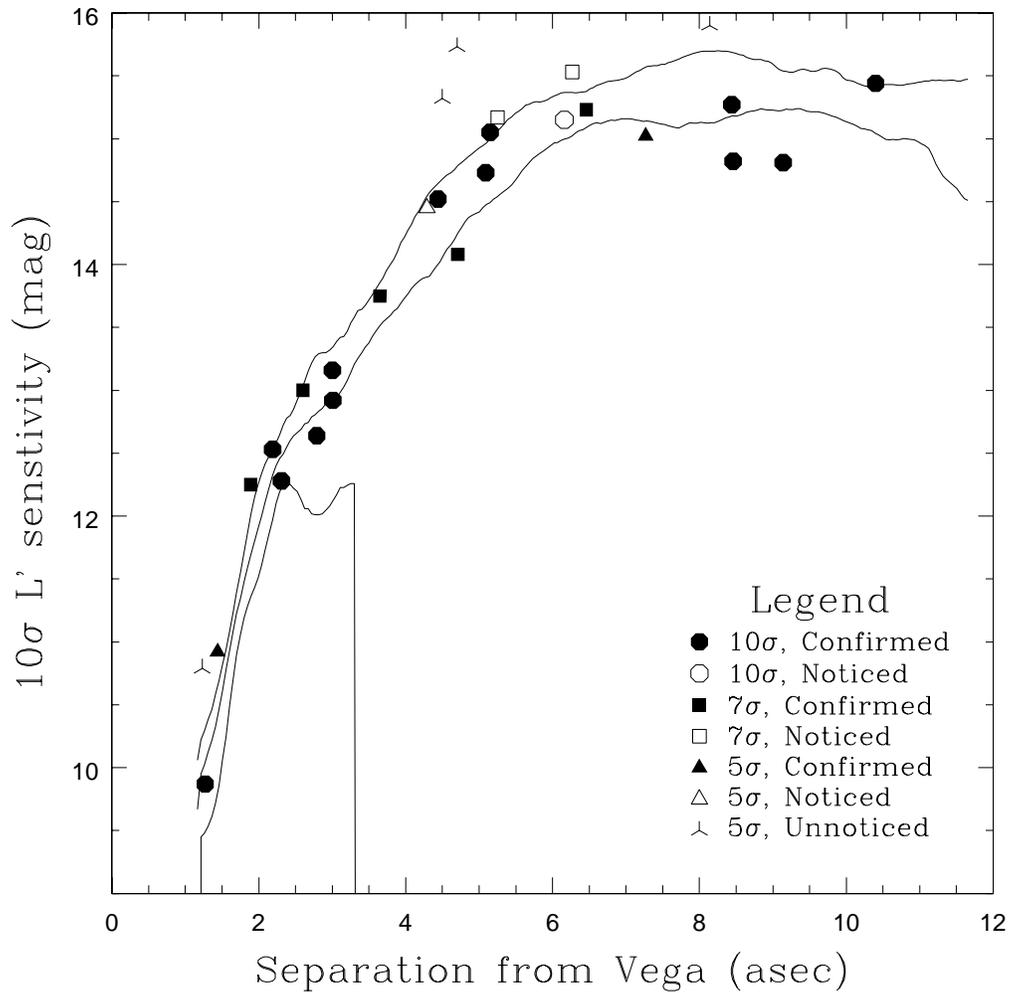


Figure 2.5  $10\sigma$  sensitivity of our Vega  $L'$  observations in magnitudes, plotted against separation in arcseconds. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

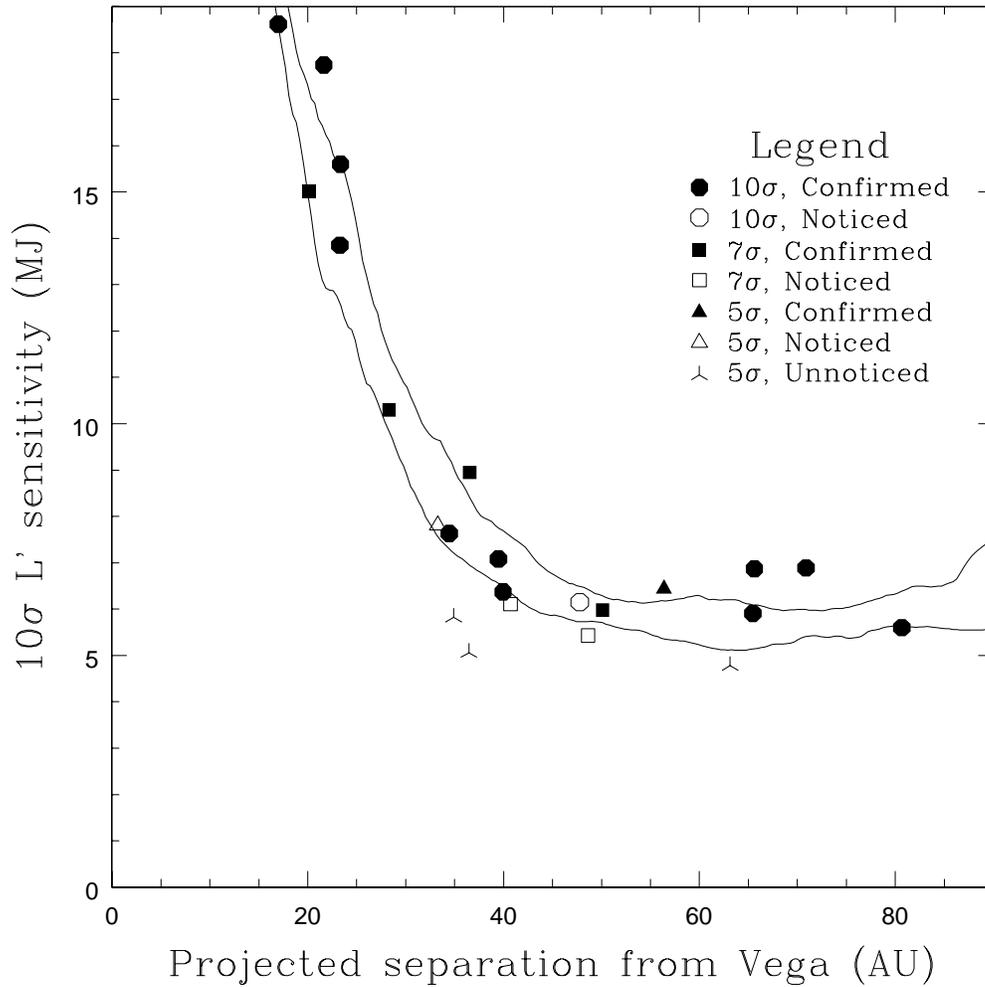


Figure 2.6  $10\sigma$  sensitivity of our Vega  $L'$  observations in terms of minimum detectable planet mass in MJ, plotted against projected separation in AU. The magnitude-mass conversion was done using the Burrows et al. (2003) models. The 50th and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

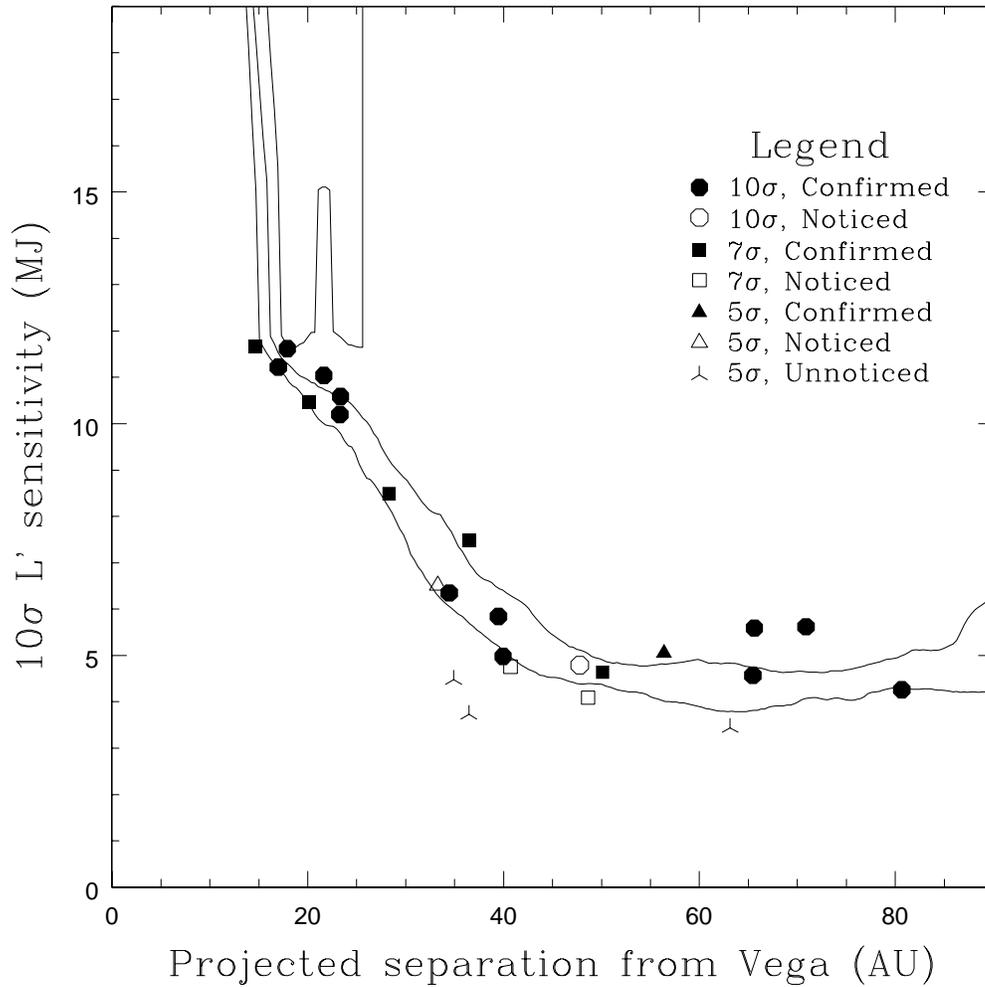


Figure 2.7  $10\sigma$  sensitivity of our Vega  $L'$  observations in terms of minimum detectable planet mass in MJ, plotted against projected separation in AU. The magnitude-mass conversion was done using the Baraffe et al. (2003) models. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

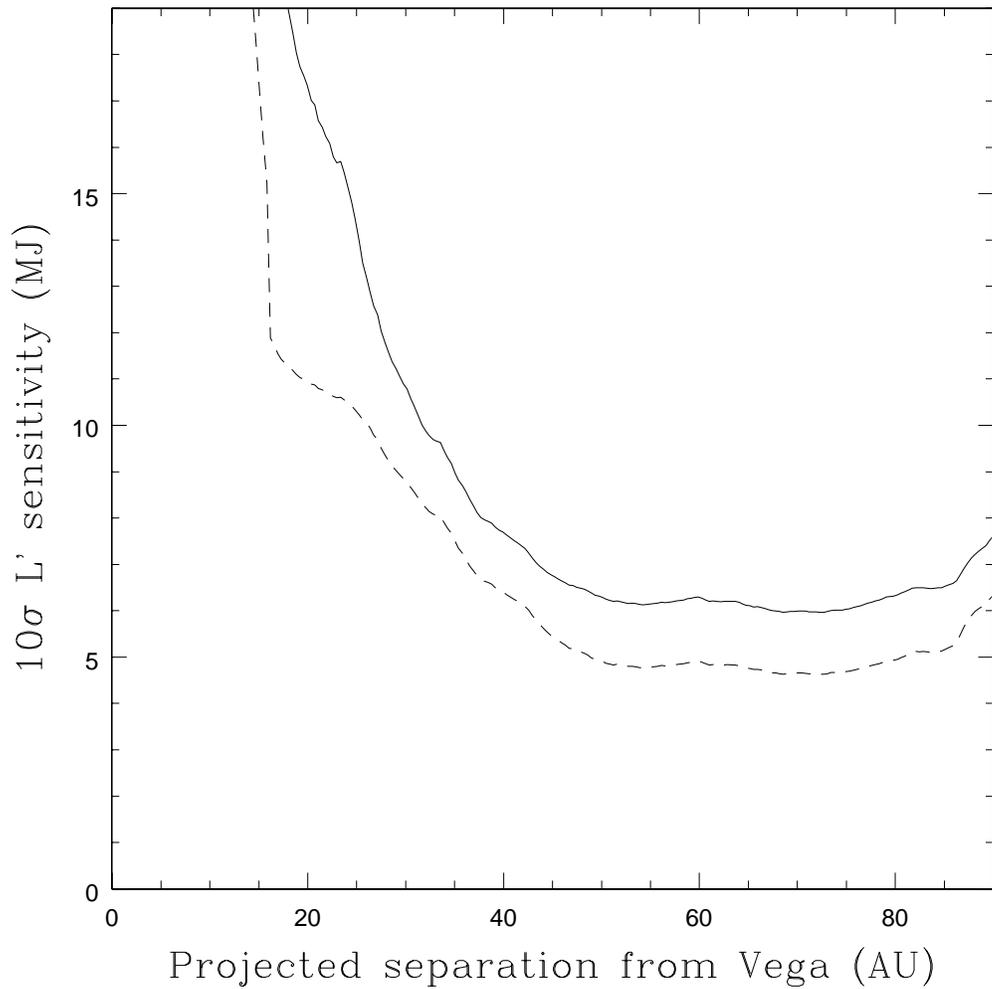


Figure 2.8  $10\sigma$  sensitivity of our Vega  $L'$  observations in terms of minimum detectable planet mass in MJ, comparing the results from the Burrows et al. (2003) (solid line) and Baraffe et al. (2003) (dashed line) planet models. The 50th percentile, or median, sensitivity curves are shown.

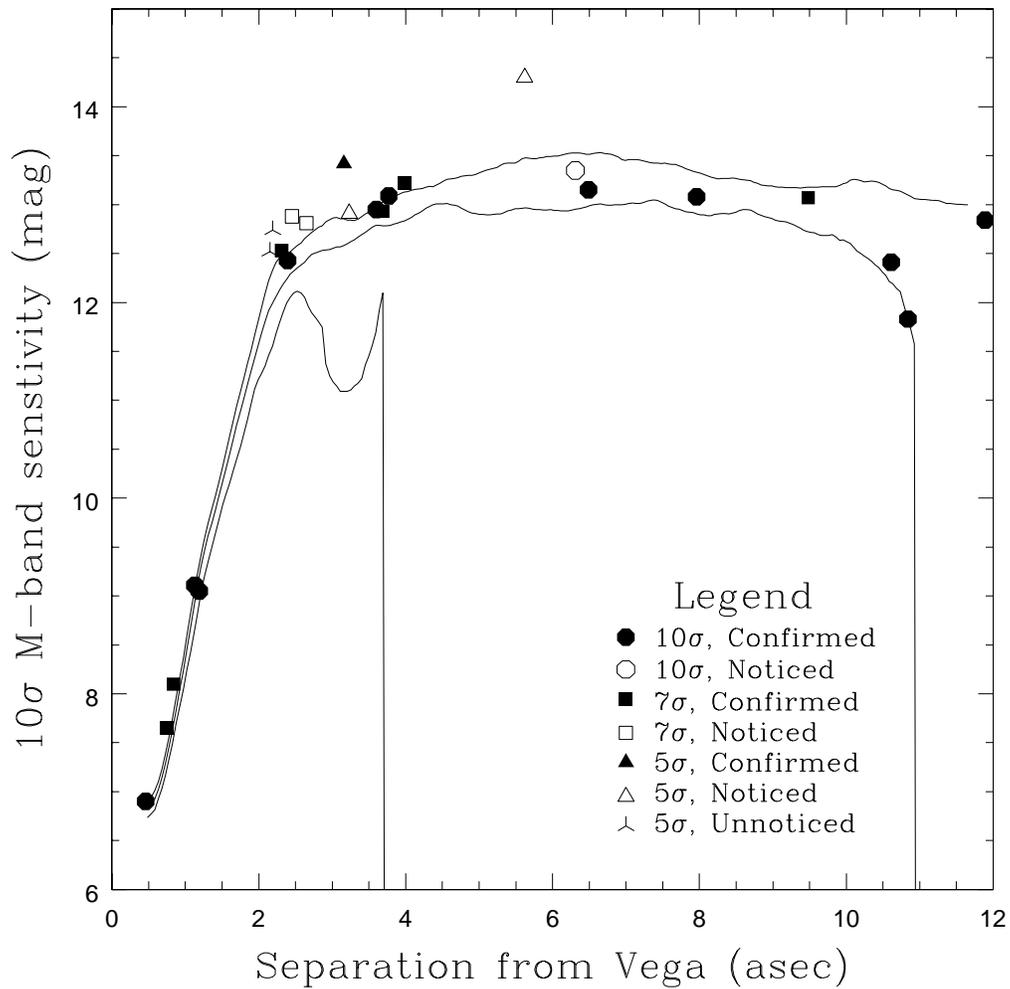


Figure 2.9  $10\sigma$  sensitivity of our Vega  $M$  band observations in magnitudes, plotted against separation in arcseconds. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

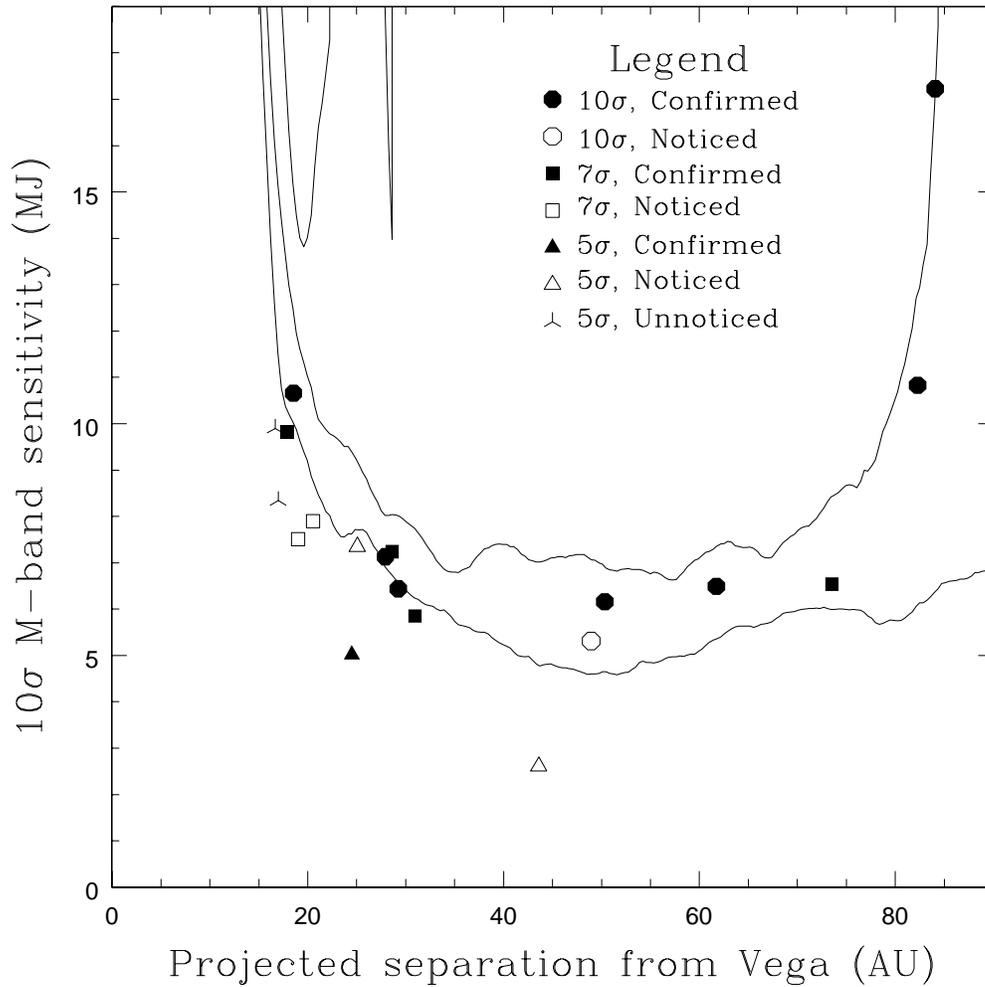


Figure 2.10  $10\sigma$  sensitivity of our Vega  $M$  band observations in terms of minimum detectable planet mass in MJ, plotted against projected separation in AU. The magnitude-mass conversion was done using the Burrows et al. (2003) models. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

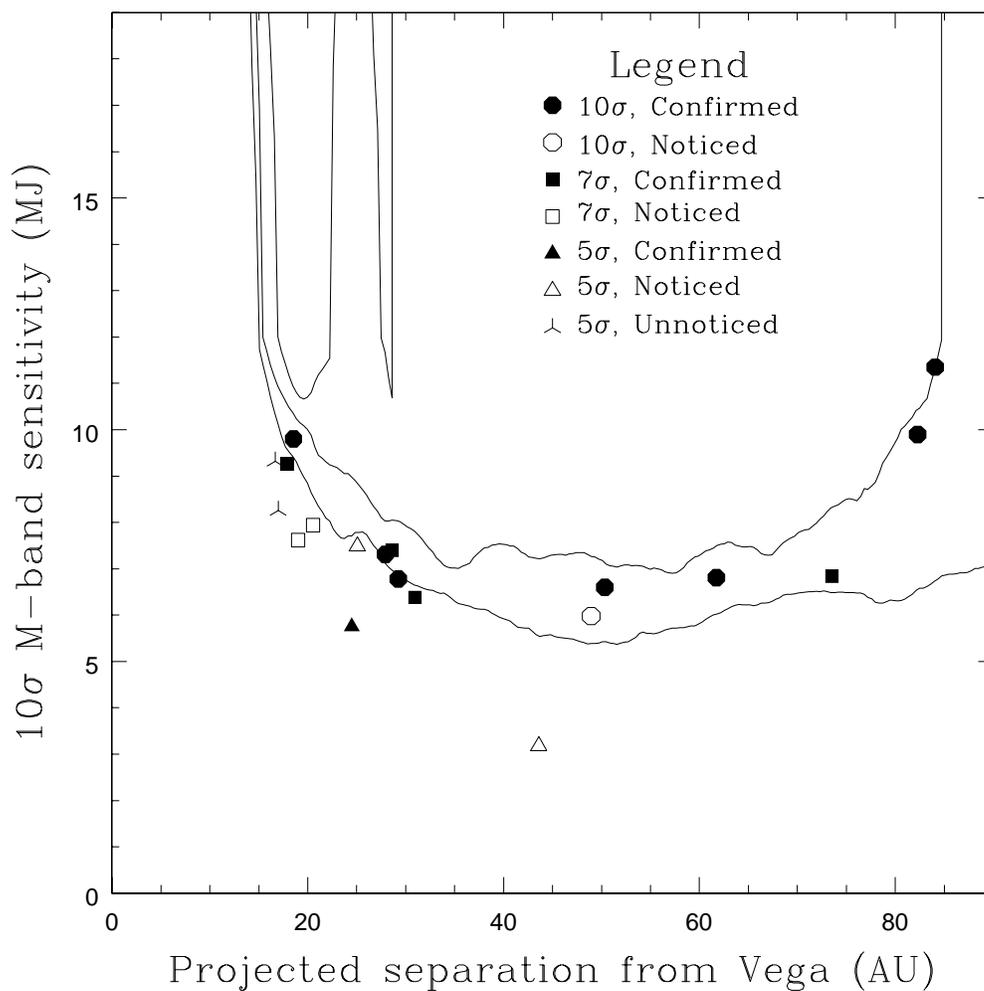


Figure 2.11  $10\sigma$  sensitivity of our Vega  $M$  band observations in terms of minimum detectable planet mass in MJ, plotted against projected separation in AU. The magnitude-mass conversion was done using the Baraffe et al. (2003) models. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

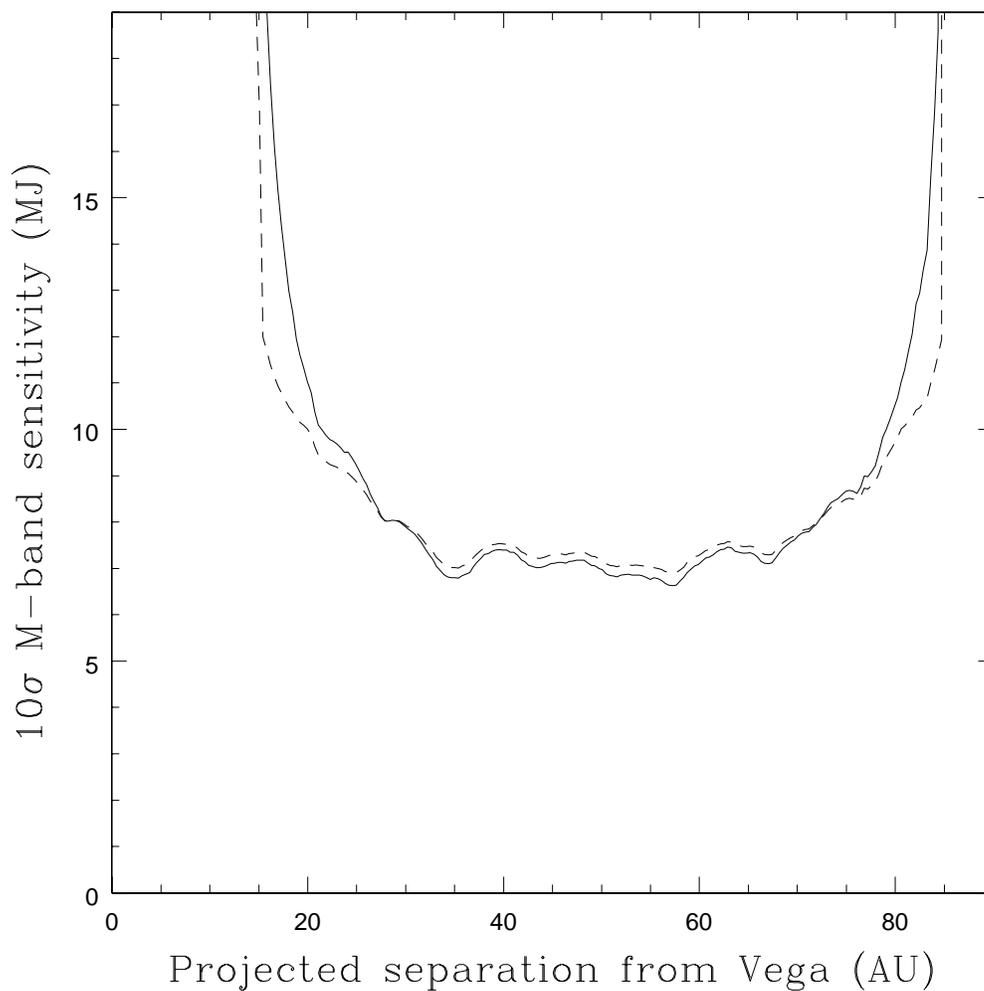


Figure 2.12  $10\sigma$  sensitivity of our Vega  $M$  band observations in terms of minimum detectable planet mass in MJ, comparing the results from the Burrows et al. (2003) (solid line) and Baraffe et al. (2003) (dashed line) planet models. The 50th percentile, or median, sensitivity curves are shown.

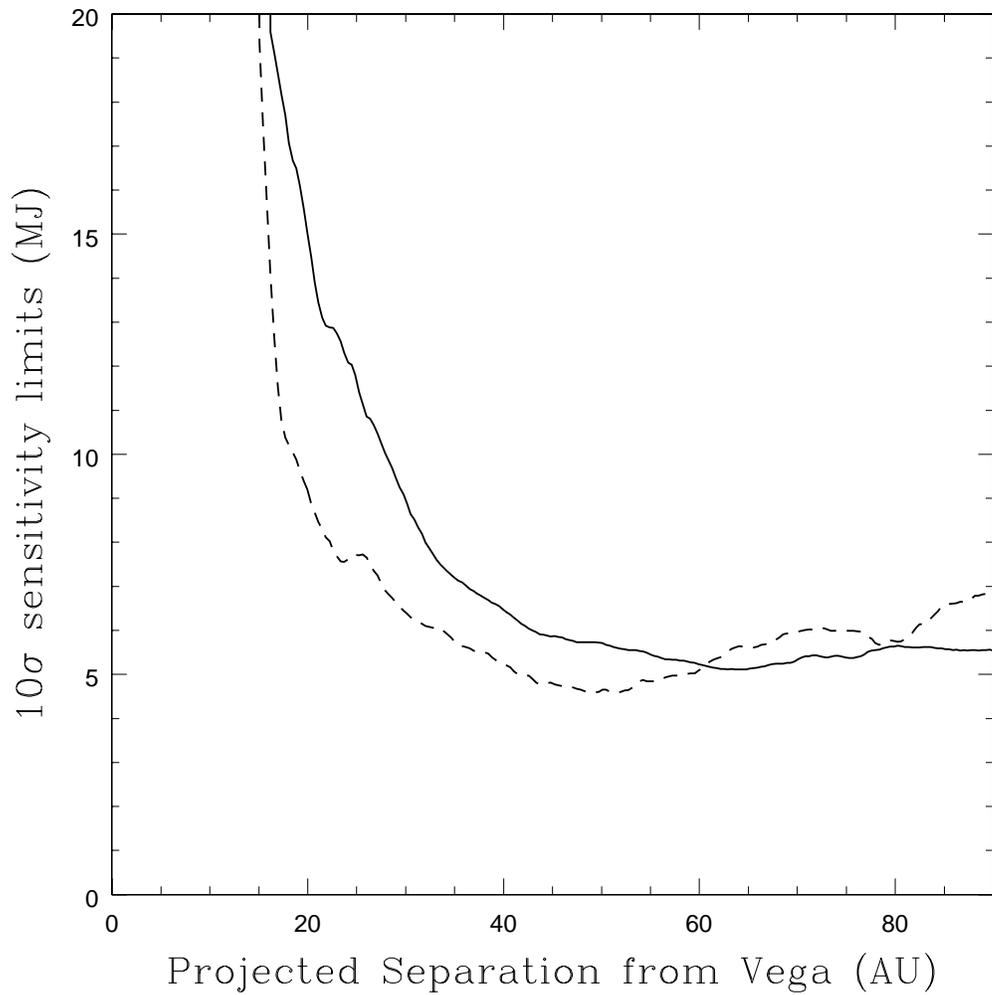


Figure 2.13  $10\sigma$  sensitivity of our Vega  $L'$  (solid line) and  $M$  band (dashed line) observations in terms of minimum detectable planet mass in MJ, based on the models of Burrows et al. (2003)

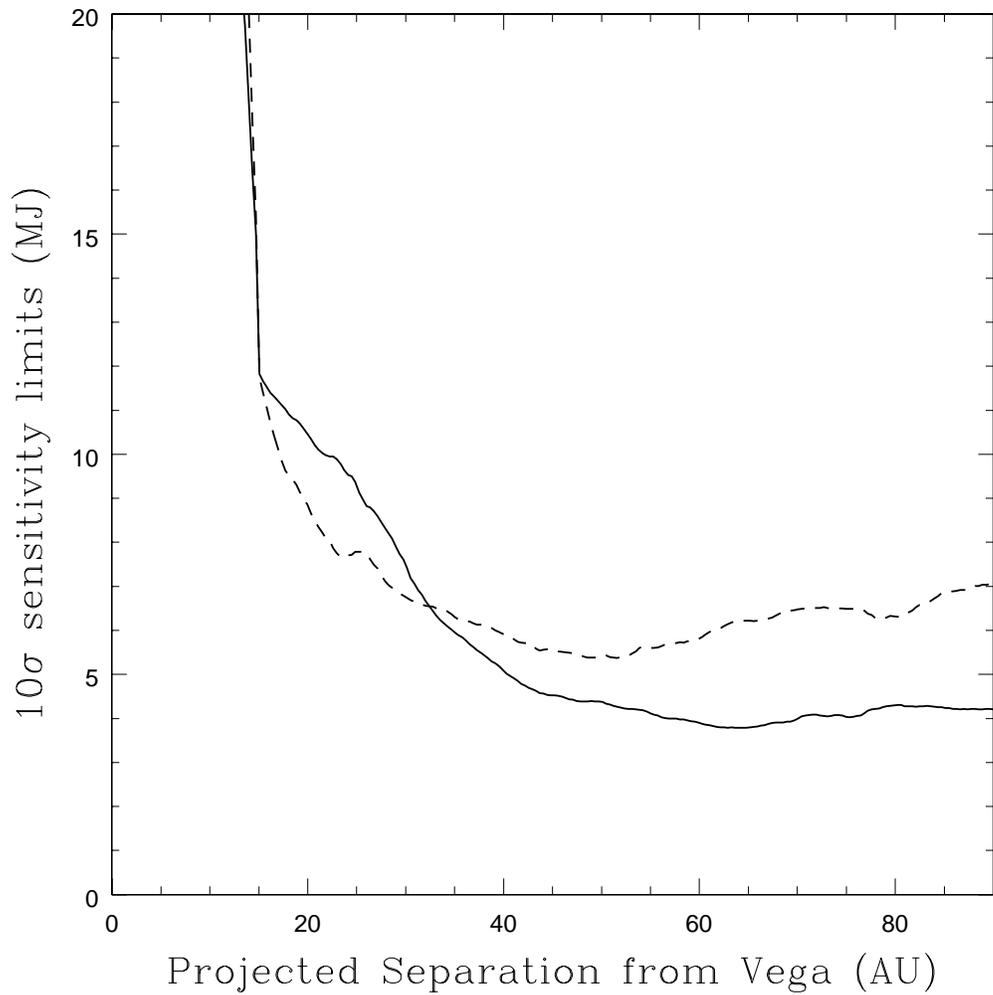


Figure 2.14  $10\sigma$  sensitivity of our Vega  $L'$  (solid line) and  $M$  band (dashed line) observations in terms of minimum detectable planet mass in MJ, based on the models of Baraffe et al. (2003)

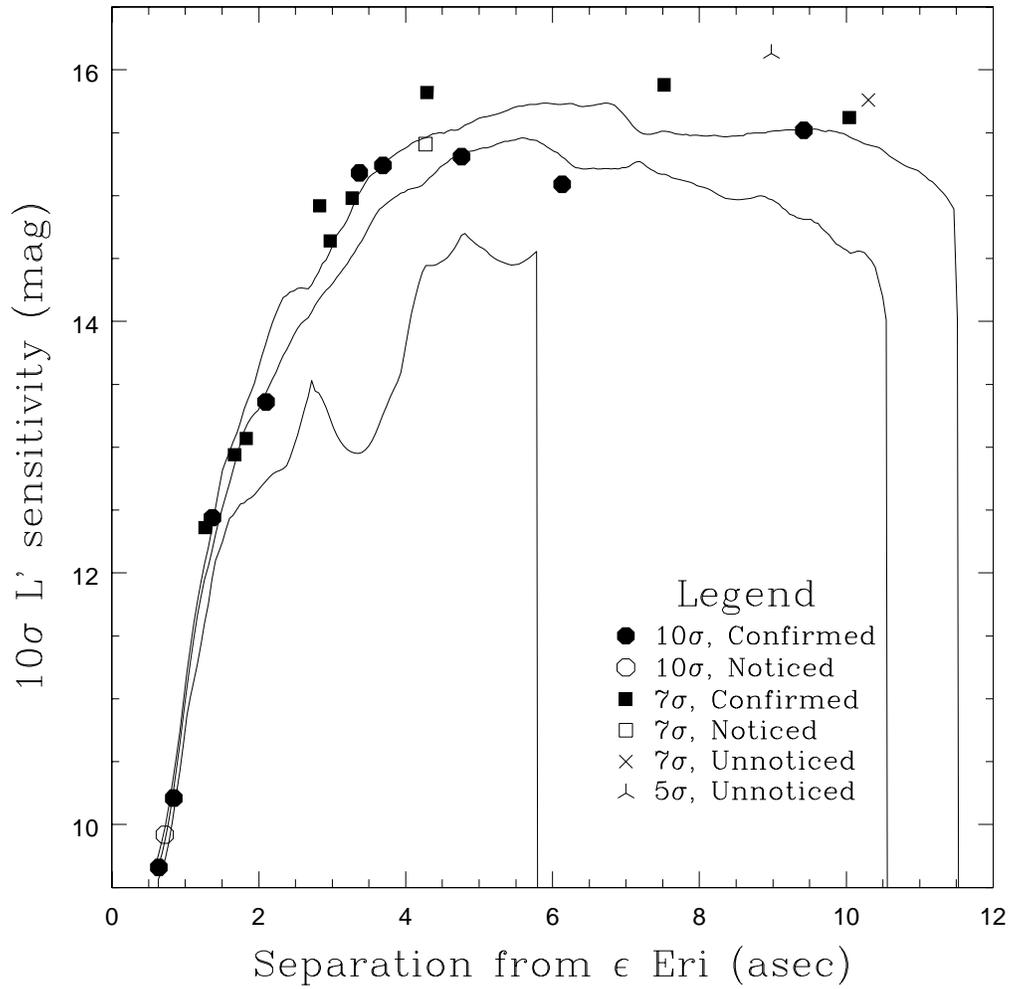


Figure 2.15  $10\sigma$  sensitivity of our  $\epsilon$  Eri  $L'$  observations in magnitudes, plotted against separation in arcseconds. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

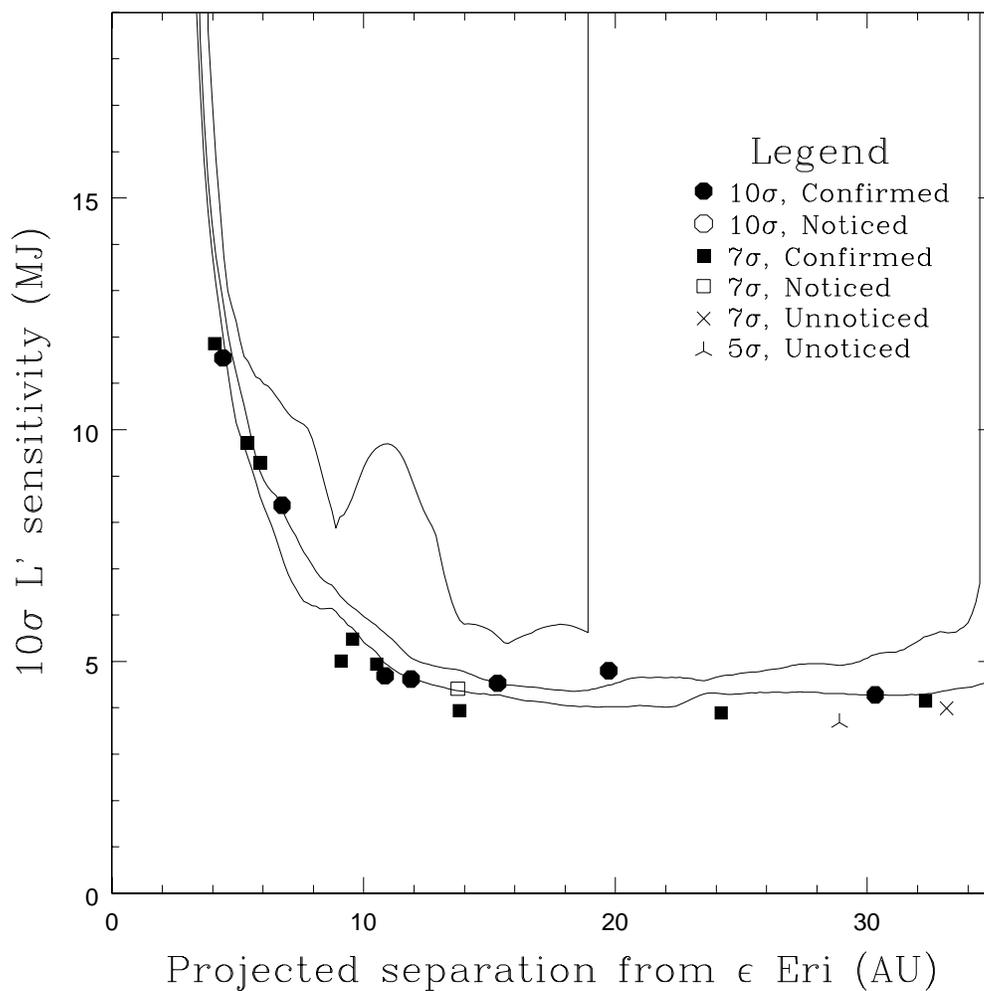


Figure 2.16  $10\sigma$  sensitivity of our  $\epsilon$  Eri  $L'$  observations in terms of minimum detectable planet mass in MJ, plotted against projected separation in AU. The magnitude-mass conversion was done using the Burrows et al. (2003) models. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

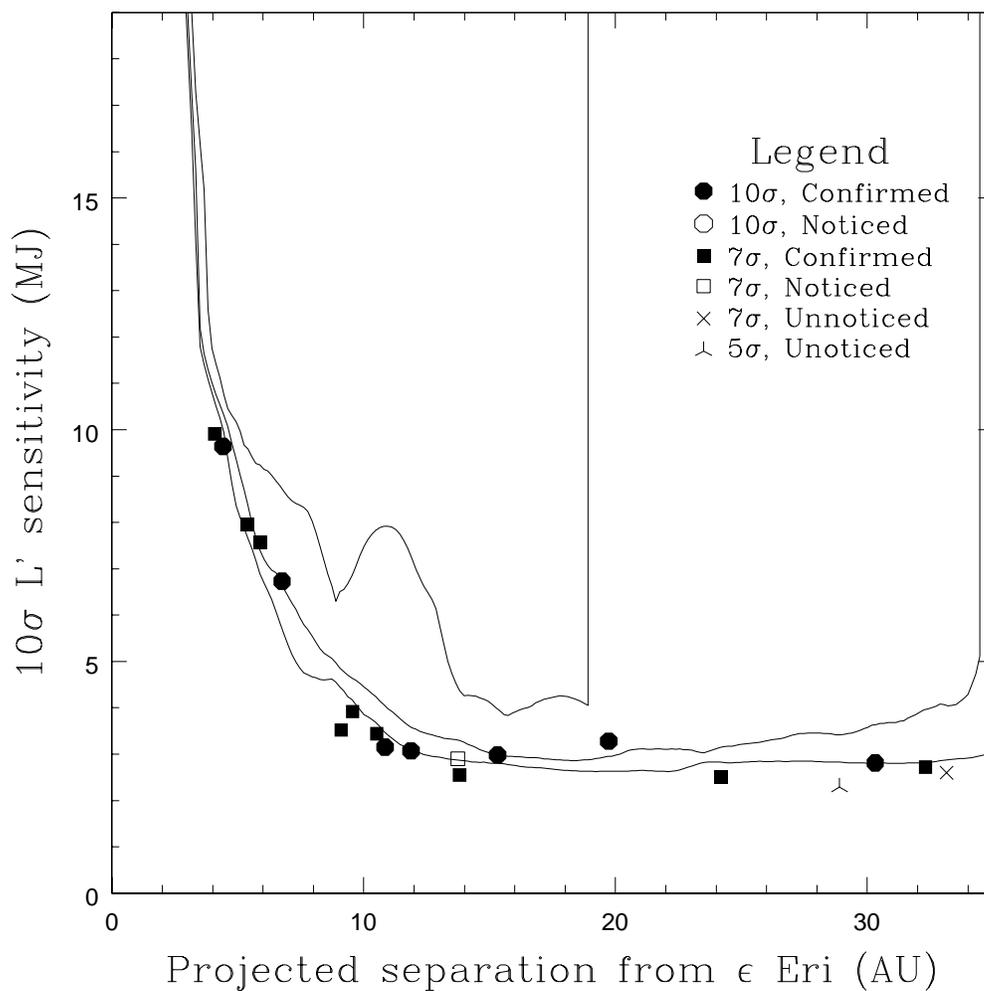


Figure 2.17 10 $\sigma$  sensitivity of our  $\epsilon$  Eri  $L'$  observations in terms of minimum detectable planet mass in MJ, plotted against projected separation in AU. The magnitude-mass conversion was done using the Baraffe et al. (2003) models. The 0th, 50th, and 90th percentile sensitivity curves are shown, along with simulated planets from the blind sensitivity test.

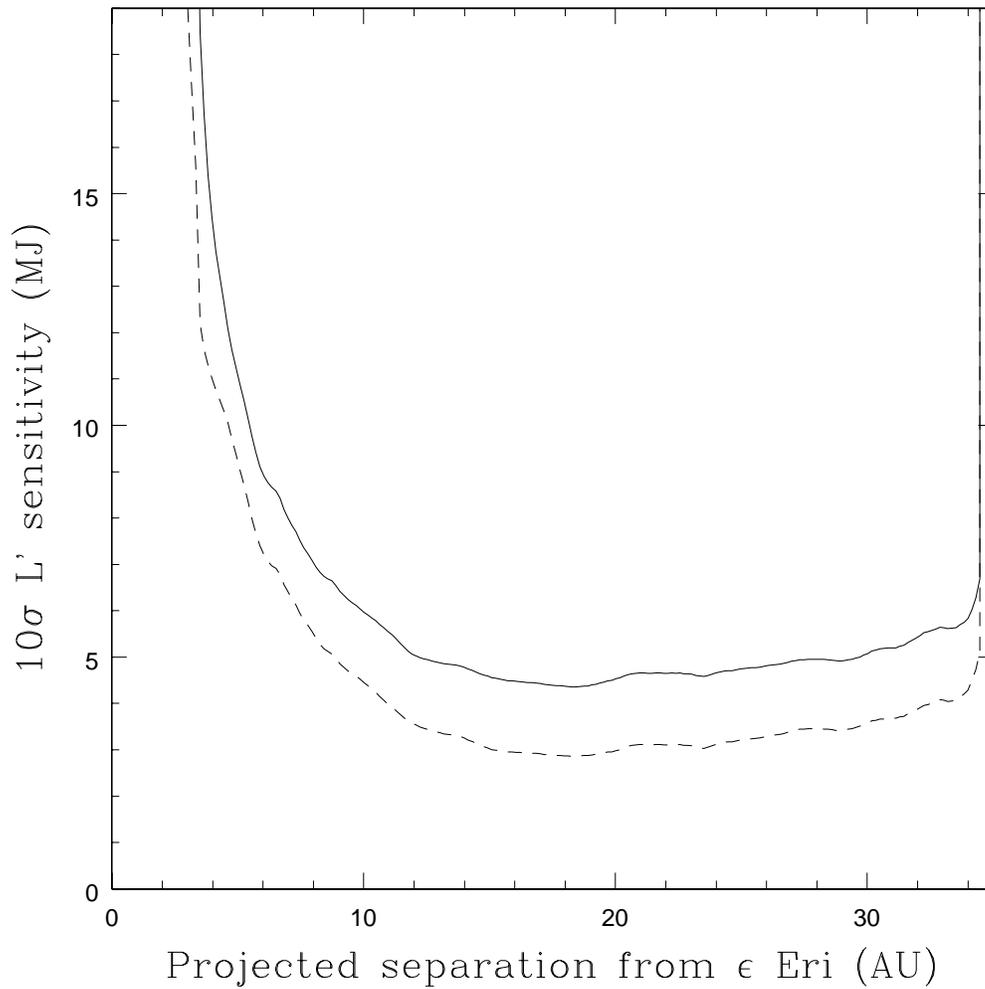


Figure 2.18  $10\sigma$  sensitivity of our  $\epsilon$  Eri  $L'$  observations in terms of minimum detectable planet mass in MJ, comparing the results from the Burrows et al. (2003) (solid line) and Baraffe et al. (2003) (dashed line) planet models. The 50th percentile, or median, sensitivity curves are shown.

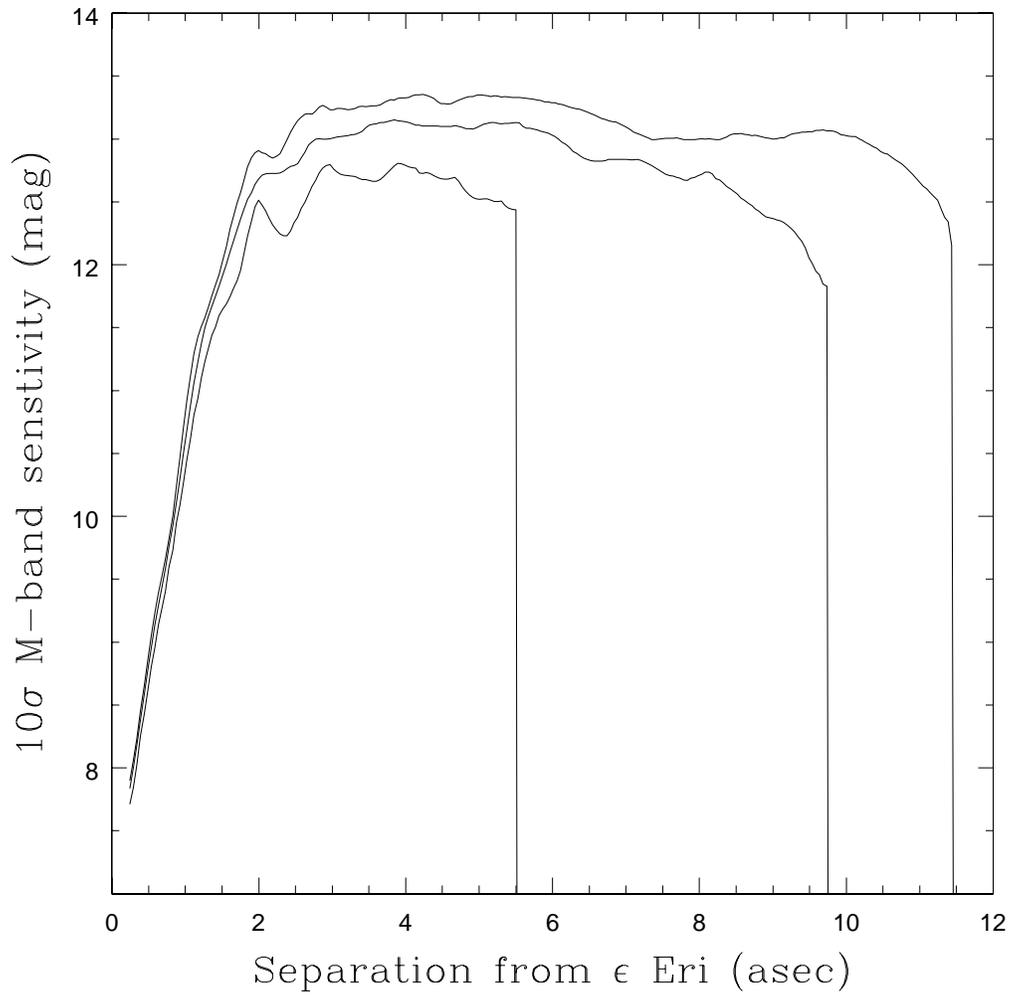


Figure 2.19  $10\sigma$  sensitivity of our  $\epsilon$  Eri  $M$  band observations in magnitudes, plotted against separation in arcseconds. The 0th, 50th, and 90th percentile sensitivity curves are shown.

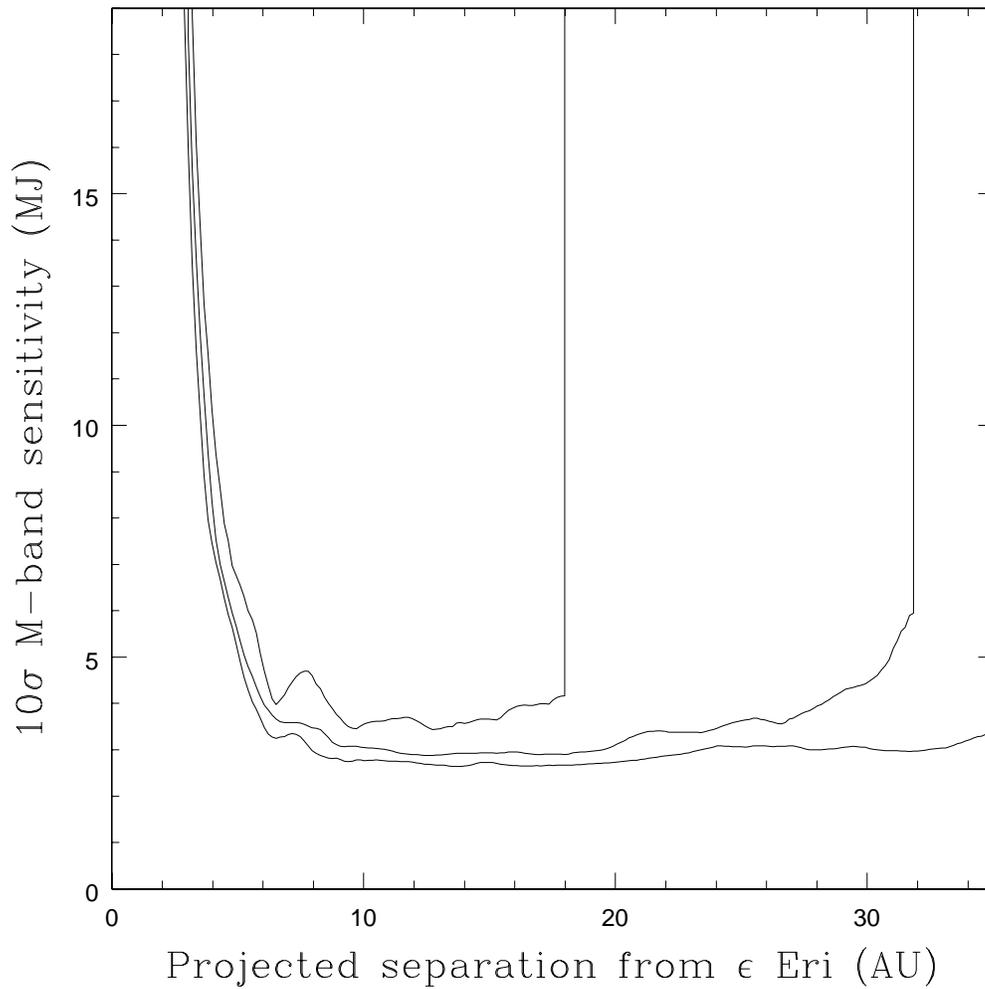


Figure 2.20  $10\sigma$  sensitivity of our  $\epsilon$  Eri  $M$  band observations in terms of minimum detectable planet mass in MJ, plotted against projected separation in AU. The magnitude-mass conversion was done using the Burrows et al. (2003) models. The 0th, 50th, and 90th percentile sensitivity curves are shown.

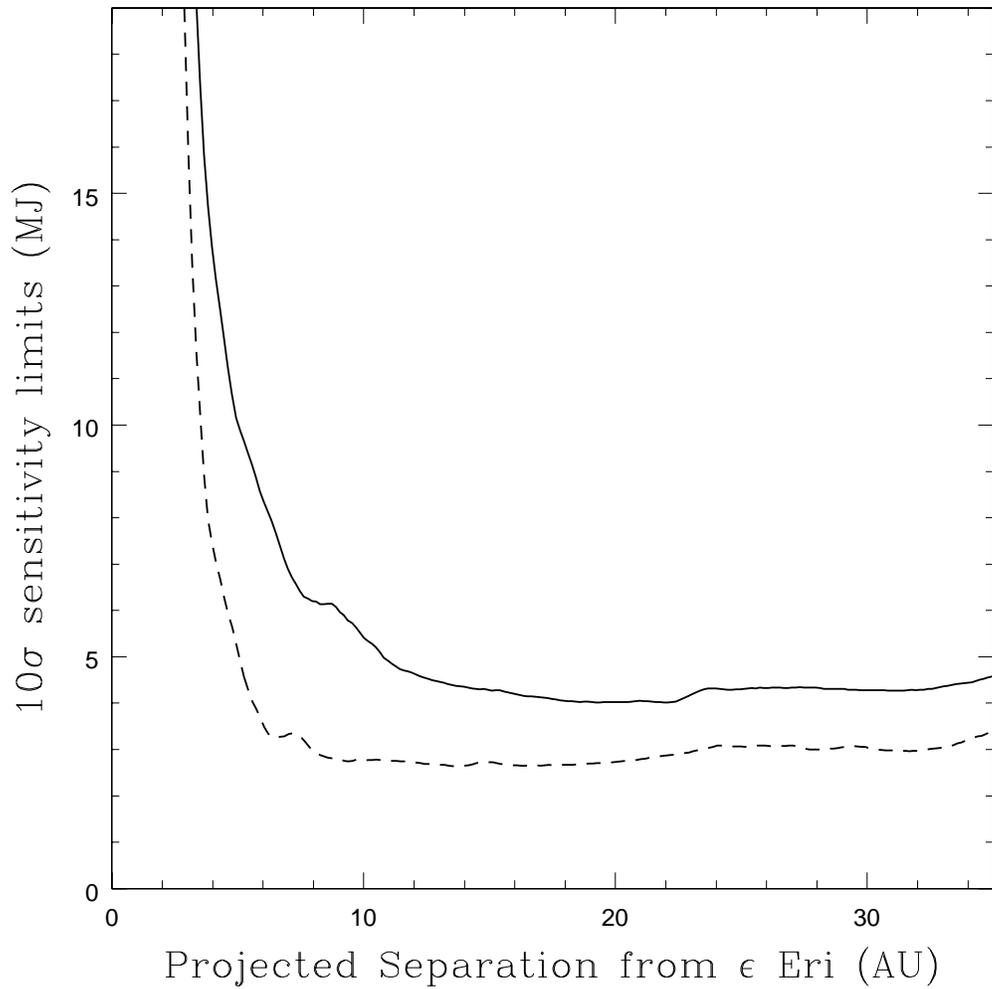


Figure 2.21  $10\sigma$  sensitivity of our  $\epsilon L'$  (solid line) and  $M$  band (dashed line) observations in terms of minimum detectable planet mass in MJ, based on the models of Burrows et al. (2003)

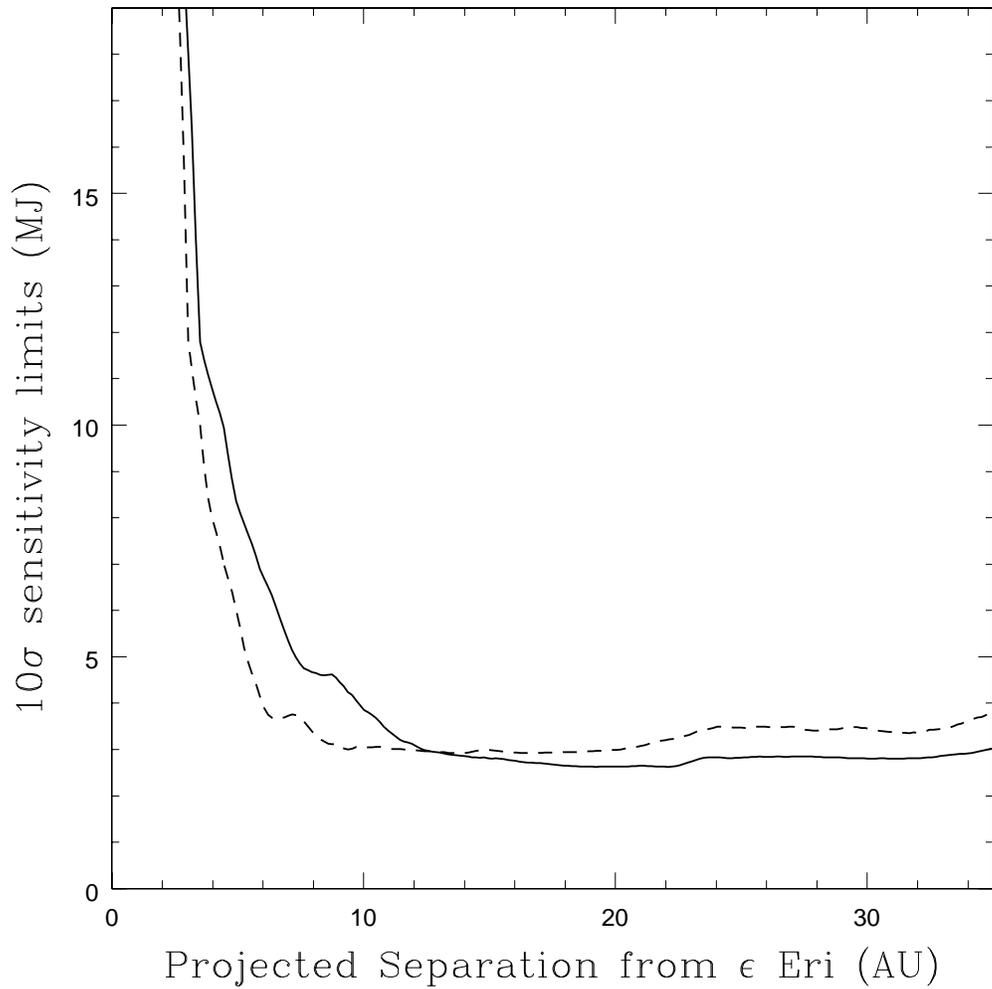


Figure 2.22  $10\sigma$  sensitivity of our  $\epsilon$  Eri  $L'$  (solid line) and  $M$  band (dashed line) observations in terms of minimum detectable planet mass in MJ, based on the models of Baraffe et al. (2003)

## 2.4 Vega: Comparison with Other Studies, and Upper Limits for Hypothetical Planets

### 2.4.1 Comparing Our Sensitivity with Other Studies

We have not attempted to compare our Vega results with an exhaustive list of all previous attempts to image planets or other faint objects around Vega. Instead, we have chosen two of the best previous results. First, the  $H$  band imaging results of Yoichi et al. (2006), and second, the narrow band,  $H$ -regime images of Marois et al. (2006). The latter almost certainly presents the most sensitive images ever acquired for faint companions at 3-10 arcsecond separations from Vega.

Before comparing our sensitivities with these other observations a word is in order about the different sensitivity estimation techniques used by the respective observers. We, as described above, used a technique that properly accounted for correlated noise. We quoted  $10\sigma$  limits without taking into account processing losses, which may mean our nominal  $10\sigma$  limits correspond to formal limits of  $7-8\sigma$  with respect to our sensitivity estimator. Our limits may still be at or above their nominal levels compared to more optimistic estimators used in other surveys.

Yoichi et al. (2006) did not calculate sensitivity limits in terms of  $\sigma$ . Instead, they calculated their sensitivities by performing numerous tests in which they placed 4 planets into their data at a fixed separation and  $\Delta$ -magnitude with respect to the primary. They set their sensitivity at each separation to the faintest  $\Delta$ -magnitude at which at least 3 of the 4 planets were recovered by their automatic detection algorithm. This is comparable to our blind sensitivity tests, except that the number and location of planets were known, and it is not clear what the false-positive rate was. In any case it would appear that the Yoichi et al. (2006) sensitivities correspond to planet brightness values at which they had at

least 75% completeness, with an unknown false-positive rate. Although they may correspond to a lower completeness and a higher false-positive rate than our  $10\sigma$  sensitivities, we have conservatively chosen to compare the Yoichi et al. (2006) sensitivity values to our own results without alteration.

Marois et al. (2006) do not explain how their quoted  $5\sigma$  sensitivity limits are obtained. We assume, however, that they used the same method as Lafrenière et al. (2007b), another planet imaging survey by a very similar set of authors, presenting observations made with the same telescope, instrument, and observing and analysis strategies. Lafrenière et al. (2007b) is the only example we know of, other than our own work, of a planet imaging survey that set  $\sigma$  limits using a sensitivity estimator that properly accounts for correlated noise. They also carefully account for processing losses, but they do not present blind sensitivity tests. Assuming that Marois et al. (2006) used the same good estimator and careful correction of processing losses, we assume that their quoted  $5\sigma$  limits are comparable to our nominal  $7\sigma$  limits. Based on this assumption we transform them to  $10\sigma$  limits for comparison with our own. We also adjust their limits by a factor of 2 (0.753 mag) in the direction of greater sensitivity, to scale from the planet-optimized narrowband filter they used to the broadband  $H$  filter. (Lafrenière et al. (2007b) estimate this correction at a factor between 1.5 and 2.5; we have used the mean value of 2.0.)

Figure 2.23 shows the sensitivities of our Vega  $L'$  and  $M$  band observations compared to those of Yoichi et al. (2006) and Marois et al. (2006). The magnitude limits, adjusted as described above, have been converted to planet masses using the theoretical planet models of Burrows et al. (2003). Figure 2.24 shows the same data with the magnitude-mass conversion performed using the models of Baraffe et al. (2003) instead. In these figures we have plotted our median sensitivity val-

ues.

Figures 2.23 and 2.24 clearly indicate that although we have obtained considerably better sensitivity to planetary-mass objects around Vega than the observations of Yoichi et al. (2006), the carefully processed narrowband observations of Marois et al. (2006) did significantly better than ours at all radii beyond 3 arcseconds, which was their approximate saturation radius. At separations between 2 and 3 arcseconds, where the other observations had no useful data, we obtained sensitivity to substellar objects ranging from very massive planets ( $\sim 10$  MJ) to low mass brown dwarfs (13-20 MJ).

We note, as before, a considerable difference in the sensitivities based on the different sets of theoretical planet models. The Burrows et al. (2003) models show  $M$  band considerably superior to  $L'$  for all separations within 5 arcseconds and indicate the Marois et al. (2006) sensitivity is better than ours by an average of about 3 MJ over the entire range for which their observations gave valid data. The Baraffe et al. (2003) models indicate that the  $L'$  sensitivity is comparable to that of the  $M$  band at small separations and considerably superior at large ones, while the Marois et al. (2006) curve bests ours by a smaller margin, about 2 MJ on average. Some of this variation might be due to small differences between the Johnson-Glass  $L'$  filter used by Baraffe et al. (2003) to produce their predicted magnitudes, and the MKO  $L'$  filter we used to integrate the Burrows et al. (2003) theoretical spectra. There are probably real differences in the underlying models as well.

To conclude, our observations attained considerably higher sensitivity than those of Yoichi et al. (2006), but the observations of Marois et al. (2006) bested ours by a consistent 1.5-3 MJ at all separations beyond 3 arcsec. Inward of 3 arcsec our images are sensitive mainly to low mass brown dwarfs, but are the

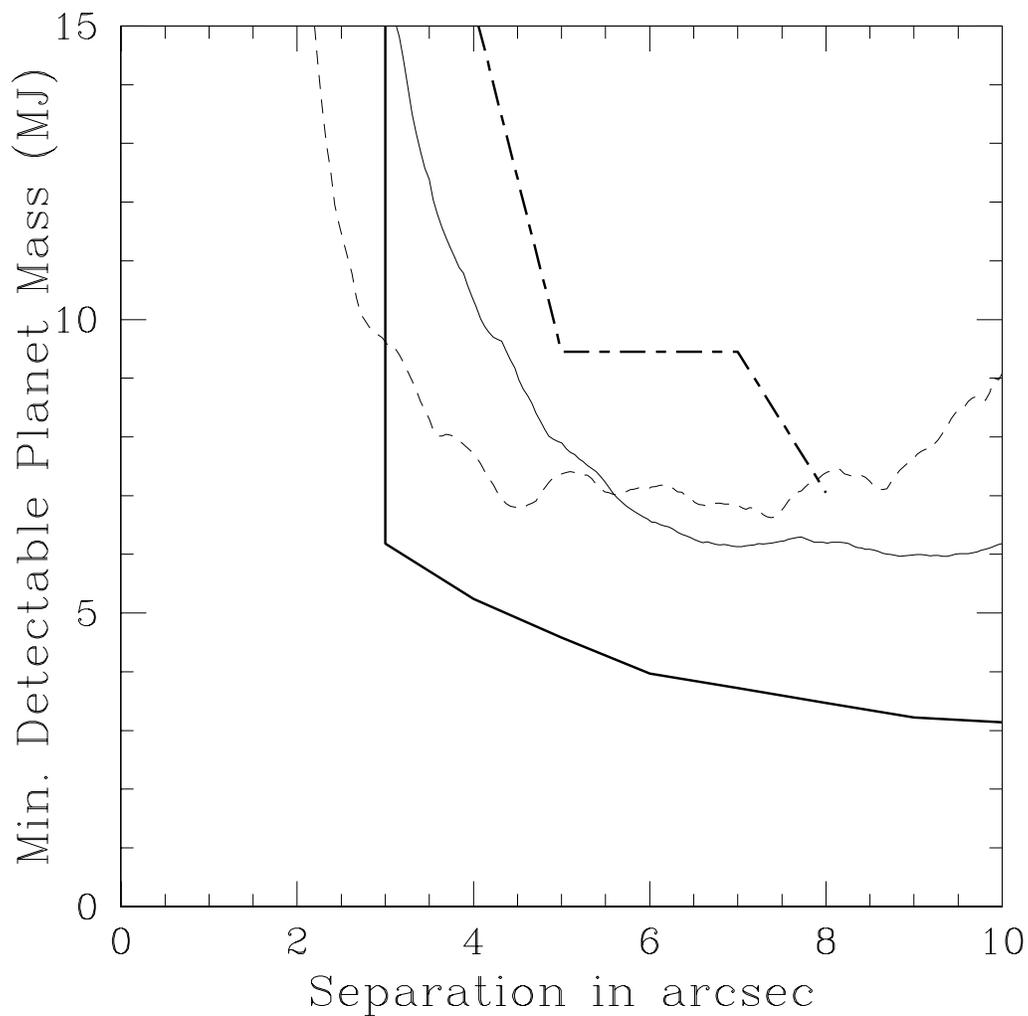


Figure 2.23 Comparison of the sensitivities obtained around Vega with different techniques. Magnitude sensitivities have been converted to planet mass limits in MJ using the theoretical models of Burrows et al. (2003). The heavy continuous line is the narrowband  $H$ -regime result from Marois et al. (2006); the heavy dot-dashed line is the  $H$  band result from Yoichi et al. (2006), the light continuous line is our  $L'$  result, and the light dashed line is our  $M$  band result.

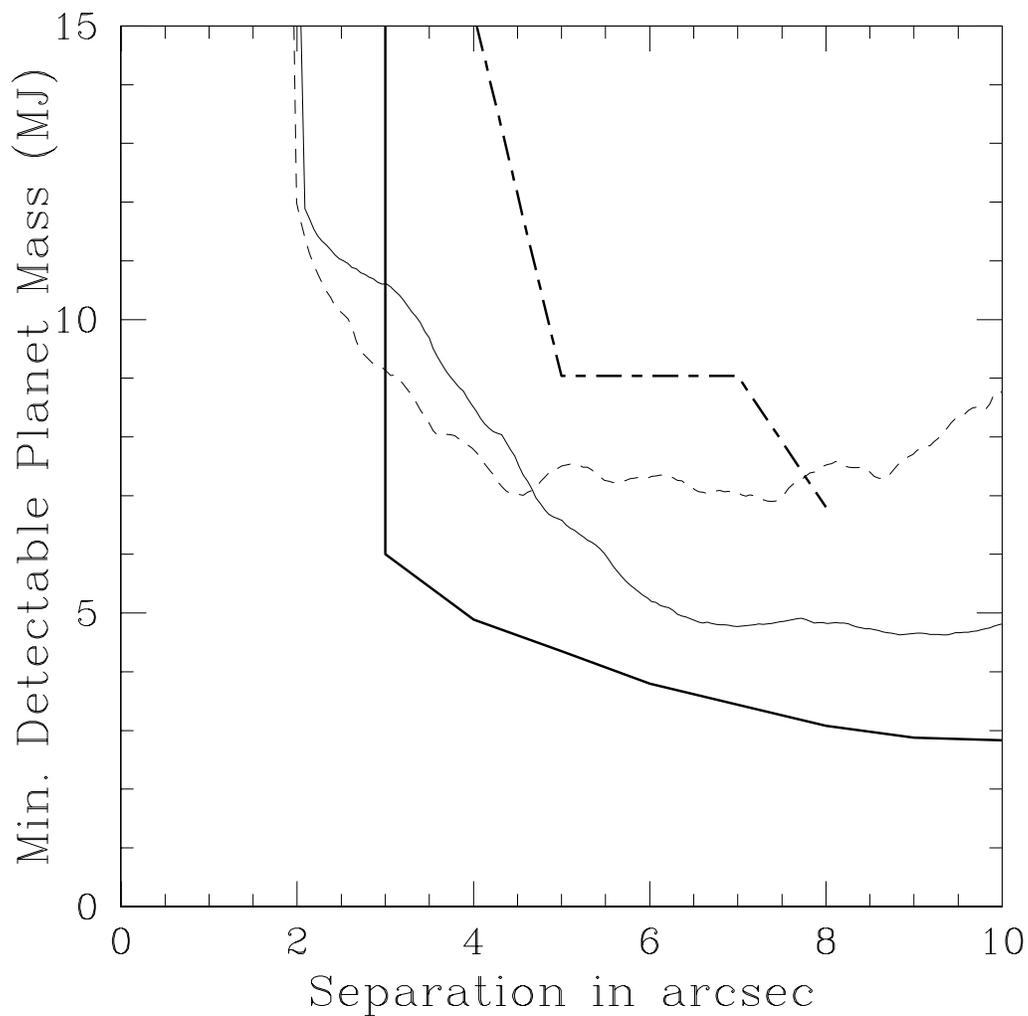


Figure 2.24 Comparison of the sensitivities obtained around Vega with different techniques. Magnitude sensitivities have been converted to planet mass limits in MJ using the theoretical models of Baraffe et al. (2003). The heavy continuous line is the narrowband  $H$ -regime result from Marois et al. (2006); the heavy dot-dashed line is the  $H$  band result from Yoichi et al. (2006), the light continuous line is our  $L$  result, and the light dashed line is our  $M$  band result.

only ones providing good sensitivity to substellar objects at all. We note that the Marois et al. (2006) observations were not more sensitive than our own by a huge margin, and that theoretical planet models are still somewhat uncertain because of the dearth of observational constraints.  $L'$  and  $M$  band observations of bright stars such as Vega make sense to diversify the investment of planet-imaging effort and hedge the overall results against the possibility that unexpected atmospheric chemistry, clouds, or evolutionary effects (see for example Marley et al. (2007)) cause planets to appear fainter in  $H$  band than current models predict. Planets could, of course, be fainter than predicted at the longer wavelengths, specifically  $M$  (Leggett et al., 2007). However, the suppression of  $M$  band flux observed by Leggett et al. (2007) applied only to objects with  $T_{\text{eff}}$  from 700-1300 K. The situation for objects cooler than 700 K is unknown. According to the Burrows et al. (2003) models, our Vega  $M$  band observations were sensitive to planets with  $T_{\text{eff}}$  below 400 K. Such objects may be too cold to have the enhanced concentrations of CO to which Leggett et al. (2007) attributed the  $M$  band flux suppression.

Because they offer better flux ratios relative to the primary star than shorter wavelengths, the  $L'$  and  $M$  bands we have used are optimal for detecting massive planets and low mass brown dwarfs at small separations from Vega and other very bright stars.

#### 2.4.2 Upper Limits at the Locations of Hypothetical Planets

Wilner (2004) presents high-resolution submillimeter observations of Vega which show two bright clumps arranged asymmetrically relative to the star. He states that is very unlikely the clumps could be background galaxies, and is essentially certain that they are concentrations of dust in the Vega system. Further, the dust could represent the remains of two different planetesimal collisions in the system, but the collisions would have to have happened fairly recently or the dust would

have dispersed. Wilner (2004) therefore concludes that the most reasonable assumption is that the clumps are dust concentrations resulting from resonant interactions between the dust and a massive planet. He shows that the observations could be explained by a 3 MJ planet in a large, eccentric orbit, which would currently be near apastron and located about 7.1 arcsec NW of the star (though the submillimeter observations were carried a few years before our imaging, a planet near apastron in such a large orbit would not move appreciably over that interval).

We chose the target position and nod direction for our Vega observations purposefully to obtain good sensitivity at the location of this hypothetical planet. The planet's location is marked on our sensitivity contour plots (Figures 2.1 and 2.2).

At approximate separation 7.1 arcsec, PA 315 degrees (due NW), our  $L'$  images of Vega gave a nominal  $10\sigma$  sensitivity of  $L' = 15.48$ , or a nominal  $7\sigma$  sensitivity at  $L' = 15.87$ . Translating these magnitudes to masses using the Burrows et al. (2003) models, and using the results of our blind sensitivity tests, we can rule out a planet at this location with a mass above 5.51 MJ with 92% confidence one more massive than 4.83 MJ with 77% confidence. Using the Baraffe et al. (2003) models instead, the masses become 4.82 and 3.49 MJ, respectively. At the same location on our  $M$  band images, we obtained nominal  $10$  and  $7\sigma$  limits of  $M = 13.67$  and  $M = 14.06$ , respectively. Using the Burrows et al. (2003) models with these magnitude limits leads us to conclude that we can rule out a planet more massive than 4.08 MJ or 2.98 MJ with 92% and 77% confidence, respectively. The Baraffe et al. (2003) models give corresponding mass limits of 4.94 MJ and 3.94 MJ. We note that no source was even suspected within an arcsecond of this location on the  $M$  band images, and therefore the confidence level for the  $7\sigma$  result is

probably higher than 77%, more like 96%, since 96% of the  $7\sigma$  planets inserted in our blind sensitivity tests were at least suspected. The upper limits we quote here are considerably better than the median sensitivities we plotted in Figures 2.23 and 2.24 because our observing strategy was optimized to give good sensitivity at this position.

We can set limits on the hypothetical planet of Wilner (2004) close to the proposed mass of 3 MJ. It would appear from Figures 2.23 and 2.24 that Marois et al. (2006) set similar or slightly lower limits, although this is not certain since they present their sensitivities only in a radially averaged sense and do not analyze them at the specific location of any hypothetical planet to set an upper limit. Observations at the  $H$ ,  $L'$ , and  $M$  bands have thus consistently set upper limits near the predicted mass of 3 MJ. A 3 MJ planet at the 0.3 Gyr age we have adopted for Vega would have  $T_{\text{eff}}$  between 300 and 400 K. No objects in this temperature range have yet been observed, so models fluxes are not observationally constrained at any wavelength. Where an upper limit from a single band would be tentative because of the uncertainties of the models, the consistent results from a range of wavelengths allow us to conclude that it is very probable no 3 MJ planet exists at the location suggested by Wilner (2004).

Wilner (2004) makes it clear that other models besides his hypothetical 3 MJ planet might explain the observed dust distribution, and that further modeling is needed to see what range of planetary orbits and masses might be capable of producing the resonant dust concentrations seen in the submillimeter. Marsh et al. (2006), for example, explain the distribution of dust they observe around Vega at 350-450  $\mu\text{m}$  wavelengths (vs 850  $\mu\text{m}$  for Wilner (2004)) by a Neptune-mass planet in a 65 AU orbit. It is not entirely clear whether their model also explains the Wilner (2004) images; however, Wyatt (2003) presents a model of a migrating

Neptune-mass planet that does match the 850  $\mu\text{m}$  images. In contrast to Wyatt (2003) and Marsh et al. (2006), Deller & Maddison (2005) present a model that explains the 850  $\mu\text{m}$  images by a 3 MJ planet in a considerably larger orbit than that suggested by Wilner (2004). It would have the same current PA as the Wilner (2004) planet (NW of the star, near PA  $315^\circ$ ), but it would be 12-13 arcsec from Vega as opposed to 7 arcsec. Our Clio observations do not obtain good sensitivity at these larger radii, though new, differently targeted Clio images could.

No current observational technique can image Neptune-mass extrasolar planets in distant orbits. The non-detections of our survey and that of Marois et al. (2006) lend some support to models explaining the Vega dust distribution using such planets rather than models such as that of Wilner (2004) in which the planet has a mass a few times that of Jupiter. However, we cannot rule out a 3 MJ planet in the more distant orbit suggested by Deller & Maddison (2005), simply because our data taking strategy was not designed to give good sensitivity at such a large separation. It is exciting that observations at  $L'$ ,  $M$  band, and the narrowband  $H$ -regime filter of Marois et al. (2006), can detect planets down to 3 MJ in the Vega system, according to theoretical planet models. More submillimeter work and orbit modeling is clearly desirable, and if it turns out that a massive planet is indeed required to account for the dust distribution, even deeper images in all 3 wavelength regimes are certainly possible.

## 2.5 $\epsilon$ Eri: Comparison with Other Studies, and Upper Limits for Hypothetical Planets

### 2.5.1 Comparing Our Sensitivity with Other Studies

As with Vega we do not attempt to compare our  $\epsilon$  Eri results with an exhaustive list of other studies, but only with a few that obtained the best sensitivity

results. We have chosen Yoichi et al. (2006), Biller et al. (2007), and Lafrenière et al. (2007b). Figures 2.25 and 2.26 show the results of the comparison.

Of these studies, the sensitivity methods of Yoichi et al. (2006) have already been discussed in Section 2.4 above, as have those of Lafrenière et al. (2007b) because we assumed Marois et al. (2006) used the same methods for their Vega data. It only remains to consider the methods of Biller et al. (2007). They use a sensitivity estimator which is based on the single-pixel RMS in 6 pixel (0.05 arc-sec, or  $1.2 \lambda/D$ ) square boxes on the images, and they quote  $5\sigma$  limits. It is not clear whether they take processing losses into account in their sensitivity calculation. In general we expect sensitivity estimators involving the single-pixel RMS to overestimate the sensitivity, as they assume independence of noise in adjacent pixels. This assumption is always violated in the speckle-dominated regions on AO images (that is, speckle noise is always spatially correlated, though the extent of the correlation depends on the details of the raw images and the type of PSF subtraction used).

The above would seem to imply that we should adjust the Biller et al. (2007)  $5\sigma$  sensitivity results by at least a factor of 2 (0.753 mag) toward decreased sensitivity in order to compare them properly against our nominal  $10\sigma$  limits. This would not include any correction for the possible overestimation of sensitivity in the presence of correlated noise. However, the ‘roll subtraction’ technique used by Biller et al. (2007) effectively creates a positive *and* a negative image of any real companion, separated by  $33^\circ$  of rotation about the primary star, and the presence of both can be used to evaluate the reality of potential sources. This doubles the data and the sensitivity should accordingly go up by  $\sqrt{2}$ . Therefore, we have considered their quoted  $5\sigma$  limits to be comparable to our  $7\sigma$  values, and scaled them to  $10\sigma$  for comparison with our  $10\sigma$  limits based on this assumption. Fur-

ther, we have adjusted the sensitivities by 0.84 mag in the sense of improving the detection limits, to convert magnitudes from the narrowband filters used in SDI (which are tuned to a predicted peak in giant planet spectra) to broadband  $H$  magnitudes.

Information we obtained by private communication from Beth Biller and Laird Close suggests that even with these various positive corrections we have underestimated the sensitivity of the Biller et al. (2007) observations relative to our own. First, their simultaneous differential imaging (SDI) technique involves two independent spectral differences. They are not necessarily equally sensitive, but in the best case this adds an additional factor of  $\sqrt{2}$  that we have not accounted for to the SDI sensitivities. Second, Beth Biller has explained to us that the Biller et al. (2007)  $5\sigma$  point-source sensitivities were calculated by comparing the single-pixel RMS noise to the brightness of the peak pixel of a PSF. This method is conservative for well-sampled data such as that of Biller et al. (2007), since it does not take into account the fact that bright pixels surrounding the peak of a PSF allow it to be detected with additional confidence. It does not overestimate the sensitivity in the presence of correlated noise. The reader should keep in mind when examining Figures 2.25 and 2.26 that we may have underestimated the Biller et al. (2007) sensitivities by a factor of  $\sqrt{2}$  (0.38 mag) or a bit more. This rather small correction will not affect any of our main conclusions.

Finally, we have converted all the magnitude sensitivities to masses using the models of Burrows et al. (2003) and Baraffe et al. (2003), and presented the results using the different model sets in Figures 2.25 and 2.26, respectively.

These figures make it clear that for  $\epsilon$  Eri, unlike for Vega, our sensitivities are significantly better than or very comparable to those of all previous observations out to a separation of 8 or more arcseconds from the star. In particular, within

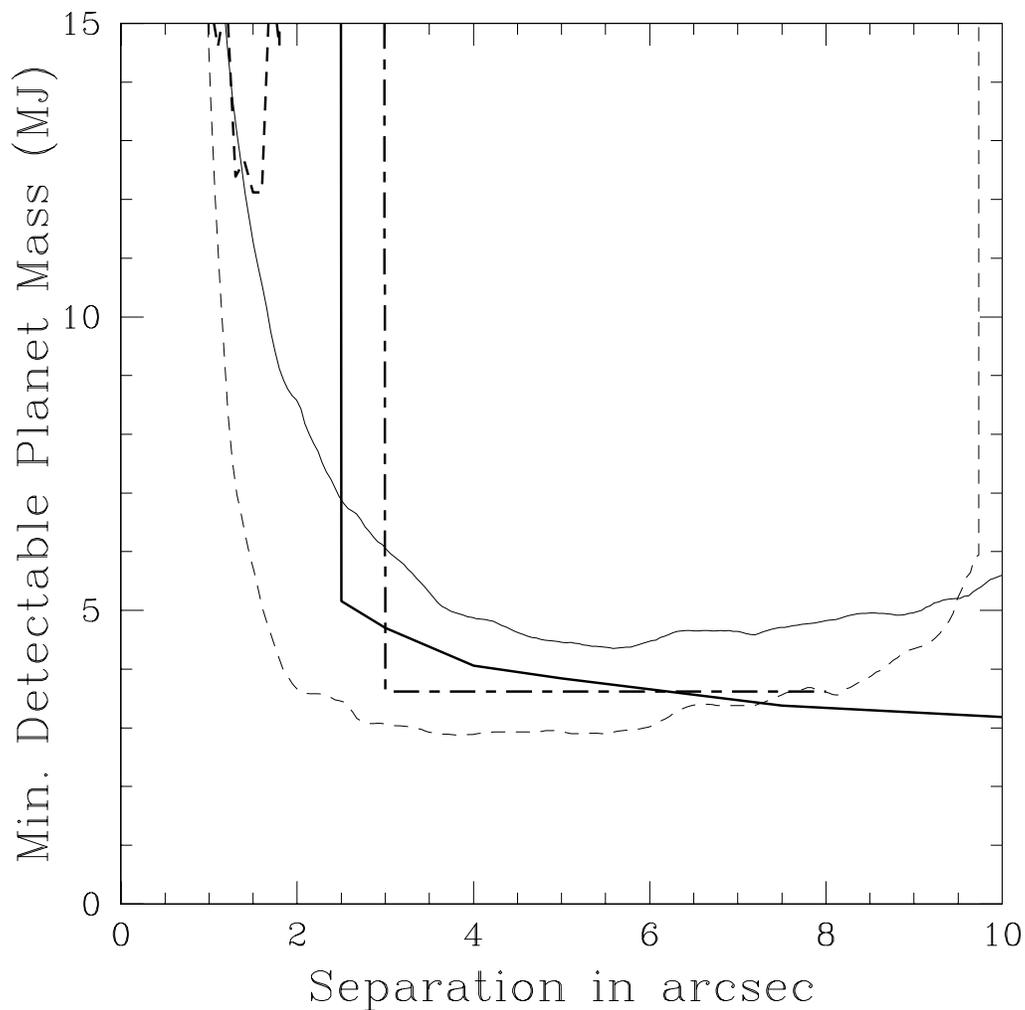


Figure 2.25 Comparison of the sensitivities obtained around  $\epsilon$  Eri with different techniques. Magnitude sensitivities have been converted to planet mass limits in MJ using the theoretical models of Burrows et al. (2003). The heavy continuous line is the narrowband  $H$ -regime result from Lafrenière et al. (2007b), the heavy dashed line is the SDI result from Biller et al. (2007), the heavy dot-dashed line is the  $H$  band result from Yoichi et al. (2006), the light continuous line is our  $L'$  result, and the light dashed line is our  $M$  band result.

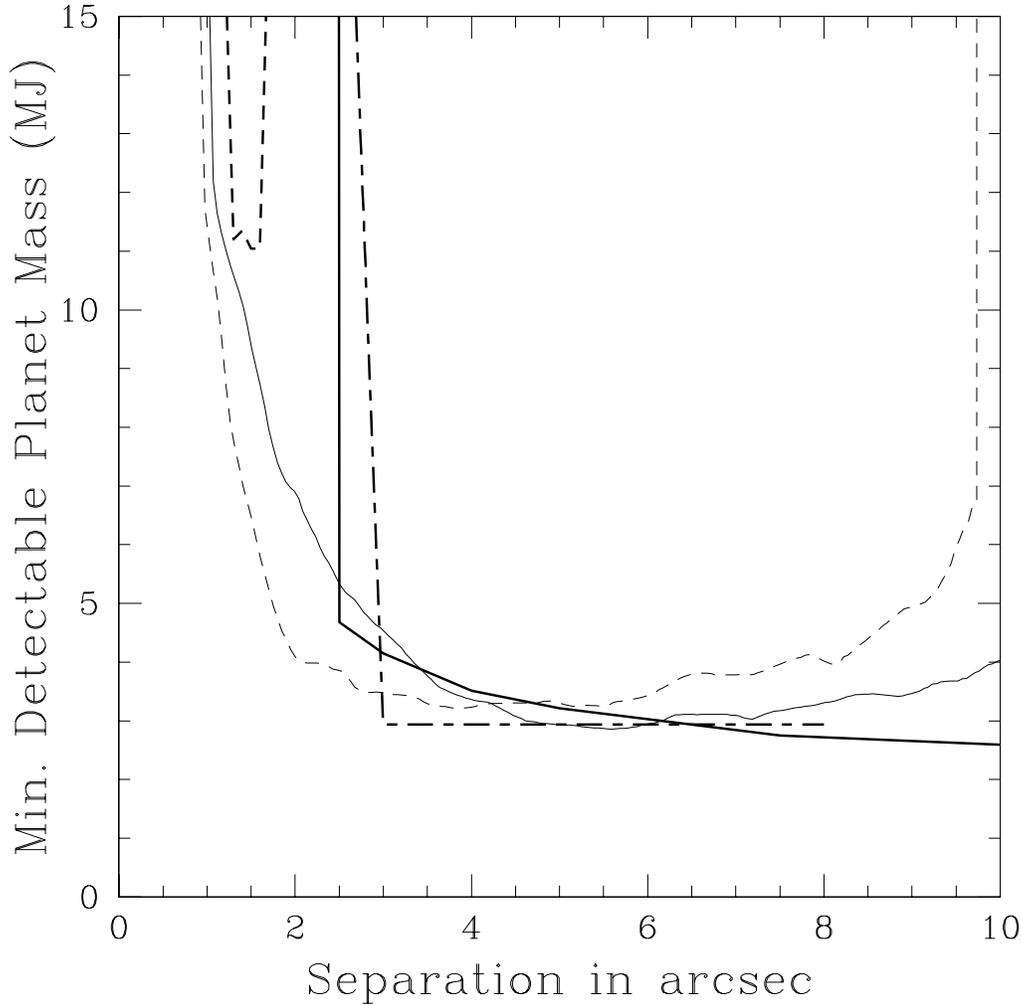


Figure 2.26 Comparison of the sensitivities obtained around  $\epsilon$  Eri with different techniques. Magnitude sensitivities have been converted to planet mass limits in MJ using the theoretical models of Baraffe et al. (2003). The heavy continuous line is the narrowband  $H$ -regime result from Lafrenière et al. (2007b), the heavy dashed line is the SDI result from Biller et al. (2007), the heavy dot-dashed line is the  $H$  band result from Yoichi et al. (2006), the light continuous line is our  $L'$  result, and the light dashed line is our  $M$  band result.

3 arcseconds we obtain significantly better sensitivity than any previous observation. Note the recurrent trend that the Burrows et al. (2003) models show the  $M$  band considerably more sensitive than  $L'$ , but the Baraffe et al. (2003) models show both as comparable. We note that the SDI method of Biller et al. (2007), which is designed to give excellent sensitivity close to bright stars, does give results comparable to our  $L'$  sensitivity, at least according to the Burrows et al. (2003) models. However, SDI is considerably bested by our  $M$  band data at all separations regardless of the model set used. This is because the planet/star flux ratio is much more favorable at  $M$  band than even in the most optimized intervals of the  $H$  band.

### 2.5.2 Upper Limits at the Locations of Hypothetical Planets

$\epsilon$  Eri has the extremely important distinction of being the one of only a few stars around which a single planet has been detected with both RV and astrometric methods (Hatzes et al., 2000; Benedict et al., 2006). This means that a complete, unique solution for the size, eccentricity, and orientation of the orbit is possible, as is a solution for the mass of the planet. Benedict et al. (2006) present such solutions, and give the mass of the planet as 1.55 MJ.

The paper is inconsistent in its presentation of the orbital solution. The orbital solution given in Table 12 is for an orbit rotated 90 degrees relative to the one shown in Figure 10. The Table 12 orbit also disagrees with the PA quoted for periastron in the appendix. We have brought this to the attention of the authors, and they inform us that the Table 12 orbit is the correct one.

At the time of our observations the orbit predicts a separation of about 0.684 arcsec. The median  $10\sigma$  sensitivity of our  $L'$  observations was  $L' = 9.77$  at 0.68 arcsec, which according to the models of Baraffe et al. (2003) corresponds to a mass of 26.84 MJ, while for the Burrows et al. (2003) models we can only say the

mass is more than 20 MJ, since the magnitude value is off the model grid. Our  $M$  band observation gave a median  $10\sigma$  sensitivity of  $M = 9.44$  at 0.681 arcsec, which corresponds to a mass of 28.04 MJ according to the Baraffe et al. (2003) models, and again to an off-grid mass above 20 MJ for the Burrows et al. (2003) models.

Since the predicted planet mass is 1.55 MJ, we do not even come close to detecting this planet. We note, however, that our observations were not timed with the idea of obtaining good sensitivity to this planet. If we had observed the planet near its apastron, at which point the separation is about 1.7 arcsec, we would have obtained median  $10\sigma$  sensitivity limits of 9.76 MJ and 7.98 MJ in the  $L'$  band, using the models of Burrows et al. (2003) and Baraffe et al. (2003) respectively, or 4.58 MJ and 5.22 MJ in the  $M$  band, again in the two different model sets. The planet would still not have been detected, unless it is far more massive than the Benedict et al. (2006) orbital solution indicates.

Could any current-technology telescope detect this planet, and if so what would be the best method?

Janson et al. (2007) applied the same SDI methodology used by Biller et al. (2007) at several different epochs for  $\epsilon$  Eri. The data from their second epoch gave them the best limit on the planet, with a  $3\sigma$  sensitivity of  $\Delta$ -magnitude 13.1 at the expected separation of the planet based on Benedict et al. (2006). As discussed above, the Biller et al. (2007) observations using the SDI method had two independent roll angles and two independent spectral differences for each observation, and the sensitivity estimation method they used was conservative. Assuming that Janson et al. (2007) used the same methodology, we will compare their  $3\sigma$  limits directly to our  $10\sigma$  limits. Note that even considering all the issues mentioned above this results in a very conservative estimation of our sensitivities relative to those of Janson et al. (2007).

We can adjust the Janson et al. (2007)  $3\sigma$  sensitivity of  $\Delta$ -magnitude 13.1 by the 0.84 mag value used before and add the  $H = 1.88$  magnitude of the star itself to get an equivalent sensitivity down to  $H = 15.8$ ; the equivalent masses are 9.6 and 9.1 MJ according to the Burrows et al. (2003) and Baraffe et al. (2003) models, respectively.

However, Janson et al. (2007) mention that the correction from the narrow-band SDI filters to  $H$  band would actually be much greater for a very cool object such as  $\epsilon$  Eri b. According to their Figure 5, the correction is about 2.2 magnitudes for the appropriate filter difference at our adopted age of 0.56 Gyr for  $\epsilon$  Eri. We cannot legitimately use this correction to obtain an upper limit mass from the Janson et al. (2007) result, because it will not apply to the minimum-mass detectable planets, which will be considerably more massive and hotter than  $\epsilon$  Eri b. However, we can use it to estimate by what factor the Janson et al. (2007) observations missed the planet.

Their sensitivity in their narrowband filter was about  $13.1 + 1.88 = 14.98$  mag, assuming that the magnitude of  $\epsilon$  Eri A is the same in the narrowband filter as in broadband  $H$ . According to the models of Burrows et al. (2003), an 0.56 Gyr-old planet of mass 1.55 MJ located 3.27 pc away has an  $H$  band magnitude of about 28.5, while the Baraffe et al. (2003) models give it  $H = 24.7$ . We correct these to narrowband magnitudes using the 2.2 mag correction, and difference them with the sensitivity of 14.98 mag. The conclusion is that the Janson et al. (2007) sensitivity was insufficient to detect the planet by 11.3 magnitudes (a factor of 34,000) under the Burrows et al. (2003) models, or 7.5 magnitudes (a factor of 1,000) under the Baraffe et al. (2003) models. Note the huge divergence in the models'  $H$  band brightness predictions for this low mass, old planet. Though the differences in filter sets used (MKO  $H$  for our integration of the Burrows et

al. (2003) models vs CIT  $H$  for the Burrows et al. (2003) models) might account for a small difference in model predictions, this very large difference (a factor of 34) is almost certainly due to large inherent differences in the models. This is especially true because the  $H$  band includes a strong peak in theoretical planet spectra Burrows et al. (2003), and therefore variations in the band should not affect the included flux as much as if it sat on a steep slope, as is the case for the  $L'$  band.

The factors quoted above indicate the SDI sensitivity would have to be increased by a factor of at least one thousand to detect the planet. Assuming we had observed the planet at apastron, by what factor would we have failed to detect it? We will consider only our  $M$  band results, as they are more sensitive to low mass planets close in to the star. Our median  $10\sigma$  sensitivity at the apastron separation of 1.7 arcsec was  $M = 12.28$ . The models of Burrows et al. (2003) give the brightness of the planet as  $M = 14.7$ , while those of Baraffe et al. (2003) give it as 15.3. This means we missed it by 2.42 mag (a factor of 9.3) under the Burrows et al. (2003) models, or 3.02 mag (a factor of 16) under the Baraffe et al. (2003) models. These much lower miss factors suggest that  $\epsilon$  Eri b might actually be detectable near apastron with ground based  $M$  band imaging. Note also that the models are much closer to agreement in the  $M$  band than at  $H$ , which suggests they are more reliable. It is likely that  $\epsilon$  Eri b is at too low a  $T_{\text{eff}}$  for its  $M$  band flux to be dimmed by the above-LTE CO concentrations suggested by Leggett et al. (2007) and Reid & Cruz (2002) to account for the suppressed  $M$  band flux observed for much hotter objects. We note also that even if the suppression of  $M$  band flux remained,  $M$  would still far better for the detection than the  $H$  band.

The next apastron of  $\epsilon$  Eri b is in 2010, and very deep  $M$  band observations with current technology might possibly succeed in imaging it then. We remind

the reader, however, that even in the best case the sensitivity of speckle-limited observations goes down only as the square root of the exposure time. Thus, even under the Burrows et al. (2003) models, an exposure  $9.3^2$  times as long as our  $M$  band integration would be required to detect the planet using Clio at the MMT. This corresponds to 82 hours of observing, or more than ten good nights. Only an exceptionally persuasive astronomer is likely to convince a TAC to assign this much time for a single target! We note that as far as we know Clio is the only currently operating instrument capable of the long, high-efficiency  $M$  band integrations required for such a project. Improved PSF subtraction techniques, or a coronagraphic option (Kenworthy et al., 2007) might reduce the required exposure for Clio to detect  $\epsilon$  Eri b. A multi-night integration would still be required.

The observation might be more practicable with the LBT, or with the next generation of giant telescopes such as the GMT, TMT, or E-ELT, provided they are equipped with adaptive secondary AO systems. It is also possible that  $\epsilon$  Eri b will be detected first from space. It might be detected at  $L'$ ,  $M$  band, or longer wavelengths using JWST, or it could be detected in reflected light at visible wavelengths by a sensitive space-based coronagraph.

Ozernoy et al. (2000) and Quillen & Thorndike (2002) suggest that the dust disk of  $\epsilon$  Eri has been sculpted by a planet of 0.1-0.2 MJ in an orbit between 40 and 65 AU in radius. Deller & Maddison (2005) agree, and prefer the model of Quillen & Thorndike (2002). Such a planet would be far too faint to detect with any telescope in the near future. However, Deller & Maddison (2005) state that an additional,  $\sim 1$  MJ planet in a smaller orbit is likely required to produce the observed clearing of the dust inside of about 30 AU (Greaves et al., 1998). The RV/astrometric planet of Hatzes et al. (2000) and Benedict et al. (2006) has too small an orbit to account for this dust clearing; Deller & Maddison (2005) suggest

orbital radii between 10 and 18 AU for the hypothetical planet. Benedict et al. (2006) mention a long-term trend in RV measurements for  $\epsilon$  Eri A that might indicate just such a planet: a  $\sim 1$  MJ object orbiting with a period longer than 50 years. Since such a planet would probably appear at least 3-4 arcsec from the star, we would likely have detected it if it had a mass of 4-5 MJ or greater, as would the Lafrenière et al. (2007b) observation. Since the mass is expected to be closer to 1 MJ, it is not surprising that our observations and those of Lafrenière et al. (2007b) did not detect it. This planet might be detected serendipitously in the course of a very long exposure to image the known RV/astrometric planet.

## 2.6 Conclusions

We have taken very deep  $L'$  and  $M$  band images of the interesting debris disk stars Vega and  $\epsilon$  Eri to search each system for orbiting planets and brown dwarfs. For both stars we obtained better sensitivity than shorter-wavelength observations at small separations from the star. Our  $\epsilon$  Eri observations were more sensitive than  $H$  band-regime observations of the same star by a greater margin than applied for our Vega images. The  $\epsilon$  Eri SDI observations presented by Biller et al. (2007) and Janson et al. (2007) are an exception to this (no SDI observations have been performed on Vega). However, even though the SDI observations yielded remarkably good sensitivity well within the saturation radii of the other  $H$  regime  $\epsilon$  Eri images, we still surpassed them at all separations because of the far better planet/star flux ratio at the  $L'$  and  $M$  bands.

The reason our  $\epsilon$  Eri observations have a greater sensitivity advantage over  $H$  regime observations than do our images of Vega is the small distance to the  $\epsilon$  Eri system. For  $\epsilon$  Eri, the sensitivity at any wavelength extends down to fainter, cooler planets than for Vega. The Burrows et al. (2003) and Baraffe et al. (2003)

models indicate the  $H - L'$  and  $H - M$  colors of cool giant planets are redder than those of hotter ones. Thus, the faintest detectable objects in the  $\epsilon$  Eri system would be considerably redder than those in the more distant Vega system, giving  $L'$  and  $M$  band observations a greater advantage for the nearer system. Another way of saying this is that the planet/star flux ratio advantage of the  $L'$  and  $M$  bands vs the  $H$  band is greatest for the nearest stars.

Planet-search observations at the  $L'$  and  $M$  bands have a considerable advantage over those in the more commonly used  $H$  band regime for  $\epsilon$  Eri and a handful of other bright, nearby stars. For slightly more distant bright stars such as Vega,  $L'$  and  $M$  band observations give markedly better results only at separations inside about 3 arcsec, and in this regime no currently employed method gives sensitivity to any but the highest mass planets. Observations in the bands we have employed are still useful on Vega, but their use tends toward a diversification of planet-search effort in case theoretical models are overpredicting planets'  $H$  band brightnesses. For nearer systems such as  $\epsilon$  Eri, by contrast,  $L'$  and  $M$  band observations clearly provide the best sensitivity at the most interesting separations, and it is the  $H$  regime images that naturally take the role of diversifying effort under the supposition that the models may overpredict planet brightness at *longer* wavelengths.

We have set a limit on the Vega planet hypothesized by Wilner (2004) that is close to the 3 MJ mass he suggested for it. It appears that Marois et al. (2006) could set a similar limit. We suspect, therefore, that no 3 MJ planet exists near the location Wilner (2004) suggested for it. More sensitive observations at  $L'$ ,  $M$  band, and  $H$  regime wavelengths would be needed to be absolutely sure. More submillimeter observations and further modeling to find the best explanation for the dust clumps will help guide future planet-imaging attempts. If a 3 MJ planet

does indeed exist at a large separation somewhere in the Vega system, it certainly could be imaged using our  $L'$  and  $M$  band observing strategy or the  $H$  regime strategy of Marois et al. (2006). A confident prediction of such a planet based on submillimeter observations and dust models would thus be extremely exciting.

$\epsilon$  Eri has a known planet detected by both radial velocity and astrometry. The planet is in a relatively small orbit, and at 1.55 MJ has quite a low mass. Still, it gets to large enough separation near apastron that it is tantalizingly close to the detection threshold of sensitive AO imaging.

We did not time our  $\epsilon$  Eri observation to catch the planet at a large separation, and as a result set limits on its mass in the range of 20-30 MJ, which are by no means interesting. Janson et al. (2007) observed it at several epochs of more promising separation using SDI, and set limits in the 9-11 MJ range.

We have explored the question of whether SDI or  $L'$  and  $M$  band imaging is the method is most likely ultimately to detect the planet. Our  $M$  band images were much more sensitive at small separations than our  $L'$  results, so we have not considered the latter. We find that the Janson et al. (2007)  $7\sigma$  sensitivity at the best epoch, where the planet was near the optimal separation for SDI imaging, was still insufficiently sensitive to detect the planet by a factor of more than 2000, even assuming the more optimistic of the two model sets. The less optimistic Burrows et al. (2003) models indicate a detection had been missed by a factor near  $10^5$ , and the large divergence in predictions indicates that  $H$  band models are very uncertain for such a low mass, relatively old (0.56 Gyr) object. **By contrast our observations, if carried out at apastron, would have missed the planet by a factor of only 9.3 or 16 under the two different model sets, respectively.**

This striking result strongly indicates that if the models of Burrows et al. (2003) and Baraffe et al. (2003) are accurate it is at  $M$  band that the planet  $\epsilon$

Eri b will first be imaged. In fact, an extremely ambitious 10-20 day observing campaign **using Clio at the MMT might detect it during the 2010 apastron passage**, since we have observed the sensitivity in the speckle-dominated regions of  $M$  band images does go up approximately as the square root of the exposure time. More advanced PSF subtraction, or coronagraphic capability in Clio (Kenworthy et al., 2007), might reduce the required exposure time to detect the planet to as little as 3-5 nights. The detection could be accomplished in a much shorter time with larger telescopes such as the LBT and GMT that are planned to be equipped with adaptive secondary AO systems. An  $L'$  and  $M$  band imager called LMIRCam is planned for the LBT (Wilson et al., 2007). Conventional AO systems probably have too high a thermal background for such deep  $M$  band images. At present, we believe **Clio with MMTAO is the only system capable of deep planet imaging integrations in the  $M$  band**. Spitzer, despite its enormously lower background and correspondingly excellent sensitivity, does not have sufficient resolution to detect objects at the separations expected for orbiting planets.

We have set a limit of 4-5 MJ for planets in distant orbits around  $\epsilon$  Eri such as the one that Benedict et al. (2006) suggest is indicated by a long term RV trend. They do not appear to think the planet is as massive as this, so it is not surprising that we have not detected it.

## CHAPTER 3

THE SURVEY: DATA ACQUISITION, DATA QUALITY, CONFIRMED AND  
SUSPECTED FAINT SOURCES

## 3.1 Introduction

We have carried out an  $L'$  and  $M$  band AO imaging survey of 50 nearby stars for orbiting extrasolar planets. The  $L'$  and  $M$  bands are longer, less used wavelength intervals (see Table 3.1 for a listing of the main astronomical IR bands) that nevertheless have considerable potential for the detection of extrasolar planets, especially with the upcoming generation of very large telescopes. The science objective of our survey is to place constraints on the statistical distributions of extrasolar planets with orbital periods too long for current radial velocity (RV) observations to detect them. In particular, we seek to determine whether a simple extrapolation to the power-law statistical distributions observed to describe known RV planets would apply to planets orbiting much farther from their parent stars. We have designed the survey with the intent of insuring that even a null result would rule out some distributions consistent with fits to the known RV planets. We did in fact obtain a null result, with no confirmed planets nor strong candidates at the end of the survey.

The constraints our null result places on the distributions of extrasolar planets orbiting at large distances from their parent stars are discussed in Chapter 4.

In the case of a planet imaging survey null result, the science hangs on the upper limits set on planet mass versus projected separation around each star surveyed. The methods used for estimating and verifying the sensitivity of the observations (from which the upper limits come) are therefore extremely critical. As planet imaging surveys of increasing sensitivity have turned up one null result

Table 3.1. Astronomical Infrared Filter Bands

Band Name	Central Wavelength ( $\mu\text{m}$ )	Band FWHM( $\mu\text{m}$ )
<i>J</i>	1.215	0.26
<i>H</i>	1.654	0.29
<i>K<sub>s</sub></i>	2.157	0.32
<i>K</i>	2.179	0.41
<i>L'</i>	3.761	0.65
<i>M</i>	4.769	0.45

Note. — The values in this table are taken directly from Table 7.5 in Cox (2000).

after another, the care with which sensitivity estimates are made and discussed has increased. However, in many cases it is still not entirely clear how much confidence may be placed in the sensitivity estimates and the planet-mass upper limits they imply. Our methods for estimating these are exceptionally careful, conservative, and rigorous. We present them here, as a contribution to the field at least as important as the science results we announce in the following Chapter.

Many of the stars that we have observed will certainly be investigated in the future in the course of other, more sensitive planet imaging surveys using the next generation of AO-equipped telescopes and IR cameras — and possibly even JWST. We have tried to provide as much information of potential use to such future workers as we can. This information divides roughly into four categories.

First, we present the sensitivity we obtained around each star in terms of magnitudes. This is important because future observers can convert the magnitude information to equivalent planet masses using new planet models and stellar age information that may be considerably superior to what is available at present. Since negative nod-images, ghosts, rays, and other artifacts caused considerable azimuthal variations in our sensitivity around each star, we present the sensitivity in the form of contour images with an astrometric grid superposed. This insures that important spatial information will not be lost, as it might be if we simply presented radial plots. Planet-imagers of the future can see at a glance what our sensitivity was at the locations of any new candidate sources. If they are using a different wavelength, our sensitivity will allow them to set a limit on the object's color — often a very powerful tool to distinguish between background stars and real planets. If they detect a bright object very close to the star, our non-detection of it at a larger projected separation may rule out the background hypothesis and provide strong evidence it is a proper motion companion.

Second, we present plots of the mass sensitivity we obtained around each star. These will allow future observers to see at a glance what limits we set, without needing to perform a magnitude–mass conversion themselves. Perhaps the designers of future surveys may even decide to exclude stars around which we obtained very good upper limits and saw nothing.

Thirdly, we present images and rough astrometry and photometry of all of the faint objects we detected around the stars we investigated. These have all been determined not to be planets. Mostly they were background stars; we also discovered a new low–mass stellar companion to one of our survey targets, and noted the previously discovered binary brown dwarf orbiting the star GJ 564. Our catalog of confirmed faint objects near the bright stars of our survey will allow future workers to avoid planet false-alarms due to background stars. The background stars can also serve as test objects for newly commissioned instruments, and may even be of some use eventually to refine the proper motions of the bright target stars.

Finally, we have catalogued a huge list of suspected sources in our data. It is certain that the vast majority of these are not real. However, if future, more sensitive surveys do turn up real objects corresponding to some of these suspected detections, our coordinates and brightness may provide a rough color, and/or aid in confirming or denying the object as a proper motion companion.

We motivate the use of the  $L'$  and  $M$  bands for planet imaging work in Section 3.2. In Section 3.3 we describe how we selected our target stars, and compare our sample to those of four other sensitive AO imaging surveys. In Section 3.4, we describe the Clio instrument with which all our survey observations were carried out. In Section 3.5, we outline our observational techniques and describe the data we obtained. Section 3.6 describes how we process the data (see Appendix A for

more detail). The crucial Section 3.7 describes how we calculate our sensitivity to faint objects in our images. We compare our methods to those used in the other surveys, and show that the sensitivities quoted in several other survey papers are probably overestimated by 0.5 to 2.0 magnitudes. The constraints these surveys place on the abundance of planets should be adjusted using rigorously verified sensitivity estimators. Since the surveys were not sensitive to planets as low in mass as was thought, the adjustments will of course be in the sense of relaxing the upper limits placed on the abundance of massive planets. We show how we have carefully calibrated and verified the accuracy of our own sensitivity estimators using blind tests with fake planets input into the raw data at random locations. In Section 3.8, we present the actual sensitivity information for each star, in the form of a magnitude contour image, and a mass sensitivity plot. In Section 3.9 we discuss how our sensitivities compare to those obtained by other surveys. In Section 3.10 we present measurements of all confirmed and suspected faint sources imaged near our target stars. We offer concluding comments in Section 3.11.

### 3.2 Motivation for an $L'$ and $M$ Band Survey

Why search for planets at the  $L'$  and  $M$  bands? These are significantly longer wavelengths than  $H$  and  $K_s$  bands more typically used for AO imaging studies, and they are difficult bands to work with for several reasons. The sky background at the  $L'$  and  $M$  bands is far brighter. Any planets must be detected against the strong photon shot noise of this bright background. Detector technology is less advanced, and the available detectors have smaller formats, higher readnoise, and higher dark current than those used for shorter wavelength observations.

**Why bother with the  $L'$  and  $M$  bands?**

**The first reason** comes from theoretical models of extrasolar planet spectra. Both long and short wavelength IR planet imaging surveys seek to detect giant planets by the IR glow that comes from their own internal heat — heat which results from the conversion of gravitational potential energy to thermal energy in the planets' initial accretion and subsequent slow contraction. As giant planets age, they cool and become fainter. However, theoretical models indicate that the flux drops more quickly at shorter wavelengths, while at longer, 'redder' IR wavelengths planets remain bright for longer (Burrows et al. (2003), Baraffe et al. (2003)), somewhat as a cooling coal glows first orange and then an increasingly deeper red as it fades. The *L'* and *M* bands may be far superior to shorter wavelengths for imaging older planets — and 'older' in this context may mean anything but very young (ie 0-100 Myr old). Older planets are interesting for several reasons.

First, they are certainly far more common in the universe. If we make a rough approximation of a constant birthrate, the mean age of solar-type stars should be half their main sequence lifetime, or about 5 Gyr. Very young stars are rare; middle-aged stars are the norm. Most very nearby stars are not very young. Planets around the closest stars may be detectable only at long wavelengths. Space-based coronagraphs might also detect some such planets by visible light reflected from their primary star; but exceptional coronagraphic performance would be required and planets in very distant orbits might be too faint to detect. Planets orbiting very nearby stars are inherently interesting both because the nearness of the stars makes a broad range of followup observations possible, and because in the very long term nearby star systems may become the first destinations for interstellar explorers from Earth.

Secondly, planetary systems up to ages of several hundred Myr may still be

undergoing substantial dynamical evolution due to planet-planet interactions (Juric & Tremaine, 2007; Gomes et al., 2005). Discovering dynamically evolving systems would be fascinating, but we also want a probe for systems old enough to have settled down into a mature stable configuration. Longer wavelength observations may provide this.

Finally, theoretical models of older planets are likely more reliable than for younger ones, as these planets are further from their unknown starting conditions and moving toward a well-understood, stable configuration such as Jupiter's. It has been suggested by Marley et al. (2007), in fact, that theoretical planet models such as those of Burrows et al. (2003) and Baraffe et al. (2003) may overpredict the brightness of young ( $< 100$  Myr) planets by orders of magnitude, while for older planets the models are more accurate.

**A second reason to use the longer wavelengths** is that models indicate that planets are brighter in these wavelengths than at shorter ones at any age (Burrows et al. (2003), Baraffe et al. (2003)). In particular, the planet/star flux ratio is more favorable at the longer wavelengths, regardless of the planet's age. This consideration favors the  $L'$  and  $M$  bands especially in the case of very bright stars, where the planet must be detected against an intense stellar halo, and choosing a wavelength that delivers a good planet/star flux ratio is extremely important.

**Thirdly**, observations at longer wavelengths are desirable simply to diversify the investment of effort in extrasolar planet searches. If observations are performed in only one wavelength regime, we risk coming to the wrong conclusions if extrasolar planets have unexpected atmospheric chemistry, clouds, or evolutionary effects that cause them to be far fainter in the chosen wavelength regime than expected. Theoretical models are still quite uncertain on these points for typical giant planets. The only planet candidates discovered so far have been distant

from their host stars, extremely young, and hot, in many cases straddling the planet/brown dwarf boundary. It is not clear that they are at all representative of the galactic population of planets in ‘normal’ orbits around solar-type stars; they may be much more akin to brown dwarfs, or they may actually be brown dwarfs. The giant planets in our own Solar System are far too cool and faint to be seen at interstellar distances, and therefore are not good examples of the kind of extra-solar planets to which current imaging surveys are sensitive. Thus there are no significant observational constraints on giant planet models so far. Models using similar assumptions such as those of Burrows et al. (2003) and Baraffe et al. (2003) still make substantially different predications about the relative usefulness of different wavelength bands, as will be seen in Section 3.7, while other work such as Marley et al. (2007) calls the shared assumptions into question. In this situation it is desirable to observe using a wide variety of techniques, including different wavelengths, to increase the likelihood of discoveries and the confidence with which null results constrain planet distributions.

**Finally**, AO systems perform better at longer wavelengths than at shorter ones. AO images taken in the  $L'$  and  $M$  bands approach the theoretically perfect images that an idealized version of the telescope would deliver in space much more closely than the images the same systems deliver at shorter wavelengths.

This list of the excellent properties of the  $L'$  and  $M$  bands may make it appear that they are superior regardless of the properties of the star system being examined. This is not the case, because of several balancing concerns.

First, the high sky background at the longer wavelengths, which has already been mentioned. Photon statistical noise (and/or read noise) is the limiting factor for long or short wavelength IR planet search observations seeking to detect planets at large angular separations from the host star. Since the thermal sky

background is far more intense at  $L'$  and  $M$ , so is the photon noise. Against this is balanced the greater brightness of planets at the longer wavelengths. Which wins out? Is the background so much fainter at the shorter wavelengths that they are superior for the background-limited case, even though the planets are also fainter? To answer this question, we must first be clear what we mean by a superior method of planet detection: given a star system of known age and distance, and an angular separation from the host star, the best method is the one that could detect the lowest mass planet in that system at that angular separation. The lower the mass of a planet, of course, the less internal energy it had to begin with, the cooler and fainter it is at any age, and the harder it is to detect.

Consider now a distant, young star system. Because it is far away, only fairly bright planets can be detected with any imaging method. However, because the system is so young even fairly low mass planets would still be hot and bright, having had little time to cool (otherwise, ie if the system were old enough that even the most massive planets had cooled below the limits of detectability at its large distance, there would be no point in observing it). Low mass planets can in principle be detected in this system, but only because they are hot. Theoretical models indicate that the near IR vs thermal IR colors (ie  $H - L'$  and  $H - M$ ) of hot extrasolar planets are not very red. Thus, **for a distant system in which only hot planets could be detected, the advantage of longer wavelengths is small or nonexistent.**

Shift focus to a very nearby star system. It is so close that intrinsically very faint planets can be detected. Theoretical models indicate that such planets have low effective temperatures and are very red in near IR vs thermal IR colors. Because of the red colors, the advantage of longer wavelength observations will be large. It is interesting to note that because low temperature objects can be de-

tected in nearby star systems and only hotter ones can be found around distant stars, the advantage of longer wavelength observations relative to shorter wavelengths depends almost solely on the **distance** to the system, **not the youth or age**. The mass of the faintest detectable objects will depend sensitively on the system age; their color will not.

Even for the nearest systems, is this advantage sufficient to cancel out the higher sky background noise? For the survey observations we present here, the answer to this question is not clear. This is in part because of an oddity in the theoretical models of extrasolar planet spectra. While models predict that planets rapidly become impossibly faint in the *Ks* band as they age and cool, they are predicted to fade much more slowly at the shorter wavelength *H* band (Burrows et al. (2003), Baraffe et al. (2003)). That is, while they are becoming ‘redder’ in terms of the *L'* and *M* bands compared to the shorter wavelengths in general (ie, redder in *L' – H* or *M – Ks* color, etc), they actually become ‘bluer’ in terms of the *Ks* band compared to its shortward neighbor *H* ( that is, in *H – Ks* color). The cause of this, in brief, is that as planets cool methane and other new molecules form in their atmospheres, and the spectral absorption lines of these molecules inhibit the ability of IR radiation to escape from the interior of the planet. The result is that the thermal glow comes out preferentially in certain ‘spectral windows’ free from strong molecular absorption (Burrows et al., 2003). The *H* band sits atop one such window; the *Ks* band sits on intense molecular absorption instead. Thus as the planets cool and new molecules form, the *Ks* band closes down and the *H* band receives an anomalous boost. The *H* band flux becomes many times what it would be for a blackbody of the planet’s temperature. The planet does still get redder with age in its *L' – H* and *M – H* color, but at a rate far less than a naive blackbody approximation would suggest. Because of this

enormous boost to the  $H$  band due to molecular absorption lines, it may be that the  $H$  band is superior to  $L'$  and  $M$  for planets sufficiently far from the star even for the nearest star systems in which the coolest, faintest, and most red planets are detectable.

Two caveats to the above may be in order. First, there may be reasons to consider the theoretical planet brightnesses in the  $H$  band regime as less robust than the predictions for longer wavelengths. This is because the  $H$  band brightnesses are dependent on more extreme deviations from blackbody spectra due to atmospheric line blanketing. In effect, the models predict that we can see into deep, high-temperature regions of extrasolar planets' atmospheres at the  $H$  band due to the unusually low opacity there. If high-altitude clouds prevent the bright radiation from deep in the planets' atmospheres from escaping freely, the  $H$  band fluxes may be much less than predicted. We note that for the very faint planet  $\epsilon$  Eri B, the Burrows et al. (2003) models, which explicitly include a cloud model, predict an  $H$  band brightness about 30 times fainter than the models of Baraffe et al. (2003). Theoretical brightnesses in the  $L'$  and  $M$  bands are elevated above blackbody values too, but to a lesser extent than the  $H$  band. We would expect these longer wavelength predictions to be more robust.

Secondly, we note that for the new generation of giant telescopes the question of which wavelength regime is optimal will certainly be resolved in favor of the  $L'$  and  $M$  bands, at least for nearby star systems. The reason is clear: the bigger telescopes will be sensitive to fainter, cooler objects — and for the nearest star systems the objects will be so cool and red that their enormously greater brightness at the longer wavelengths relative to  $H$  band will overwhelm the increased sky noise. This issue is discussed in more detail in Chapter 5. **It is very likely that the first mature planetary systems to be directly imaged will be discovered by**

**the next generation of giant telescopes, using the  $L'$  or  $M$  band and some of the methods we have pioneered in this survey.** The other main contenders for this achievement will be space telescopes. JWST will offer background-limited sensitivity at  $L'$ ,  $M$ , and longer wavelengths far in excess of even the largest ground based telescopes, but it will not perform well in the contrast-limited regime close to bright stars. A 20-30 meter ground-based telescope such as the GMT, operating in the  $M$  band, may outperform it at the separations where giant planets are most likely to be found. Space-based visible-light coronagraphs may also image planets of very nearby stars in reflected starlight, but exceptional coronagraphic performance will be required.

Planets close to the star must be detected against the bright halo of the stellar image. In this case the background limit is not relevant, and the planet/star flux ratio becomes very important. The brighter the star, the more urgent the need to use a band that delivers a good planet/star flux ratio. There are still competing considerations between the  $L'$  and  $H$  bands: the diffraction core of  $H$  band images is smaller in angular terms, yielding greater potential maximum resolution. However, the  $L'$  images more closely approach theoretical perfection, meaning they are cleaner, with fewer speckles and less glare. The question, however, is not whether the  $L'$  and  $M$  bands are more promising wavelengths than the  $H$  band for the nearest, brightest stars, but exactly how bright and nearby the star must be before the  $L'$  advantage is clear.

We will discuss exactly where this breakpoint occurs, and also make an attempt to answer the question of whether  $L'$  or  $H$  is better for background-limited observations of the nearest stars with current telescopes, in Section 3.7. We note in closing that although we have framed the above discussion mainly in terms of the  $H$  band versus the  $L'$  band, two of the most ambitious recent planet imaging

surveys (Biller et al. (2007) and Lafrenière et al. (2007b)) have been carried out using narrowband slices within the  $H$  band. The discussion above remains broadly applicable. The narrowband methods, and how the sensitivities they obtained compare with ours, will also be discussed in Section 3.7.

The reader may well wonder how the  $L'$  and  $M$  bands compare with one another. The answer is that the difference between  $L'$  and  $M$  is analogous to the difference between  $H$  and  $L'$ . The  $M$  band delivers an even better planet/star contrast ratio than  $L'$ , and AO systems deliver even better images at  $M$  band, but the sky background is even higher. Thus, while  $L'$  observations are superior to  $H$  band for bright, nearby stars,  $M$  band observations are superior to  $L'$  for very bright, nearby stars. Accordingly, we have carried out most of our survey in the  $L'$  band, but have observed 3 of the most interesting, brightest, nearest star systems in the  $M$  band as well.

In conclusion, the  $L'$  and  $M$  bands are certainly superior to shorter wavelength observations for detecting mature planetary systems around some subset of bright, nearby stars. As larger, more sensitive telescopes are built, these bands will become increasingly attractive and may well result in some of the first images of planets in mature, stable solar systems like our own.

### 3.3 Our Survey Sample

In Section 3.2 we discussed the reasons for carrying out an  $L'$  and  $M$  band imaging survey. The conclusion was that these wavelengths are superior to the more commonly used shorter IR wavelengths for stars that are nearby and bright. Since at any wavelength younger planets are brighter and easier to detect, young star systems are preferred. However, while shorter wavelength surveys are strongly driven to prioritize the youth of star systems first and their nearness second, our

better sensitivity to old planets allows us to reverse these priorities and observe some of the very nearest stars regardless of age.

Thus we have selected **nearby, bright, young stars** for our survey. Early estimates of the sensitivities of the Clio instrument, and the statistical distributions of extrasolar planets (based on the known RV planets), led us to believe that if we surveyed 50 stars we would expect to detect 4-6 planets. This expected number of detections is sufficiently large that a null detection would be very surprising and statistically significant, indicating that the estimated distribution could be ruled out with high confidence. This would be an interesting result, indicating that the distributions of the RV planets do not hold when extrapolated to the larger orbital radii to which our survey was sensitive.

RV surveys have focused mainly on F, G, and K type stars, and this spectral type range fits in nicely with the optimal targets for an  $L'$  and  $M$  band survey: it leaves out the A-type and hotter stars, which are mostly distant, and the nearby M-stars, which are all faint. It is precisely the FGK stars that are both nearby and fairly bright. We therefore set out to observe 50 of the nearest young FGK stars.

We knew that telescope scheduling and weather issues could render observations of some of our intended targets impossible, and so in order to complete a survey of 50 stars it would be necessary to have a sample list of more than 50. Our key criterion in choosing this larger list was to choose the nearest FGK stars around which, based on the theoretical models of Burrows et al. (2003) and Baraffe et al. (2003), we could detect planets down to 10 Jupiter Masses (MJ) or below. This practically meant that stars within 5 parsecs were potential targets up to ages of several Gyr, while at larger distances we would consider only fairly young stars. We set out initially to investigate only FGK stars within 25pc of the sun, in order to make our sample comparable in spectral type to the samples of

the RV surveys and to focus on the nearest stars at which the  $L'$  and  $M$  bands offer the largest advantage over shorter wavelengths. In the end we included a few M-stars and a few stars slightly beyond 25pc. We did this because it was difficult to come up with a sufficient number of targets using the more stringent constraints, and because some M-stars or stars slightly beyond 25pc were very interesting (that is, we predicted a minimum detectable mass much less than 10 MJ for planets in these systems). The stars of our sample that we actually observed are presented in Tables 3.2, 3.3, and 3.4. Promising sample targets we did not observe are discussed in Chapter 5.

Table 3.2. Ages and Age References for Observed Survey Targets

Star Name	Age #1 Gyr	Age #1 Ref	Age #2 Gyr	Age #2 Ref	Age Adop- ted
GJ 5	0.11	Fischer	0.2	Bryden	0.155
HD 1405	0.1-0.2	W03a	0.03-0.08	Lopez	0.1
$\tau$ Ceti	4.4-12	Lachaume	...	...	5.0
GJ 117	0.1	W03a	0.03	Fischer	0.1
$\epsilon$ Eri	0.56	Fischer	...	...	0.56
GJ 159	0.03-0.01	FEPS	...	...	0.1
GJ 166BC	2	...	...	...	2
HD 29391	0.01-0.03	Z01	...	...	0.1
GJ 211	0.52	Fischer	...	...	0.52
GJ 216A	0.3	UMMG K03>M01	0.4-0.6	K03	0.44
BD+20 1790	0.06-0.3	FEPS	...	...	0.18
GJ 278C	0.1-0.3	CMG Kar>BYN98	...	...	0.2
GJ 282A	0.49	Fischer	0.4-0.6	K03	0.5
GJ 311	0.3	M01	0.1-0.3	FEPS	0.24
HD 77407AB	0.05	W03b	...	...	0.1
HD 78141	0.1-0.2	W03a	...	...	0.15
GJ 349	0.37	Fischer	...	...	0.37
GJ 355	0.1	W03a	0.05-0.15	M01	0.1
GJ 354.1A	0.1	W03a	0.02-0.15	M01	0.1
GJ 380	2	...	...	...	2
GJ 410	0.3	UMMG K03>M01	0.4-0.6	0.37	
HD 96064AB	0.1-0.2	W03	...	...	0.15
GJ 450	1	RASS	...	...	1

Table 3.2—Continued

Star Name	Age #1	Age #1	Age #2	Age #2	Age
	Gyr	Ref	Gyr	Ref	Adop- ted
BD+60 1417	0.1-0.2	W03a	...	...	0.15
HD 113449	0.1-0.2	W03a	...	...	0.15
GJ 505AB	0.79	Fischer	...	...	0.79
GJ 519 0.3	UMMG K03>M01	0.4-0.6	K03	0.37	
GJ 3860	0.28	Fischer	0.4-0.6	K03	0.28
GJ 564	0.1-0.2	W03a	...	...	0.15
GJ 3876	2	...	...	...	2
$\xi$ Boo AB	0.43	Fischer	0.1	W03a	0.29
HD 139813	0.1-0.2	W03a	...	...	0.15
GJ 625	0.4-0.6	K03a	...	...	0.5
GJ 659A	1	Favata	...	...	1
GJ 659B	1	Favata	...	...	1
GJ 702AB	2	...	...	...	2
GJ 820A	2	...	...	...	2
GJ 820B	2	...	...	...	2
BD+48 3686	0.1-0.2	W03a	...	...	0.15

Table 3.2—Continued

Star Name	Age #1 Gyr	Age #1 Ref	Age #2 Gyr	Age #2 Ref	Age Adop- ted
GJ 879	0.1-0.3	BYN98	...	...	0.2
HD 220140AB	0.025-0.15	M01	...	...	0.1
GJ 896AB	<0.3	BYN98	...	...	0.3

Note. — References: Bryden is Bryden et al. (2006), BYN98 is Barrado y Navascués (1998), Favata is Favata et al. (1998), Fischer is Fischer (1998), K03 is King et al. (2003), Kar is Karatas et al. (2005), Lachaume is Lachaume et al. (1999) Lopez is Lopez-Santiago et al. (2006), M01 is Montes et al. (2001), W03a is Wichmann et al. (2003), W03b is Wichmann & Schmitt (2003), and Z01 is Zuckerman et al. (2001). FEPS means the age estimate is from the FEPS target list, which we received through private communication from Michael Meyer in 2006. RASS means that the Hünsch et al. (1998) catalog reports a ROSAT detection at a flux level that suggests an age of 1 Gyr or less. Sometimes ages are assigned based on one paper’s assignment of a star to the Ursa Major Moving Group (UMMG) or the Castor Moving Group (CMG), and another paper’s age for that group. This is indicated by G P1>P2, where G is the group acronym, P1 is the assigning paper, and P2 gives the quoted value for the group age. Adopted ages are usually averages of quoted values; not all values are weighted equally. The 2 Gyr ages with no references given are for nearby stars that have, so far as we know, no indication of unusual youth or age. We expect the average age of stars in the solar neighborhood to be about 2 Gyr.

Table 3.2 gives the different age measurements available to us for each of the sample stars we observed, along with our final adopted ages. As can be seen, some of our stars have average estimated ages below 100 Myr. We have chosen to approximate these ages as 100 Myr. There are several reasons for this. First, the readily available models from Baraffe et al. (2003) and Burrows et al. (2003) do not give the type of observables we need for planets younger than 100 Myr. Second, setting the ages of these stars slightly older than they are thought to be fits in with our generally conservative approach to the volatile subject of extrasolar planet searches, and ensures that our survey results do not hang on just a few very young stars and will not be invalidated if the age estimates are revised upward.

Table 3.3. Distances, Ages, and Spectral Types of Observed Survey Targets

Star	Age(Gyr)	Dist(pc)	Type	Binary?
GJ 5	0.155	14.25	K0Ve	No.
HD 1405	0.1	30	K2V	No.
$\tau$ Ceti	10	3.50	G8Vp	No.
GJ 117	0.1	8.31	K2V	No.
$\epsilon$ Eri	0.56	3.27	K2V	No.
GJ 159	0.1	18.12	F6V	No.
GJ 166BC	2	4.83	DA + dM4.5e	Yes.
HD 29391	0.1	14.71	F0V	No.
GJ 211	0.52	12.09	K1Ve	No.
GJ 216A	0.44	8.01	F6V	No.
BD+20 1790	0.18	24	K3	No.
GJ 278C	0.2	14.64	M0.5Ve	No.
GJ 282A	0.5	13.46	K2Ve	No.
GJ 311	0.24	13.85	G1V	No.
HD 77407AB	0.1	30.08	G0V	Yes.
HD 78141	0.15	21.4	K0	No.
GJ 349	0.37	11.29	K3Ve	No.
GJ 355	0.1	19.23	K0	No.
GJ 354.1A	0.1	18.87	dG9	No.
GJ 380	2	4.69	K2Ve	No.

Table 3.3—Continued

Star	Age(Gyr)	Dist(pc)	Type	Binary?
GJ 410	0.37	11	dM2e	No.
HD 96064AB	0.15	24.63	G5V + M3V	Yes.
GJ 450	1	8.1	M1Ve	No.
BD+60 1417	0.15	17.7	K0	No.
HD 113449	0.15	22.1	G5V	No.
GJ 505AB	0.79	11.9	K2V + M0.5V	Yes.
GJ 519	0.37	9.81	dM1	No.
GJ 3860	0.28	14.93	K0	No.
GJ 564	0.15	17.94	G2V	No.
GJ 3876	2	43.3	F9IV	No.
ξ Boo AB	0.29	6.71	G8V + K4V	Yes.
HD 139813	0.15	21.7	G5	No.
GJ 625	0.5	6.28	dM2	No.
GJ 659A	1	20.2	K8V	No.
GJ 659B	1	20.2	dK8	No.
GJ 702AB	2	5.03	K0V + K4V	Yes.
GJ 820A	2	3.46	K5V	No.
GJ 820B	2	3.46	K7V	No.
BD+48 3686	0.15	23.6	K0	No.
GJ 879	0.2	7.81	K5Ve	No.

Table 3.3—Continued

Star	Age(Gyr)	Dist(pc)	Type	Binary?
HD 220140AB	0.1	19.74	G9V	Yes.
GJ 896AB	0.3	6.58	M3.5 + M4.5	Yes.

Finally, setting the ages conservatively hedges our results to some extent against the possibility suggested in Marley et al. (2007) that young massive planets may be far fainter than work such as Baraffe et al. (2003) and Burrows et al. (2003) predicts, because much of the gravitational potential energy of the accreting material may get radiated away in an accretion shock and thus never get deposited in the planet's interior. Figure 4 in Marley et al. (2007) shows that in this accretion scenario planets start out at much lower luminosities than predicted by 'hot start' models such as those in Baraffe et al. (2003) and Burrows et al. (2003), but over time the predictions converge. By 100 Myr, the differences are less than an order of magnitude for planets less massive than 10 MJ, and are negligible for planets of 4 MJ and lower masses.

Table 3.3 gives the final adopted ages, distances, and spectral types for each of our stars, and notes which stars were binaries. By 'binary' in the context of this table, we mean a binary with a sufficiently small angular separation that both components fit in the same Clio field and could be observed simultaneously. Table 3.4 gives the celestial coordinates and magnitudes in the  $V$ ,  $H$ ,  $K$ , and  $L'$  bands. Most of the celestial coordinates, spectral types, distances, and magnitudes are from the Simbad website. In general Simbad coordinates and distances for bright stars such as these are from Perryman et al. (1997), and magnitudes for the  $H$  and  $K$  bands are from Cutri et al. (2003) The  $L'$  magnitudes, and  $H$ ,  $K$ , and  $L'$  magnitudes in cases where only a  $V$  magnitude was available from Simbad, are calculated using the spectral types and standard IR colors from Table 7.6 of Cox (2000).

Table 3.4. Positions and Magnitudes of Observed Survey Targets

Star	RA	DEC	V	H	K	L'
GJ 5	00:06:36.80	29:01:17.40	6.13	4.69	4.31	4.25
HD 1405	00:18:20.90	30:57:22.00	8.60	6.51	6.39	6.32
$\tau$ Ceti	01:44:04.10	-15:56:14.90	3.50	1.77	1.70	1.65
GJ 117	02:52:32.10	-12:46:11.00	6.00	4.23	4.17	4.11
$\epsilon$ Eri	03:32:55.80	-09:27:29.70	3.73	1.88	1.78	1.72
GJ 159	04:02:36.70	-00:16:08.10	5.38	4.34	4.18	4.14
GJ 166B	04:15:21.50	-07:39:22.30	9.50	...	...	...
GJ 166C	04:15:21.50	-07:39:22.30	11.17	5.75	5.45	5.05
HD 29391	04:37:36.10	-02:28:24.80	5.22	4.77	4.54	4.51
GJ 211	05:41:20.30	53:28:51.80	6.23	3.99	4.27	4.21
GJ 216A	05:44:27.80	-22:26:54.20	3.60	2.47	2.42	2.38
BD+20 1790	07:23:43.60	20:24:58.70	9.93	7.61	7.51	7.42
GJ 278C	07:34:37.40	31:52:09.80	9.07	5.42	5.24	5.05
GJ 282A	07:39:59.30	-03:35:51.00	7.20	5.06	4.89	4.82
GJ 311	08:39:11.70	65:01:15.30	5.65	4.28	4.17	4.12
HD 77407A	09:03:27.10	37:50:27.50	7.10	5.53	5.44	5.39
HD 77407B	09:03:27.10	37:50:27.50	...	...	...	...
HD 78141	09:07:18.10	22:52:21.60	7.99	5.92	5.78	5.72
GJ 349	09:29:54.80	05:39:18.50	7.22	5.00	4.79	4.70
GJ 355	09:32:25.60	-11:11:04.70	7.80	5.60	5.45	5.39

Table 3.4—Continued

Star	RA	DEC	V	H	K	L'
GJ 354.1A	09:32:43.80	26:59:18.70	7.01	5.24	5.12	5.06
GJ 380	10:11:22.10	49:27:15.30	6.61	3.93	2.96	2.89
GJ 410	11:02:38.30	21:58:01.70	9.69	5.90	5.69	5.46
HD 96064A	11:04:41.50	-04:13:15.90	7.64	5.90	5.80	5.75
HD 96064B	11:04:41.50	-04:13:15.90	...	...	...	...
GJ 450	11:51:07.30	35:16:19.30	9.78	5.83	5.61	5.40
BD+60 1417	12:43:33.30	60:00:52.70	9.40	7.36	7.29	7.23
HD 113449	13:03:49.70	-05:09:42.50	7.69	5.67	5.51	5.46
GJ 505A	13:16:51.10	17:01:01.90	6.52	4.58	4.38	4.31
GJ 505B	13:16:51.10	17:01:01.90	9.80	5.98	5.75	5.43
GJ 519	13:37:28.80	35:43:03.90	9.07	5.66	5.49	5.28
GJ 3860	14:36:00.60	09:44:47.50	7.51	5.63	5.55	5.49
GJ 564	14:50:15.80	23:54:42.60	5.88	4.47	4.42	4.37
GJ 3876	14:50:20.40	82:30:43.00	5.64	4.19	3.92	3.87
ξ Boo A	14:51:23.40	19:06:01.70	4.55	2.82	2.75	2.70
ξ Boo B	14:51:23.40	19:06:01.70	6.97	4.45	4.34	4.24
HD 139813	15:29:23.60	80:27:01.00	7.31	5.56	5.46	5.41
GJ 625	16:25:24.60	54:18:14.80	10.40	6.06	5.83	5.60
GJ 659A	17:10:10.50	54:29:39.80	8.80	6.23	6.12	5.97
GJ 659B	17:10:12.40	54:29:24.50	9.29	6.13	5.97	5.83

Table 3.4—Continued

Star	RA	DEC	V	H	K	L'
GJ 702A	18:05:27.30	02:30:00.40	4.20	2.32	2.24	2.18
GJ 702B	18:05:27.30	02:30:00.40	6.00	3.48	3.37	3.27
GJ 820A	21:06:53.90	38:44:57.90	5.21	2.47	2.36	2.25
GJ 820B	21:06:55.30	38:44:31.40	6.03	3.02	2.87	2.74
BD+48 3686	22:20:07.00	49:30:11.80	8.57	6.58	6.51	6.45
GJ 879	22:56:24.10	-31:33:56.00	6.48	3.80	3.81	3.70
HD 220140A	23:19:26.60	79:00:12.70	7.54	5.74	5.66	5.60
HD 220140B	23:19:26.60	79:00:12.70	...	...	...	...
GJ 896A	23:31:52.20	19:56:14.10	9.95	5.24	4.99	4.64
GJ 896B	23:31:52.20	19:56:14.10	12.40	6.98	6.68	6.28

Note. — Coordinates are epoch J2000.0 and are mostly from Perryman et al. (1997). *H* and *K* magnitudes are from Cutri et al. (2003), or else calculated from Simbad website spectral types and *V* magnitudes using Table 7.6 of Cox (2000). *L'* magnitudes are similarly calculated from either *V* or *K* values.

The survey was planned mostly in 2004-5 (Freed et al., 2004), and the Clio instrument was commissioned on the MMT in June 2005. However, problems with the MMT, weather, and issues with Clio itself delayed the actual start of survey observations until April 2006. By May 2007 we had observed 50 stars and were ready to conclude the survey. In the roughly 3 years between the survey's initial planning and final completion, many RV planets were discovered and the power law fits to the mass and orbital semimajor axis distributions of the RV planets were significantly revised. For example, Freed et al. (2004) describe the early stages of survey planning in 2004. They note that the mass distribution of RV planets known at that time could be fit with a power law of  $dn/dM \propto M^{-0.7}$ , and the semimajor axis distribution could be fit with a rising power law of  $dn/da \propto a^{0.7}$ . From these power laws, with the semimajor axis distribution truncated at 50 AU, Freed et al. (2004) predict a Clio survey of 80 stars should detect 12-18 planets.

The current picture is very different. Since more RV planets have been discovered at lower masses, the best-fit power law slope of the mass distribution has been adjusted down, predicting fewer massive planets. Complete samples of RV planets are now available (Fischer & Valenti, 2005), so secure normalizations of the statistical distributions are possible. Also, as always happens with astronomical instruments, Clio has turned out to be somewhat less sensitive than we had initially predicted. Therefore, the initial possibility discussed in Freed et al. (2004) that an 80-star Clio survey could find 12-18 planets (corresponding to 7-11 planets in a 50 star survey) is no longer realistic. The expected number of planets in the final analysis is more like 0.3-0.5. This number is based on mass and semimajor axis power laws with slopes of -1.44 and -0.35, respectively, with the distribution truncated at 40 AU and normalized based on Fischer & Valenti

(2005) (see Chapter 4). To get an expected number of planets of 7-11 we would need to observe at least 700 stars! Even if we had known this in the planning stages of our survey, we could not have selected even 100 stars as promising as the 50 we did in fact observe. There are simply not enough nearby, young FGK stars. In Chapter 5 we discuss how future surveys at the  $L'$  and  $M$  bands may be designed so that the expected number of planets detected will be well above unity.

In Section 3.2 above, we concluded that the ideal targets for an  $L'$  and  $M$  band imaging survey are nearby, bright, and young, but that youth need not be prioritized as highly as it is for selecting targets of shorter wavelength surveys. Our sample should thus consist in general of brighter, nearer, and older stars than those of other surveys. In Figures 3.1 through 3.3 we present histograms of the star magnitude, distance, and age distributions for the samples of our survey and 4 others: Masciadri et al. (2005), Kasper et al. (2007), Biller et al. (2007), and Lafrenière et al. (2007b). Masciadri et al. (2005) used the VLT/NACO system to survey young stars for extrasolar planets at the  $H$  and  $K_s$  bands; Kasper et al. (2007) used the VLT/NACO to carry out the only  $L'$  planet imaging survey undertaken before the commissioning of Clío; Biller et al. (2007) used a narrowband  $\text{CH}_4$  spectral differential imaging (SDI) method at  $H$  band wavelengths to search for planets at the VLT and MMT; and Lafrenière et al. (2007b) searched for planets using an optimized narrowband filter in the  $H$  band wavelength regime at the Gemini North telescope. It is clear that our target list was chosen well to fit the niche of  $L'$  and  $M$  band observations that we have described above. The Kasper et al. (2007) survey sample does not seem to be well tuned to the niche of the  $L'$  band, however Kasper et al. (2007) did obtain excellent sensitivity to low-mass planets around the stars surveyed, and the work is a fine example of using the  $L'$

band to diversify observing effort in the search for extrasolar planets.

### 3.4 The Clio Instrument

Clio is an  $L'$  and  $M$  band optimized AO imaging camera for the MMT. It was developed by Phil Hinz, Melanie Freed, Andy Breuninger, and Suresh Sivaraman explicitly to take advantage of the very low thermal background delivered by the MMT adaptive secondary AO system to perform AO imaging surveys for extrasolar planets.

Clio uses a 320x256 pixel InSb detector. On the MMT the pixel scale is about 48.55 mas/pixel, yielding a field of view of 15.5x12.4 arcsec. The key advantages of this detector are its fast readout time ( $\sim 59.6$ msec) and its deep full well capacity ( $3.4 \times 10^6 e^-$ ). Because of these characteristics the detector is able to handle very high flux levels and still be read full-frame without saturating. These capabilities are greatly enhanced by the remarkable software Andy Breuninger wrote to run the detector, which allows it to be used in integrate-while-reading (IWR) mode with high ( $\geq 90\%$ ) duty cycle even with short frame times such as 100 msec. Clio is able to handle the very bright sky flux in the broadband  $M$  filter. It is the only AO imaging camera of which we are aware that can take efficient, long integrations at  $M$ . Most cameras are limited to the narrower  $M'$  band, where the flux is more manageable; however  $M$  band is substantially better than  $M'$  for planet searches because the wider bandwidth includes more photons from the planet and allows detection in a shorter amount of time (Freed et al., 2004). At the shorter wavelength  $L'$  band we have used for most of our observations, the sky background is less bright but Clio's deep full well capacity allows us to use a long frame time of typically 1500-2000 msec to minimize the read noise hit and obtain a duty cycle well above 90%.

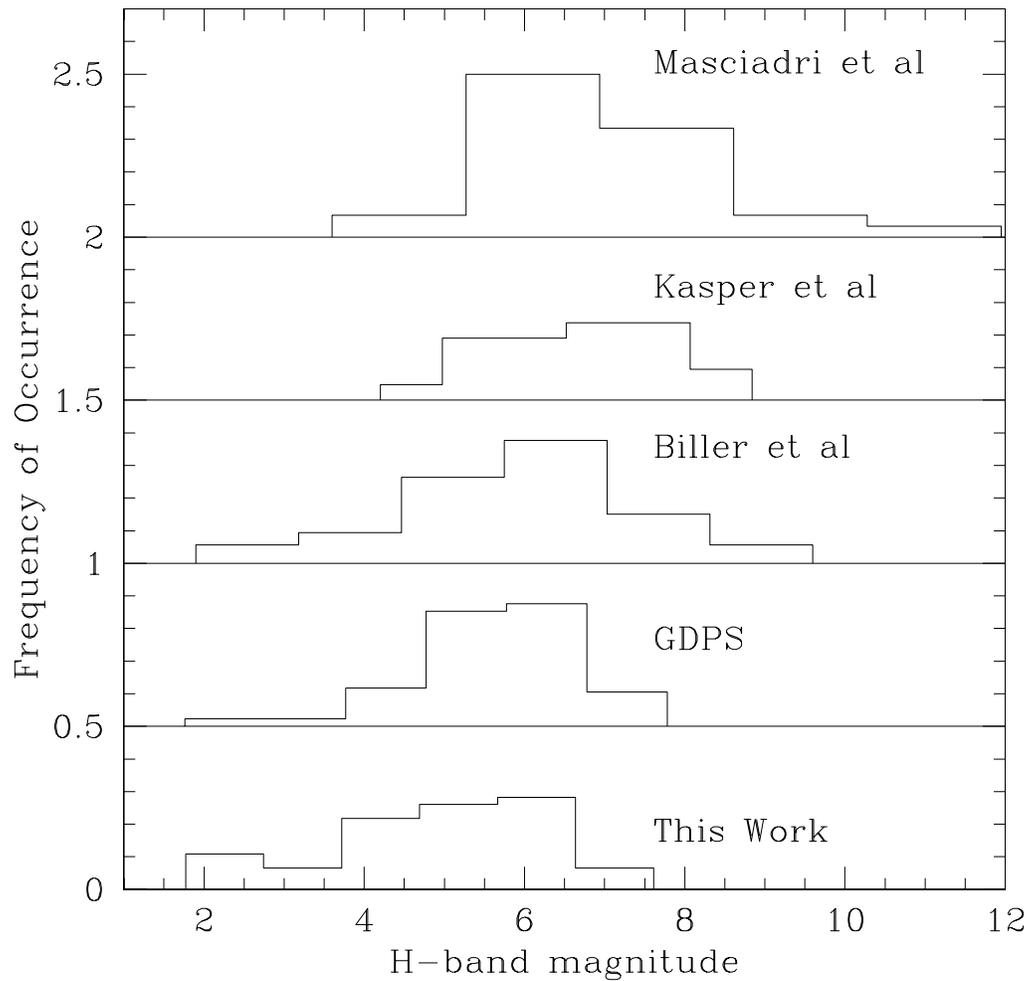


Figure 3.1 Stellar magnitude histograms for the target lists of various planet imaging surveys.  $H$  band magnitudes are used except for the Kasper et al. (2007) survey, for which we have used  $L'$  because  $H$  band magnitudes were not tabulated.  $H - L'$  is usually small. Our target list is the most heavily weighted toward bright stars.

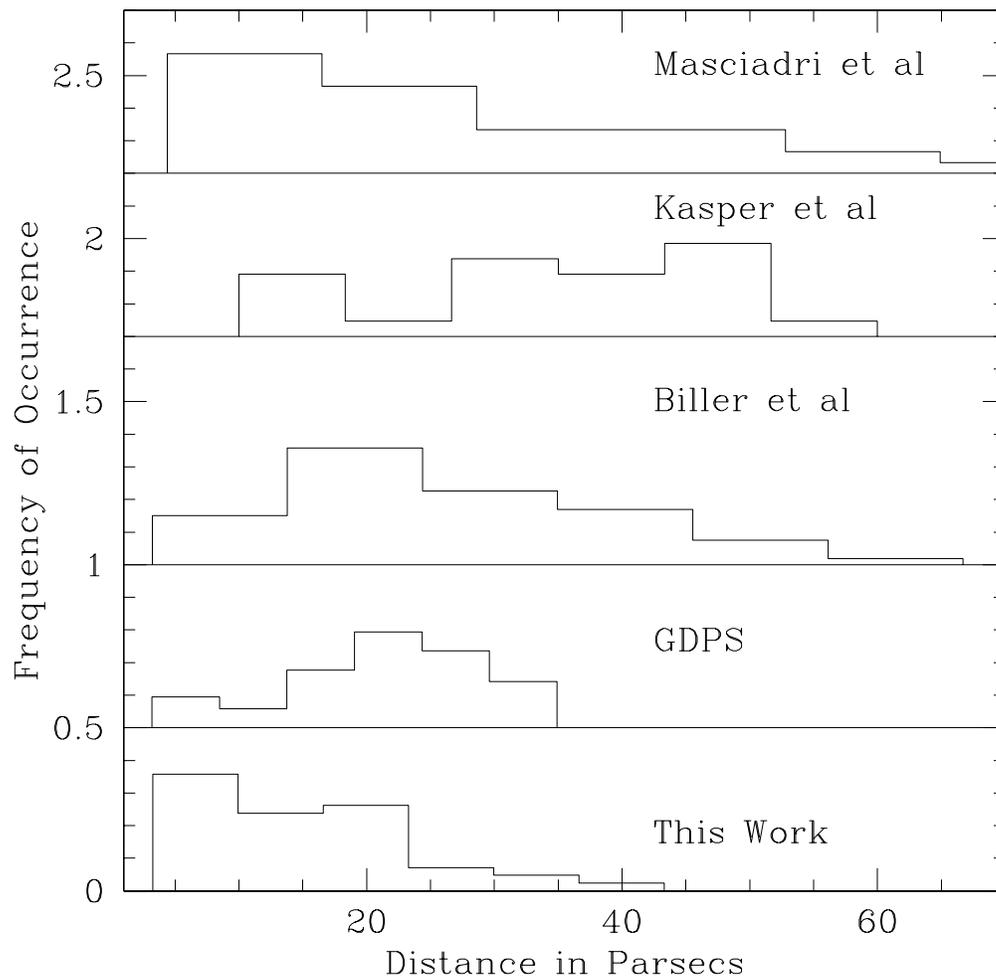


Figure 3.2 Stellar distance histograms for the target lists of various planet imaging surveys. One star at 150 pc has been left out of the Biller et al. (2007) sample to keep the histogram comparable with the others. Our target list is the most focused on nearby stars.

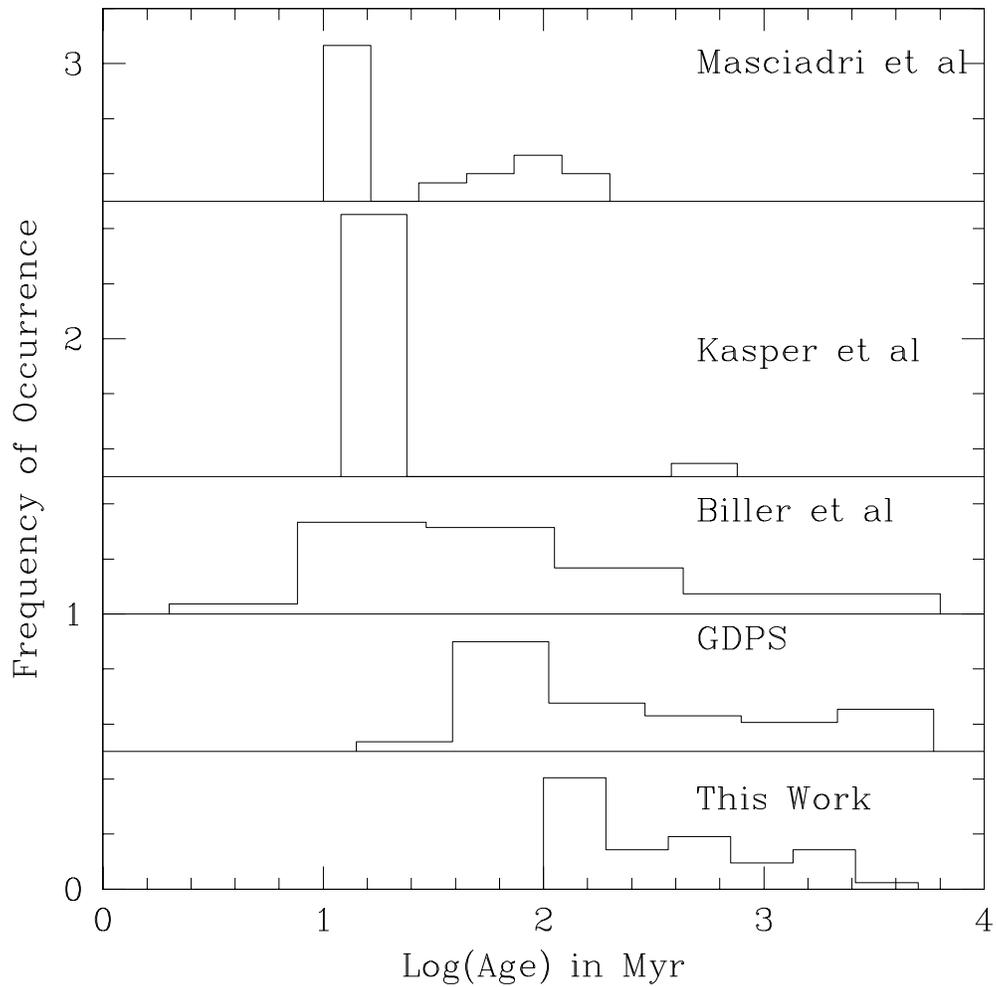


Figure 3.3 Stellar age histograms for the target lists of various planet imaging surveys. Our target list is the most weighted toward old stars.

Clio's disadvantages are high read noise ( $\sim 700 e^-$ ; compare with the full well of  $3.4 \times 10^6 e^-$ ) and high dark current. We reduce the dark current by cooling Clio to about 55K using solid nitrogen. At this temperature it takes 6-10 seconds for the detector to saturate on dark current. Running the detector colder results in decreased pixel stability and more image noise. The high dark current and read noise severely decrease Clio's sensitivity at the  $H$  and  $K$  bands where the sky is relatively dark, but at the  $L'$  and  $M$  bands for which the instrument was designed their effect is negligible. This is because the bright sky dominates the dark current, and if the exposure is set by the sky, the photon shot noise of the bright background dominates the read noise. Clio is thus strongly specialized for the  $L'$  and  $M$  bands, but impressively capable within them. The CONICA camera on the VLT may slightly outperform Clio at  $L'$  (Kasper et al., 2007); at  $M$  band Clio's capability is currently unique.

### 3.5 Observations

For each star in our sample we sought to acquire at least one hour of cumulative integration at the  $L'$  band. In most cases we achieved this. For a few of our brightest nearby targets we acquired  $M$  band integrations as well. If possible we desired to observe the star through transit. In this way our observations would span the greatest possible range of parallactic rotation for the star. This was very helpful for eliminating artifacts, as will be explained in the next section.

After acquiring a target with MMTAO, we would determine a long 'science' exposure time based solely on the sky background, chosen so that the sky background flux filled 60 – 80% of the detector full-well capacity. This insured that the observations were background-limited rather than readnoise limited. Close to the star, of course, the sensitivity is limited by the stellar speckles. Our typical

exposures were 1200-2000 msec at  $L'$  and 150-250 msec at  $M$  band. In normal operation Clio coadds several individual frames and saves them as a single FITS image. We ordinarily set the coadd number between 10 and 20 for  $L'$  or 50 to 200 for  $M$  band, where the exposures are shorter. The one exception was the star GJ 380, for which we took FITS data-cubes rather than coadding the images, and then sliced them post-observation and processed the frames individually. This has potential advantages for special observations, but for routine survey work we found the resulting data volume and processing runtimes excessive. For our normal observations using coadding rather than data cubes, each FITS image output by Clio typically corresponded to 20-30 seconds of integration

In addition to the long, 'science' exposure set by the sky background, we would if possible determine a shorter, 'PSF-monitoring' exposure that would yield unsaturated images of the primary star. We would intersperse such shorter exposures among the longer ones that made up our science integration. They were invaluable for estimating the true point spread function (PSF) that applied to the science data, as well as providing a photometric reality check on the sky conditions. The PSF could not in general be obtained from the science exposures because the core of the star image was saturated.

We took our data using the standard IR imaging technique of nodding, in which a sequence of images is taken in one position, the telescopes is moved ('nodded') slightly, and then another image sequence is acquired. Images taken at one position can then be subtracted from images taken at the other position. Any real celestial objects leave both bright and dark images, but artifacts of the bright sky interacting with the telescope and the detector vanish. This is a powerful technique and is practically indispensable for  $L'$  and  $M$  band observations.

We typically took 5 or 10 coadded images at each nod position in each given

nod cycle. This meant that a given nod cycle lasted about 5 minutes. This seemed to be fast enough that alterations in the sky background did not introduce appreciable noise into our data. This is in sharp contrast with the situation for ground-based observations at much longer wavelengths such as the  $10\ \mu\text{m}$  *N* band, where a special ‘chopping’ mirror must be used to carry out a procedure analagous to nodding on a timescale of seconds or less.

Good Clio science data sets consisted of 15 or more nod cycles, corresponding to around 150 FITS images that, when coadded, would produce a master image with a cumulative exposure of over an hour. We would typically intersperse 2-4 nod cycles of shorter, ‘PSF monitoring’ exposures evenly among the longer, science image nod cycles.

After the end of an imaging sequence we would take a short series of dark frames at the same integration and coadd settings as the science data, and then another with setting matching the PSF monitoring data. It was often possible to do this while the telescope was slewing to the next target, and the AO system was locking on it. We considered it important to take darks as close to the science observations as possible because of variations in the bias level or dark current with slow, subtle changes in detector temperature.

Figures 3.4, 3.5, and 3.6 illustrate how nod subtraction works. In Figure 3.4, we show an image from nod position # 1, in which the star is on the left hand side of the image. In Figure 3.5, from the other nod position, the star is on the right. We subtracted the Figure 3.5 image from the Figure 3.4 image to produce Figure 3.6, in which detector and skyglow artifacts are enormously reduced, so the image looks clean away from the star despite being stretched 20 times more than the images in Figures 3.4 and 3.5. In a real Clio data processing sequence, another image would be made in which Figure 3.4 was subtracted from Figure

3.5, producing a frame with the positive stellar image on the right. After further processing, the two images would be shifted so that the positive stellar images overlaid, and then coadded. In this way science data from both nod positions is used; we do not have ‘throw-away’ sky subtraction images. A fully reduced Clio data set consists of a registered and coadded stack of 100 or more Figure 3.6 type images.

The black negative star images do, of course, greatly reduce the sensitivity in some parts of the image. We have a processing method that reduces or eliminates this problem in some data sets (see next section). For all data sets we have taken care to nod the telescope far enough that the region near the star, where planets are most likely to be found, is not affected by the negative nod-subtraction images.

Table 3.5 shows the date on which each of our target stars was observed, the values to which we set the Clio single-frame integration time and the coadds, and the number of coadded FITS images we acquired. Note that the true single-frame integration for Clio is the nominal integration plus about 59.6 msec. For short nominal integration times this can make an important difference. Note also the enormously greater number of images for GJ 380, in which we did not use coadds. Processing frames singly could potentially result in slightly better resolution. However, since the data volume and processing times were effectively increased about  $15\times$ , we found this to be impractical for regular survey observations. We have used it extensively on specialty targets such as close binaries, especially in poor seeing when the ability to cull out individual sharp frames can be essential.

Table 3.6 shows the total cumulative exposure time, the mean airmass, and the total parallactic rotation for each of our target observations. The parallactic



Figure 3.4 Raw Clío image of  $\epsilon$  Eri from nod position 1. Note the strong column striping from the detector, and other artifacts.

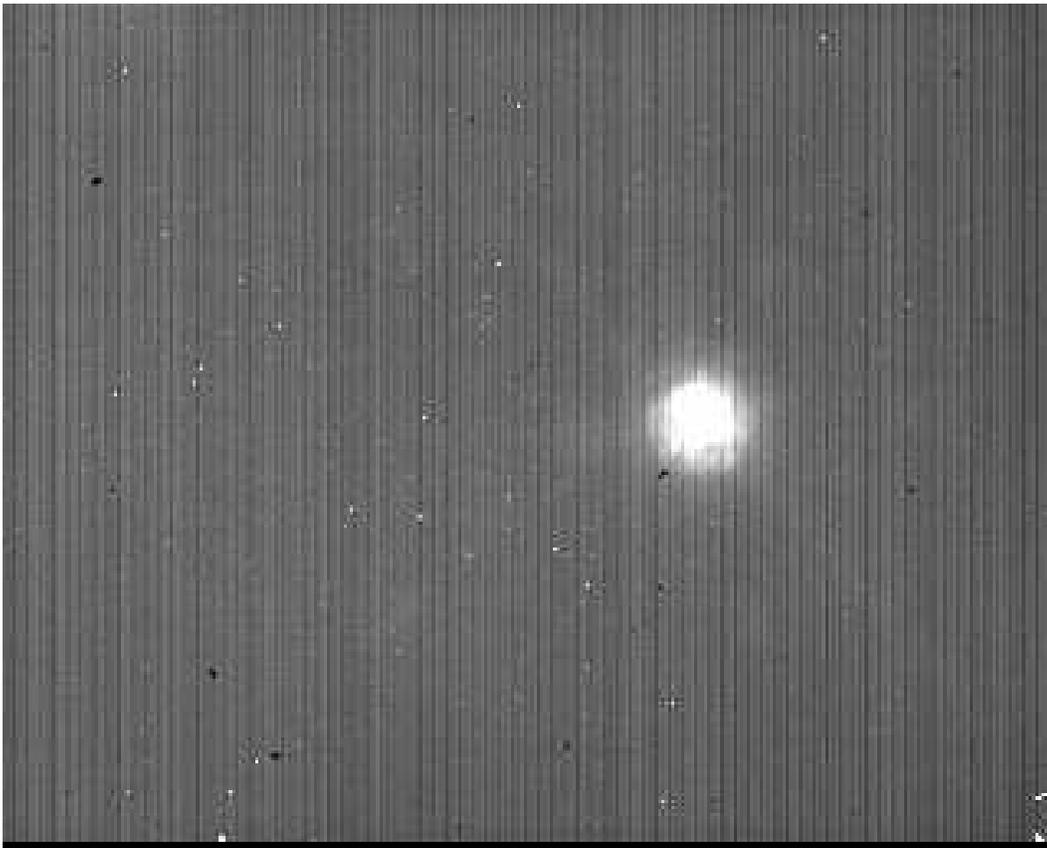


Figure 3.5 Raw Clío image of  $\epsilon$  Eri from nod position 2. Note the strong column striping from the detector, and other artifacts.

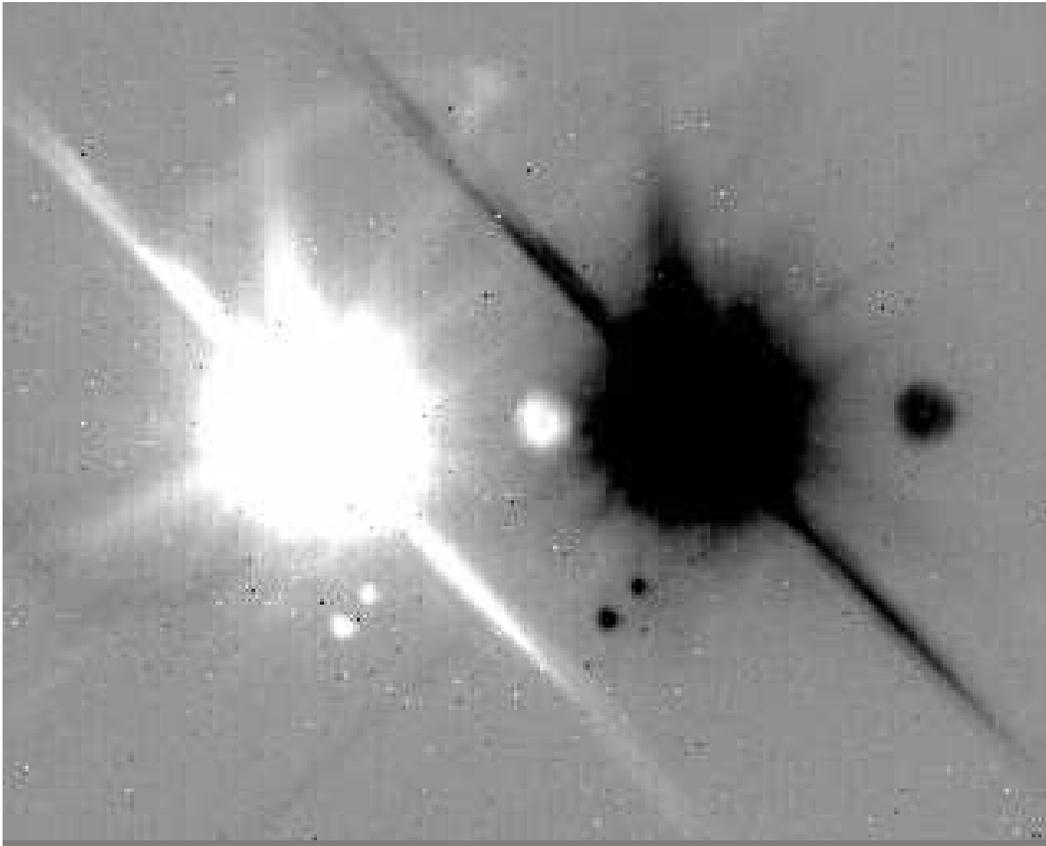


Figure 3.6 Nod-subtracted image of  $\epsilon$  Eri made from Figures 3.4 and 3.5. Many detector and sky-glow related artifacts have vanished, revealing faint structure in the outer halo of this very bright star.. This image is stretched 20 times more than the previous two figures, yet it looks cleaner apart from the bright artifacts from the star itself.

rotation was important to determine how effective our PSF subtraction method would be (see Section 3.6).

Finally, in Table 3.7 we give the photometric calibrations for each Clio observing run. Our calibration is based on  $L'$  and  $M$  band standard stars from Legget et al. (2003). In most cases we used early A-stars of about  $L' = 7$ . The magnitude choice simply gave good SNR against the sky without saturating the images. We chose A-stars to insure the magnitude differences from one band to another would be very small so that, for example, we could get an approximate  $Ks$  calibration using the  $L'$  magnitude from Legget et al. (2003), or an approximate  $M$  band calibration using their  $M'$  value. The standard number we quote is for photometric calibrations is the ADU/sec expected within a 30 pixel (1.46 asec) radius aperture from a star of 10th magnitude. For the  $M$  band we used a 10 pixel (0.486 asec) radius aperture instead. These apertures were chosen because they were large enough to give consistent results across a wide range of seeing conditions, but small enough not to include excessive sky noise relative to the good signal from these fairly bright calibration stars. The calibration accuracy is perhaps 5% or better for the  $L'$  and  $Ks$  bands, and about 10% for the  $M$  and  $M'$  bands.

Table 3.5. Observations of Science Targets: Basic Parameters

Star	Date Obs	Band	Clio int(msec)	Coadds	# Images
GJ 659A	11/04/06	L'	2000	10	90
GJ 354.1A	12/04/06, 2006	L'	2000	10	232
GJ 450	12/04/06	L'	2000	10	260
GJ 625	12/04/06	L'	2000	10	208
GJ 349	13/04/06	L'	2000	10	240
GJ 564	13/04/06	L'	2000	10	193
GJ 3876	13/04/06	L'	2000	25	68
GJ 3860	09/06/06	L'	1500	15	170
HD 139813	09/06/06	L'	1200	20	148
GJ 702AB <sup>a</sup>	09/06/06	L'	1200	20	95
GJ 820A	09/06/06	L'	1200	20	133
BD+60 1417	10/06/06	L'	1200	20	160
$\xi$ Boo AB <sup>a</sup>	10/06/06	L'	1200	20	157
GJ 820B	10/06/06	L'	1500	15	140
GJ 519	10/06/06	L'	1500	15	180
BD+48 3686	11/06/06	L'	1200	20	130
$\xi$ Boo AB <sup>a</sup>	11/06/06	M	100	100	260
GJ 505AB <sup>a</sup>	12/06/06	L'	1200	20	149
GJ 659B	12/06/06	L'	1200	20	170
GJ 820A	12/06/06	M	100	100	176

Table 3.5—Continued

Star	Date Obs	Band	Clio int(msec)	Coadds	# Images
GJ 896AB <sup>a</sup>	13/07/06	L'	1500	20	105
$\epsilon$ Eri	09/09/06	M	130	100	180
GJ 5	11/09/06	L'	1500	15	210
$\epsilon$ Eri	11/09/06	L'	1500	15	184
GJ 117	01/12/06	L'	1500	15	139
GJ 211	01/12/06	L'	1500	15	170
GJ 282A	01/12/06	L'	1500	15	190
HD 1405	02/12/06	L'	1500	15	98
GJ 159	02/12/06	L'	1500	15	180
GJ 216A	02/12/06	L'	1500	15	158
GJ 278C	02/12/06	L'	1500	15	132
GJ 355	02/12/06	L'	1500	15	159
GJ 879	03/12/06	L'	1500	15	54
HD 220140AB <sup>a</sup>	03/12/06	L'	1500	15	170
GJ 166BC <sup>a</sup>	03/12/06	L'	1500	15	149
GJ 311	03/12/06	L'	1500	15	90
GJ 410	03/12/06	L'	1500	15	100
$\tau$ Ceti	04/01/07	L'	1700	15	160
HD 29391	04/01/07	L'	1700	15	200
BD+20 1790	04/01/07	L'	1700	15	188

Table 3.5—Continued

Star	Date Obs	Band	Clio int(msec)	Coadds	# Images
HD 96064AB <sup>a</sup>	04/01/07	L'	1700	15	180
HD 77407AB <sup>a</sup>	05/01/07	L'	1700	15	79
HD 78141 <sup>b</sup>	11/04/07	L'	1700	15	203
HD 113449	11/04/07	L'	1500	15	190
GJ 380	30/04/07	L'	1500	1	2066

<sup>a</sup>These stars were sufficiently close binaries that both stars appeared on the same Clio images, and meaningful sensitivity to substellar objects could be obtained around both.

<sup>b</sup>A small fraction of the images of this star were accidentally taken with a 1500 msec rather than a 1700 msec nominal integration time.

Table 3.6. Observations of Science Targets: Data Acquired

Star	Band	Exposure(sec)	Mean Airmass	Rotation
GJ 659A	L'	1853.64	1.113	15.80
GJ 354.1A	L'	4778.27	1.032	130.75
GJ 450	L'	5354.96	1.031	110.37
GJ 625	L'	4283.97	1.117	45.65
GJ 349	L'	4943.04	1.178	40.61
GJ 564	L'	3975.03	1.036	70.70
GJ 3876	L'	3501.32	1.601	27.23
GJ3860	L'	3976.98	1.086	47.09
HD139813	L'	3728.42	1.529	30.15
GJ 702AB <sup>a</sup>	L'	2393.24	1.149	25.50
GJ 820A	L'	3350.54	1.012	101.25
BD+60 1417	L'	4030.72	1.153	37.65
$\xi$ Boo AB <sup>a</sup>	L'	3955.14	1.047	71.20
GJ 820B	L'	3275.16	1.012	103.68
GJ 519	L'	4210.92	1.011	139.97
BD+48 3686	L'	3274.96	1.074	35.97
$\xi$ Boo AB <sup>a</sup>	M	4149.60	1.060	46.142
GJ 505AB <sup>a</sup>	L'	3753.61	1.070	45.30
GJ 659B	L'	4282.64	1.112	43.93
GJ 820A	M	2808.96	1.025	44.24

Table 3.6—Continued

Star	Band	Exposure(sec)	Mean Airmass	Rotation
GJ 896AB <sup>a</sup>	L'	3275.16	1.026	66.49
$\epsilon$ Eri	M	3412.80	1.334	23.406
GJ 5	L'	4912.74	1.011	146.98
$\epsilon$ Eri	L'	4304.50	1.342	36.92
GJ 117	L'	3251.77	1.463	34.05
GJ 211	L'	3976.98	1.097	50.12
GJ 282A	L'	4444.86	1.281	30.28
HD 1405 <sup>b</sup>	L'	2292.61	1.036	162.97
GJ 159	L'	4210.92	1.189	37.65
GJ 216A	L'	3696.25	1.739	30.10
GJ 278C <sup>b</sup>	L'	3088.01	1.017	170.627
GJ 355	L'	3719.65	1.380	25.74
GJ 879	L'	1263.28	2.232	11.68
HD 220140AB <sup>a</sup>	L'	3976.98	1.494	14.14
GJ 166BC <sup>a</sup>	L'	3485.71	1.301	28.66
GJ 311	L'	2105.46	1.201	26.23
GJ 410	L'	2339.40	1.026	34.26
$\tau$ Ceti	L'	4223.04	1.535	37.03
HD 29391	L'	5278.80	1.227	39.63
BD+20 1790	L'	4962.07	1.068	47.94

Table 3.6—Continued

Star	Band	Exposure(sec)	Mean Airmass	Rotation
HD 96064 AB <sup>a</sup>	L'	4750.92	1.252	41.74
HD 77407 AB <sup>a</sup>	L'	2085.13	1.008	95.44
HD 78141 <sup>c</sup>	L'	5297.98	1.022	109.11
HD 113449	L'	4444.86	1.263	35.36
GJ 380	L'	3222.13	1.341	20.58

<sup>a</sup>These stars were sufficiently close binaries that both stars appeared on the same Clio images, and meaningful sensitivity to substellar objects could be obtained around both.

<sup>b</sup>Though the rotation on this star is very large, difficulties arise because the star transited very near the zenith and almost all the rotation happened in a short span of time during which observations were not possible. PSF subtraction had to be performed on a subset of the data with equal numbers of images on each side of transit.

<sup>c</sup>A small fraction of the images of this star were accidentally taken with a 1500 msec rather than a 1700 msec nominal integration time. The total exposure time has been corrected accordingly.

Table 3.7. Photometric Calibrations for Clío Observing Runs

Run Dates	Star	Band	Phot. Cal.	Uncert. (mag)	Mean Airmass
April 11-13, 2006	HD 162208	L'	16825	0.016	1.011
April 11-13, 2006	HD 162208	M	4317	0.046	1.011
June 9-12, 2006	HD 162208	Ks	27231	0.044	1.020
June 9-12, 2006	HD 162208	L'	13135	0.033	1.032
June 9-12, 2006	HD 162208	M	3044	0.057	1.023
July 12-13, 2006	HD 203856	L'	14204	0.021	1.043
September 9-11, 2006	HD 203856	Ks	26715	0.022	1.032
September 9-11, 2006	HD 203856	L'	14258	0.020	1.026
September 9-11, 2006	HD 203856	M	3148	0.053	1.029
December 1-3, 2006,	HD 22686	Ks	26396	0.017	1.260
December 1-3, 2006,	HD 22686	L'	14917	0.021	1.216
December 1-3, 2006,	HD 22686	M	3654	0.060	1.239
January 4-5, 2007	HD 22686	Ks	30768	0.042	1.145
January 4-5, 2007	HD 22686	L'	15614	0.044	1.143
January 4-5, 2007	HD 22686	M	3580	0.093	1.144
April 11, 2007	HD 106965*	Ks	29540	0.021	1.030
April 11, 2007	HD 106965*	L'	16041	0.036	1.030
April 11, 2007	HD 106965*	M	4043	0.091	1.030
April 30, 2007	HD 136754	L'	14101	0.048	1.009

Table 3.7—Continued

Run Dates	Star	Band	Phot. Cal.	Uncert. (mag)	Mean Airmass
-----------	------	------	------------	---------------	--------------

Note. — The photometric calibrations quoted here for the  $L'$  and  $Ks$  bands are the ADU/sec from a 10th magnitude stars within a 30.0 pixel (1.46 asec) aperture. For the  $M$  band the ADU/sec within a 10.0 pixel (0.486 asec) aperture is quoted.

### 3.6 Data Processing

We give here a brief description of the Clio data processing techniques. They are explained more fully in Appendix A.

To process Clio images, we first paired each image with a nod subtraction partner taken in the opposite nod position. In this pairing, every image served as the ‘science’ image exactly once, and most images also served as the ‘subtraction image’ once or twice. As we describe the processing we will draw a sharp distinction between the science and subtraction images, but we remind the reader that in general each image got to serve each role.

Since the sky background does slowly vary, a subtraction image should come as close to the science image as possible in the observing sequence. However, if possible no subtraction image should be used for two different science images. The restriction is desirable to keep the noise in nod-subtracted images as independent as possible. If the same subtraction image were used for many science

images, the noise on that image would be repeated many times in the subtracted frames, and would amplify in the final stack. This effect would be somewhat mitigated by parallactic rotation and image jitter from frame to frame, but the noise would still be needlessly increased. Our algorithm for making the nod subtraction pairs was to take as the subtraction image for a given science image the image in the opposite beam that came closest to it in the observing sequence, and had not been used as the subtraction image for a different science frame.

In some cases the numbers of images in each beam were not equal because some images had to be rejected due to bad seeing, AO loop failures, or other problems. It would not then be possible to avoid re-using some of the subtraction images. To deal with this case, we chose for each data set a maximum allowed distance in the observing sequence between the science image and the subtraction image. We typically set this distance to 13 images, so that, for example, image 83 could be used as the subtraction image for science image 70, but image 84 could not be because it came too far away in the sequence. If no unused subtraction images were available within the specified range of the science image, our algorithm was permitted to re-use a previously used subtraction image. However, if all subtraction images in the specified range had already been used twice, our algorithm would simply report a failure to pair images, and the user would have to set the maximum allowed sequence distance between the science and subtraction images to a larger value. For data sets taken under difficult conditions, in which many images were missing, we sometimes had to set the maximum sequence distance to 30 or even 50 images to obtain a successful pairing.

Once the pairings were complete, our processing pipeline divided each raw science image and its subtraction partner by the number of coadds, thus normalizing each image to a single frame. We then subtracted a similarly normalized

master dark frame from each image, and divided by a normalized flat. Finally, we applied hot-pixel fixing and bad column fixing to the images. Severe column patterns in the images at this point usually kept the hot pixel fixing and bad column fixing algorithms from finding many pixels, or any columns, that stood strongly enough above the noise to be fixed at this stage.

The next step was to perform nod subtraction. Before performing the subtraction, we scaled the subtraction image to match the science image in a user-defined region of clean sky. This partially compensated for any slight variation in sky brightness between the taking of the science image and the subtraction image. It also insured that the mean sky level on the nod-subtracted image would be zero. Note that if it were not for this scaling, the dark subtraction would have no effect, because the dark frame would be subtracted from each image and then the images would be subtracted from each other, cancelling the dark. We would, of course, not have used a dark in this case. However, since we did perform the scaling, we believed the subtraction of a dark frame prior to scaling was desirable, so that what was scaled would be actual sky background, not sky background plus dark counts and bias.

After nod subtraction the same pixel fixing and column fixing algorithms are applied again, with much more effect since the column variations in the image largely vanish with the nod subtraction, leaving any anomalous pixels or columns very obvious above the greatly reduced background noise.

Some residual column variations remain in the subtracted image. Often these are not uniform across whole columns — for example, a column may have anomalously high counts at the bottom and anomalously low ones at the top. To deal with these, we have developed a complicated algorithm (affectionately known as colfudge) which divides the image into 36x36 pixel blocks, attempts to find the

average deviation of each column in each block from the mean sky background, and then adds to each pixel an interpolated value meant to correct the column variation. This algorithm is powerful, but it does introduce artifacts and dim faint point sources slightly. It does not dim brighter point sources because it is designed to reject statistically deviant pixels from its column averages, so that sources bright enough to be rejected from the averages are not dimmed.

Although it has its problems, 'colfudge' (or, some similar algorithm) is indispensable in dealing with Clio data. By optimizing the parameters of colfudge we have significantly reduced the artifacts it introduced. See Appendix A for further information, and a description of an alternative algorithm due to Matt Kenworthy that performs the same function as colfudge with greater efficiency.

Nod-subtracted, hot pixel fixed, and 'colfudged' images were zero padded, shifted, and rotated to prepare for the final stack. All images were zero padded from their initial 256x320 size to 500x500 to avoid data loss in shift and rotation.

The rotation placed celestial north up on the images with an accuracy of about 0.2 degrees. Since we did not use the instrument rotator, a different rotation was required for each image: the parallactic angle plus a constant offset determined by the way the instrument was bolted onto the detector. We determined the offset to the required 0.2 degree accuracy by observing known binary stars. Note that an error in this offset produces a small systematic error in position angle astrometry, but does not introduce rotational blurring into the images. The parallactic angle was calculated very accurately from the UT time stamp of the images and the RA and DEC of the star. This is somewhat better than using the value written to the image header because it can be corrected for the duration of the exposure. We found that the clean, symmetrical stellar images produced by the MMT AO system at the  $L'$  band could be centroided so well, even if saturated, that the par-

allactic rotation of bright binary stars over just a 30 sec coadd sequence could be detected with great significance even for stars not close to the zenith and therefore not rotating rapidly: the internal precision of our parallactic angle calculation is far better than the accuracy with which we could obtain the constant offset. We do not expect that sufficient parallactic rotation occurred during any Clio coadd sequence to appreciably blur the science images.

We shifted our images based on a simple, iterative centroid on the primary star, with the shifts set to place the star at pixel coordinates exactly 250,250 on the zero padded 500x500 image. AO images are so clean and symmetrical at  $L'$  and  $M$  band that this simple centroiding procedure is sufficient even for the most strongly saturated of our survey targets. This was clearly demonstrated by two bright background stars detected in our images of the star GJ 820A, our second brightest object (after  $\epsilon$  Eri). The PSFs of the background stars were as tight as expected for perfect registration, though they had been registered only based on centroiding the heavily saturated primary star. An even more stringent test was our observations of the white dwarf companion to the star Procyon. Procyon is significantly brighter at  $L'$  than even  $\epsilon$  Eri, and yet the white dwarf showed a sharp PSF.

Our shift, rotation, and zero padding were all accomplished in a single bicubic spline operation, using bicubic spline routines from Press et al. (1992). Some clever handling was required to fold in the zero padding to this procedure. It was necessary to do so, however, because performing a bicubic spline on an already zero padded image would have caused severe artifacts at the zero pad boundary. We chose to shift and rotate in a single bicubic spline operation to maintain the highest image sharpness possible, after tests we performed showed that PSFs were noticeably broadened in bilinear interpolations, and even in bicubic spline

operations if the shift and rotation were performed separately so that two interpolations were required.

Note that because of the nodding and the shifts and rotations, and because our processing took pains to avoid significant loss of data at the image edges, our final stacked images contain data over a considerably larger area than the 258x320 pixel (12.5x15.5 arcsec) field of a single Clio frame. A corollary of this is at large separations from the star the ‘coverage’, or fraction of the total images that contributed data at a given location, is less than 100%, and the background noise goes up accordingly. This is why our sensitivity curves (see Sections 3.8 and 3.9) usually show the best sensitivity at moderate separations from the star rather than at the image edge.

We stacked our processed images to make a master image for each processing method using a creeping mean combine. This method of image stacking uses a single parameter, the rejection fraction, which we set to 20% for our standard master images. The mean of each given pixel through the image stack is computed; then the most deviant value is rejected and the mean is computed again. The procedure of rejecting the worst value and then recomputing the mean is continued until the required fraction of data points have been rejected; then the value of the given pixel in the master, combined image is set to the mean over the unrejected data points. Suresh Sivanandam greatly improved the speed of our processing pipeline by showing me how a creeping mean could be computed in  $N \log(N)$  time, rather than my previous  $N^2$  implementation.

We chose the creeping mean algorithm over the more standard median with sigma clipping because of its power to cause large deviant chunks of data to disappear without a trace, provided the fraction of deviant data is less than the creeping mean rejection fraction. The final mean is computed over a data set

in which the deviant data simply do not appear. Median combines, by contrast, are slightly biased by large amounts of deviant data.

Since we observe with the MMT instrument rotator off, ghosts, rays, and other artifacts in our images rotate slowly with respect to real celestial sources. The intense artifacts from the negative nod subtraction images also rotate. This creates exactly the situation where the creeping mean performs best: when a large, high-amplitude artifact rotates through a region, it makes the pixels there strongly deviant. However, provided the parallactic rotation was sufficient, a sufficient fraction of images will exist with good data in that region that the creeping mean will reject the artifact. Figure 3.7 compares a 20% rejection creeping mean with a median stack with a single iteration  $5\sigma$  sigma clip. Along the bottom of the images are simulated real planets present in 100% of the data. At top left is a simulated bright star image, also present in 100% of the frames. The other objects are simulated ghosts. In order from left to right and top to bottom they are present in 50%, 20%, 19%, 18%, 15%, 10%, and 5% of the data. As expected, neither the creeping mean nor the median stack helps against the 50% ghost; if half the data is bad, we're out of luck. For the other ghosts the creeping mean clearly does a better job. The median sigma clip removes the cores but is more strongly biased by the remaining bright haloes. Because data has literally been thrown away and was not then available to average over, the creeping mean has a slightly higher background noise level than the median. This is a small price to pay for effective removal of ghosts.

We comment, however, that side-by-side tests of the creeping mean against the median on real Clio data do not necessarily show a strong, clear advantage to the creeping mean. The median does better in some areas, the creeping mean in others. What is clear is that the creeping mean is a reasonable alternative to

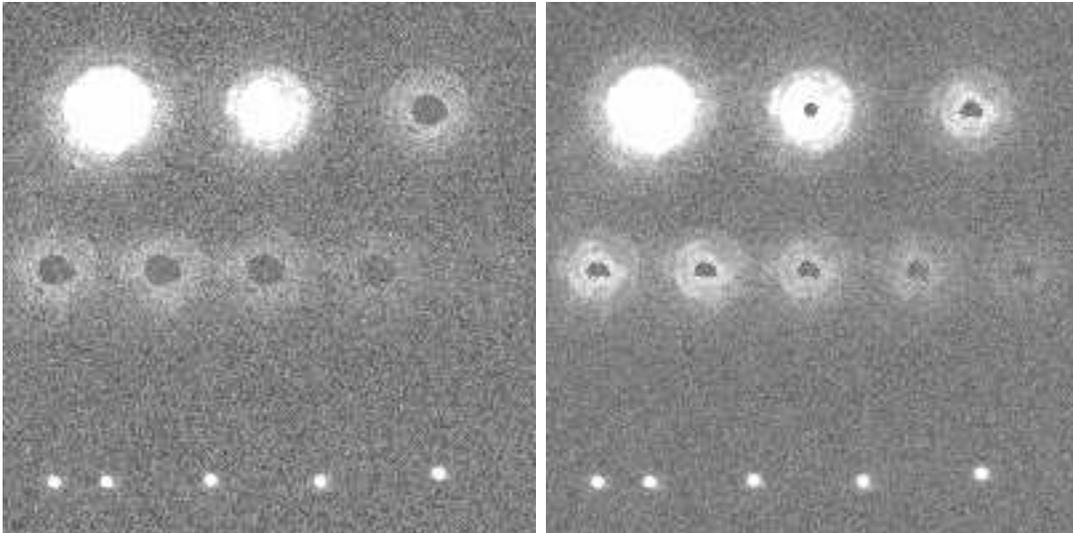


Figure 3.7 Comparison of different image stacking methods. On the left, a 20% rejection creeping mean. On the right, a median combine with a single-iteration  $5\sigma$  sigma-clip. The faint images at bottom are simulated planets; the bright upper left image is a simulated star. In order from left to right and top to bottom the other artifacts simulate ghosts present in 50%, 20%, 19%, 18%, 15%, 10%, and 5% of the data. The creeping mean clearly removes most of the ghosts more effectively, while the median removes only the inner cores and is more strongly biased by the halo regions. The sky noise is somewhat lower in the median combined image.

the median, and has worked well for us. Observers wanting the best possible sensitivity for heavily haunted or artifact-ridden images would do well to test the creeping mean against more conventional image combination methods to see if it might be the best choice for their particular data set. A creeping mean combine with an extreme rejection fraction of 50%, in particular, sometimes seems to work magic on very bad data sets, at the cost of increased background noise since half the data has been rejected.

Our implementation of the creeping mean algorithm ignored pixels that were equal to zero. This was very important, because the rotation in our data sets made it inevitable that valid data regions on some images would overlay zero padded regions on others. It was essential to keep the zero pads from being averaged in to the real data. Since our sky backgrounds have mean zero, zero is a possible data value. However, the creeping mean stack is performed on fully processed images with FITS floating point pixel values. Valid data pixel values of exactly zero are vanishingly improbable.

After the creeping mean stack, we unsharp masked the final image to remove the bright stellar halo. This also helped with dark haloes left over from the negative nod subtraction images. Our procedure for unsharp masking was to create a version of the image convolved with a gaussian kernel of  $\sigma = 5\text{pix}$ , and then subtracted this from the original image. The FWHM of this gaussian was 11.8 pixels, as opposed to about 3 pixels for a typical PSF, so the unsharp masking had no appreciable effect on point sources. A few innovations in our masking strategy substantially improved our results. First, we did not use Fourier methods in the convolution, as these would have caused artifacts at the zero pad boundaries. Rather, we simply implemented the mathematical definition of a convolution on a pixel by pixel basis. That is, we set each pixel in the convolved image equal

to the integral over all the pixels in a region surrounding that pixel on the original image times the normalized gaussian kernel centered on the pixel. This was much slower than a Fourier convolution, but it gave us very useful flexibility. First, we defined the region over which the integration was performed as a disk with radius set so that the gaussian value at its edge was only 1% of the central value. Second, we rejected zero-valued pixels in the main image from consideration in the integral. This prevented artifacts at the edge of the zero pad. Second, we performed several iterations of high-rejection sigma-clipping on the set of pixels contained in the disk on the original image. Only pixels that did not get rejected were included in the integral. This meant that sufficiently bright point sources were rejected from the convolved image. Therefore they were totally unaffected by the unsharp masking operation, rather than being only slightly affected. This innovation also greatly reduced the tendency for dark halos to form around bright point sources as a result of the unsharp masking.

The unsharp masking was the final step in our image processing pipeline.

The above describes our baseline processing method. There were six specializations of this method, which we called the 'b', 'c', 'd', 'e', 'x', and 'y' processing strategies, with the baseline method itself called 'a'.

In the 'b' processing method a master PSF image is constructed and subtracted from each science image after the nod subtraction step. The master PSF image is made from a stack of images analyzed in the same way as regular science images up to the final stack, except that they are not rotated to account for parallactic rotation. This means that celestial objects rotate about the central star on these images, but artifacts fixed to the telescope or detector do not rotate. Thus, while in the baseline processing, the creeping mean stack rejects ghosts, rays, and other artifacts, in the construction of the master PSF image the creeping mean rejects

any real sources, creating an image with only the rays, ghosts, and other artifacts. The most important artifacts present in this image are not ghosts or rays but persistent speckles in the inner part of the stellar PSF. They are presumably due to imperfections in the primary mirror, perhaps due to gravitational flexure, that are too high in spatial frequency to be corrected by the AO system, but smooth and low enough in spatial frequency that they create coherent speckles rather than merely scattering the light. These speckles evolve slowly, but for a typical Clio integration time of 1.5-2 hours on a single target the speckles remain constant enough that they can be effectively subtracted.

In order to avoid subtracting real sources, we make two master PSF images, one by stacking the first half of the science data, and the other by stacking the second half. We subtract the second master PSF from images taken in the first half of the observing sequence, and vice-versa. For good data sets with sufficient rotation, the 'b' processing method offers substantially better sensitivity close to the star than the baseline 'a' method. We pay careful attention to the amount of parallactic rotation present in each data set, to avoid calculating sensitivities from PSF subtracted data so close to the star that real sources would be subtracted away. For many stars, the parallactic rotation was large enough that this was not a concern.

The 'c' data reduction method is designed to produce usefully PSF subtracted images for data sets with insufficient parallactic rotation. A model of the PSF is constructed for each individual image by averaging arcs at constant radius from the central star. Typically the length of these arcs is set to 20-30 pixels. Creeping mean averaging of the arcs is used to reject real point sources from the averaged PSF and thus minimize the extent to which the PSF subtraction would dim a real source. The model PSF generated from each individual science image is

subtracted from that image prior to the final stack. This reduction method can be quite powerful, but was not usually used for Clio data analysis. We felt that except in special circumstances the 'b' reduction method was more powerful.

The 'd' reduction method was developed to combat a problem we sometimes experienced with creeping mean combined images of bright stars. If the star had an extended halo, there could be parts of the image that were very bright from the stellar halo on images from one nod position, but dark because of the negative stellar halo on images from the other position. This would create a strong bimodality within the final image stack for pixels in certain regions. The creeping mean would randomly select one mode of the bimodal distribution, and reject pixels from the other mode. The chosen mode could be different for neighboring pixels, so that one pixel would be very bright in the final stacked image, while a nearby one would be dark. The data quality in the affected regions was totally destroyed by this intense 'bimodality noise'. We note that it is likely a median/sigma clip stack would have experienced a similar problem; this is not an exclusive weakness of the creeping mean. We solved the problem by unsharp masking each individual image before the final stack. This removed the stellar haloes and insured the stack was performed on images with good, zero-mean sky values in most regions. The improvements in the data quality were enormous.

The 'e' data reduction method combines the 'b' and 'd' methods: the master PSF image is subtracted, and then the pre-stack unsharp masking is performed. This was helpful for data sets where regions in the PSF subtracted images were affected with bimodality noise.

The 'x' data reduction method uses a variant on nod subtraction that avoids the dark negative images. A single master sky images is made for the whole data set, and this is subtracted from each individual science image, rather than

the assigned subtraction image. The master sky image is made by taking the science/subtraction image pairs and slicing each image along the perpendicular bisector of the line between the position of the star in the science image and the position of the star in the subtraction image. The halves with the stars are then discarded, and the halves without the stars are joined, after they have been scaled to match in a region of blank sky in a manner identical to the way we scale the subtraction image in baseline processing. After a starless sky image has been made from each science/subtraction image pair, the sky images are scaled to match each other and then stacked with a 50% rejection creeping mean. The usefulness of the 'x' processing method varies enormously from one data set to another, presumably depending on how stable the sky background was. In all cases the 'x' processed image has a noisier sky background than the 'a' image. However, in some cases the differences are very slight, and the absence of the dark nod subtraction artifacts on the 'x' image seems almost miraculous, delivering good sensitivity in regions of the image previously destroyed by nod artifacts. In other cases the 'x' image is hopelessly full of intense column noise and other artifacts.

The 'y' image reduction method is a combination of the 'x' and 'd' methods, in which the images are unsharp masked after the subtraction of the master sky image but before the final stack.

In our sample we had some binary stars of near equal brightness. For these, we used a scaled version of the PSF of each star to subtract the other. We did this prior to the final stack, on an image-by-image basis. The resulting PSF subtraction was generally substantially better than that obtained by any other method. The baseline version of this we called the 'f' image processing method; a version that combined pre-stack unsharp masking was called 'g'.

The final images produced by each method were unsharp masked after the final stack, even if prestack unsharp masking had been applied.

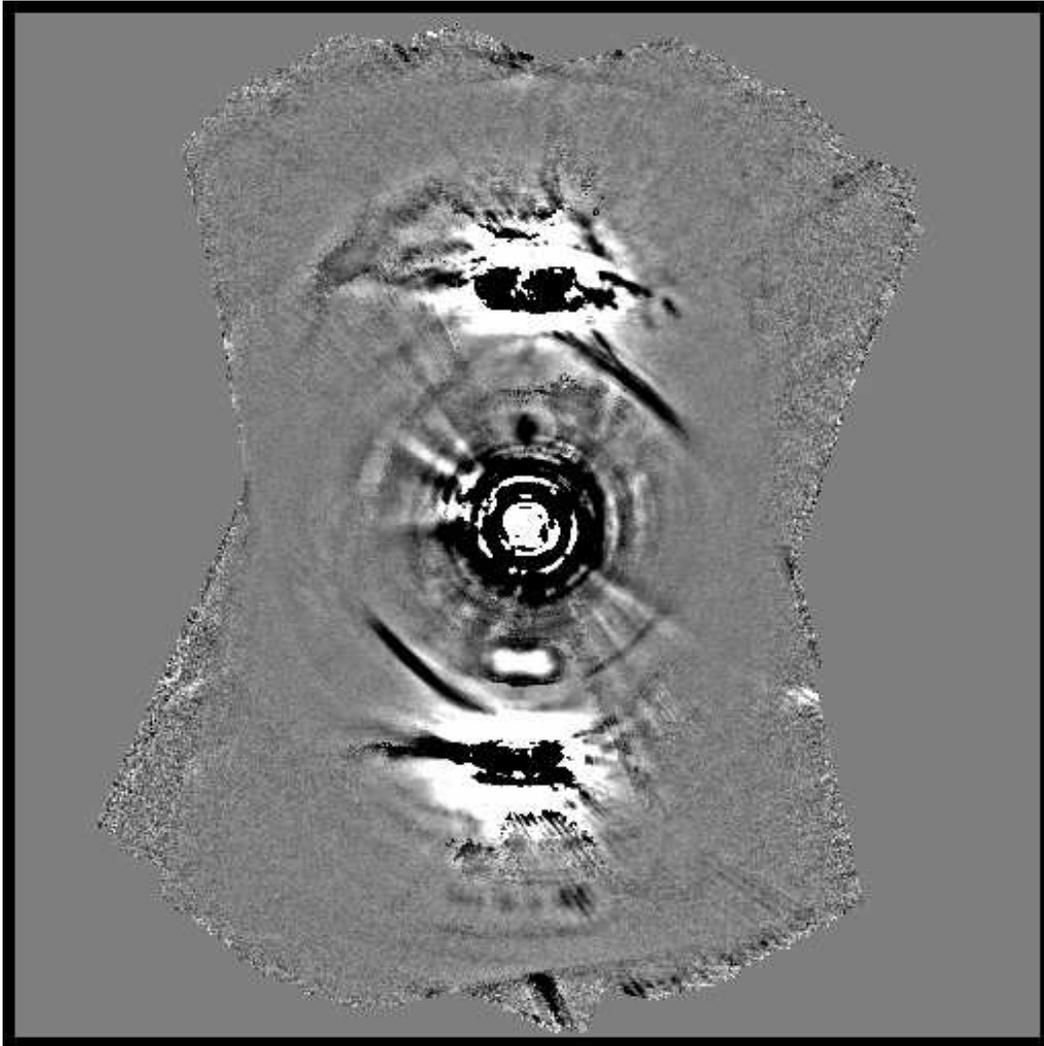


Figure 3.8  $L'$  Image of  $\epsilon$  Eri, Processed with Method 'a' (baseline processing). Compare with the 'd' method image (Figure 3.9) to see the effect of prestack unsharp masking, or with the 'x' image (Figure 3.10) to see the effect of the master sky subtraction image.

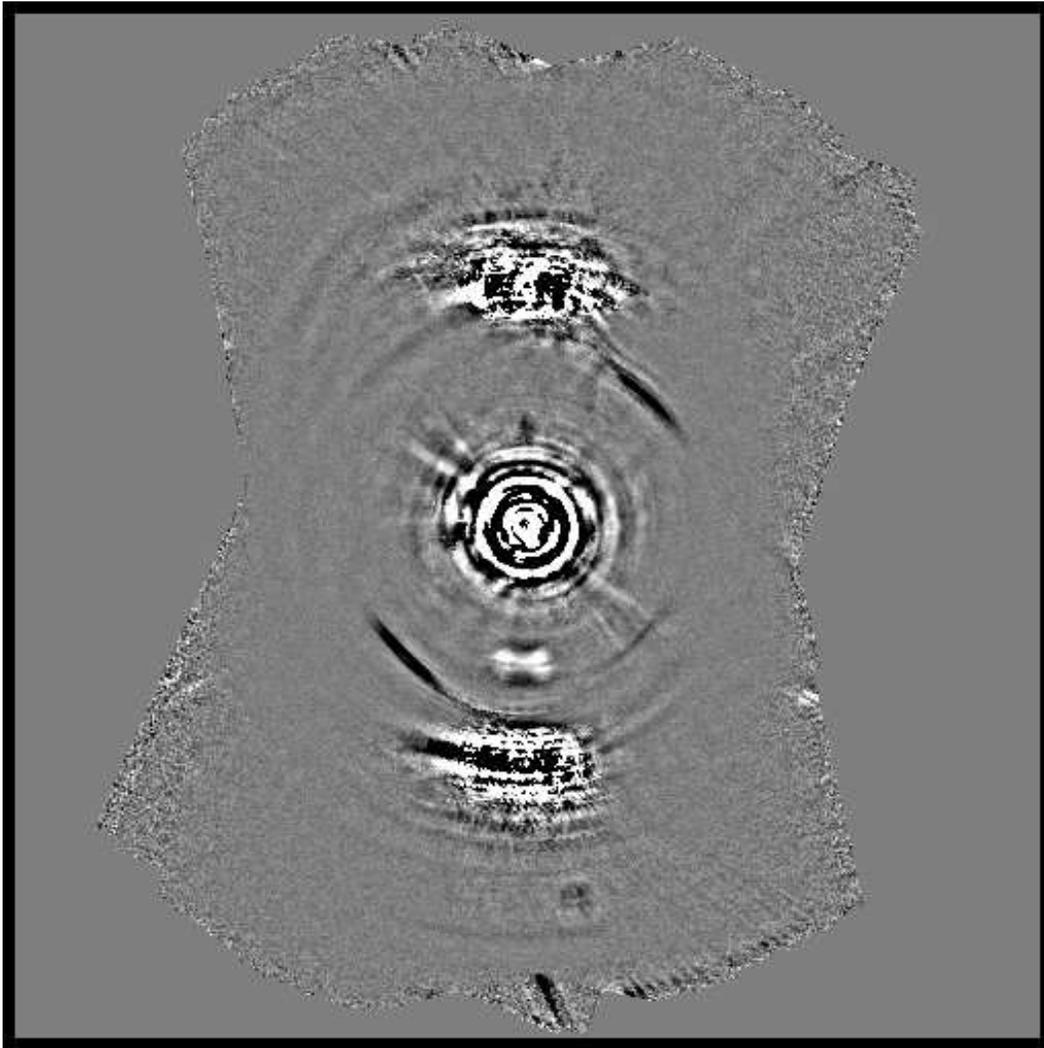


Figure 3.9  $L'$  Image of  $\epsilon$  Eri, Processed with Method 'd' (pre-stack unsharp masking). Compare with the 'y' method image (Figure 3.11) to see the effect of the master sky subtraction image.

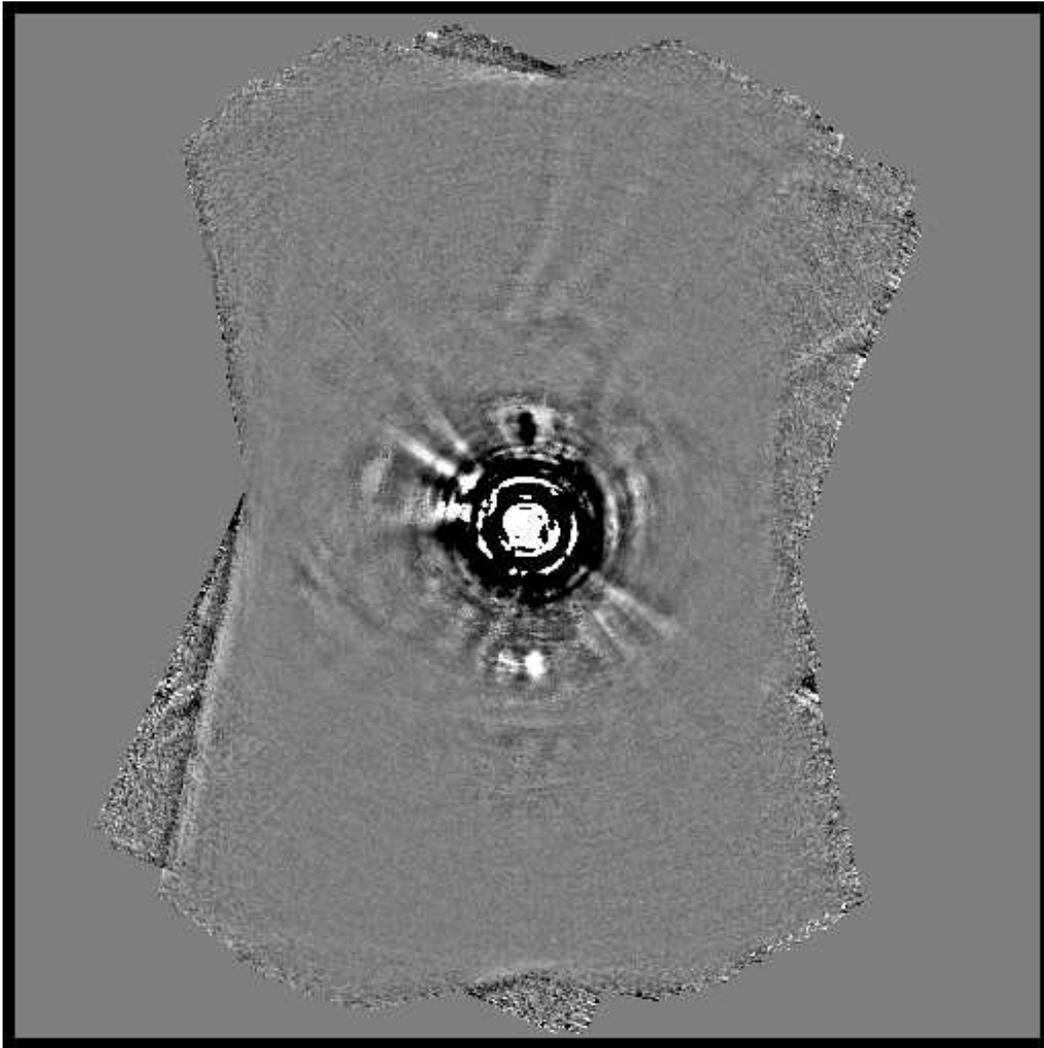


Figure 3.10  $L'$  Image of  $\epsilon$  Eri, Processed with Method 'x' (use of a master sky image to eliminate negative star images from nod subtraction). The master sky image subtraction worked very well for this data set.

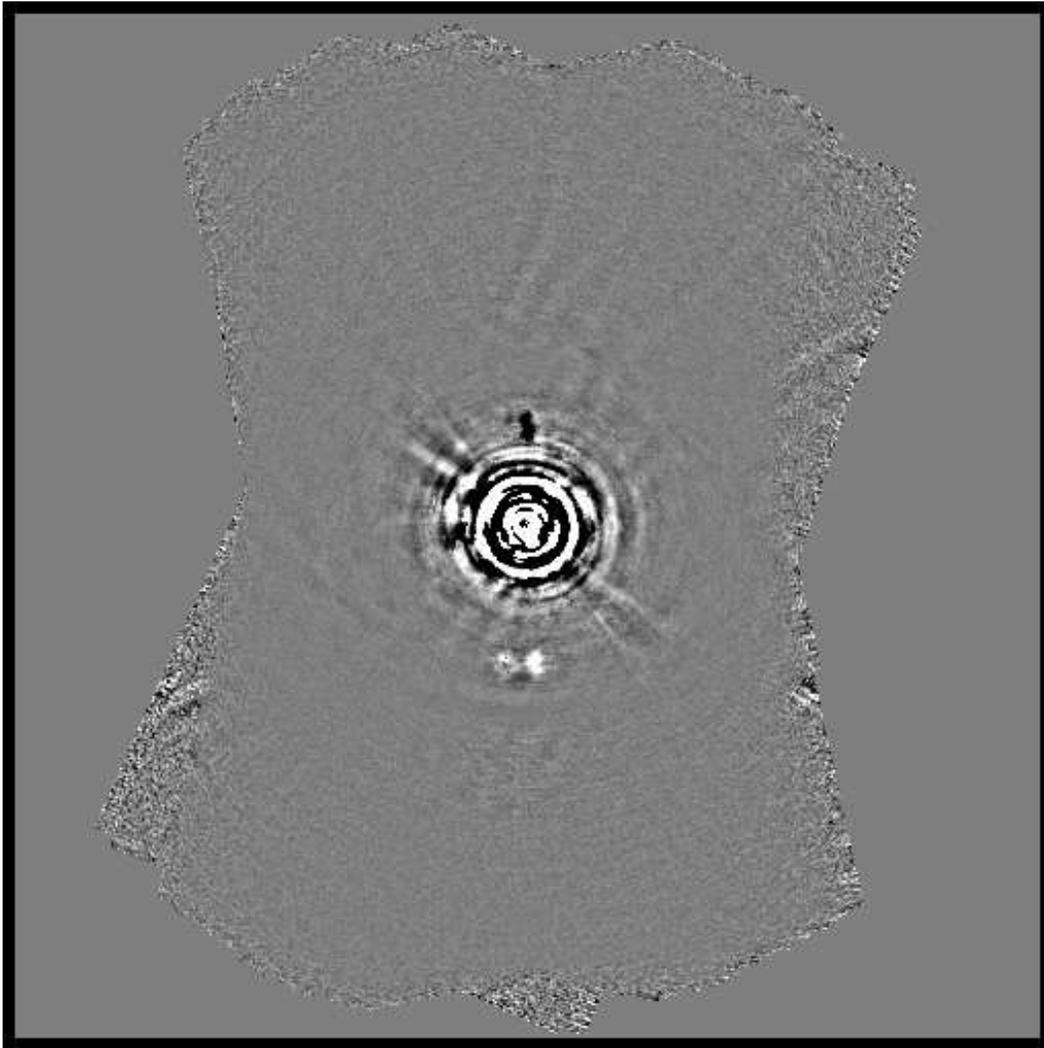


Figure 3.11  $L'$  image of  $\epsilon$  Eri, Processed with Method 'y' (use of a master sky image and pre-stack unsharp masking). The master sky image subtraction worked very well for this data set.

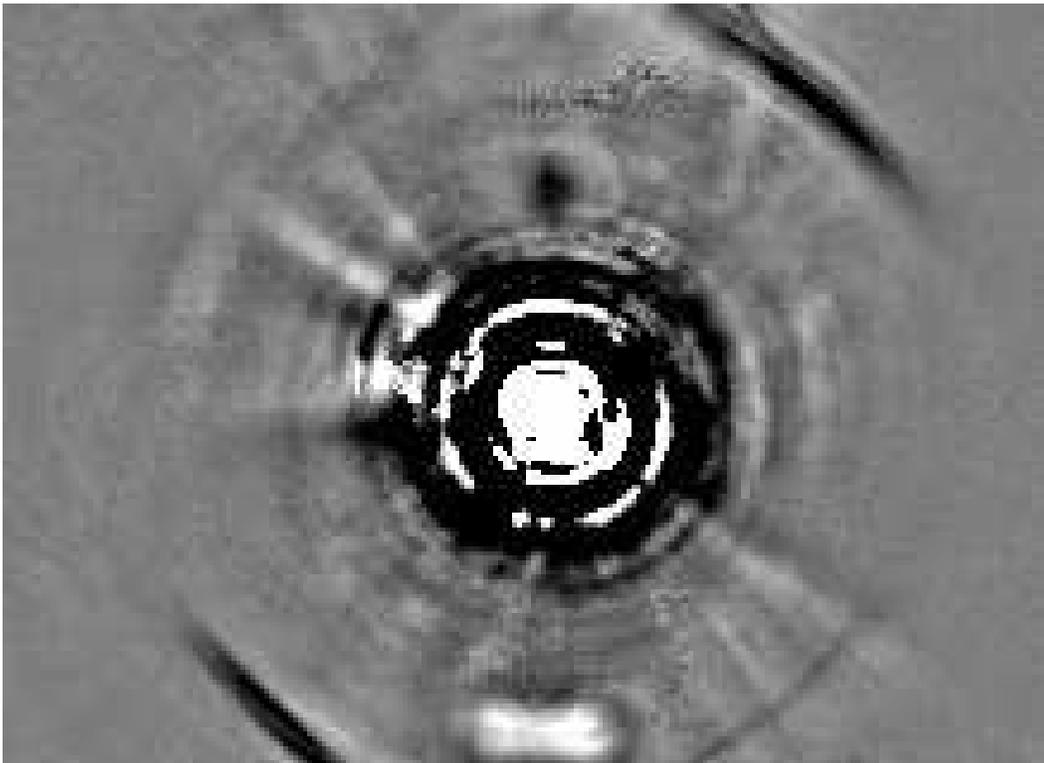


Figure 3.12  $L'$  Image of  $\epsilon$  Eri, Processed with Method 'a' (baseline processing). This is a closer view to be compared with the 'b' image (Figure 3.13) to see the effect of PSF subtraction.

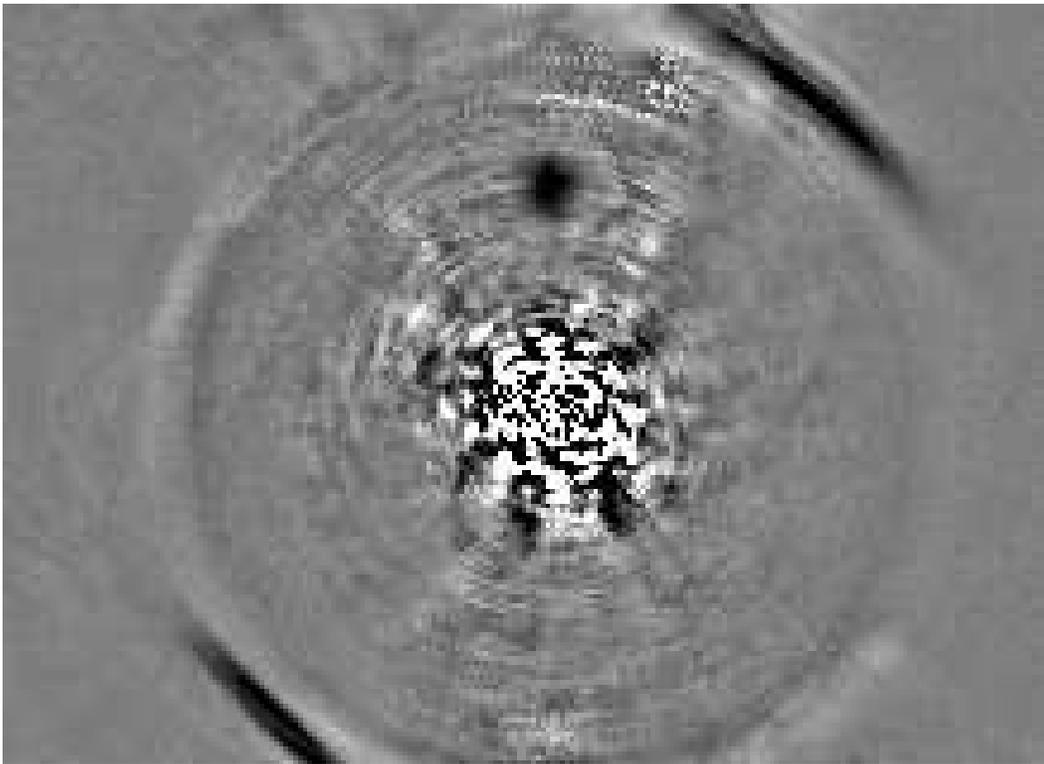


Figure 3.13  $L'$  Image of  $\epsilon$  Eri, Processed with Method 'b' (PSF subtraction using parallactic rotation to construct a master PSF based on images of the science target). Compare with the 'e' method image (Figure 3.14 to see the effect of prestack unsharp masking on PSF subtracted images).

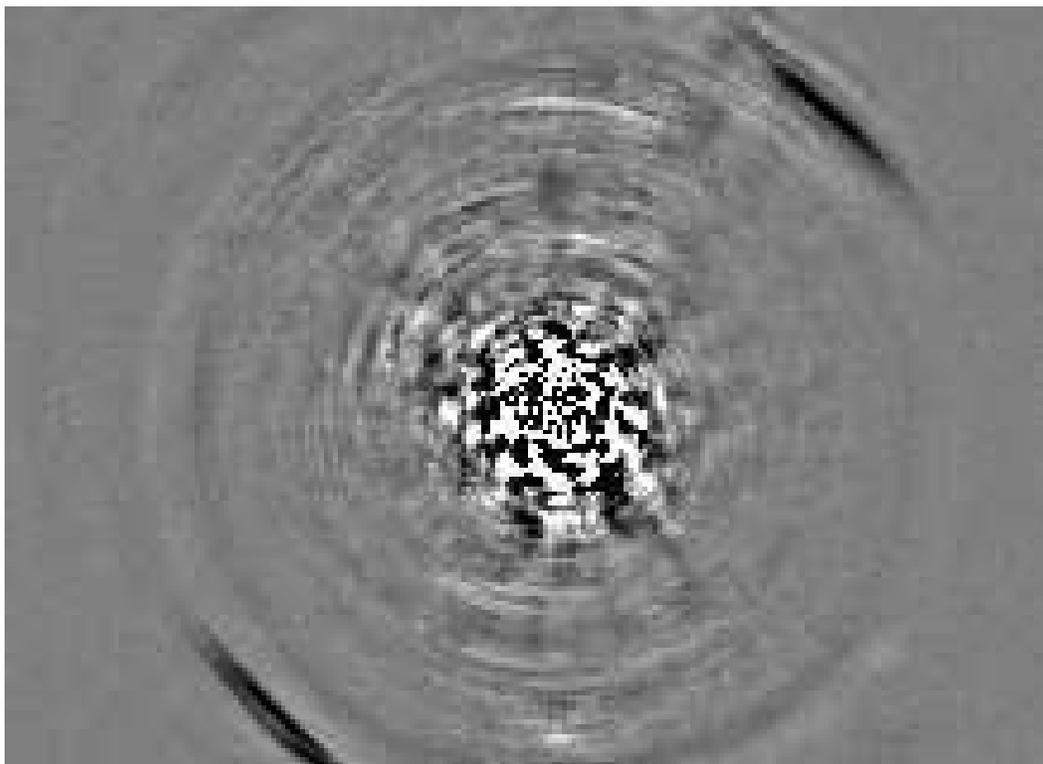


Figure 3.14  $L'$  Image of  $\epsilon$  Eri, Processed with Method 'e' (PSF subtraction and pre-stack unsharp masking).

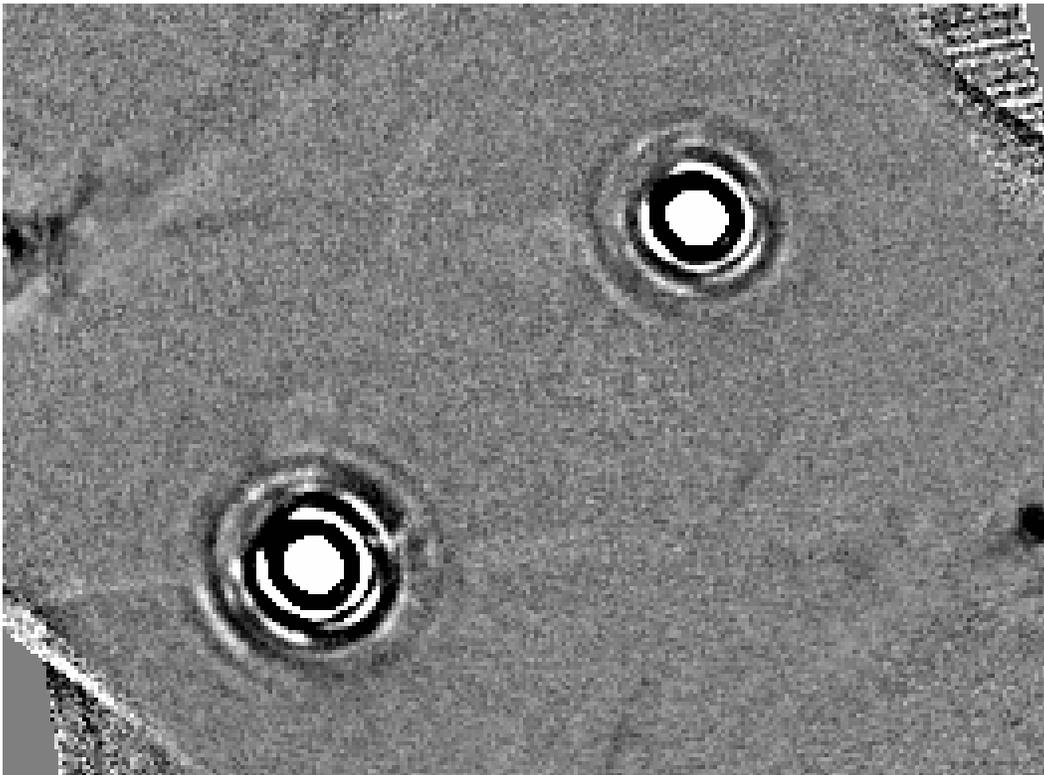


Figure 3.15 *M* Band Image of  $\xi$  Boo AB, Processed with Method 'a' (baseline processing). Compare to the 'f' method image (Figure 3.16) to see the effect of binary star PSF subtraction.

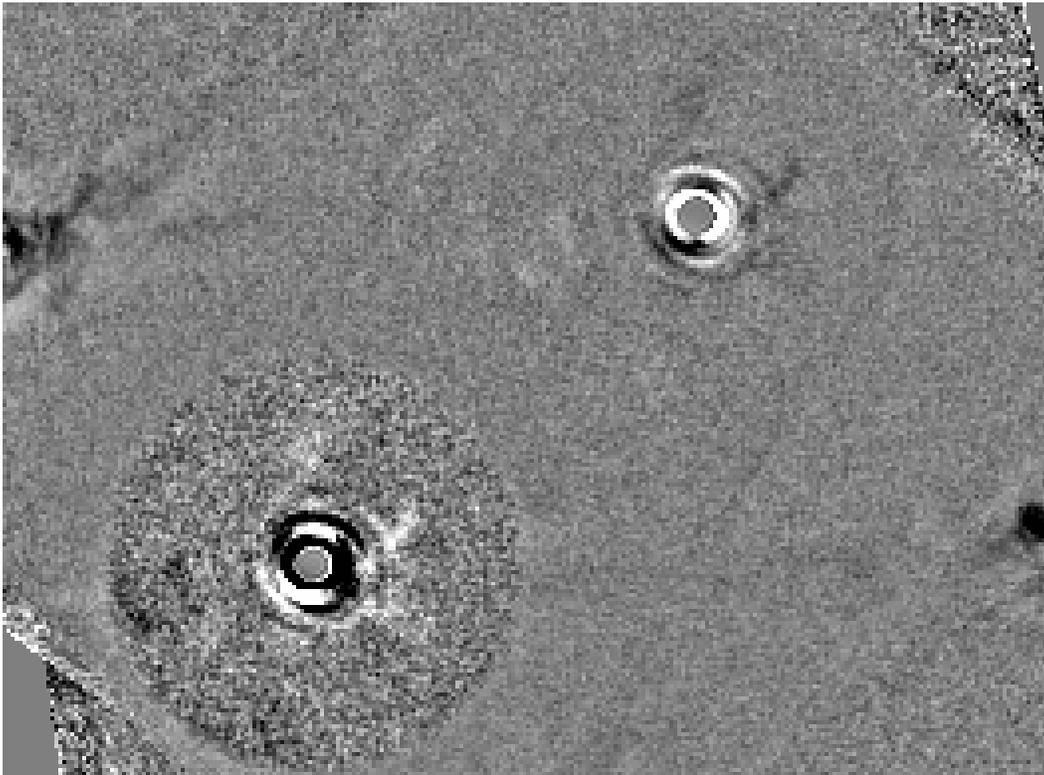


Figure 3.16 *M* Band Image of  $\xi$  Boo AB, Processed with Method 'f' (subtraction of each component of a binary using the PSF from the other).

### 3.7 Sensitivity Analysis

Sensitivity estimation begins where the image processing ends. For most of our data sets, the image processing ended with final images produced by the 'a', 'b', 'd', 'e', 'x', and 'y' processing methods. In some cases the 'x' and 'y' images were too low quality to be useful, so these were not included in the sensitivity analysis. In other cases, there was insufficient rotation for the standard PSF subtraction to work, so the 'b' and 'e' methods were not used. In some of these cases we made a 'c' image to compensate for this deficiency. Finally, for binary stars of near-equal brightness we made 'f' and 'g' images in addition to all the others.

#### 3.7.1 Three Sensitivity Estimators

So far, no unambiguous imaging detections of planets orbiting normal stars have been made. Our survey, and other comparable surveys such as Masciadri et al. (2005), Kasper et al. (2007), Biller et al. (2007), and Lafrenière et al. (2007b) have all arrived at null results. The science content of a survey null result is always some form of an upper limit on the occurrence of extrasolar planets. The accuracy of this upper limit depends entirely on the accuracy of the method chosen to estimate the sensitivity of the survey images to faint planets.

The importance of choosing an accurate sensitivity method can hardly be overestimated. **All of the science** of a null result planet imaging survey hinges on the sensitivity estimation method. If the method is not accurate, neither will be the science result of the survey.

There are two key choices to be made in selecting a sensitivity estimation method. First, what method will be used to determine point-source sensitivity from the pixel statistics of a given region of the image? Second, in what regions will the sensitivity be calculated? **The first question is far more important than**

**the second.** Once a good sensitivity estimation method has been adopted, it will give reasonable results within any sufficiently large, sensibly defined region.

We will discuss 3 different answers to the first question; that is, three different ways to calculate the point-source sensitivity in a given region of the image. Answers to the second question are discussed in the sixth paragraph of Section 3.7.2.

**Method 1:** The simplest solution is to calculate the RMS of the pixel values within the assigned region and translate this to a point source sensitivity by some means. If the pixel-to-pixel noise is gaussian and spatially uncorrelated (that is, the noise is independent from pixel to pixel), a good way to convert the pixel RMS to point source sensitivity is to define an aperture diameter near the FWHM of the PSF, and multiply the pixel RMS by the square root of the number of pixels contained within this aperture. This gives the expected RMS of the sum of pixels within any randomly chosen aperture of the selected diameter on the image.

To rephrase this in equation form, if  $\sigma_{pix}$  is the RMS of the pixel values within the given region and  $r$  is the radius of our chosen aperture for point source detection (likely about half the FWHM), the  $1\sigma$  point source sensitivity  $\sigma_{PSF}$  is:

$$\sigma_{PSF} = \sigma_{pix} \sqrt{\pi r^2} = \sigma_{pix} r \sqrt{\pi}. \quad (3.1)$$

In our case an aperture correction must then be applied, because we calculate the sensitivity in magnitudes based on photometry of standard stars with an aperture of radius 30 pixels (1.46 asec). This large aperture was chosen to reliably include all the flux even in conditions of poor seeing. The optimal radius  $r$  for detecting faint point sources against the noise is much smaller; typically  $r = 2.5$  pixels, or 0.12 arcsec. We perform aperture photometry on an unsaturated PSF image using the 30 pixel radius and the much smaller radius  $r$ . We then

aperture-correct the calculated  $\sigma_{PSF}$  by multiplying it by the ratio of the flux contained in the large aperture to that in the smaller aperture. If the sensitivity is desired in the form of  $\Delta$ -magnitudes relative to the central star, rather than actual photometric magnitudes, the calculated  $\sigma_{PSF}$  may simply be ratioed to the exposure-scaled flux on an unsaturated PSF image within the narrow sensitivity aperture, and no aperture correction is needed. The  $\sigma_{PSF}$  value is finally multiplied by the chosen sensitivity threshold, ie 5 for a  $5\sigma$  sensitivity limit, and then converted to magnitudes or  $\Delta$ -magnitudes as desired.

Method 1 is not valid if the pixel variations are spatially correlated (that is, if the noise is not independent for adjacent pixels). This is because Equation 3.1 assumes the noise is independent from pixel to pixel. It predicts  $\sigma_{PSF}$ , the RMS variation of the sum over an aperture, from  $\sigma_{pix}$ , the pixel-to-pixel RMS, under the assumption that the pixels are independent. This assumption is violated if there is spatially correlated noise such as speckle-noise close to a star or detector pattern noise. An aperture randomly placed over an image with correlated noise may fall on a clump of bright or dark pixels. The variation in the sum over the aperture will thus be much more than would be predicted from the RMS of single pixel values under the assumption of uncorrelated noise. If the noise is spatially correlated, the assumption on which Method 1 is based is mathematically invalid. The method will produce unreliable results under such conditions no matter the size of the region over which it is applied. In Section 3.7.3 we will consider the extent to which the sensitivity is overestimated in such cases.

**Method 2:** Method 2 is a very simple solution to the problem Method 1 has when correlated noise is present. Rather than calculating the single-pixel RMS and predicting the RMS over an aperture, one simply sums up actual apertures on the image, and measures the RMS variation of the sums. The discussion of

the possible need for an aperture correction under Method 1 applies here as well. The final RMS, as before, is multiplied by the desired sensitivity in sigma and converted into magnitudes,  $\Delta$ -magnitudes, or whatever form is desired. If correlated noise is present it increases the variation of the aperture sums and is naturally included in the measurement.

**Method 3:** Method 3 is an extension of Method 2. Method 2 can be thought of as an aperture photometry method: the RMS of the summed flux in apertures laid down on the image is calculated. The other standard way of doing photometry, of course, is PSF fitting. Method 3 is a PSF fitting sensitivity method. It requires a good model or measurement of the expected PSF for faint sources in the data set. A PSF is fit to the noise of the image at various locations within the sensitivity region, and the RMS of the fit amplitudes is calculated. For our implementation of this method we forced the PSF to be centered on a given pixel, and did a simple 2-parameter least square fit within a radius of 6 pixels. The parameters of the fit were a constant sky background level and the amplitude of the PSF. Once the RMS of the best-fit PSF amplitudes has been calculated, it may be multiplied by the desired sigma value for the sensitivity measurement just as for the other methods. If the PSF has been properly normalized, no aperture correction is necessary for this method.

### 3.7.2 Our Choice of Estimator

We rejected Method 1 because of its inability to account for correlated noise, which is always strong in the speckle-dominated areas of AO planet search images.

We implemented Methods 2 and 3. They produced very similar, almost interchangeable results in general. In the end we decided somewhat subjectively that Method 3, the most sophisticated, gave the most reliable results. It seemed

to make cleaner-looking sensitivity maps, and not to be biased by large ghosts and other diffuse features as Method 2 sometimes was. We carried out all of our final sensitivity analysis using Method 3. However, we used both Method 2 and Method 3 for identifying possible sources in our data.

In implementing Method 2, we set the size of the aperture to a different optimum value for each of our data sets. This value was determined from the unsaturated PSF from the short, PSF-monitoring exposures for that star. The optimum detection aperture maximizes the ratio of enclosed flux to enclosed noise. While the behavior of enclosed noise in the speckle-dominated regime is unknown, in the background-limited regime we expect the noise to go up as the square root of the number of pixels enclosed within the aperture. We determined this value experimentally for each PSF by the crude expedient of summing the flux in a finely spaced range of apertures centered on the PSF core, ratioing this sum to the square root of the number of enclosed pixels, and picking the aperture radius that maximized the ratio. This optimal value was usually between 1.9 and 2.9 pixels (0.09 and 0.14 asec), depending on the seeing. This yields aperture diameters slightly larger than the PSF FWHM, and it seems a very reasonable choice for detecting point sources against either background or speckle-dominated noise. As discussed above, we applied aperture corrections derived from the PSF image to scale flux from this small aperture to the 30 pixel (1.46 asec) radius aperture used for our photometric standard stars (10 pixels, or 0.49 asec, was used for *M* band data), and then used our photometric calibration to convert the sensitivity to magnitudes. Our aperture corrections were usually factors of 2 to 3, so neglecting them would have rendered our sensitivity estimates extremely inaccurate.

One further question remains in the implementation of Method 2, and that is how the different places the aperture would be placed on the image were cho-

sen. We elected to place it centered on every pixel in the image, creating a new image that was effectively an aperture photometry map of the old. The RMS variation of the apertures was then easily calculated in any region of our choice on the aperture photometry map. We note that adjacent pixels in this aperture photometry map are not independent. This non-independence does not affect the RMS calculation, as long as the RMS is calculated over a sufficiently large region. ‘Sufficiently large’ in this context means a region large enough to contain several PSFs without significant overlap.

We implemented Method 3 in the same way; fitting a PSF centered at every pixel in the image, to create a PSF-fitting photometry map of the image. Our rather large choice of a 6.0 pixel fitting radius for the PSF may need some explanation. We felt that it was important to include a sky background in the least-square PSF fit, so that the fits could ride easily over diffuse ghosts, remaining unsubtracted portions of the stellar halo, etc. We found that with background subtraction the fitting radius needed to be substantially larger than the FWHM to insure that the background was sufficiently well sampled to be accurately fit.

The only remaining questions are over what regions we calculated our final RMS values, and what sigma value we chose for our sensitivities. For areas far from the star, we calculated the sensitivity in a disk of radius 8 pixels (0.39 arcsec, or about  $3 \lambda/D$ ). This was large enough to span many PSF-sizes, but small enough to sample the local noise properties of the image in a reasonable way. For areas within 60 pixels (3 arcsec) of the star we felt that the sensitivity might vary radially too quickly for our standard 16 pixel-wide disk to adequately sample it (within 8 pixels of the star core this was obviously true!). Therefore, within 60 pixels of the star we calculated the sensitivity not within a disk but along an azimuthal arc 45 pixels (2.2 arcsec) long and 1 pixel wide, at a fixed radius from the

star. We centered disks or arcs on every pixel of the aperture photometry and PSF photometry maps, producing sensitivity maps of the entire image based on both Methods 2 and 3. We chose to calculate  $10\sigma$  sensitivities, based on blind planet detection experiments that indicated our completeness at  $7\sigma$  or  $5\sigma$  was substantially below 100%.

In order to make a master sensitivity image based on all the final processed images available for a given star (ie the master images produced by the ‘a’, ‘b’, ‘d’, etc. processing methods), we first smoothed the sensitivity map from each master image slightly by convolving it with a gaussian kernel with  $\sigma = 3.0\text{pix}$ . This diminished the ragged appearance resulting from the hard edges of the 8.0 pixel radius sensitivity disk. We then minimum-combined the sensitivity maps from all the processing methods used on each star to produce a single master sensitivity image giving the best sensitivity obtained by any processing method at each point in our image. We made master sensitivity images using both Method 2 and Method 3, but performed all of our final analyses using the Method 3 images only. Figures 3.20 through 3.64 present our final sensitivity maps for every target, transformed into magnitude contour images.

### 3.7.3 Testing the Three Estimators

To verify that our choice of estimator and our implementation thereof were accurate and effective, we performed tests of Methods 1, 2, and 3 on two different images from our survey.

First, we tested the estimators on different regions of a PSF-subtracted (‘e’ method) image of 61 Cygni A, the second brightest star in our survey (after  $\epsilon$  Eri). This image has large areas close to the star that are dominated by residual speckle noise, which has strong spatial correlations. Far from the star are areas of clean sky where the noise should be largely uncorrelated; however, very close

examination suggests residual faint rays and other artifacts from the star may have introduced some correlation even here.

We tried out each estimator in the 4 quadrants of an inner and outer annulus close to the star. All 8 regions are dominated by speckle noise. We also calculated the sensitivity in 4 regions of clean sky far from the star. Figure 3.17 shows the image with the regions in which we tested the estimators outlined in white. Table 3.8 gives the sensitivity results in magnitudes. Averaged over the 8 speckle-dominated regions, the Method 1 estimator reports sensitivity 1.09 magnitudes better than Method 3 and 0.77 magnitudes better than Method 2. For the blank sky regions, Method 1 estimates 0.37 magnitudes better sensitivity than Method 3 and 0.25 magnitudes better sensitivity than Method 2. As expected, the Method 1 estimator appears to overestimate the sensitivity in the presence of correlated noise.

It was interesting that even in the blank sky regions Method 1 estimated better sensitivity than the other methods. Do the estimators inherently differ, or was there low-level correlated noise in the sky background? Also, Method 2 estimated better sensitivity than Method 3 in the speckle-dominated regime. Could Method 3 be estimating lower sensitivity because the PSF fitting radius extends closer to the star and samples higher noise there? To answer these questions, we performed another set of tests on 'd' method images of a much fainter star, GJ 625. Here there were image regions where sky noise appeared completely uncorrelated even under the closest examination, but there were also regions slightly affected by nod subtraction artifacts. The latter had significant correlated noise, without the steep gradient of noise experienced close to the star. We defined 5 boxes of clean sky, and 5 boxes of strong correlated noise from the nod subtraction regions. Figure 3.18 shows the image, and Table 3.9 gives the sensitivity

results from each estimator. For the 5 clean sky regions, Method 1 estimated, on average, a sensitivity 0.006 magnitudes worse than Method 2 and an 0.128 magnitudes better than Method 3. For the correlated noise regions, Method 1 gave an average sensitivity estimate 0.75 magnitudes better than Method 2 and 0.92 magnitudes better than Method 3.

It is not clear whether our implementation of the PSF fitting inherent in Method 3 indeed causes it to underestimate the sensitivity slightly (perhaps 0.3 magnitudes) in the presence of a strong gradient in the noise, such as is found close to the star, or whether it consistently delivers results 0.2-0.3 magnitudes more pessimistic than Method 2 in the presence of correlated noise. What is clear is that all 3 methods agree well in regions of truly uncorrelated sky noise, but **Method 1 consistently overestimates the sensitivity by nearly a magnitude in regions of strong correlated noise**. Even when the correlated noise is weak, as in the sky regions of the 61 Cygni image, the Method 1 estimator appears to be affected at the 0.3-0.4 magnitude level. Differences between Methods 2 and 3 are far less significant. Finally, it is clear that in Method 3 we have chosen an accurate but also a conservative sensitivity estimator.

#### 3.7.4 Blind Sensitivity Tests

Since having an accurate sensitivity estimate is so important to the scientific validity of a planet imaging survey null-result, we tested ours to make sure the rather complicated sensitivity estimation methods described above had produced valid results, and to determine the appropriate sigma values to quote for our final sensitivity. In these experiments, simulated planets were inserted at random locations into the raw data for a given star. The flux of each simulated planet was scaled to 5, 7, or  $10\sigma$  significance using the (Method 3) final sensitivity map for that data set. The data was then processed in exactly the same way as

Table 3.8. Different Methods of Sensitivity Estimation Tested on GJ 820A

Annulus	Sector	Aperture	PSF Fitting	Single Pix	Sing. - PSF
		Sensitivity	Sensitivity	Sensitivity	Difference
		( $10\sigma, L'$ mag)	( $10\sigma, L'$ mag)	( $10\sigma, L'$ mag)	( $L'$ mag)
Inner	NE	11.43	11.28	12.23	0.95
Inner	SE	11.1	10.77	11.95	1.18
Inner	SW	11.42	10.98	12.18	1.20
Inner	NW	11.41	11.01	12.18	1.17
Second	NE	12.87	12.61	13.61	1.00
Second	SE	12.72	12.4	13.46	1.06
Second	SW	12.67	12.35	13.45	1.10
Second	NW	13.02	12.65	13.74	1.09
Outer Arc	1	15.49	15.39	15.73	0.35
Outer Arc	2	15.54	15.41	15.74	0.32
Outer Arc	3	15.51	15.39	15.74	0.35
Outer Arc	4	15.52	15.41	15.87	0.46

Note. — The inner and second annuli are placed in the contrast-limited, speckle noise dominated region close to the star, while the outer arc is placed over largely clean sky. The single pixel method overestimates the sensitivity in the inner, speckle-dominated regions because it does not account well for correlated noise.

Table 3.9. Different Methods of Sensitivity Estimation Tested on GJ 625

Box #	Aperture Sensitivity ( $10\sigma, L'$ mag)	PSF Fitting Sensitivity ( $10\sigma, L'$ mag)	Single Pix Sensitivity ( $10\sigma, L'$ mag)	Sing. - PSF Difference ( $L'$ mag)
1	15.5	15.31	15.49	0.18
2	15.55	15.46	15.44	-0.02
3	15.54	15.38	15.52	0.14
4	15.25	15.15	15.32	0.17
5	14.73	14.6	14.77	0.17
6	13.19	13.07	14.11	0.92
7	13.66	13.42	14.27	0.61
8	13.4	13.28	14.2	0.80
9	13.25	13.04	13.94	0.69
10	13.44	13.3	14.16	0.72

Note. — The first 5 boxes are positioned over very clean, background limited sky. All sensitivity estimators agree for them. The last 5 boxes are over artifact-plagued regions of the image. The single pixel method overestimates the sensitivity in these regions by neglecting pixel correlations.

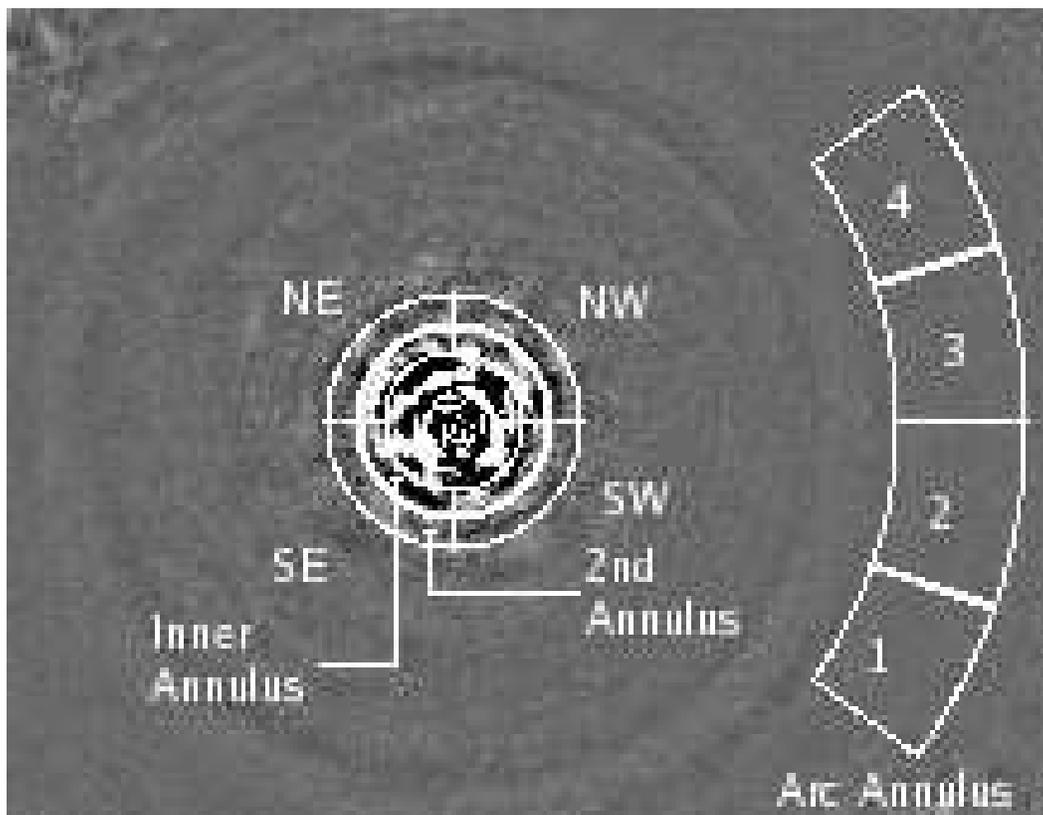


Figure 3.17 Fully processed, PSF-subtracted  $L'$  image of GJ 820A, with regions used to test different sensitivity calculation methods outlined in white and labeled. The sensitivity results are in Table 3.8. The conclusion is that Methods 2 and 3 agree well with Method 1 in the background-limited case, but in the contrast-limited, speckle-dominated regions near the star the Method 1 overestimates the sensitivity by neglecting pixel-to-pixel correlations.

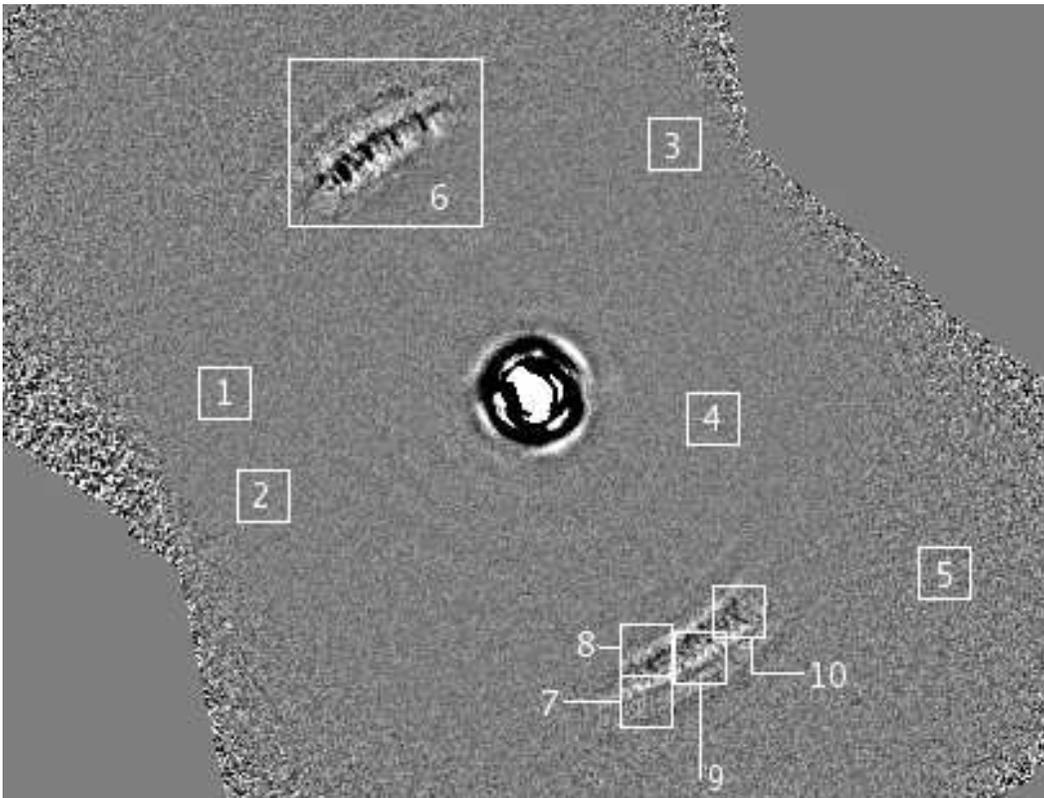


Figure 3.18 Fully processed  $L'$  image of GJ 625, with regions used to test different sensitivity calculation methods outlined in white and labeled. The sensitivity results are in Table 3.9. The conclusion is that Methods 2 and 3 agree well with Method 1 in the background-limited case, but in the areas affected by nod-subtraction artifacts, where the noise is strongly correlated from pixel-to-pixel, Method 1 overestimates the sensitivity by nearly a magnitude.

the real, unaltered science data for that star, and planets were sought in the fully processed images by the same combination of manual and automatic methods used for the real images.

The simulated planets were inserted into the raw, completely unprocessed frames using the PSF from short, PSF monitoring exposures appropriate to the data set. They were put in the correct positions on each raw image using stellar centroids and image rotations from a log file output by our image processing pipeline. The PSF image was normalized based on the sensitivity map and shifted to the correct location with sub-pixel accuracy by bicubic spline interpolation. The random positions of the planets were chosen from a distribution that placed them close to the star with higher probability than far from it, so that we would adequately sample the sensitivity close to the star. The numbers of planets at each significance level was random, as well as the location, so the user trying to recover the simulated planets from the processed data would not know when/whether they were all found.

The program that we used to place the simulated planets in the raw data wrote a log file detailing the position and brightness of each planet. The procedure for the blind test was for the user to attempt to recover simulated planets using exactly the same methods we applied to unaltered data sets to search for real planets. At the end, a list of confirmed and suspected planets was constructed. Only then did the user open the planet insertion log and see what planets had been inserted and how many had been confirmed or suspected.

Tables 3.10 through 3.14 give the results of these simulations. In these tables, a designation of 'Confirmed' for a simulated planet means it was bright and consistent enough through the data set that a real source of similar appearance could have been confidently detected. Further telescope time could be allocated to ob-

serve such sources with very little concern that the source would turn out not to be a real celestial object. ‘Noticed’ in these tables means flagged as a possible source, either manually or automatically. ‘Not confirmable’ means the question of whether the source was ‘real’ or due to an artifact could not be confidently resolved. ‘Denied’ means the apparent source seems to be fully explained by an artifact such as a stellar ghost or speckle.

Early in our survey, we did a blind test using our excellent data set for the rather faint star GJ 450. In this test a random number of planets between 15 and 20 was to be inserted into the data, all at a significance of exactly  $10\sigma$ . The planet insertion code input 17 planets, and all were recovered with no false positives. Table 3.10 gives the properties of the planets in this simulation.

Later, we performed blind tests on both  $L'$  and  $M$  band images of the very bright star Vega, which is not properly one of our survey targets but still provides a good example of the sensitivity obtained in the difficult case of an extremely bright star. In each simulation random numbers of planets were inserted at exact significance levels of 10, 7, and  $5\sigma$ , based on our master Method 3 sensitivity map. For the Vega  $L'$  test 13 of 14 planets at  $10\sigma$ , 5 of 7 planets at  $7\sigma$ , and 2 of 7 planets at  $5\sigma$  were confirmed. For the  $M$  band test 11 of 12 planets at  $10\sigma$ , 6 of 8 planets at  $7\sigma$ , and 1 of 5 planets at  $5\sigma$  were confirmed. See Tables 3.11 and 3.12 for further details.

We performed a similar test on the  $L'$  data from  $\epsilon$  Eri. In this case there was a glitch in the planet insertion code involving a mirror-flip of the sensitivity map relative to the real data, so the planets did not end up being inserted at exactly their assigned sensitivities. The experiment was nonetheless satisfactory and useful. For comparison with the other blind tests we assigned all planets brighter than  $8.37\sigma$  to the  $10\sigma$  class, all planets between  $5.92\sigma$  and  $8.37\sigma$  to the  $7\sigma$  class,

and all planets below  $5.92\sigma$  to the  $5\sigma$  class. The boundaries chosen are the logarithmic averages of 10 and 7, and 7 and 5, respectively. The final result was that 9 of 10 planets at  $10\sigma$ , 9 of 11 at  $7\sigma$ , and 0 of 1 at  $5\sigma$  were confirmed. See Table 3.13 for details.

Because  $\epsilon$  Eri and Vega are both very bright stars, and the GJ 450 test included only  $10\sigma$  planets, we performed one further test on a star of more typical brightness for our survey: HD 29391. In this simulation 12 of 12 planets at  $10\sigma$ , 5 of 10 planets at  $7\sigma$ , and 1 of 5 planets at  $5\sigma$  were confirmed. See Table 3.14 for details. An 'e' method master image from the HD 29391 test is shown in Figure 3.19.

Table 3.10. GJ 450  $L'$ -band fake planet experiment.

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
0.51	12.53	>20	26.99	10.00	Confirmed
0.56	13.32	>20	19.67	10.00	Confirmed
0.95	15.35	11.26	9.44	10.00	Confirmed
1.14	15.60	10.54	8.50	10.00	Confirmed
1.27	15.96	9.51	7.35	10.00	Confirmed
1.58	16.06	9.21	7.07	10.00	Confirmed
1.90	16.51	7.93	5.81	10.00	Confirmed
2.50	16.59	7.73	5.66	10.00	Confirmed
2.69	16.57	7.78	5.69	10.00	Confirmed
2.91	16.38	8.29	6.10	10.00	Confirmed
2.98	16.60	7.70	5.64	10.00	Confirmed
3.71	16.51	7.93	5.81	10.00	Confirmed
3.90	16.59	7.73	5.66	10.00	Confirmed

Table 3.10—Continued

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
3.93	16.62	7.65	5.60	10.00	Confirmed
5.02	16.49	7.98	5.85	10.00	Confirmed
6.52	16.43	8.15	5.97	10.00	Confirmed
6.53	16.27	8.61	6.44	10.00	Confirmed

Note. — All of the input planets were confirmed. Planet magnitude to mass conversion carried out by interpolation of the mass vs. magnitude tables at different ages provided in Burrows et al. (2003) and Baraffe et al. (2003). See Section 3.10 for details on how sources are investigated and classified as confirmed or otherwise.

Table 3.11. Vega L'-band fake planet experiment.

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
1.27	9.87	>20 MJ	39.42 MJ	10.00 $\sigma$	Confirmed
2.19	12.53	18.62 MJ	11.22 MJ	10.00 $\sigma$	Confirmed
2.31	12.28	>20 MJ	11.62 MJ	10.00 $\sigma$	Confirmed
2.79	12.64	17.74 MJ	11.04 MJ	10.00 $\sigma$	Confirmed
3.00	13.16	13.85 MJ	10.20 MJ	10.00 $\sigma$	Confirmed
3.01	12.92	15.60 MJ	10.59 MJ	10.00 $\sigma$	Confirmed
4.44	14.52	7.63 MJ	6.35 MJ	10.00 $\sigma$	Confirmed
5.09	14.73	7.08 MJ	5.84 MJ	10.00 $\sigma$	Confirmed
5.15	15.05	6.37 MJ	4.98 MJ	10.00 $\sigma$	Confirmed
6.16	15.15	6.15 MJ	4.79 MJ	10.00 $\sigma$	Noticed and denied
8.44	15.27	5.91 MJ	4.57 MJ	10.00 $\sigma$	Confirmed
8.46	14.82	6.87 MJ	5.59 MJ	10.00 $\sigma$	Confirmed
9.14	14.81	6.89 MJ	5.62 MJ	10.00 $\sigma$	Confirmed
10.40	15.44	5.60 MJ	4.26 MJ	10.00 $\sigma$	Confirmed
1.89	12.25	>20 MJ	11.67 MJ	7.00 $\sigma$	Confirmed
2.60	13.00	15.01 MJ	10.46 MJ	7.00 $\sigma$	Confirmed
3.65	13.75	10.30 MJ	8.49 MJ	7.00 $\sigma$	Confirmed
4.71	14.08	8.95 MJ	7.49 MJ	7.00 $\sigma$	Confirmed
5.25	15.17	6.11 MJ	4.76 MJ	7.00 $\sigma$	Noticed; not confirmable

Table 3.11—Continued

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
6.27	15.53	5.43 MJ	4.09 MJ	$7.00\sigma$	Noticed and denied
6.46	15.23	5.98 MJ	4.64 MJ	$7.00\sigma$	Confirmed
1.23	10.79	>20 MJ	27.23 MJ	$5.00\sigma$	Unnoticed
1.44	10.92	>20 MJ	25.96 MJ	$5.00\sigma$	Confirmed
4.29	14.45	7.81 MJ	6.51 MJ	$5.00\sigma$	Noticed; not confirmable
4.50	15.32	5.82 MJ	4.48 MJ	$5.00\sigma$	Unnoticed
4.70	15.73	5.06 MJ	3.73 MJ	$5.00\sigma$	Unnoticed
7.27	15.02	6.44 MJ	5.05 MJ	$5.00\sigma$	Confirmed
8.14	15.90	4.78 MJ	3.44 MJ	$5.00\sigma$	Unnoticed

Note. — Planets confirmed: 13/14 at  $10\sigma$ ; 5/7 at  $7\sigma$ ; 2/7 at  $5\sigma$ . Planets noticed: 14/14 at  $10\sigma$ ; 7/7 at  $7\sigma$ ; 3/7 at  $5\sigma$ . Planet magnitude to mass conversion carried out by interpolation of the mass vs. magnitude tables at different ages provided in Baraffe et al. (2003). See Section 3.10 for details on how sources are investigated and classified as confirmed or otherwise.

Table 3.12. Vega M-band fake planet experiment.

Sep (asec)	<i>M</i> Band Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
0.46	6.90	> 20 MJ	>100 MJ	10.00 $\sigma$	Confirmed
1.13	9.11	> 20 MJ	74.59 MJ	10.00 $\sigma$	Confirmed
1.19	9.05	> 20 MJ	76.72 MJ	10.00 $\sigma$	Confirmed
2.39	12.43	10.66 MJ	9.80 MJ	10.00 $\sigma$	Confirmed
3.60	12.95	7.13 MJ	7.31 MJ	10.00 $\sigma$	Confirmed
3.77	13.09	6.44 MJ	6.78 MJ	10.00 $\sigma$	Confirmed
6.31	13.35	5.31 MJ	5.98 MJ	10.00 $\sigma$	Noticed; not confirmable
6.49	13.15	6.16 MJ	6.60 MJ	10.00 $\sigma$	Confirmed
7.96	13.08	6.49 MJ	6.81 MJ	10.00 $\sigma$	Confirmed
10.61	12.41	10.83 MJ	9.90 MJ	10.00 $\sigma$	Confirmed
10.84	11.83	17.23 MJ	11.35 MJ	10.00 $\sigma$	Confirmed
11.89	12.84	7.74 MJ	7.80 MJ	10.00 $\sigma$	Confirmed
0.74	7.65	> 20 MJ	>100 MJ	7.00 $\sigma$	Confirmed
0.84	8.10	> 20 MJ	>100 MJ	7.00 $\sigma$	Confirmed
2.31	12.53	9.82 MJ	9.27 MJ	7.00 $\sigma$	Confirmed
2.45	12.88	7.51 MJ	7.62 MJ	7.00 $\sigma$	Noticed; not confirmable
2.65	12.81	7.90 MJ	7.94 MJ	7.00 $\sigma$	Noticed; not confirmable
3.69	12.93	7.24 MJ	7.40 MJ	7.00 $\sigma$	Confirmed
3.99	13.22	5.85 MJ	6.38 MJ	7.00 $\sigma$	Confirmed

Table 3.12—Continued

Sep (asec)	<i>M</i> Band Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
9.48	13.07	6.54 MJ	6.84 MJ	7.00 $\sigma$	Confirmed
2.15	12.52	9.90 MJ	9.32 MJ	5.00 $\sigma$	Unnoticed
2.19	12.74	8.34 MJ	8.26 MJ	5.00 $\sigma$	Unnoticed
3.16	13.42	5.02 MJ	5.75 MJ	5.00 $\sigma$	Confirmed
3.23	12.91	7.35 MJ	7.49 MJ	5.00 $\sigma$	Noticed; denied
5.62	14.30	2.61 MJ	3.18 MJ	5.00 $\sigma$	Noticed; not confirmable

Note. — Planets confirmed: 11/12 at 10 $\sigma$ ; 6/8 at 7 $\sigma$ ; 1/5 at 5 $\sigma$ . Planets noticed: 12/12 at 10 $\sigma$ ; 8/8 at 7 $\sigma$ ; 3/5 at 5 $\sigma$ . Planet magnitude to mass conversion carried out by interpolation of the mass vs. magnitude tables at different ages provided in Baraffe et al. (2003). See Section 3.10 for details on how sources are investigated and classified as confirmed or otherwise.

Table 3.13.  $\epsilon$  Eridani L'-band fake planet experiment.

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
0.64	9.66	>20 MJ	27.79 MJ	10.48 $\sigma$	Confirmed
0.72	9.92	>20 MJ	25.39 MJ	10.08 $\sigma$	Noticed; not confirmable
0.84	10.21	>20 MJ	22.70 MJ	8.64 $\sigma$	Confirmed
1.37	12.44	11.55 MJ	9.64 MJ	8.53 $\sigma$	Confirmed
2.1	13.36	8.37 MJ	6.73 MJ	11.16 $\sigma$	Confirmed
3.37	15.18	4.69 MJ	3.15 MJ	8.40 $\sigma$	Confirmed
3.69	15.24	4.62 MJ	3.07 MJ	10.19 $\sigma$	Confirmed
4.76	15.31	4.53 MJ	2.98 MJ	12.10 $\sigma$	Confirmed
6.13	15.09	4.80 MJ	3.28 MJ	11.01 $\sigma$	Confirmed
9.42	15.52	4.28 MJ	2.81 MJ	8.57 $\sigma$	Confirmed
1.27	12.36	11.86 MJ	9.91 MJ	7.83 $\sigma$	Confirmed
1.67	12.94	9.72 MJ	7.95 MJ	7.39 $\sigma$	Confirmed
1.83	13.07	9.29 MJ	7.56 MJ	6.33 $\sigma$	Confirmed
2.83	14.92	5.01 MJ	3.52 MJ	6.52 $\sigma$	Confirmed
2.97	14.64	5.48 MJ	3.92 MJ	6.30 $\sigma$	Confirmed
3.27	14.98	4.94 MJ	3.44 MJ	7.58 $\sigma$	Confirmed
4.27	15.41	4.41 MJ	2.90 MJ	6.85 $\sigma$	Noticed; not confirmable
4.29	15.82	3.94 MJ	2.55 MJ	7.90 $\sigma$	Confirmed
7.52	15.88	3.89 MJ	2.50 MJ	6.32 $\sigma$	Confirmed

Table 3.13—Continued

Sep (asec)	$L'$ Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
10.04	15.62	4.15 MJ	2.72 MJ	$8.25\sigma$	Confirmed
10.3	15.76	3.99 MJ	2.60 MJ	$6.06\sigma$	Unnoticed
8.98	16.13	3.68 MJ	2.29 MJ	$3.90\sigma$	Unnoticed

Note. — Planets confirmed: 9/10 at  $10\sigma$ ; 9/11 at  $7\sigma$ ; and 0/1 at  $5\sigma$ . Planets noticed: 10/10 at  $10\sigma$ ; 10/11 at  $7\sigma$ ; and 0/1 at  $5\sigma$ . Planet magnitude to mass conversion carried out by interpolation of the mass vs. magnitude tables at different ages provided in Baraffe et al. (2003). See Section 3.10 for details on how sources are investigated and classified as confirmed or otherwise.

Table 3.14. HD 29391  $L'$ -band fake planet experiment.

Sep (asec)	$M$ Band Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
0.42	11.59	>20	16.26	10.00	Confirmed
0.76	12.56	16.85	10.86	10.00	Confirmed
1.23	15.35	4.97	4.14	10.00	Confirmed
2.06	15.90	3.92	3.23	10.00	Confirmed
2.27	16.10	3.63	2.93	10.00	Confirmed
3.26	14.58	6.95	5.81	10.00	Confirmed
3.60	15.77	4.15	3.44	10.00	Confirmed
4.29	15.48	4.72	3.91	10.00	Confirmed
4.41	16.22	3.46	2.76	10.00	Confirmed
5.31	16.21	3.47	2.77	10.00	Confirmed
8.92	16.15	3.56	2.86	10.00	Confirmed
10.69	16.15	3.56	2.86	10.00	Confirmed
1.25	15.17	5.40	4.48	7.00	Confirmed
1.86	16.32	3.31	2.62	7.00	Confirmed
2.00	16.47	3.09	2.42	7.00	Unnoticed
2.69	16.54	2.99	2.32	7.00	Unnoticed
2.92	16.61	2.93	2.23	7.00	Noticed; not confirmable
3.29	16.47	3.09	2.42	7.00	Confirmed
4.69	15.83	4.03	3.35	7.00	Noticed; not confirmable

The total statistics from all 5 blind tests are that 62 of 65 planets were recovered at  $10\sigma$ , 25 of 36 at  $7\sigma$ , and 4 of 18 at  $5\sigma$ . In percentages we have 95% completeness at  $10\sigma$ , 69% completeness at  $7\sigma$ , and 22% completeness at  $5\sigma$ . We note, however, that the three  $10\sigma$  planets that were not confirmed were from the three data sets involving very bright stars. All of the  $10\sigma$  planets in the GJ 450 and HD 29391 data sets were recovered. For simplicity, we have adopted a completeness level of 100% at  $10\sigma$ , a conservative 50% at  $7\sigma$ , and 0% at  $5\sigma$ , and we have elected to quote  $10\sigma$  sensitivities as our standard, since it is only at this level that we know we have obtained near 100% completeness.

In all the blind tests only one object that had not been inserted was confirmed, and that was a much more tentative confirmation than for most of the ‘real’ simulated planets. It was found in the Vega  $M$  band data, very close in, at low significance against the intense speckle noise but so bright it would have to be a low-mass star rather than a planet or brown dwarf if it were real. It seems unlikely that such a massive object would have evaded astrometric or RV detection, and the source also does not seem to appear in some of our later  $M$  band Vega images that ought to be sensitive enough at least to suggest it. In any case it appears that our false-positive rate for confirmed objects is very low. In fact, in the whole survey we confirmed only one object which turned out on further analysis probably to be unreal – and this was caught by reanalysis of the data, before any telescope time had been expended following it up.

**Why were lower significance planets not detected?** If the random noise background in an image is truly gaussian, there is only an 0.13% chance of a spurious noise burst appearing as a source detected at more than  $3.0\sigma$ . Why, then, could we not confirm  $3\sigma$  sources? Why not even  $5\sigma$  sources?

One part of the answer is that because some stages of our processing remove

Table 3.14—Continued

Sep (asec)	<i>M</i> Band Mag	Mass (Bur)	Mass (Bar)	Det. Sig.	Status
5.72	16.38	3.22	2.54	7.00	Confirmed
6.28	15.97	3.82	3.12	7.00	Noticed; not confirmable
10.53	15.94	3.86	3.17	7.00	Confirmed
1.19	15.39	4.89	4.07	5.00	Confirmed
1.93	16.77	2.78	2.01	5.00	Noticed; not confirmable
5.76	16.57	2.97	2.28	5.00	Noticed; not confirmable
6.68	16.25	3.41	2.72	5.00	Unnoticed
7.70	16.18	3.51	2.82	5.00	Unnoticed

Note. — Planets confirmed: 12/12 at  $10\sigma$ ; 5/10 at  $7\sigma$ ; 1/5 at  $5\sigma$ . Planets noticed: 12/12 at  $10\sigma$ ; 8/10 at  $7\sigma$ ; 3/5 at  $5\sigma$ . Planet magnitude to mass conversion carried out by interpolation of the mass vs. magnitude tables at different ages provided in Burrows et al. (2003) and Baraffe et al. (2003). See Section 3.10 for details on how sources are investigated and classified as confirmed or otherwise.

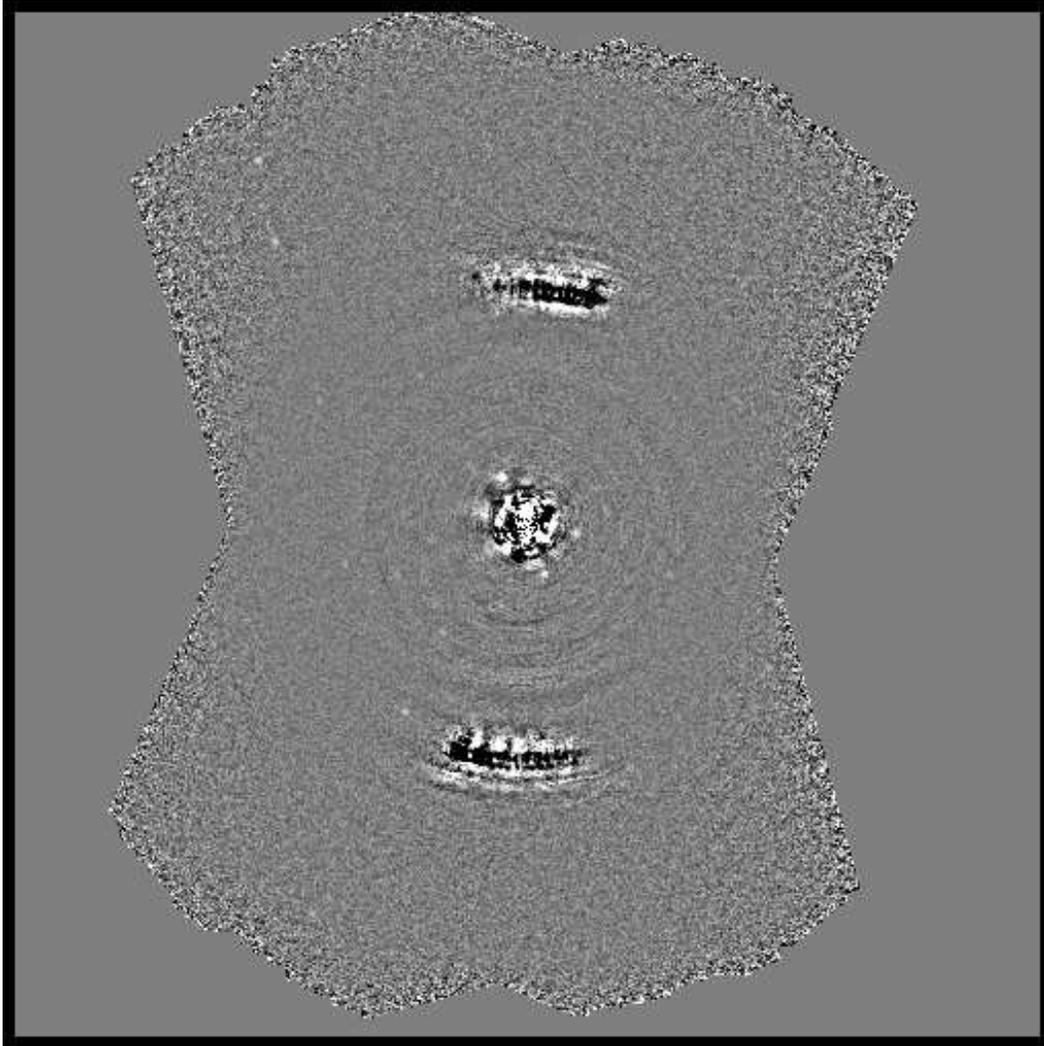


Figure 3.19 Fully processed 'e' method master image from the blind sensitivity test on HD 29391. Identifying the simulated planets is left as an exercise for the reader. The contrast stretch chosen here makes visible all the planets except the innermost and brightest ones, which are lost in speckle noise that saturates the image display range. In this data set there are 12 planets of  $10\sigma$  significance, 10 at  $7\sigma$ , and 5 at  $5\sigma$

flux from faint point sources, the formal significance of the 10, 7, and  $5\sigma$  planets in our blind tests is probably somewhat less than their nominal values. Typically about 20% of the flux is lost, with higher losses possible in some cases. See Section 3.7.5 for a detailed analysis of this issue.

Another part of the answer, which Phil Hinz first brought to my attention, is that an 0.13% false-positive probability is actually unacceptably high. A typical image from our survey has data over perhaps 120,000 pixels. If we take the diameter of a PSF to be 3.2 pixels (a fairly typical FWHM for our unsaturated stellar images), about 15,000 different PSFs could fit in our image. That is, each image samples the distribution of spurious PSFs drawn from the noise 15,000 times. Even if the whole image had only pure, gaussian noise, we would expect  $0.0013 \times 15,000 = 19.5$  spurious  $3\sigma$  detections per image. We need the threshold for useful confirmation of a planet to result in far less than 1 false positive per image, because effective followup of faint sources is so costly in telescope time (for a source near the detection limit, a followup observation must use an exposure at least as long, in conditions at least as good, as the original observation). If we wanted to set a limit so that we would expect only two false positives in our entire 50-star survey, under the assumption of gaussian noise statistics, we would need to choose the significance that yields a one-sided gaussian tail probability of  $1/(15000 \times 25) = 2.67 \times 10^{-6}$ . This corresponds to approximately a  $5\sigma$  limit.

However, **the statistics of planet-search images are never gaussian**. On the contrary, the distribution always has a long tail because of speckles, ghosts, diffraction rays, and detector pattern noise artifacts. Thus **it is to be expected that the reasonable high-completeness detection limit for planet search surveys will be well above  $5\sigma$** . Further complicating this issue is that fairly small regions of the image are usually used to calculate the RMS used in sensitivity estimates. This

is as it should be, since due to the stellar speckles, there is always a strong radial dependence to the noise, and ghosts, rays, variable pattern noise, and other things can introduce azimuthal sensitivity variations as well. However, using such small regions compounds our error of approximating a long-tailed distribution with a gaussian with the further problem of calculating the RMS in a region too small to measure the long tails. That is, **we are calculating the RMS of a distribution with a long tail toward low probability, high brightness events in a region small enough that it usually will not include such ‘tail events’**. Yet it is these very tail events, or gaussian-breaking bright artifacts, that produce unacceptable numbers of false-positive planet detections if the detection threshold is set too low. Thus **our test result, that only planets with a significance well above  $5\sigma$  are confirmed with high completeness if the detection threshold is set high enough to avoid an unacceptable number of false positives, is not surprising.**

In this context, it is interesting to consider what fraction of the planets were ‘noticed’ (ie suspected) in the blind tests. The percentages of planets that were noticed (and either confirmed or found too faint to be securely confirmed) over the five blind tests are 100% at  $10\sigma$ , 92% at  $7\sigma$ , and 50% at  $5\sigma$ . This suggests that real planets well below our quoted sensitivities would have been suspected in our data. However, very many unreal objects are suspected too. In case one of our suspected detections turns out to be a real object of some interest, we have included a table of suspected sources in Section 3.10, for the use of future observers who may confirm objects we only suspected.

It is not at all clear that if other observers conducting planet imaging surveys adopt our exact sensitivity estimation method, the same detection rates (that is, completeness levels) for sources at 5, 7, and  $10\sigma$  significance will apply. Statistics differ between data sets produced by different instruments at different

wavelengths. There are clever methods for distinguishing between real and fake sources at the same significance (see Section 3.10 for examples), and these can increase completeness levels for low-significance sources without increasing the false-positive rate.

However, **Careful blind tests such as those we describe here are the only sure way to validate sensitivity estimators and map completeness vs significance level** for the complex, intensively processed data sets typical of planet imaging surveys. As far as we know **we are the first to attempt such tests, which should be standard for all future planet imaging surveys**. Without them, the completeness of the survey at different significance levels cannot be known, and the sensitivities from different surveys cannot be accurately compared.

### 3.7.5 One Caveat to Our Sensitivity Estimates

There is one effect of potential significance that we have neglected in computing our formal  $10\sigma$  sensitivity maps: the degree to which some of the steps in our processing pipeline may dim real point sources. This is potentially relevant in comparing our sensitivity measurements to those of other surveys, assuming the methods were sufficiently comparable. For example, if our processing dims point sources by 30%, our nominal  $10\sigma$  limits are really  $7\sigma$  based on our sensitivity estimation method, though they may well be  $10\sigma$  or even more based on a more optimistic method. Note well that this consideration **does not in any way affect the validity of the completeness vs nominal significance level that we derived from the blind sensitivity tests**. This is because the simulated planets were inserted into the raw data and therefore were subject to exactly the same processing any real planet images would have been.

There are three stages of our processing that may reasonably be suspected of dimming inserted planets. They are the hot pixel fixing, which might mistake the

cores of planetary PSFs for hot pixels, the unsharp masking, which will slightly dim a real point source by subtracting from it a blurred version of itself, and the 'colfudge' correction, which could cause dark streaks to emanate from a planetary image as the brightness of the image itself is considered a potential column artifact that must be corrected.

The effects are listed here in order of increasing likelihood to dim real sources substantially. The hot pixel fixing algorithm considers every pixel in the image in turn to see if it is bad. It computes the mean over a 3x3 pixel box centered on the pixel under consideration but not including the pixel itself. It then computes the single-pixel RMS over a 5x5 box centered on the pixel under consideration, again excluding the pixel itself. If the pixel differs from the 3x3 mean by more than 5 times the RMS in the 5x5 box (ie if it is more than  $5\sigma$  deviant), it is replaced by the value of the 3x3 mean. Since our data are well sampled, it is essentially certain that the wings of any PSF image will keep the 5x5 RMS far too high for the central pixel to be rejected. Thus, the hot pixel fixing should have no effect on the brightness of real point sources. The unsharp masking could cause only a slight dimming. However, the colfudge artifacts can be strong and might have a larger effect.

We note that in both the unsharp masking and the colfudge routine we attempt to reject point sources from the calculations of the corrections to be applied to the image, so that point sources will not be dimmed (see Section 3.6 for the details of the unsharp masking algorithm or Appendix A for colfudge). These precautions unfortunately do not work for very faint point sources that are below the noise level in individual frames and can be seen only on the stacked master image. The unsharp masking and colfudge algorithms cannot detect them statistically and therefore cannot reject them, so the corrections applied are

biased and the point sources slightly dimmed.

We performed several experiments to see how severe the dimming was. First, we constructed 1000 simulated images, ran them through various parts of our processing, and stacked them to get final images on which the flux loss from point sources could be measured. The images were made by creating a background of pure gaussian noise and then adding in bright and faint point sources based on a real stellar PSF from our science data. Stacked but unprocessed versions of these images are shown in Figure 3.7, where they serve to illustrate the performance of creeping mean and median combines. To prove that the hot pixel fixing really had as negligible an effect as we supposed, we used the PSF of the star GJ 5 for the inserted point sources. This star was observed in above-average seeing and is one of the sharper PSFs we measured. This PSF, if any, should have its peak pixels rejected by the hot pixel fixing. The results of this flux loss experiment are presented in Table 3.15. The 'Faint Sources' from Table 3.15 correspond to the five simulated planets that appear along the bottom of the images in Figure 3.7; the loss from 'Bright Sources' in the table is measured from the bright simulated star at the upper left corner of the Figure 3.7 images.

The loss from hot pixel fixing is zero, and that from post-stack unsharp masking is less than 3%. However, colfudge takes out 14.6% of the flux, and combining colfudge with post-stack unsharp masking (equivalent to our 'a', 'b', 'c', 'f' and 'x' processing methods) removes 16.9% of the flux. Adding pre-stack unsharp masking as used in our 'd', 'e', 'g', and 'y' processing methods takes out a few more percent for a total loss of 21.8% on faint sources. Note that the flux loss for bright sources in this final case is still only 3.4%, indicating that our methods of statistically rejecting point sources to avoid dimming them are fairly effective.

This test was fairly realistic, but it used simulated, rather than real, images,

Table 3.15. Flux Loss in Processing from Simulated GJ 5 Experiment

Processing Applied	Flux Loss (Faint Sources)	Flux Loss (Bright Sources)
Hot Pixel Fixing	0.0%	0.0%
Final Unsharp Masking	2.7%	1.2%
Colfudge	14.6%	-0.2%
Hot Pixel Fixing, Final Unsharp Masking, and Colfudge	16.9%	2.6%
All of the Above With Pre-Stack Unsharp Masking	21.8%	3.4%

Note. — ‘Bright’ in this context means bright enough to stand out above the noise on a single frame; ‘faint’ means anything fainter than this. For a real data set with typically 100-200 frames and a background limit of  $L' = 16.0$ , a ‘bright’ source that stands out at the  $7\sigma$  level on one frame appears at around the  $90\sigma$  level on the full stack. The corresponding magnitude is  $L' = 13.5$ ; this therefore approximately marks the bright/faint boundary.

and no rotation was involved. Since colfudge artifacts appear as dark streaks out from a point source, it seemed possible that rotation would change the loss from colfudge either up or down. Also, the GJ 5 PSF was very sharp, and a less sharp PSF might be more affected. Accordingly, we tried another experiment in which ten simulated point sources were placed in the raw data for star GJ 354.1A, using the PSF from unsaturated images of that star. This was one of our blurriest PSFs, with a FWHM of about 3.7 pixels as opposed to 3.0 for GJ 5. The simulated sources we inserted were all faint enough not to be rejected by the unsharp masking and colfudge statistical rejection methods on individual images, but bright enough to be detected with high significance on the final image. The measured average dimming was 18.5% for the 'a' method reduction and 28.4% for the 'd' method.

We see that for the 'a' method reduction, a blurry PSF in a data set with rotation lost only a little more flux than a sharp one in our unrotating test (indeed, the difference is barely above the scatter in our photometry of faint sources against the noise). Since we would expect a blurry source to be more dimmed by colfudge than a sharp one, while it is not clear what effect rotation would have, it seems reasonable to ascribe all of the small difference between the 'a' script flux loss in the GJ 5 and GJ 354.1A experiments to the blurriness of the GJ 354.1A PSF. The implication is that rotation has little or no effect, so the GJ 5 simulated image results are probably applicable to real rotated data sets with sharp PSFs.

Under the 'd' method reduction the flux loss from the GJ 354.1A PSF was significantly more than for the GJ 5 PSF. We conclude that pre-stack unsharp masking affects blurry sources considerably more severely, which is not surprising. Since our final sensitivity maps are calculated by choosing the best sensitivity obtained at a given location in any of the individual sensitivity maps from the

different processing methods, and the ‘d’ method images are usually cleaner, the ‘d’ images are probably weighted higher than the ‘a’ in our final sensitivity results. Thus the degree of flux loss in the ‘d’ images is unfortunate. However, we note that most of our data sets had sharper PSFs than GJ 354.1A, so the flux losses in processing should in general be less severe.

Our PSF subtraction methods remove some further flux from very close-in sources by self-subtraction. However, we have in general allowed sensitivities to be computed from the PSF-subtracted images (that is the ‘b’ and ‘e’ method images) only beyond a minimum radius set by the requirement that parallactic rotation move a source through an arc of length at least 6 pixels, (ie, about twice the FWHM) throughout the observing sequence. Dimming should not be severe at this radius, and will drop to negligible values at slightly larger radii. Also, we note that the self-subtraction dims sources precisely in the region where our Method 3 sensitivity estimator appears to be most pessimistic relative to the other reasonable estimator, Method 2. This probably compensates the self-subtraction to some extent.

To conclude, our processing dims sources by a bit more than 25% in the worst cases. In these cases our 10, 7, and  $5\sigma$  sources will really be a bit less than 7.5, 5.25, and  $3.75\sigma$ , formally, according to our sensitivity estimator. **They may still be equal to or higher than their nominal significance values based on more optimistic sensitivity estimators.** In many cases our processing losses are likely less than 25%. Since it is not clear that the sensitivity estimates from other surveys have taken into account their own processing losses, which may be equal to or greater than ours, we have retained our nominal significance values when comparing our sensitivities to those of other surveys in Subsection 3.9. We note, again, that **the accuracy of our completeness vs nominal source significance**

values from the blind sensitivity tests is not at all affected by flux losses in processing. Also, again, truly accurate comparison of the sensitivity of one survey with another requires the results of blind sensitivity tests from both. Only then can the sensitivities be compared at nominal significance levels for which both surveys had the same completeness. The comparison is then no longer at the mercy of flux losses in processing or differences in the sensitivity estimators.

The photometry of faint apparent companions in Table 3.10.2 has been corrected for flux loss assuming an average loss of 0.2 magnitudes, or about 20%.

### 3.7.6 Sensitivity Estimators of Other Surveys

Published and submitted papers presenting the results of other planet imaging surveys devote surprisingly little space to the crucial issue of sensitivity estimation. In most cases the method is not explained in sufficient detail for other workers to duplicate it for comparison with their own methods. Here, we describe and comment on the methods used in the four surveys we have compared with our own.

**Masciadri et al. (2005):** Since only a pixel-to-pixel RMS is given, it would appear that the method is used is analogous to our Method 1 (see Subsection 3.7.1 above), applied over “...a box of  $d \times d$  pixels (with typical values of  $d$  of 4-6 pixels, which corresponds to the FWHM of the nonsaturated PSFs)...” Note that 6 pixels is 0.08 arcsec, or about  $2 \lambda/D$  at  $H$  band for the VLT NACO camera used for the Masciadri et al. (2005) observations. The method for converting the single-pixel RMS to a point source sensitivity is not explained, so we cannot be sure how close to our Method 1 it was. As sensitivities are presented in  $\Delta$ -magnitudes it is probable that no aperture correction was needed; none is mentioned. Processing losses were likewise not mentioned, though based on the filtering methods described they probably were not negligible.

Method 1 does not correctly estimate the sensitivity in the presence of correlated noise, and the images presented show strong correlated speckle noise near the star, as we would expect. Even at large distances from the star correlated pattern noise seems to persist, as is not uncommon for astronomical IR detectors. Another problem is that a 4-6 pixel (0.05-0.08 arcsec) square box is a very small region indeed over which to calculate the sensitivity (contrast our 8-pixel (0.4 arcsec) radius disk or 45 pixel (2.2 arcsec) long arc). Calculating the sensitivity in too small a region biases the measurement toward over-optimistic results since (as described in Subsection 3.7.4) outlying points are less likely to be included. Finally  $5\sigma$  limits are quoted. We have found that Method 1 overestimates the sensitivity by about 1 magnitude in regions of strong correlated noise, and that 100% completeness is attained only near  $10\sigma$ . Adding the factor of 2 (0.8 mag) from the  $5/10\sigma$  conversion to the potential 1 magnitude error of Method 1 leads to the conclusion that the sensitivity could be overestimated by as much as 1.8 magnitudes in the speckle dominated areas of the image. As we see below, there is evidence the overestimation is not this severe.

Masciadri et al. (2005) do not record a blind sensitivity test, but they do present in Figure 4 an image with 2 simulated planets inserted, one at 0.7 arcsec and the other at 1.0 arcsec. The inner one, based on the plot included in Figure 4, has a significance only slightly above  $5\sigma$ . That it appears as clearly as it does probably indicates the sensitivity is not overestimated by as much as the 1.8 magnitudes we suggest above, at least at this separation from the star. However, if the bright arrows that highlight the inserted planets are removed, it becomes clear that the inner one would not be confirmed in any detection scheme with low false-alarm probability. It is certain that the detection limits in this paper are overly optimistic, though by how much could only be determined with a blind sensitivity

test.

We have reduced the sensitivities from this paper by 0.753 mag, or a factor of 2, to compare them with our own nominal  $10\sigma$  limits in Section 3.9. Note that this is only the amount required to convert  $5\sigma$  sensitivities to  $10\sigma$  ones. It does not include any correction for the likely error of a Method 1 sensitivity estimator, and as such is a conservative adjustment.

We comment that although we have strongly criticised the sensitivity estimates in this paper, it remains a useful and significant work early in the history of AO planet imaging surveys. It is certain that the true sensitivity of many of the images would have been sufficient to detect planetary mass objects, had any been present at sufficient separations from the survey targets.

**Kasper et al. (2007):** Again, the sensitivity calculation section of this paper records only the pixel-to-pixel RMS, leaving us to assume that the method used is analogous to our Method 1. The sensitivities are presented in the form of  $\Delta$ -magnitudes, so it is probable that no aperture correction was needed; none is discussed. The RMS calculation was performed not in a small box but in a 4 pixel-wide annulus centered on the primary star. This large region would yield a better, less optimistic sensitivity estimate even with Method 1, although, as we have mentioned above, in our opinion the best option is to use an accurate estimator in a smaller region.

The authors mention that one of their processing steps produced a flux loss of 15% from PSF cores, but it is not clear this was the only processing flux loss nor that they made any correction for it in their sensitivity estimates. They choose to quote  $5\sigma$  limits, but they do recognize that the distribution is non-gaussian and go to some length using a pixel brightness histogram and Chebyshev's inequality to explain why  $5\sigma$  is a good choice. It is not clear that this discussion

takes into account either the fact that speckle noise is correlated noise or the false-positive problems and related statistical issues we have discussed in Subsection 3.7.4 above.

They do not perform a blind sensitivity test, but do present in Figure 3 an image with three fake planets in a line, quite similar to Figure 4 in Masciadri et al. (2005). The text claims that all 3 planets are detected with at least  $5\sigma$  significance, but, once again, it does not appear that the innermost one, at least, stands out remarkably above the speckle noise. It would not have been confirmed in a blind test.

Again, we note that although the sensitivity estimates are overly optimistic, the survey did attain sensitivity even to fairly low mass planets around the stars surveyed, and represents an important contribution to the field.

**Biller et al. (2007):** Once again, the authors quote only a pixel-to-pixel RMS in their discussion of the sensitivity, and do not explain how this is converted to a point-source sensitivity. This led us initially to assume they must have used a Method 1-type estimator. However, Beth Biller (2007, private communication) has explained to us that this is not correct. Instead, the sensitivity is calculated by comparing the pixel-to-pixel RMS to the peak pixel value expected for a real PSF. The  $5\sigma$  limits quoted correspond to the brightness of a point source whose peak pixel would stand above the mean by 5 times the pixel-to-pixel RMS. This is a very conservative sensitivity estimator for well-sampled data such as that of Biller et al. (2007). It supposes a PSF must be detected based only on the peak pixel, but actually the surrounding pixels over a diameter approximately equal to the FWHM will contribute to the detection. The method will in general underestimate the sensitivity, though in the presence of correlated noise the underestimation will be less severe. The sensitivity is calculated in  $\Delta$ -magnitudes based

on scaling from the peak pixel of an unsaturated image of the primary star, so no aperture correction is needed. Processing flux losses are not mentioned, though they probably exist.

The authors consider calculating the single-pixel RMS either within a 6 pixel (0.05 arcsec) square box or in an annulus centered on the primary star as was done in Kasper et al. (2007). They choose to use the box rather than an annulus for a reason with which we entirely agree: the point source sensitivity at a given location on an image is determined by the noise properties of the immediately surrounding region, not those in a full annulus around the primary at the same radius. Six pixels or 0.05 arcsec is only  $1.2 \lambda/D$  for the SDI/NACO camera on the VLT. This is perhaps a bit too small, as it can only contain about one PSF. Because of the conservative sensitivity estimator, this is most likely not a problem.

This survey has two unique properties that make the sensitivity estimator even more conservative. All the data sets were taken at two different instrument rotator angles separated by 33 degrees, and the images from the two positions were subtracted in the final stages of processing. Thus any real celestial object should appear as a pair of sharp PSF images, one dark and one bright, separated by 33 degrees of rotation about the primary star. This effectively doubles the data available. A well designed search algorithm or a well-trained human analyst could probably detect such dark/bright pairs at a significance a factor of  $\sqrt{2}$  lower than the minimum significance a single bright object would have to have to be detected in the same data set. Additionally, images are made using two independent differences of narrowband CH<sub>4</sub> filters. The two filter differences are not necessarily equally sensitive, but when they are they contribute an additional factor of  $\sqrt{2}$  to the sensitivity.

The authors present in Figure 3 the traditional ‘line of planets’ image, with 3

simulated planets at different separations. The innermost is obviously not confirmable, and the authors do not claim it is confirmed. They state that the outer two, both of which are clearly confirmable, have significance greater than  $10\sigma$ . No test object near the  $5\sigma$  detection limit is presented.

Because of the special properties mentioned above, we have only worsened the sensitivities from this paper by 0.365 mag, or a factor of 1.4, to compare them with our own nominal  $10\sigma$  limits in the figures discussed in Section 3.9. This is roughly the amount required to convert  $7\sigma$  sensitivities to  $10\sigma$  ones, and is equivalent to the assumption that nominal  $7\sigma$  sensitivities from Biller et al. (2007) are comparable to our nominal  $10\sigma$  limits. Because of the conservative estimator and the possibility of 2 factors of  $\sqrt{2}$  from the roll angles and then independent wavelength differences, it would most likely have been better to consider their  $5\sigma$  limits directly comparable to our  $10\sigma$  ones. The reader should keep in mind when examining the Section 3.9 figures that we may have underestimated the Biller et al. (2007) sensitivities by a factor of  $\sqrt{2}$  (0.38 mag) or a bit more. This rather small correction will not affect any of our main conclusions.

The Biller et al. (2007) paper represents a remarkable contribution to the field. The authors present significant new ideas in instrument design, observation strategy, and data analysis. Sensitivity to planets exceptionally close to their host stars was obtained.

We note that the sensitivity discussion sections of Masciadri et al. (2005), Kasper et al. (2007), and Biller et al. (2007) all consider the pixel-to-pixel RMS and explain over what areas it was calculated, but do not explain how it was converted to a point source sensitivity. It seemed most natural to us in every case to suppose that a method analogous to our Method 1 had been used, and we raised concerns that this had resulted in overestimated sensitivity in areas of the images affected

with correlated noise. In the case of Biller et al. (2007) it turns out that instead a very conservative estimator has been used, which is not in fact analogous to Method 1. We do not know if this is the case for either of the other papers. If it is, the concerns we have mentioned above about the estimator being overly optimistic do not apply. **The sensitivity discussion in each paper is insufficiently detailed, leaving us uncertain about important aspects of the estimators.**

**Lafrenière et al. (2007b):** This paper represents by far the most careful and accurate sensitivity analysis prior to our own work. A method analogous to our Method 2 is used in annular regions centered on the primary star. Thus **the paper is the first in the field of planet imaging surveys to account properly for correlated noise.** The large, annular regions used will even give somewhat conservative results relative to our own for the same nominal  $\sigma$ . Flux losses in processing are carefully analyzed and corrected in the final sensitivity estimate. As with the other papers an aperture correction is not discussed but the  $\Delta$ -magnitude form of the sensitivities suggests none was needed. The authors quote  $5\sigma$  limits and cite Lafrenière et al. (2007a) to support their claim that the advanced LOCI PSF subtraction algorithm they have used produces true gaussian noise and thus a  $5\sigma$  limit is appropriate. In fact, the noise histograms in Lafrenière et al. (2007a) do show a distribution that deviates remarkably little from a pure gaussian, indicating that the LOCI algorithm is indeed powerful. However, some deviations do exist, and the claim that  $5\sigma$  limits really apply needs to be substantiated by a blind sensitivity test, which neither Lafrenière et al. (2007b) nor Lafrenière et al. (2007a) provide. This lack is the only weakness in this remarkable pair of papers presenting the most sensitive planet imaging survey to date, combined with the most careful and thorough sensitivity analysis prior to our own.

Because of the good estimator used in this paper, the large annular region over

which it was used, and the good accounting for processing flux losses, we have assumed that the quoted  $5\sigma$  sensitivities are more equivalent to our nominal  $7\sigma$  sensitivities. Therefore we have only worsened the sensitivities from this paper by 0.36-0.39 magnitudes, or a factor of about 1.4, to compare them with our own nominal  $10\sigma$  limits in Section 3.9.

### 3.8 Sensitivity Obtained Around Each Target

We have described above in detail how we calculated our sensitivities; in this section, we present the sensitivity results. First, we present the nominal  $10\sigma$  Method 3 sensitivity maps, converted to magnitude contour images. Each image has celestial North up, the primary star in the exact center, and an astrometric grid of  $2\times 2$  arcsec squares is overlaid. Other observers who detect a source near one of our survey stars can use these images and the astrometric grid to determine what sensitivity we obtained at the location of their source. They can also search our list of suspected detections in Section 3.10 to see if we noticed a corresponding source.

After the contour maps, we present plots of the minimum mass planet we could detect vs the projected separation in AU for each of our targets. Since our sensitivities showed considerable azimuthal variations for reasons such as the negative nod subtraction images, we have plotted the 0th, 50th, and 90th percentile sensitivities at each radius. The 0th percentile is of course the worst sensitivity obtained at a given radius, the 50th percentile is the median, while only 10% of the image at a give radius had better sensitivity than the 90th percentile. Note that because we observed with the instrument rotator off, and shifted and rotated the images in processing to register celestial sources, the data regions in our images are not square or round, and can have rather odd shapes. This is

why the different percentile sensitivities we plot reach the edge of the valid data region and drop to zero sensitivity at different projected radii from the star.

We have converted the three sensitivity curves to planet masses using both sets of theoretical models. The Burrows et al. (2003) results are plotted as solid lines, and the Baraffe et al. (2003) results as dotted lines. Note that the Baraffe et al. (2003) models indicate better sensitivities in the  $L'$  data sets, than the Burrows et al. (2003) models, but often comparable or worse sensitivities in the  $M$  band data sets.

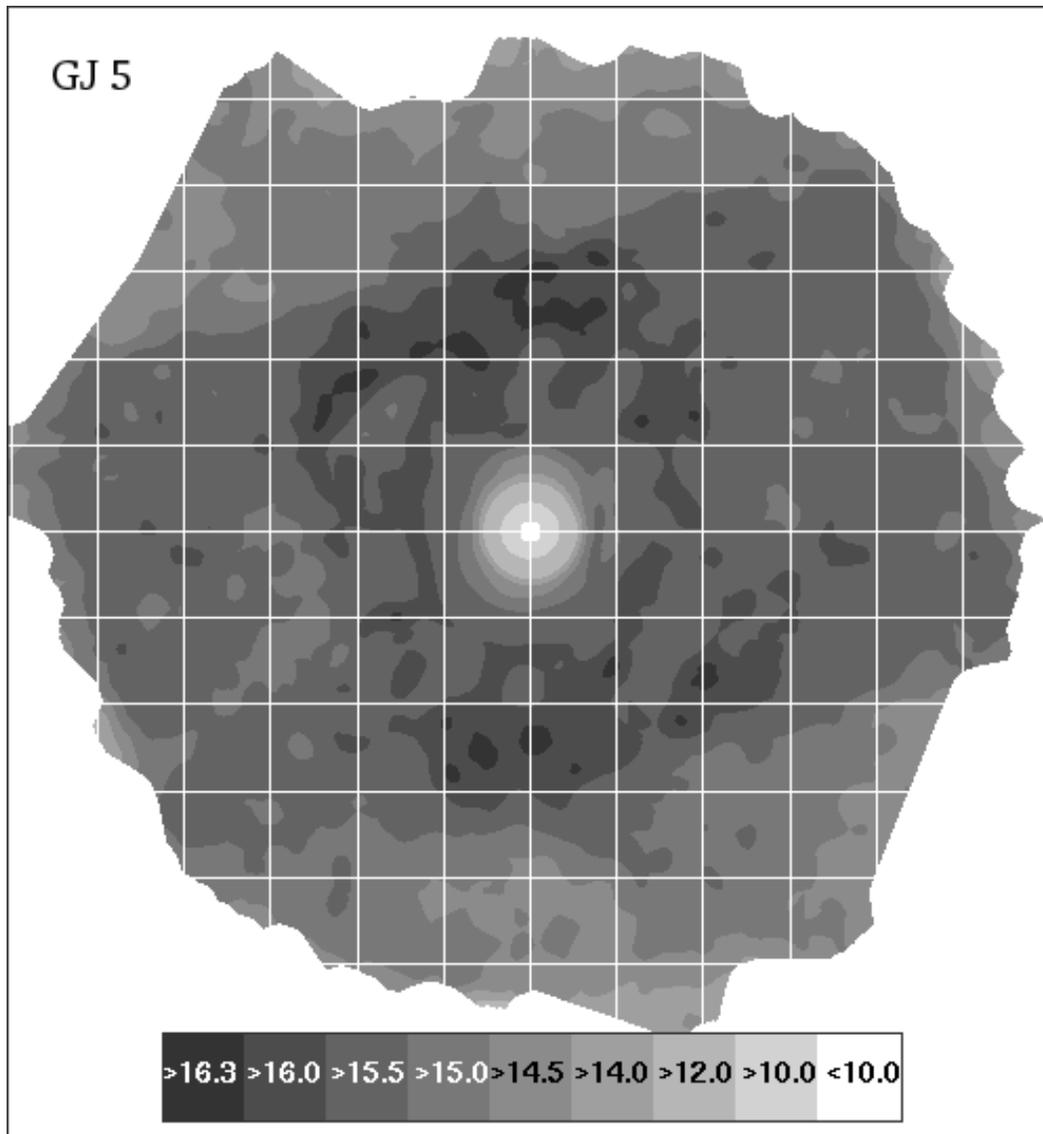


Figure 3.20 Sensitivity contours for our  $L'$  images of star GJ 5. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

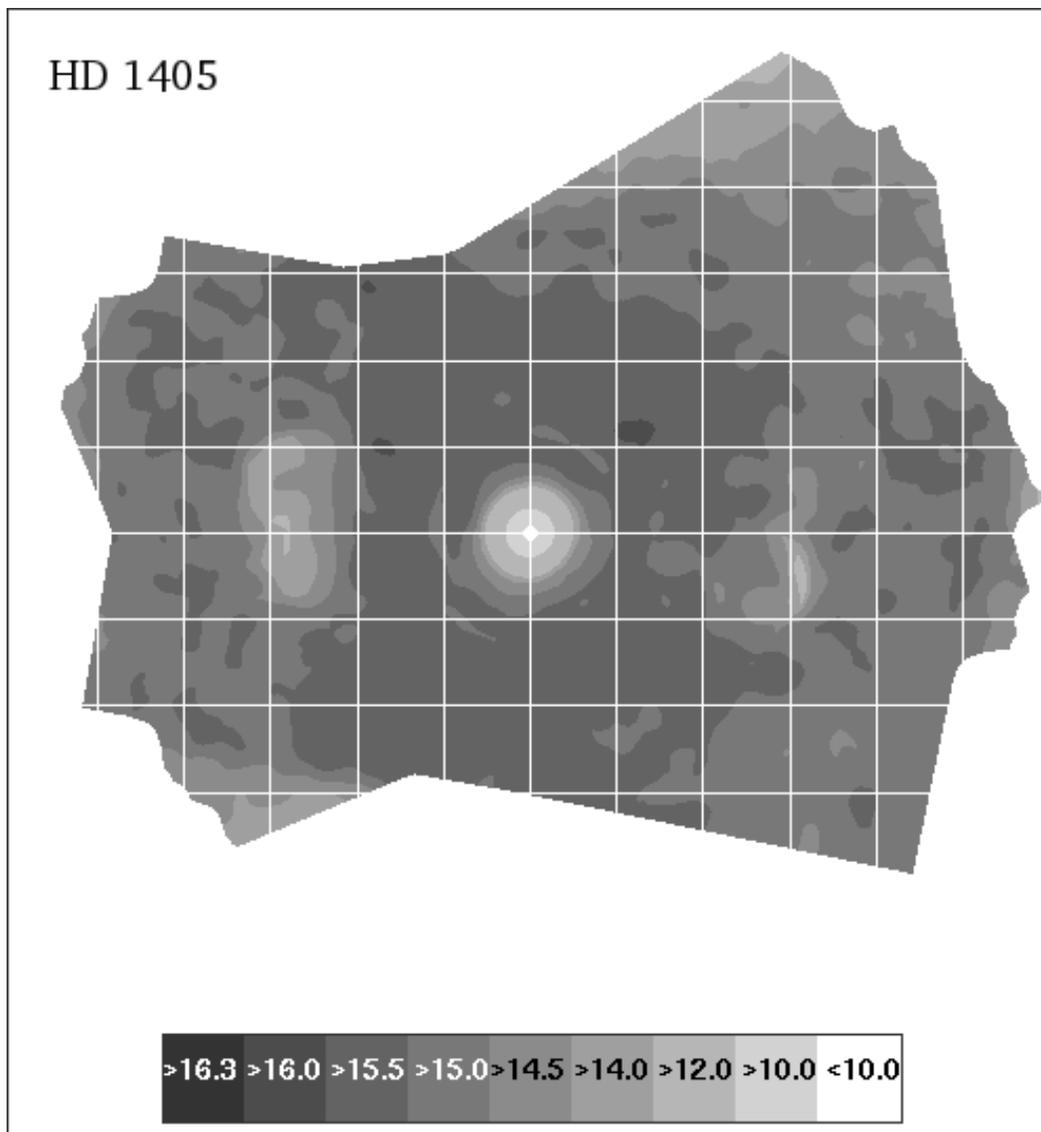


Figure 3.21 Sensitivity Contours for our  $L'$  images of star HD 1405. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

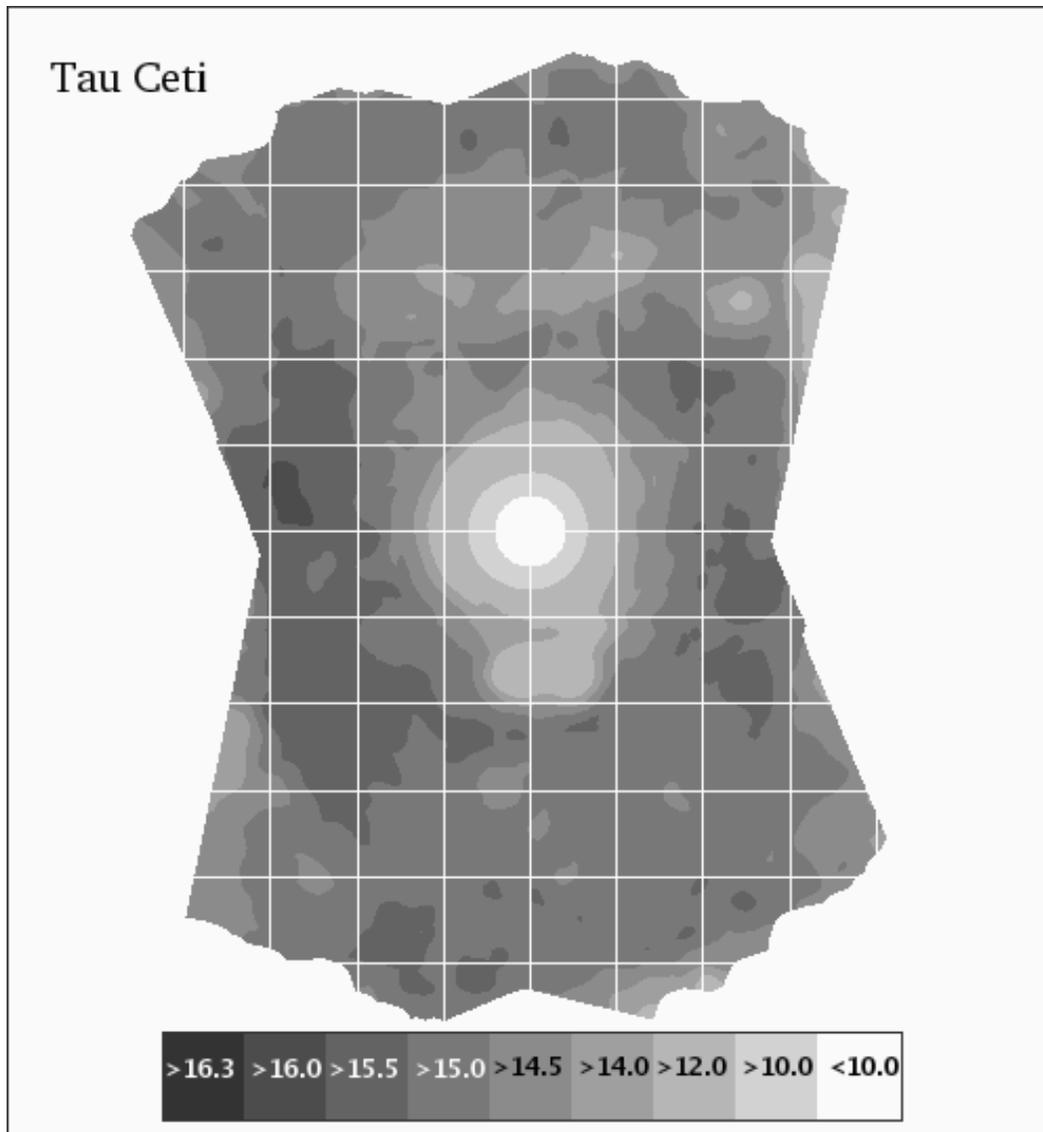


Figure 3.22 Sensitivity Contours for our  $L'$  images of  $\tau$  Ceti. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.0$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.5$  but not as faint as  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

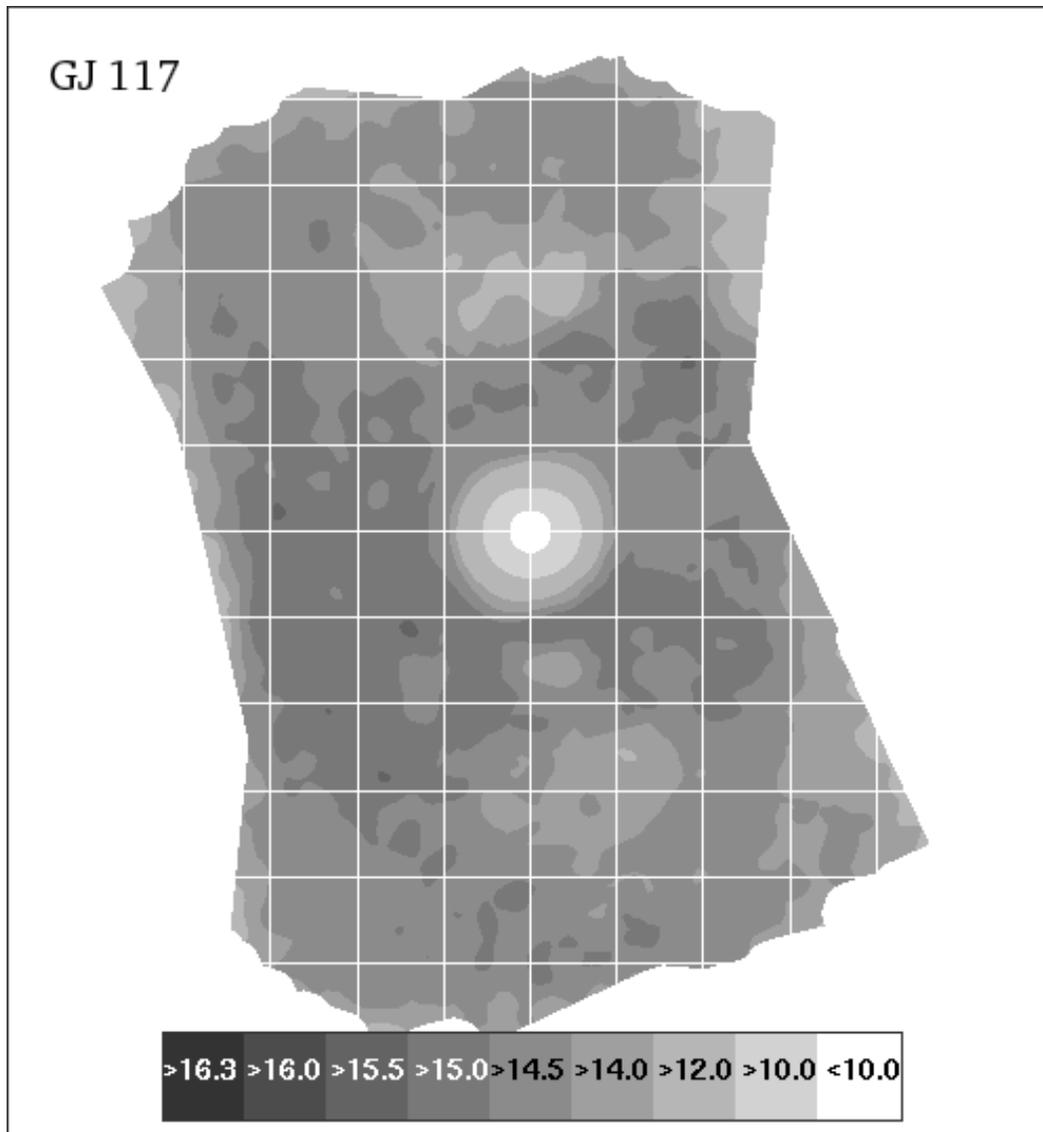


Figure 3.23 Sensitivity Contours for our  $L'$  images of GJ 117. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.5$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.0$  but not as faint as  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

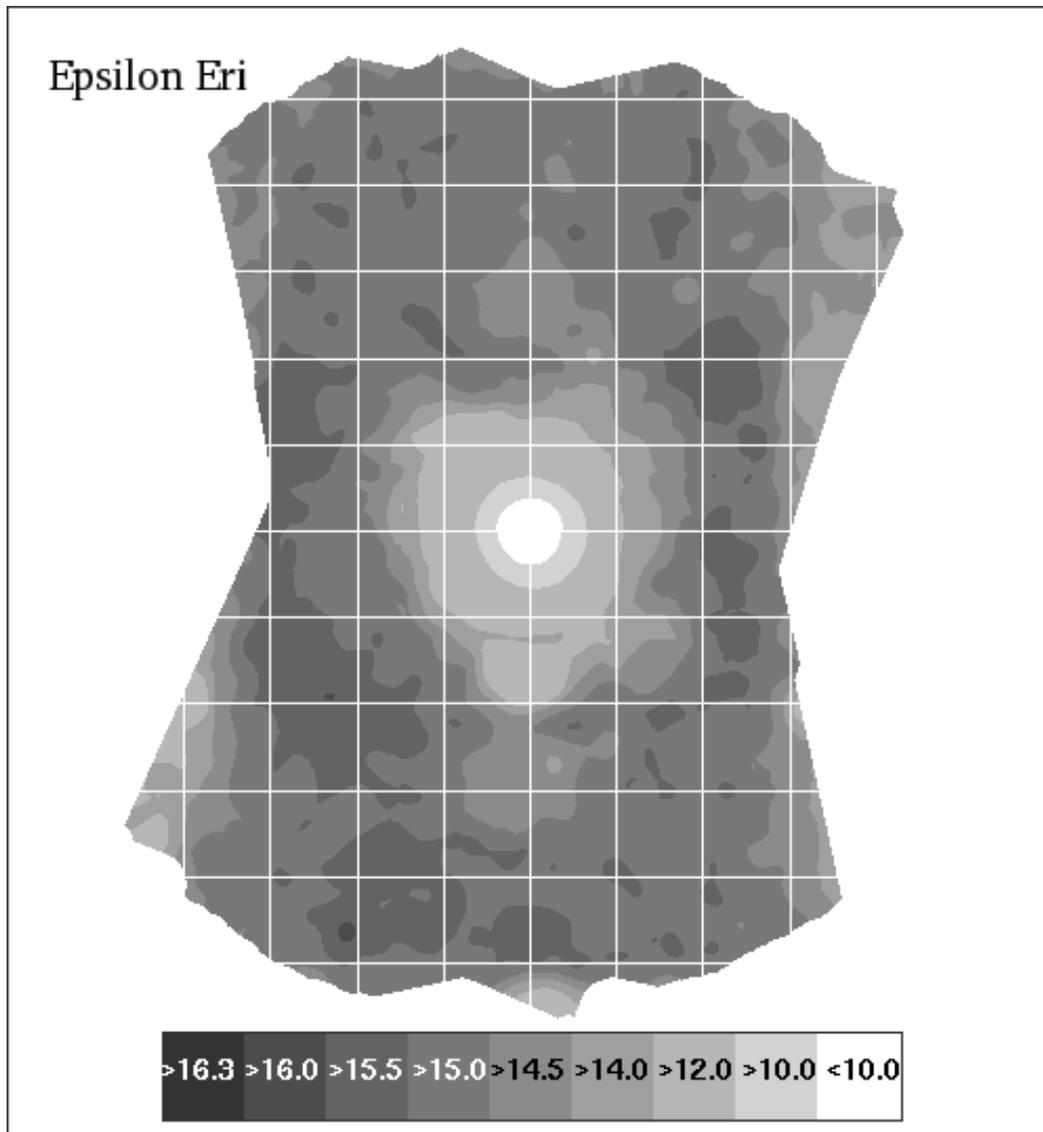


Figure 3.24 Sensitivity Contours for our  $L'$  images of  $\epsilon$  Eri. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.0$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.5$  but not as faint as  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

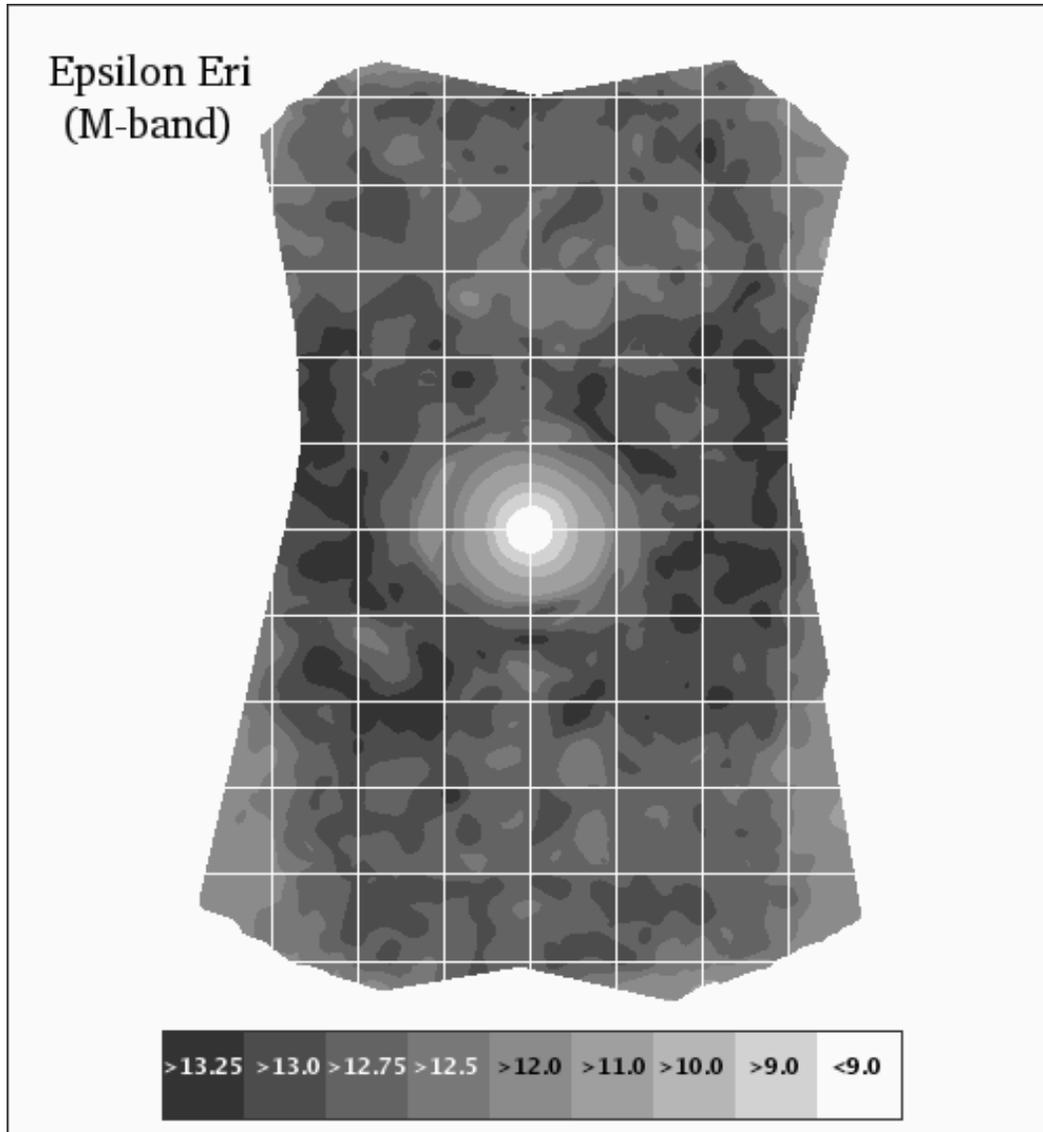


Figure 3.25 Sensitivity Contours for our  $M$  band images of  $\epsilon$  Eri. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $M = 13.25$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

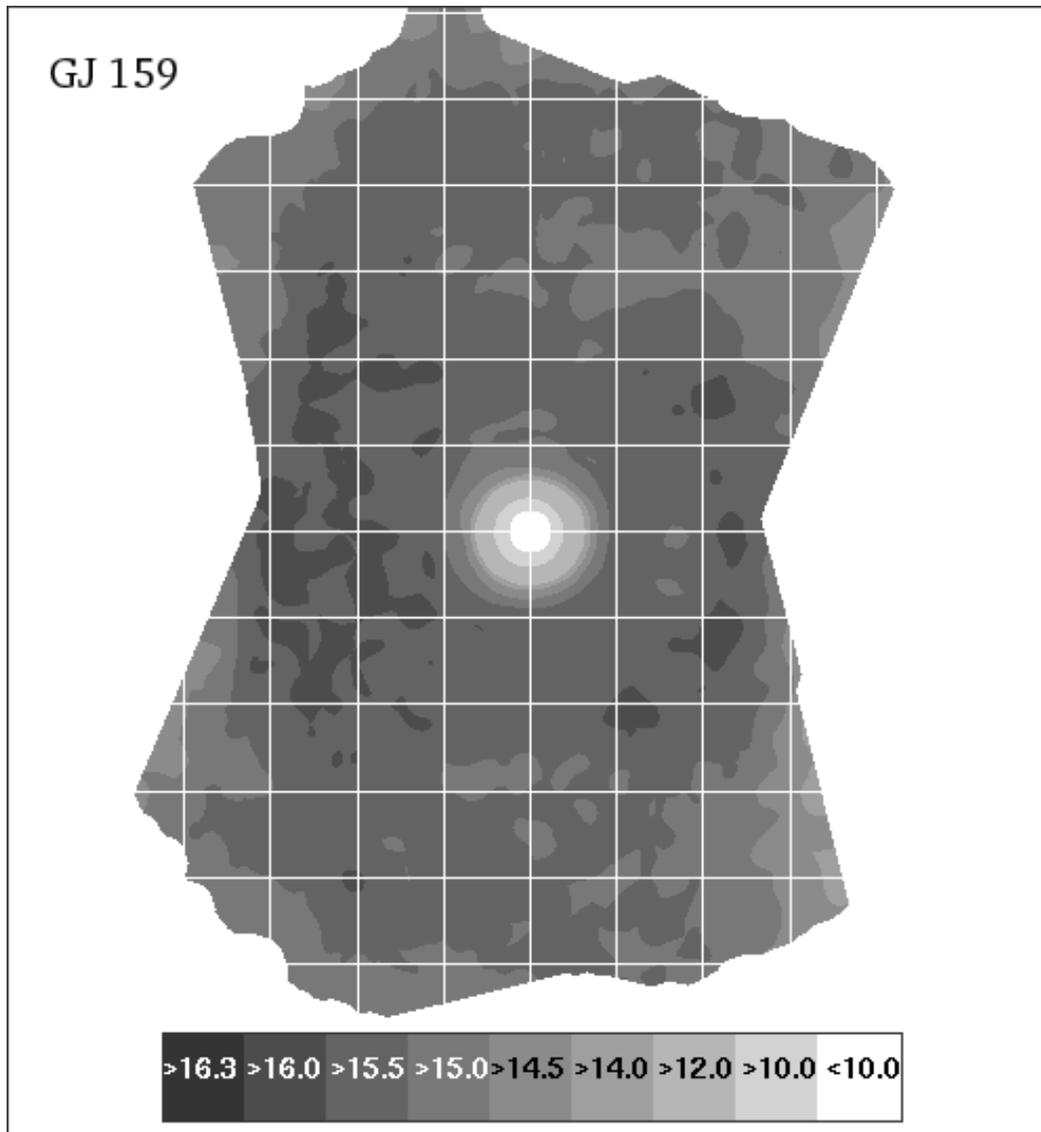


Figure 3.26 Sensitivity Contours for our  $L'$  images of GJ 159. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

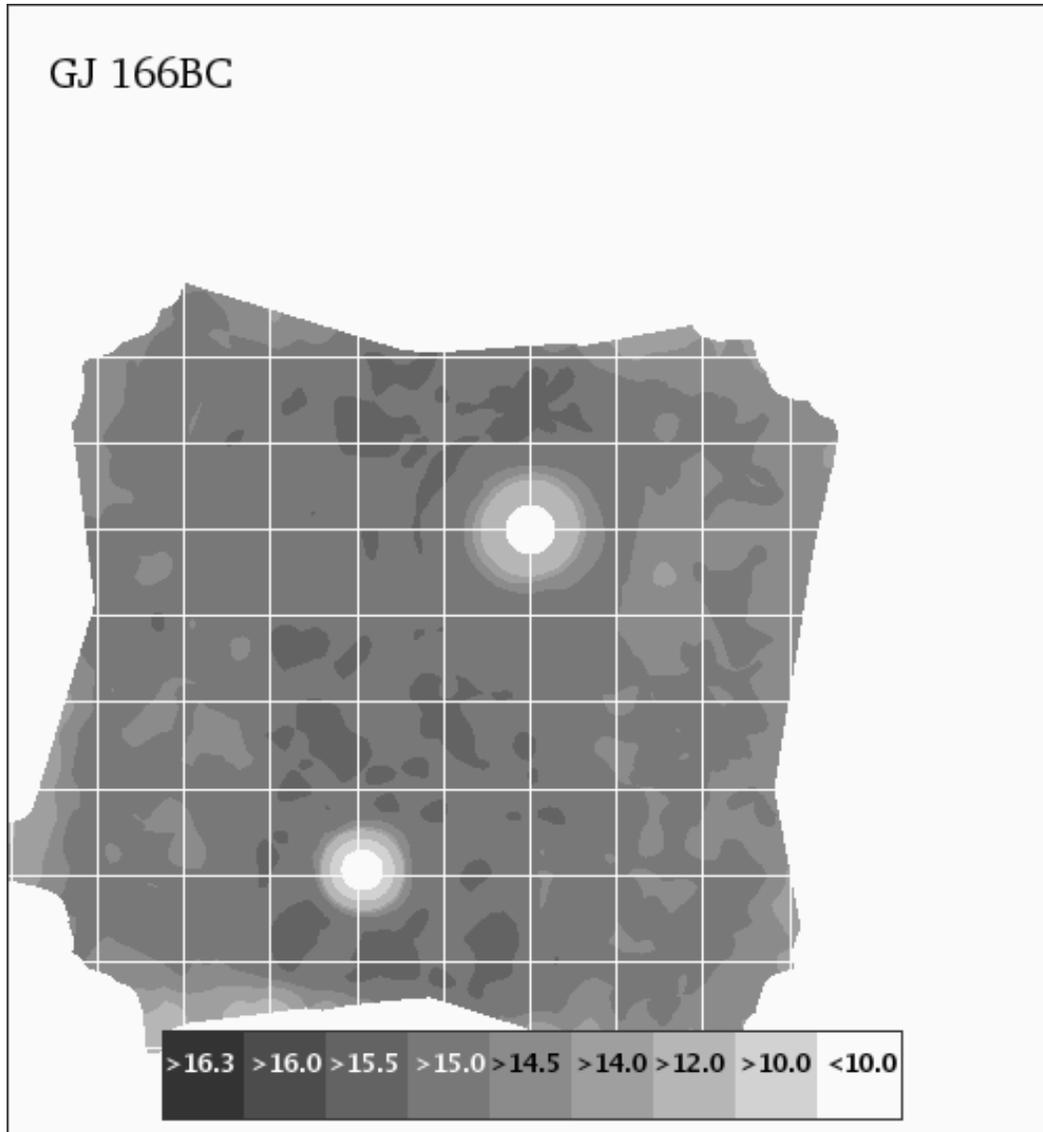


Figure 3.27 Sensitivity Contours for our  $L'$  images of GJ 166BC. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

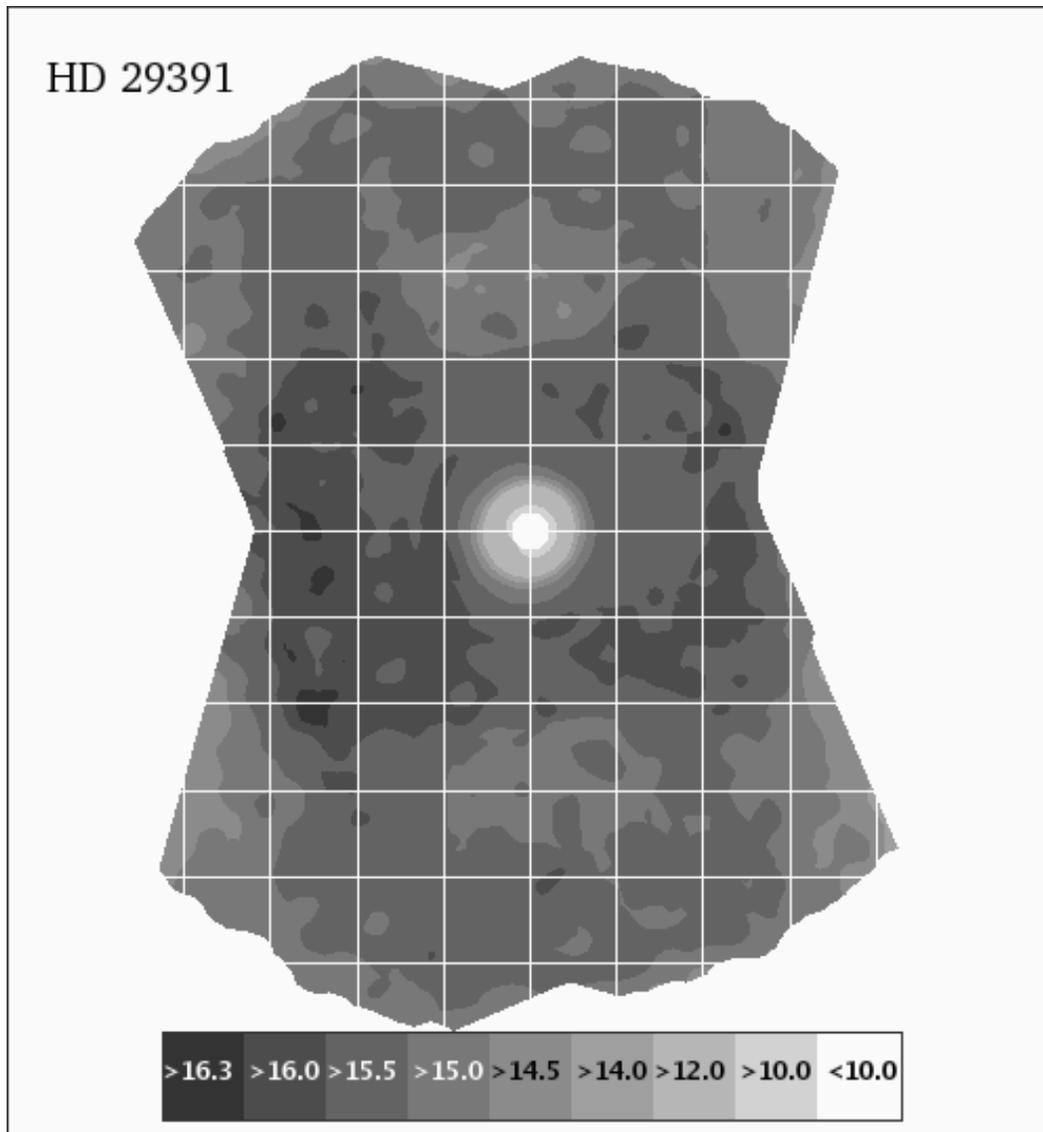


Figure 3.28 Sensitivity Contours for our  $L'$  images of HD 29391. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.3$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

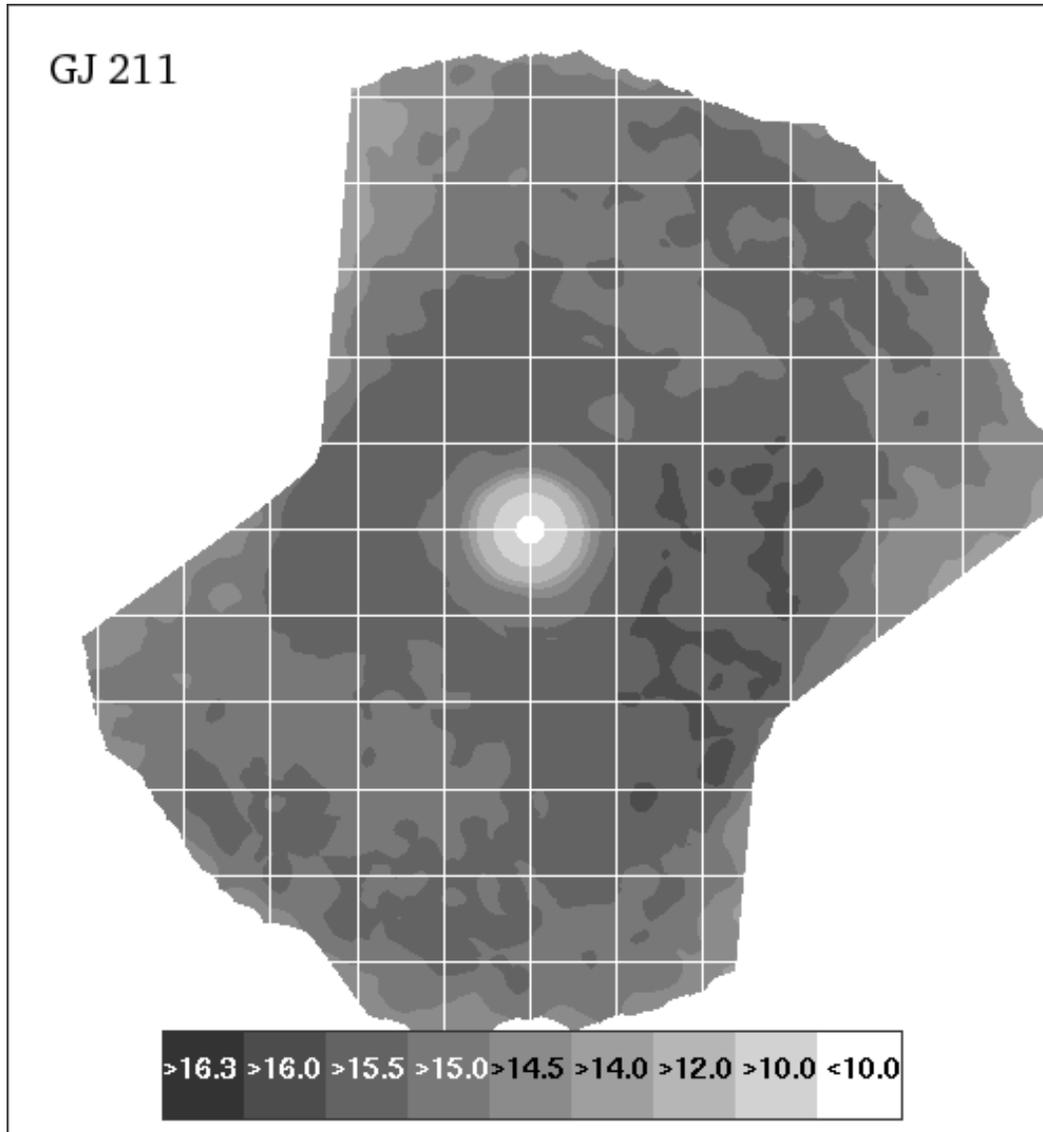


Figure 3.29 Sensitivity Contours for our  $L'$  images of GJ 211. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.0$ . These are 10 nominal  $\sigma$  sensitivities. North is up; the grid squares are 2x2 arcsec.

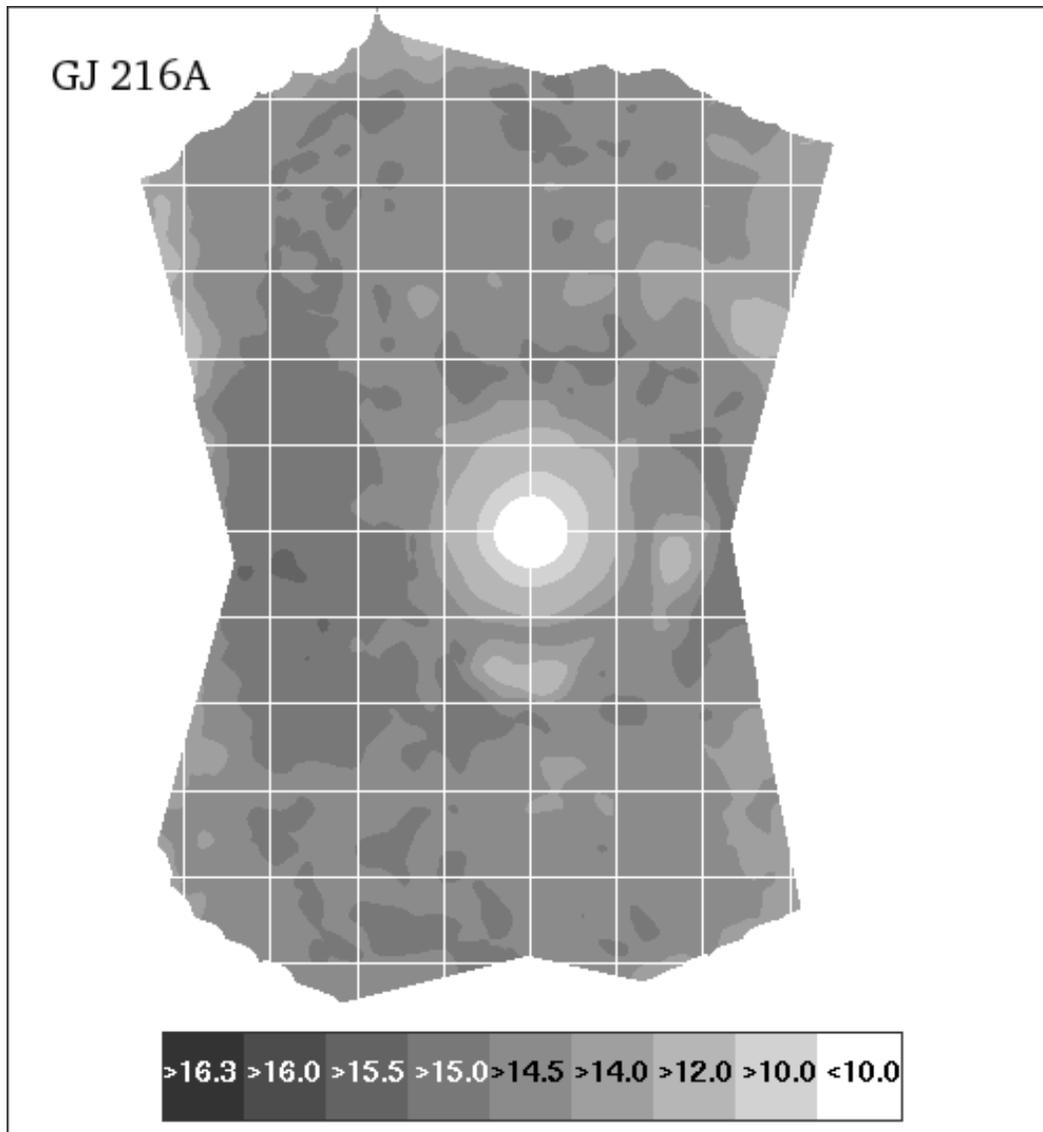


Figure 3.30 Sensitivity Contours for our  $L'$  images of GJ 216A. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.5$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.0$  but not as faint as  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

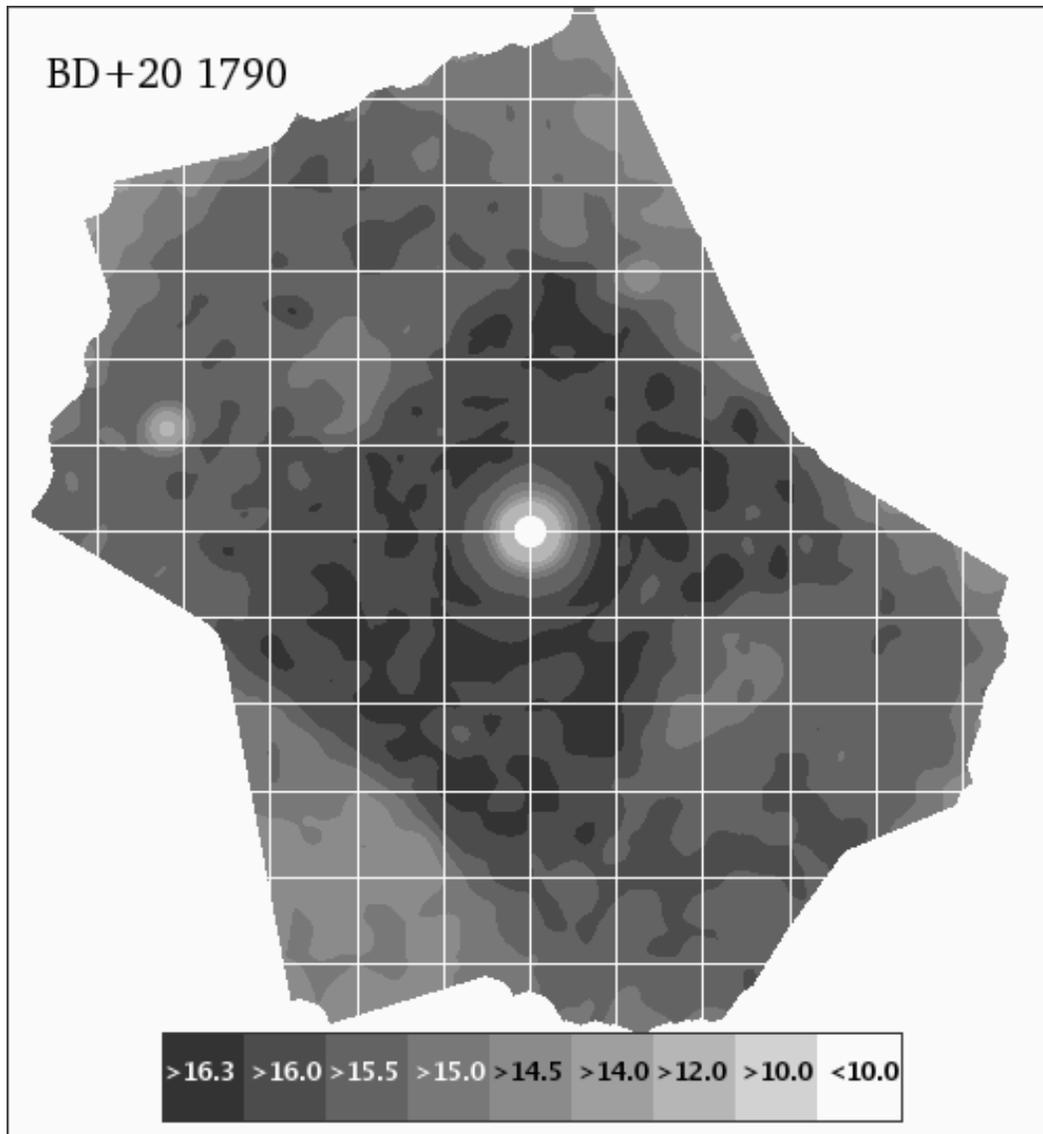


Figure 3.31 Sensitivity Contours for our  $L'$  images of BD+20 1790. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

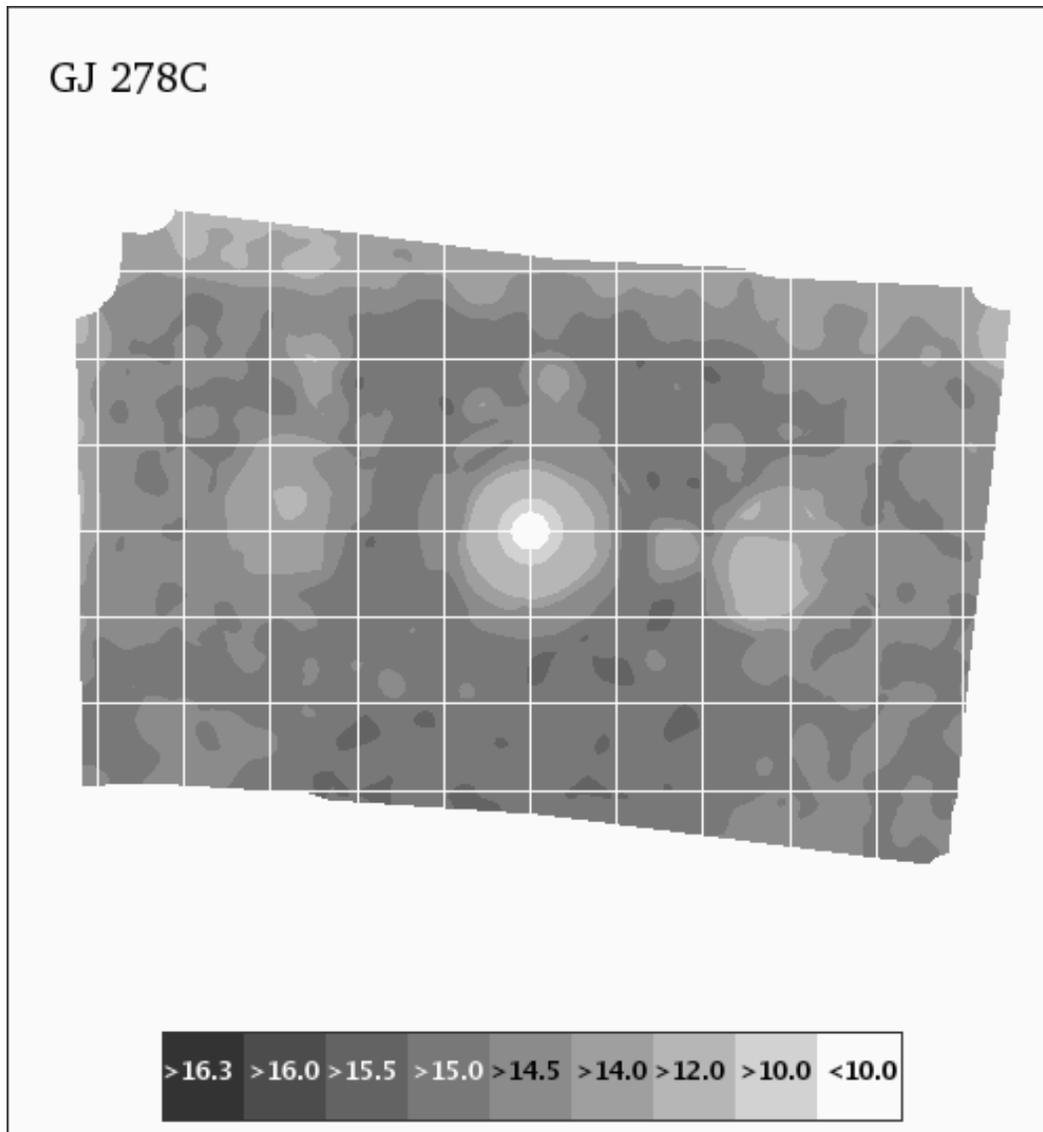


Figure 3.32 Sensitivity Contours for our  $L'$  images of GJ 278C. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.5$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.0$  but not as faint as  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

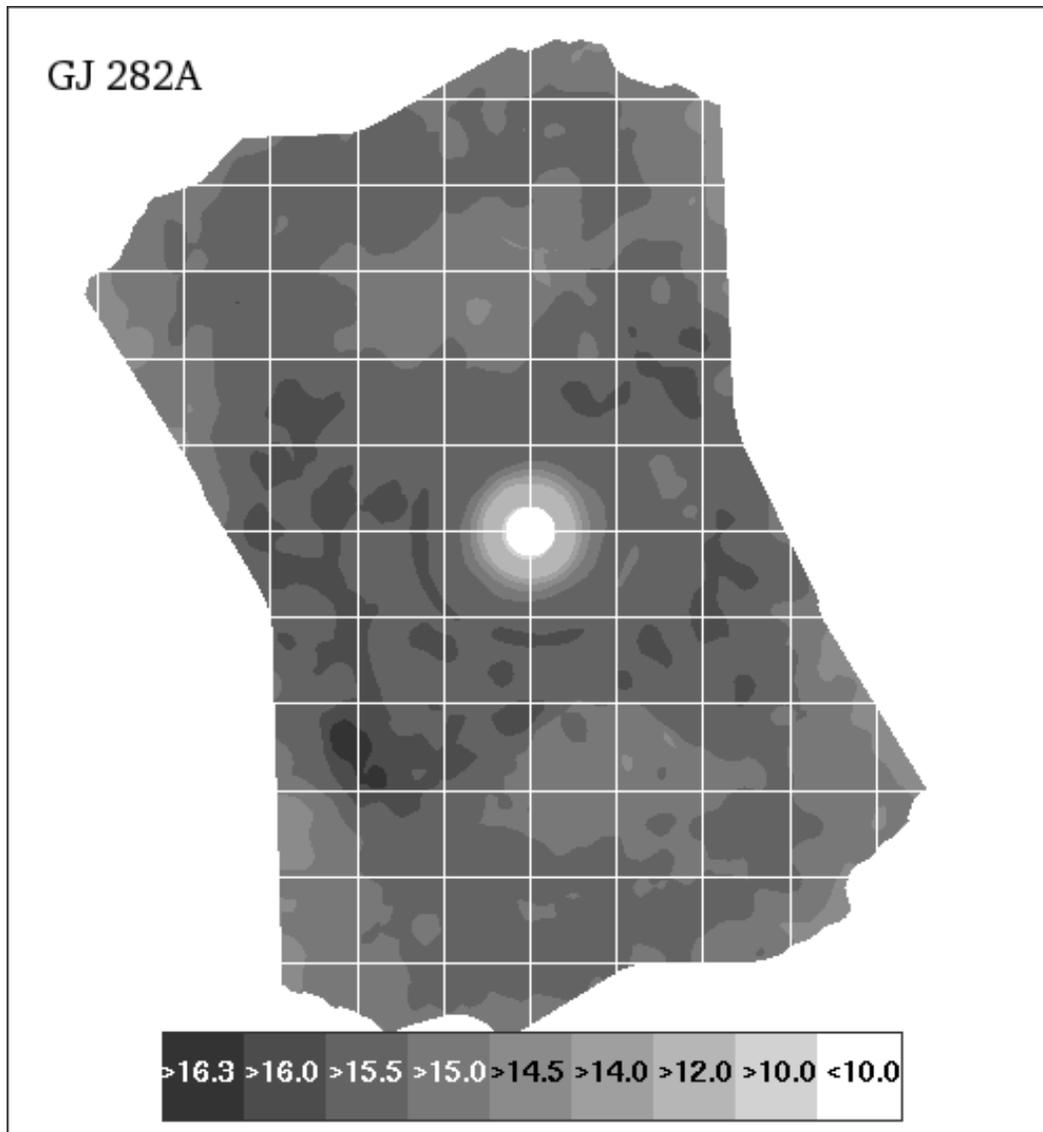


Figure 3.33 Sensitivity Contours for our  $L'$  images of GJ282A. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.3$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

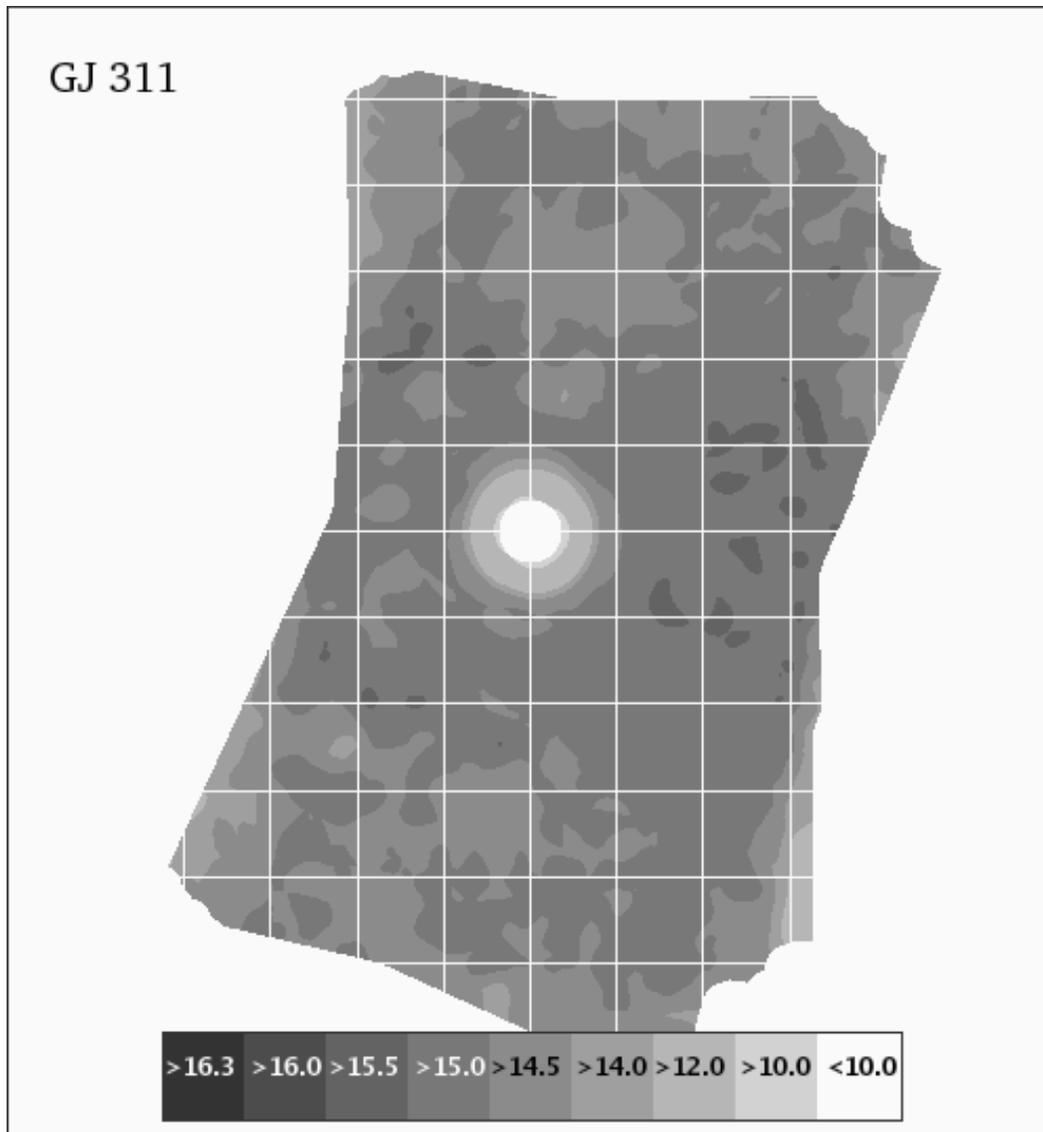


Figure 3.34 Sensitivity Contours for our  $L'$  images of GJ 311. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.5$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.0$  but not as faint as  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

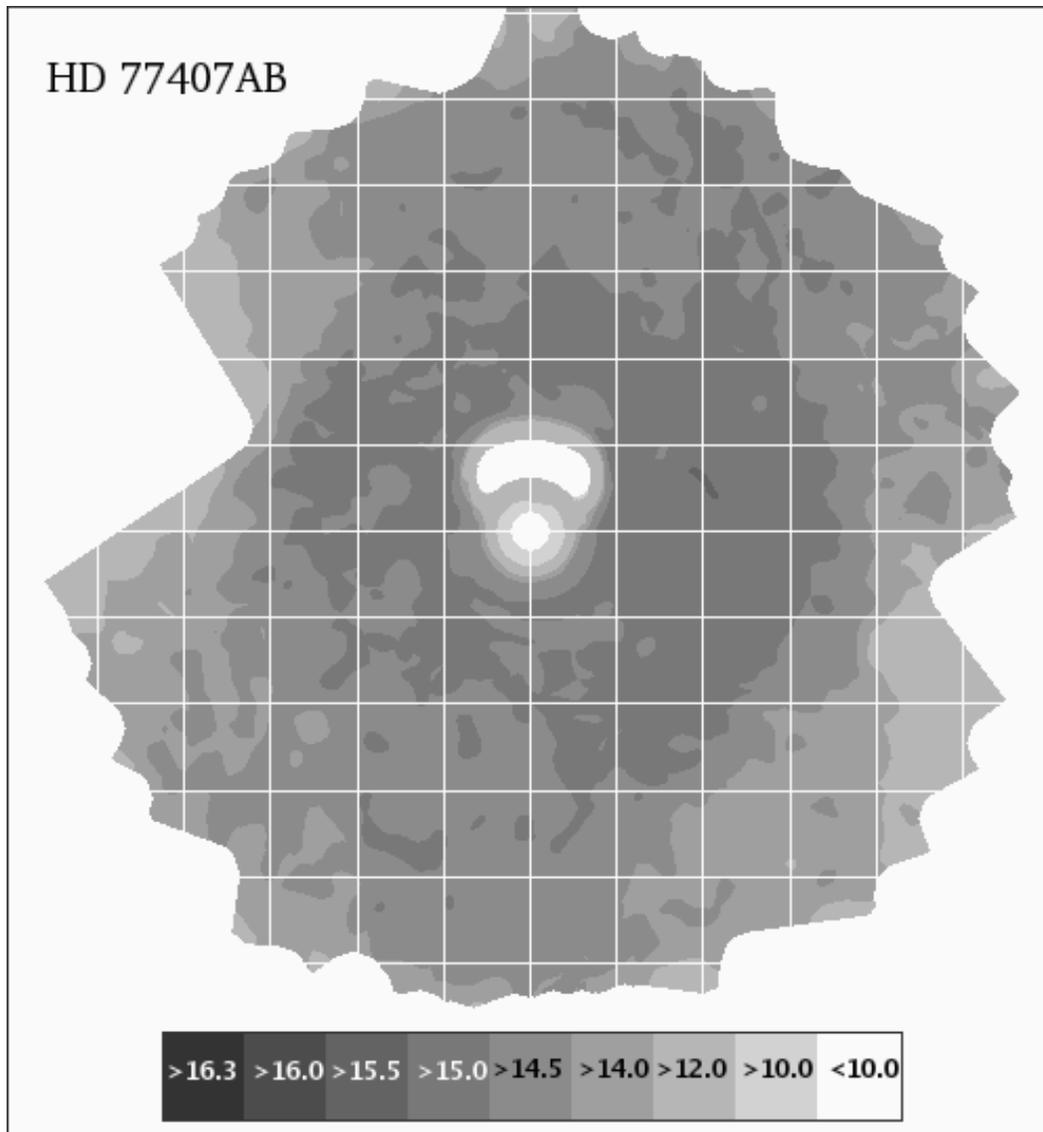


Figure 3.35 Sensitivity Contours for our  $L'$  images of HD 77407AB. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.0$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 14.5$  but not as faint as  $L' = 15.0$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

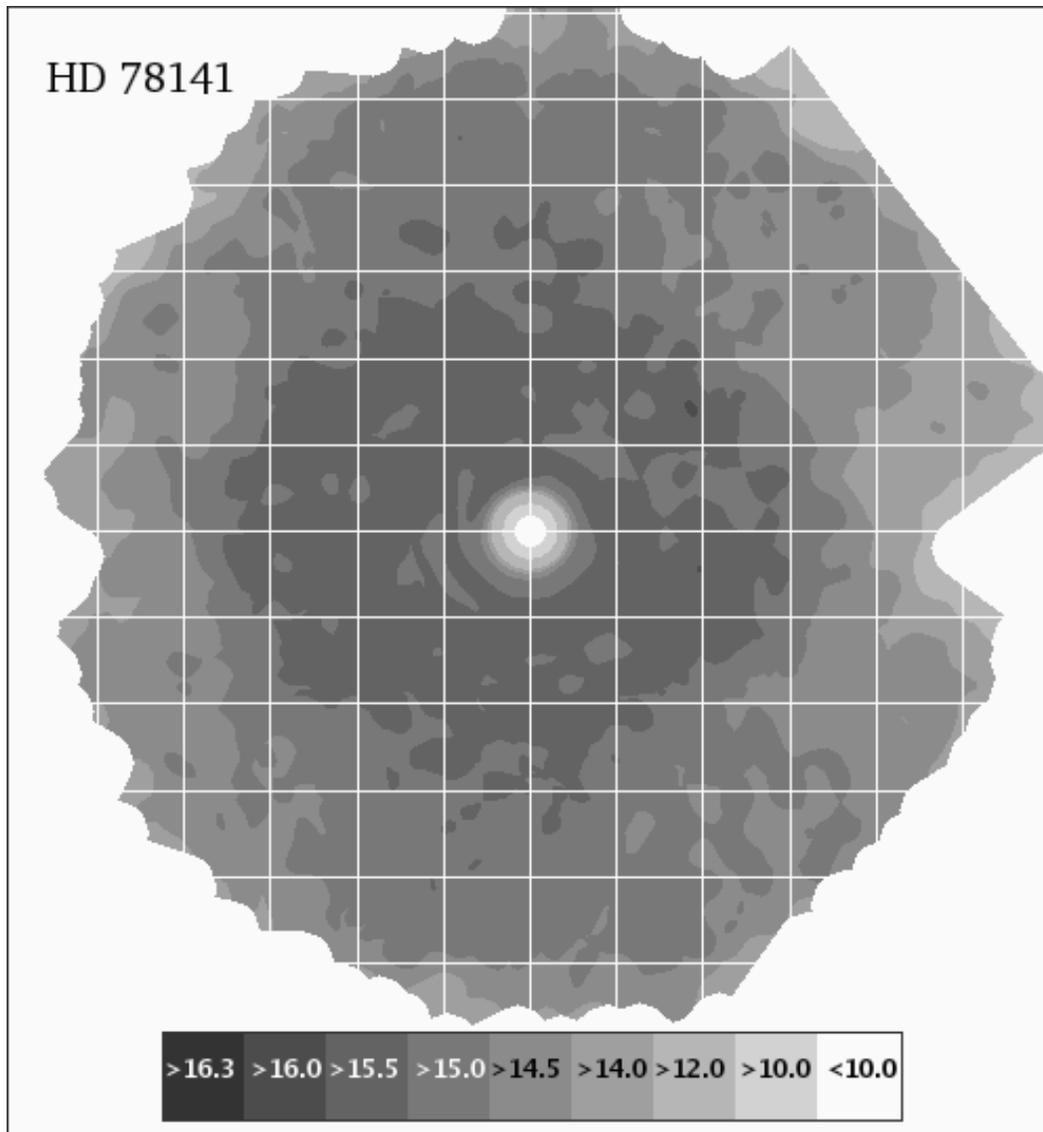


Figure 3.36 Sensitivity Contours for our  $L'$  images of HD 78141. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.0$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 15.5$  but not as faint as  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

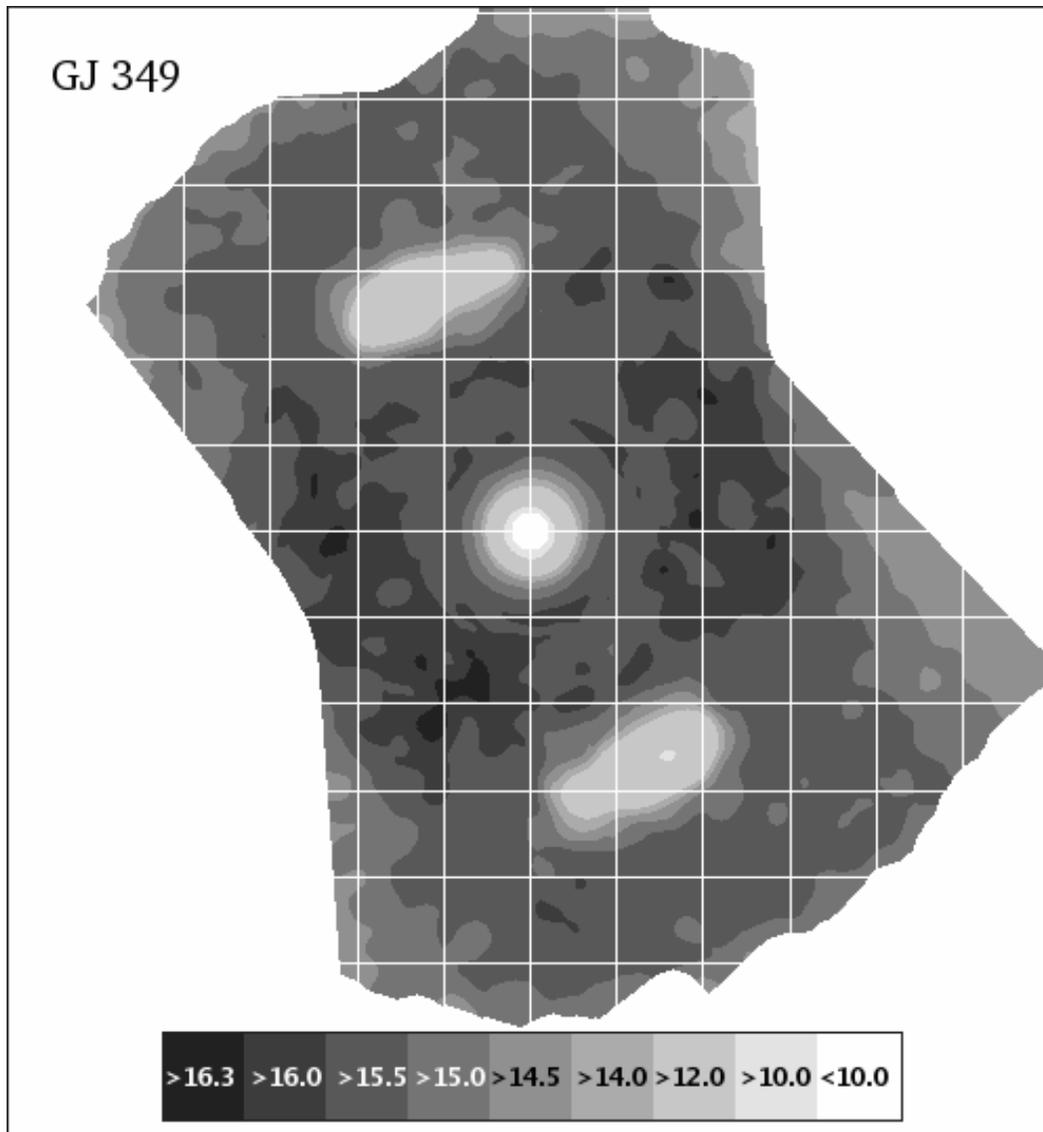


Figure 3.37 Sensitivity Contours for our  $L'$  images of GJ 349. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.3$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

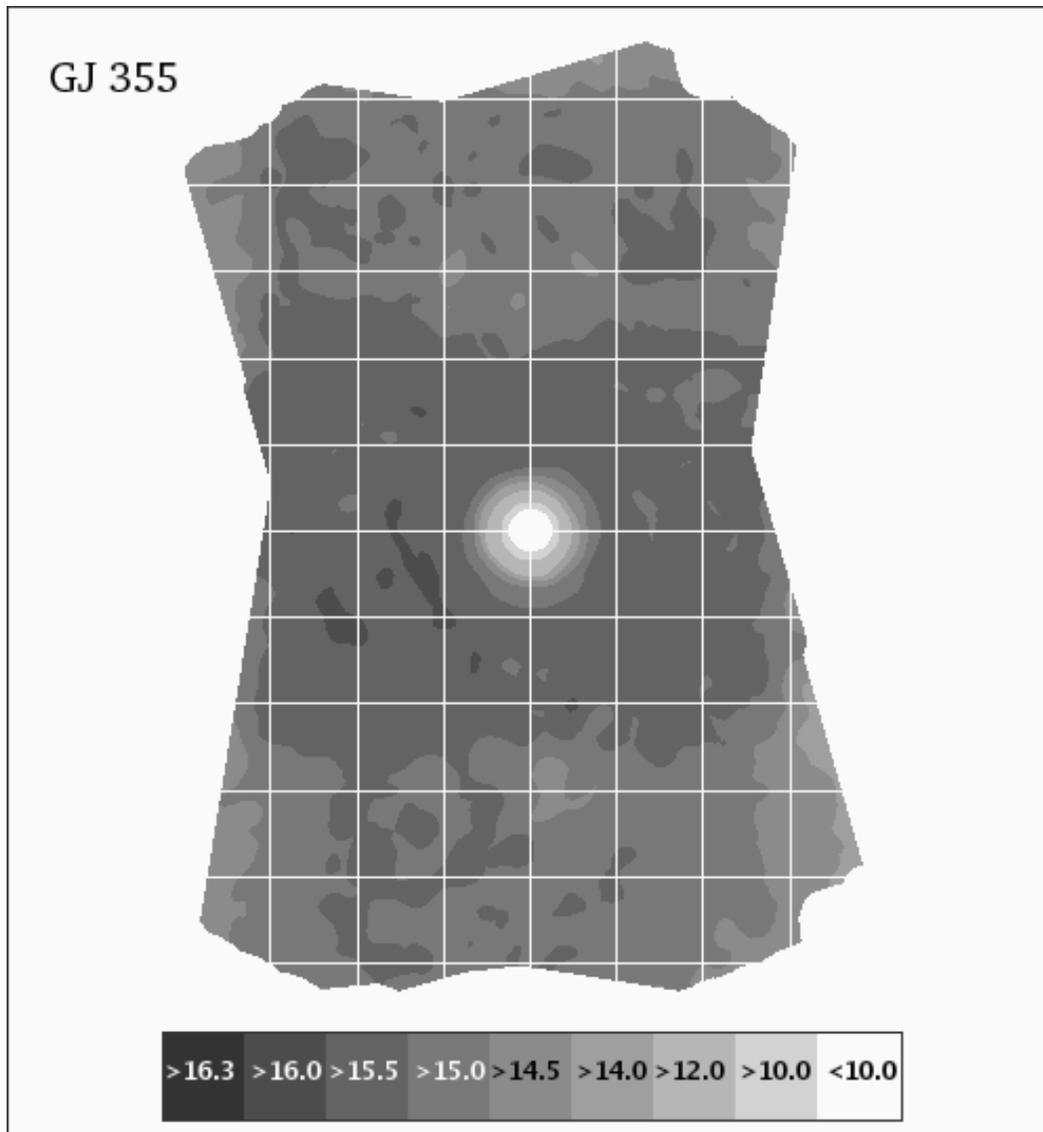


Figure 3.38 Sensitivity Contours for our  $L'$  images of GJ 355. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.0$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.5$  but not as faint as  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

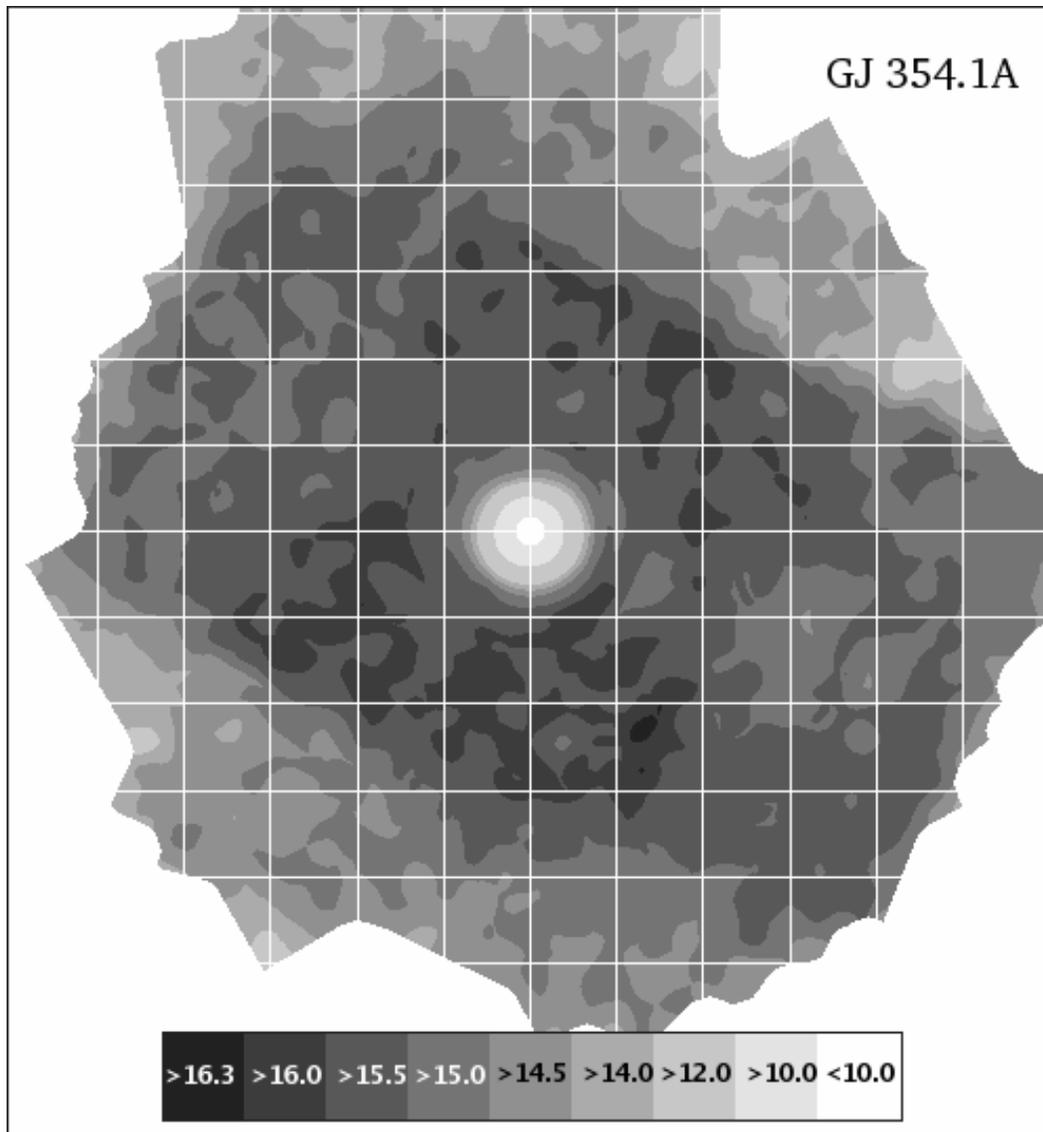


Figure 3.39 Sensitivity Contours for our  $L'$  images of GJ 354.1A. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.3$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

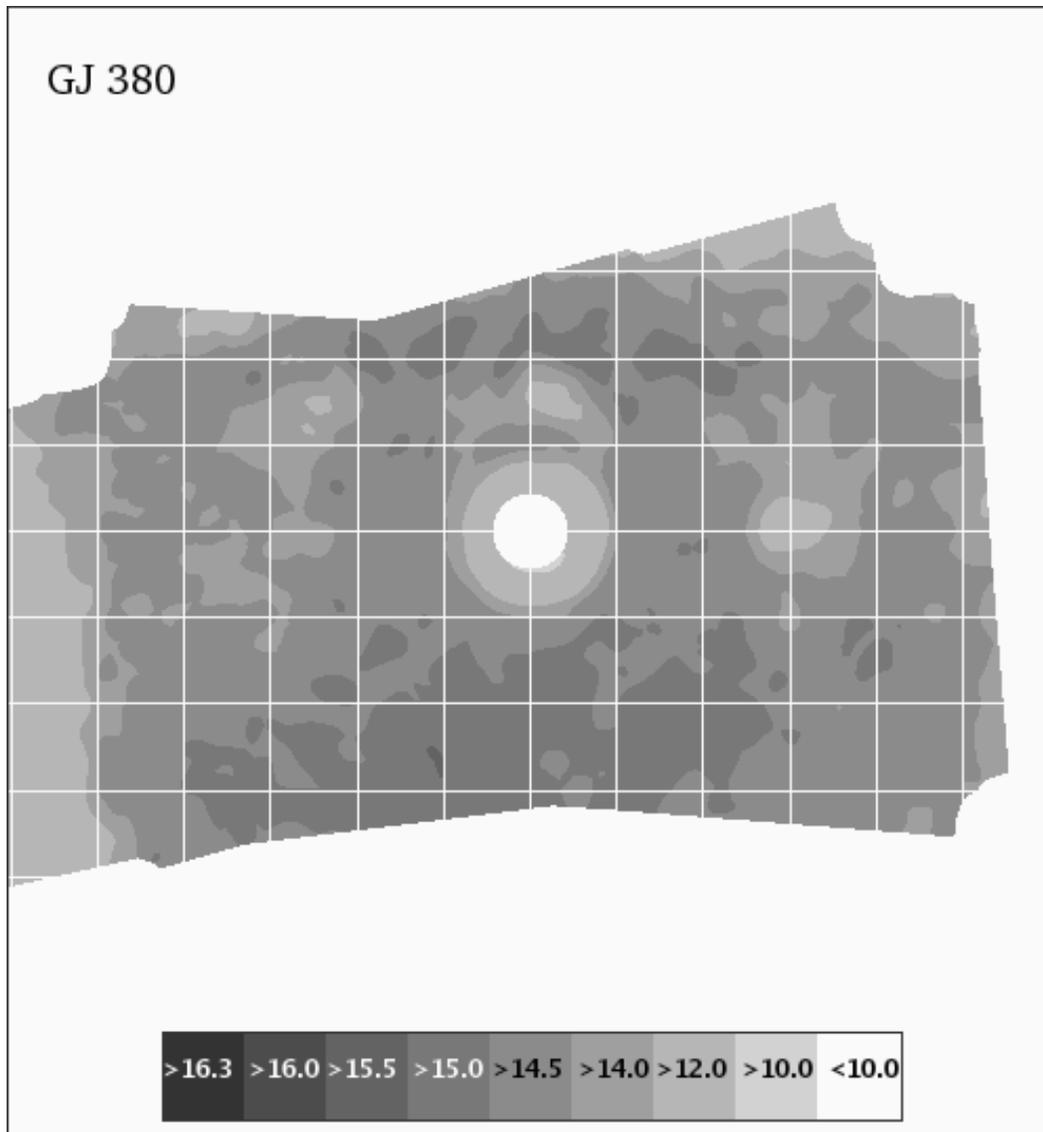


Figure 3.40 Sensitivity Contours for our  $L'$  images of GJ 380. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.5$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 15.0$  but not as faint as  $L' = 15,5$ . These are nominal  $10 \sigma$  sensitivities. North is up; the grid squares are  $2 \times 2$  arcsec.

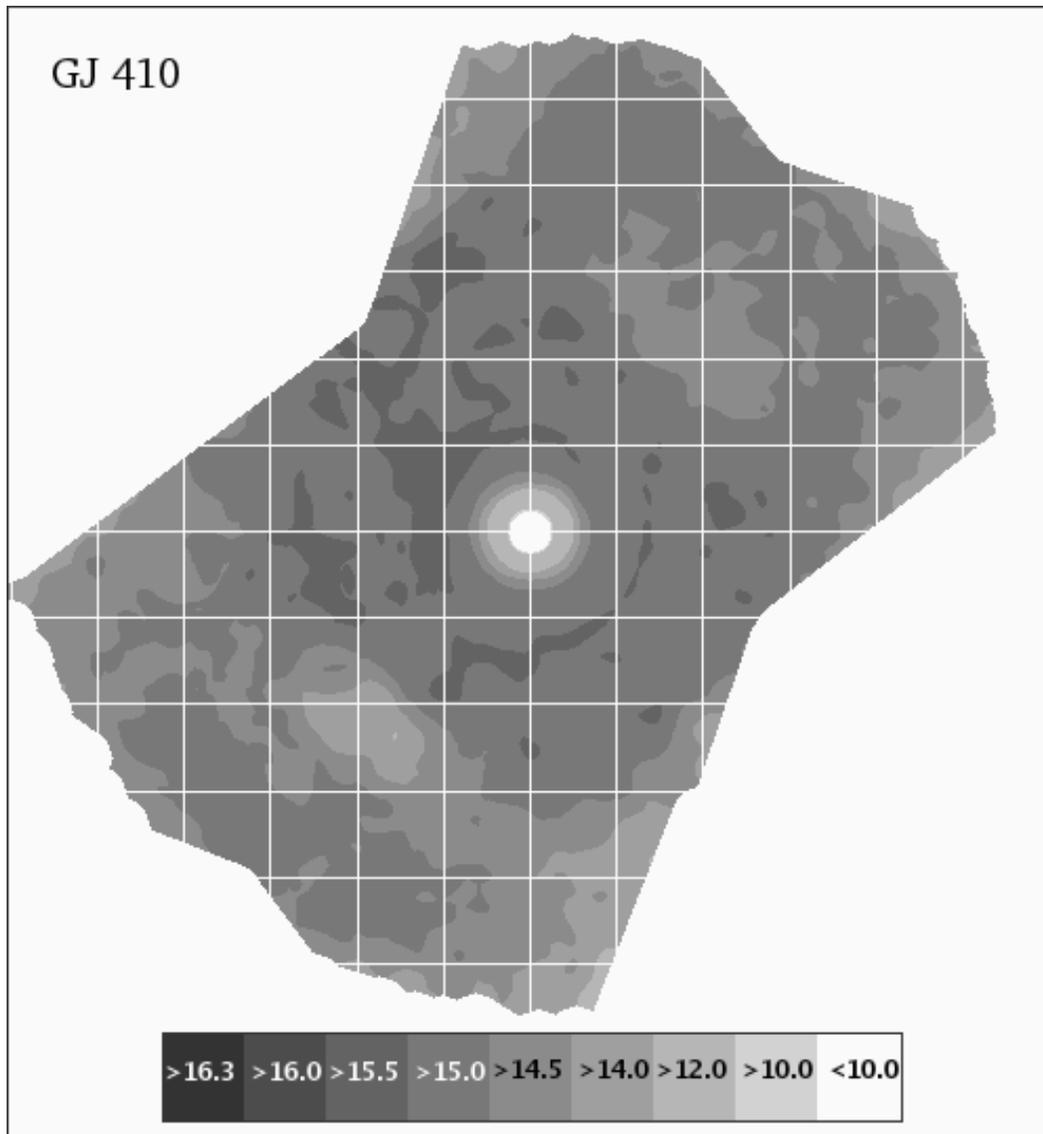


Figure 3.41 Sensitivity Contours for our  $L'$  images of GJ 410. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

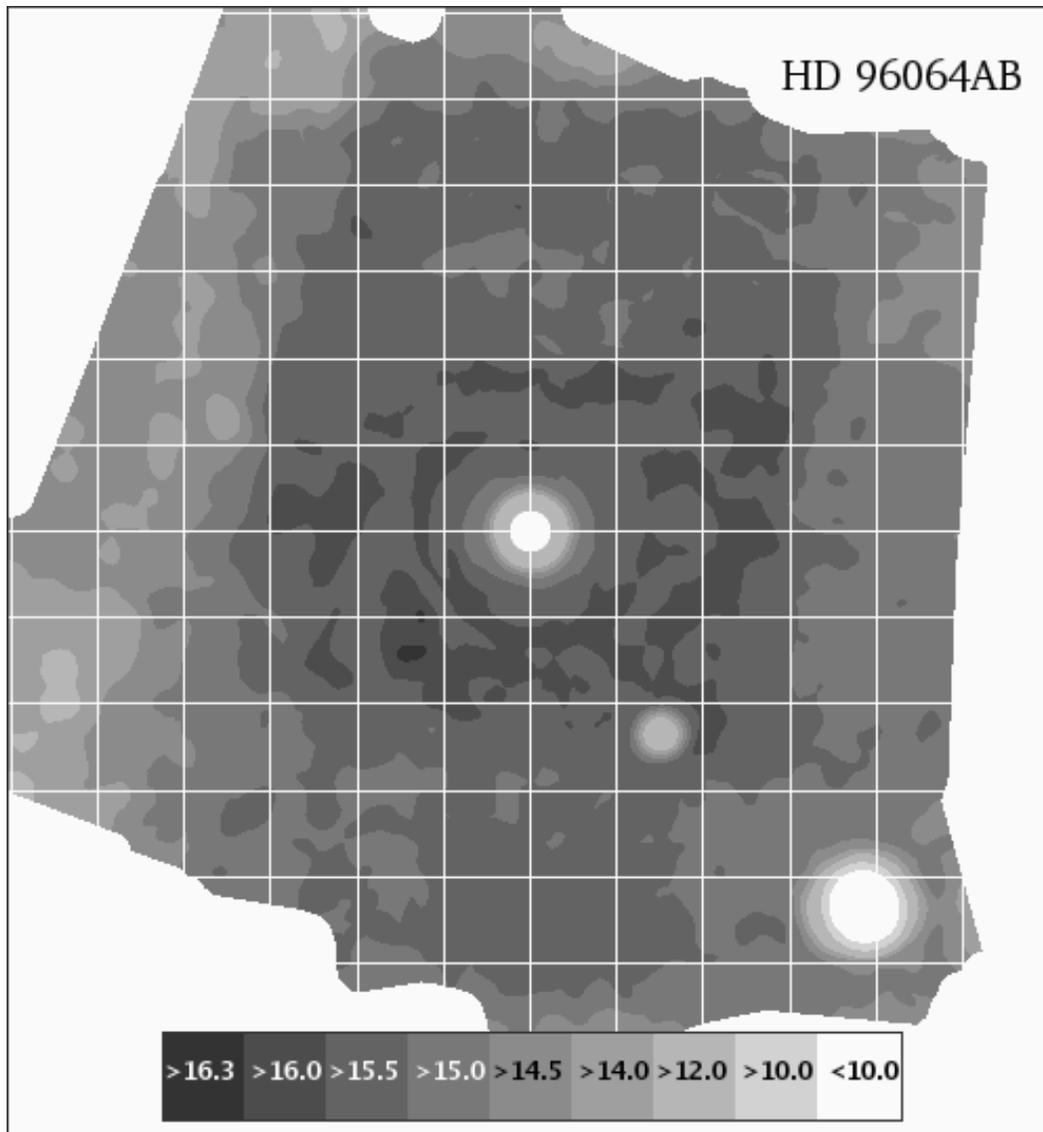


Figure 3.42 Sensitivity Contours for our  $L'$  images of HD 96064AB. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.3$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2\times 2$  arcsec.

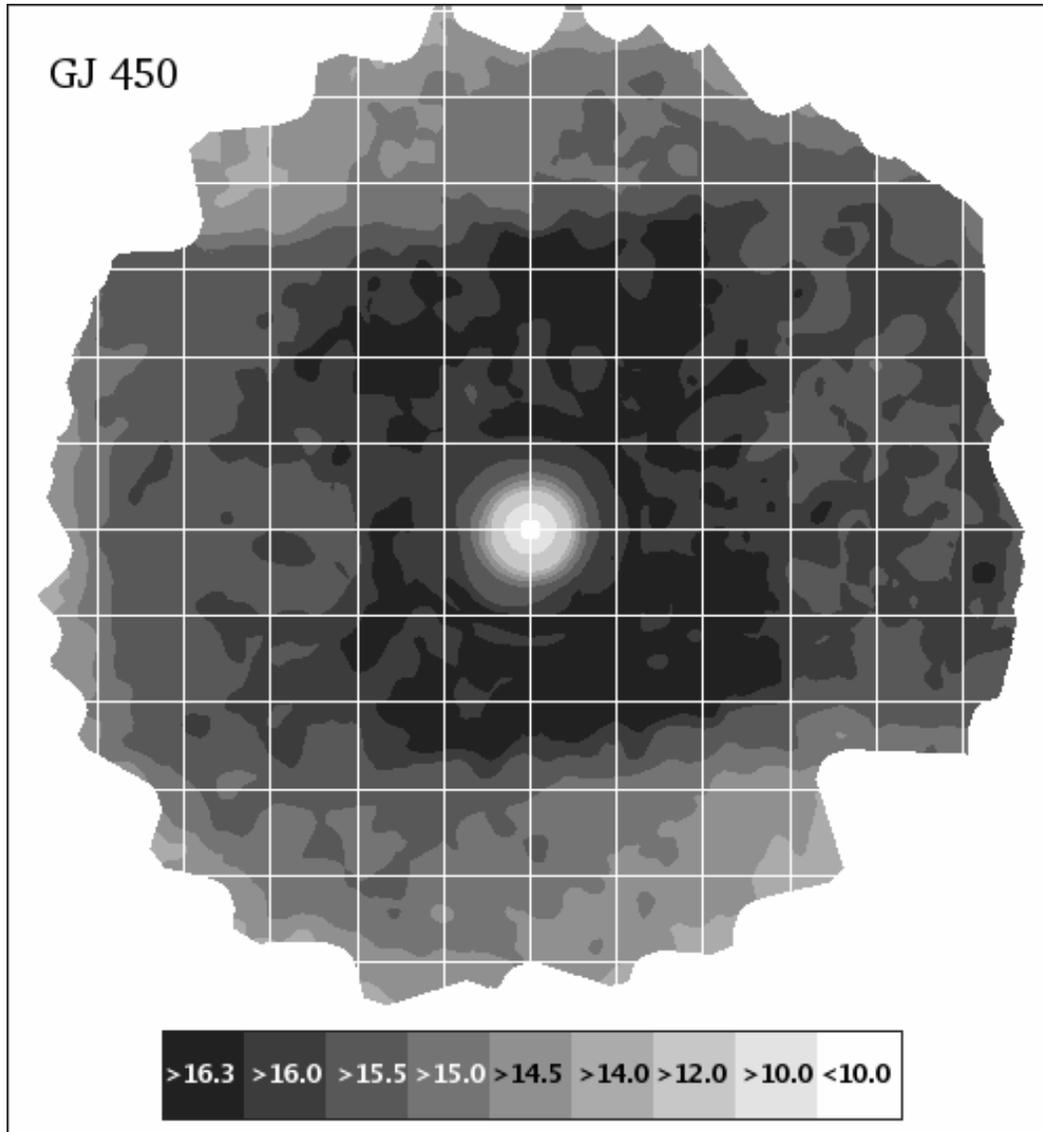


Figure 3.43 Sensitivity Contours for our  $L'$  images of GJ 450. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

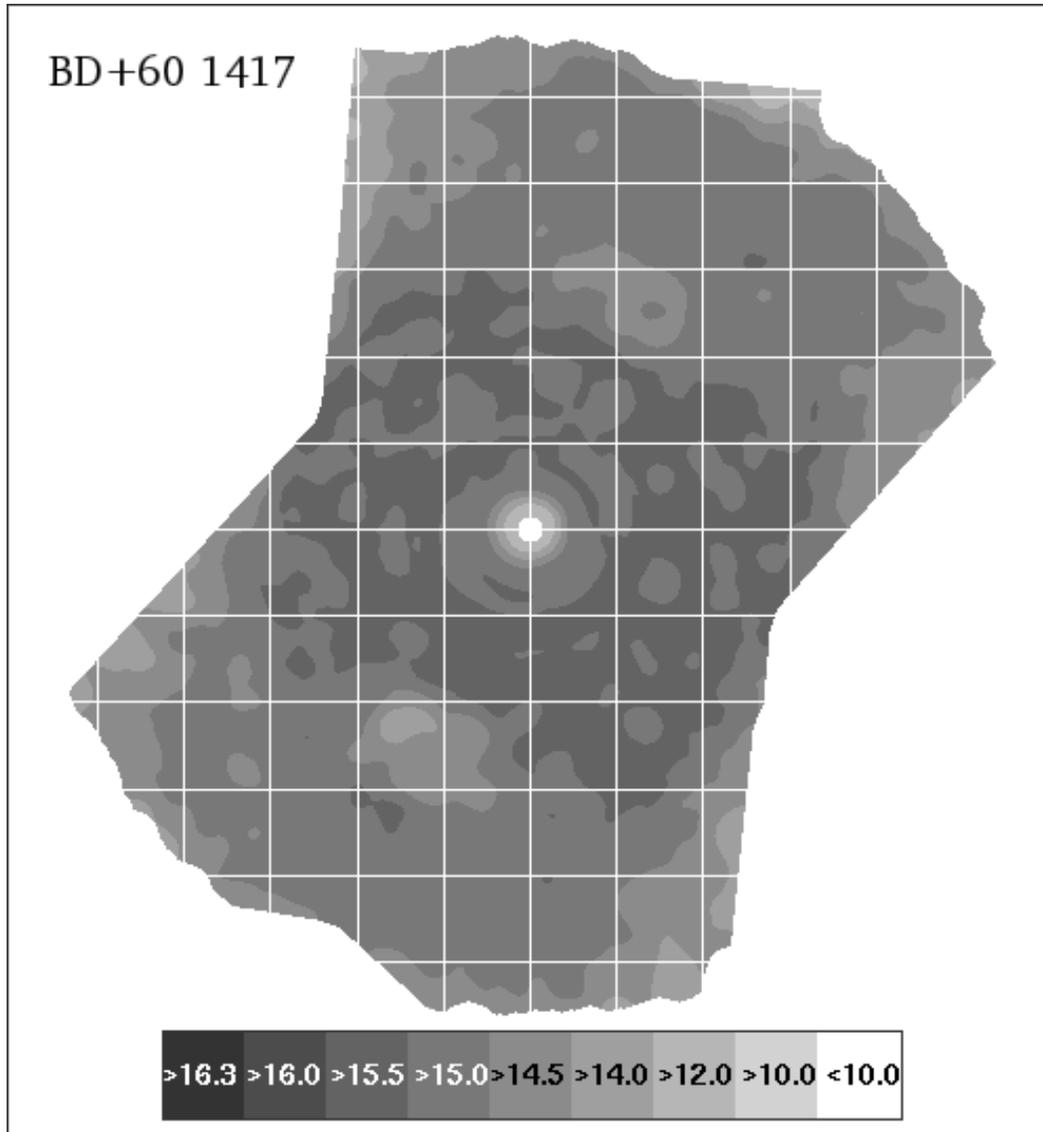


Figure 3.44 Sensitivity Contours for our  $L'$  images of BD+60 1417. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

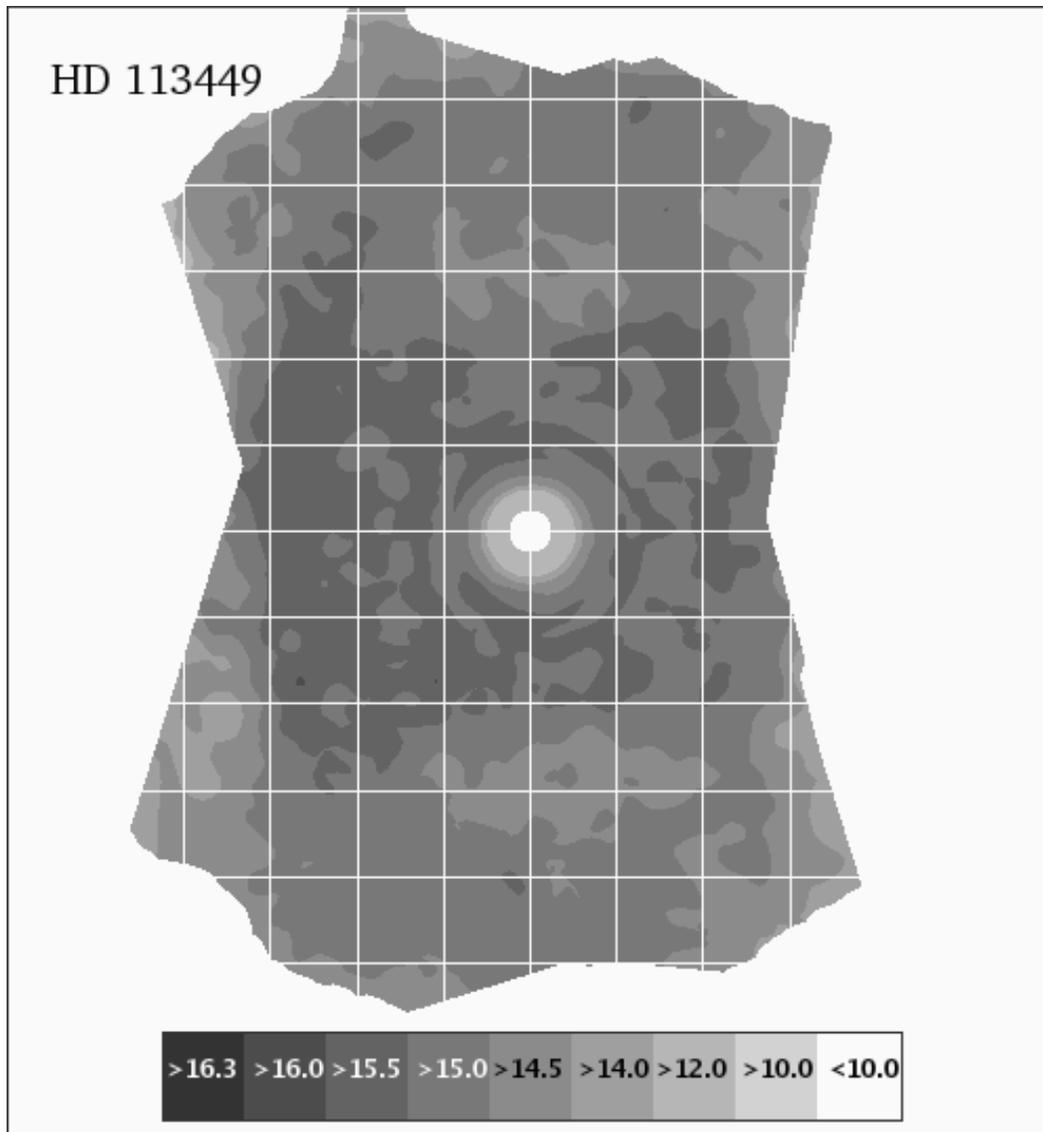


Figure 3.45 Sensitivity Contours for our  $L'$  images of HD 113449. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.0$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 15.5$  but not as faint as  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

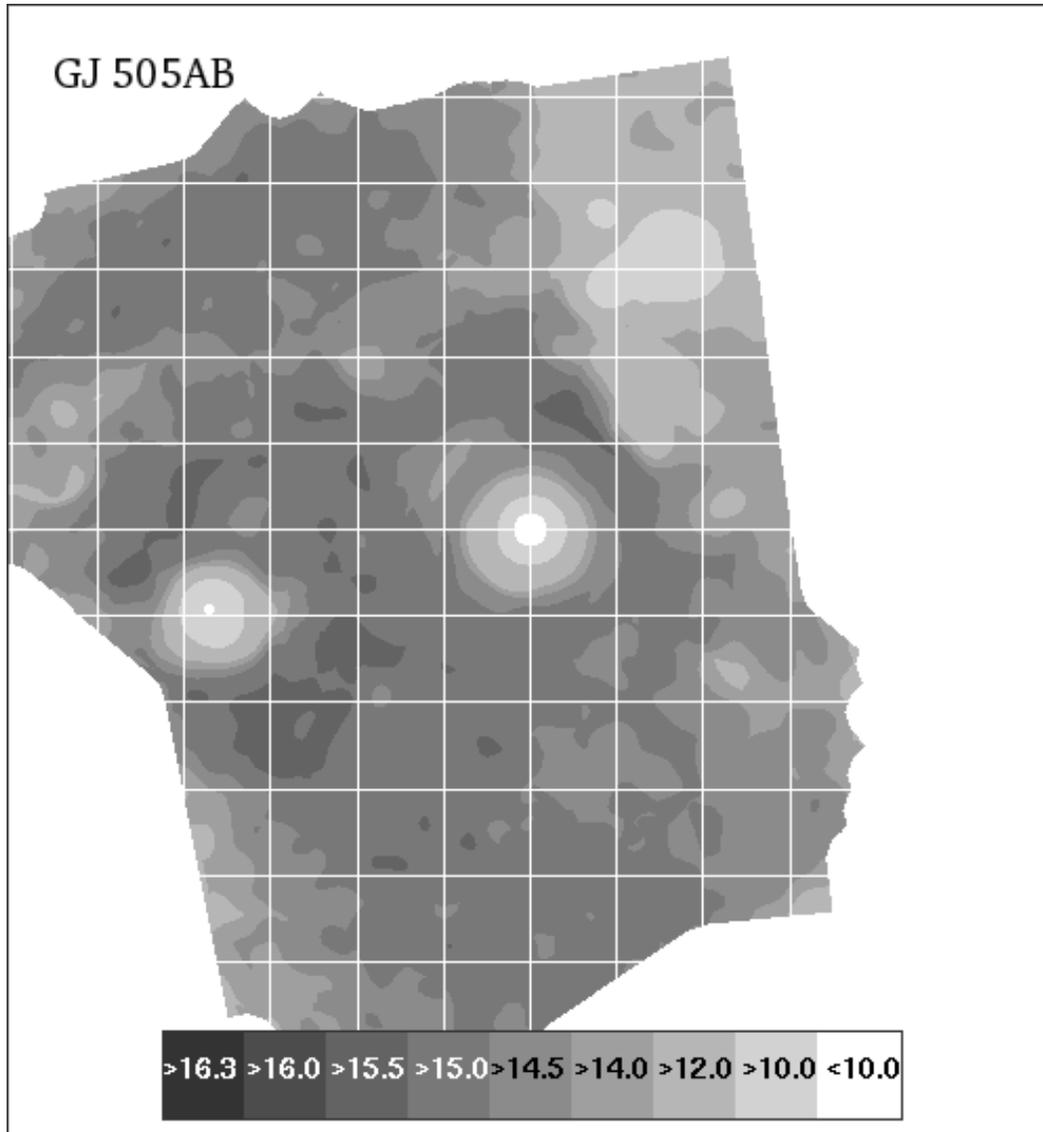


Figure 3.46 Sensitivity Contours for our  $L'$  images of GJ 505AB. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

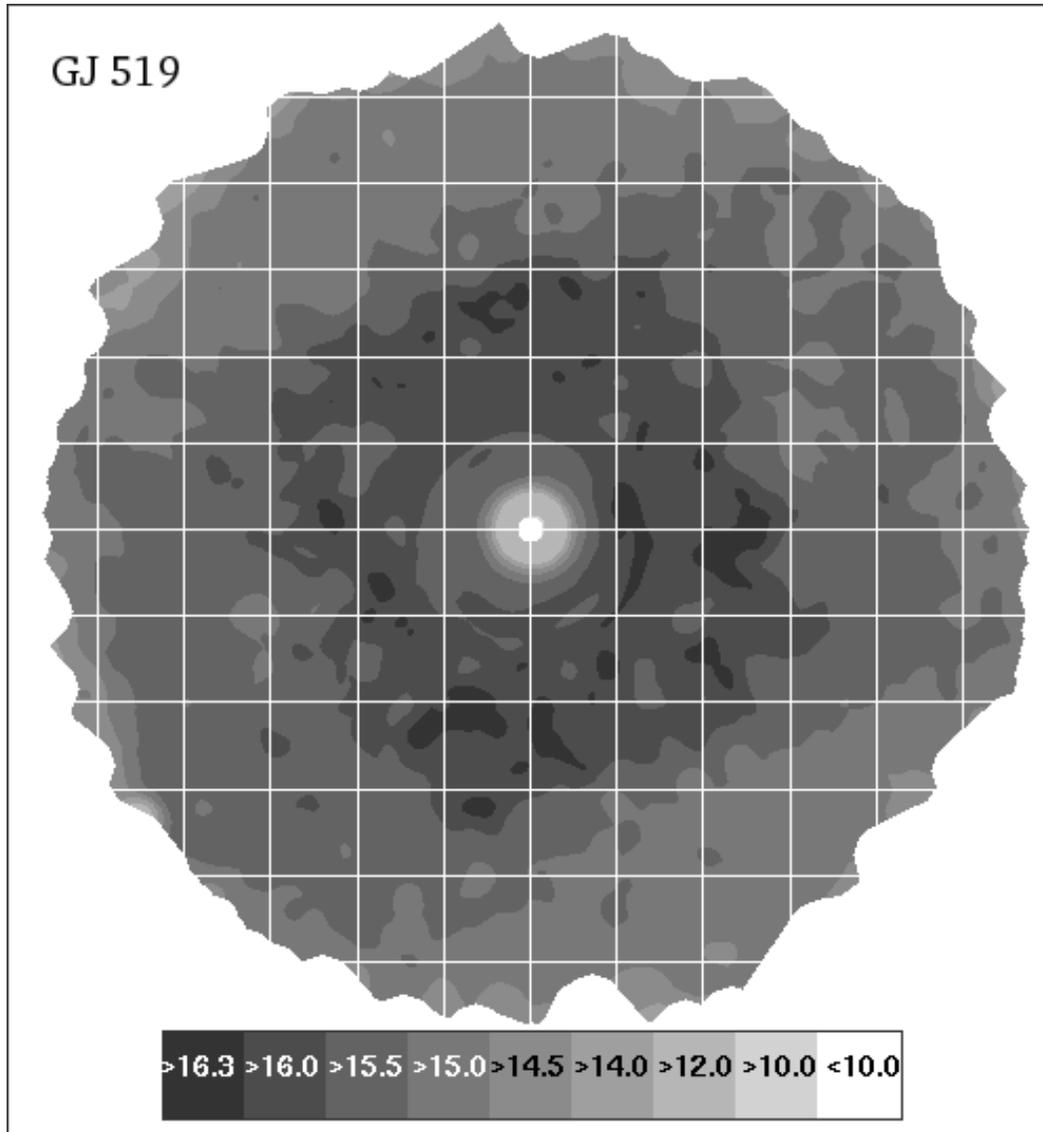


Figure 3.47 Sensitivity Contours for our  $L'$  images of GJ 519. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

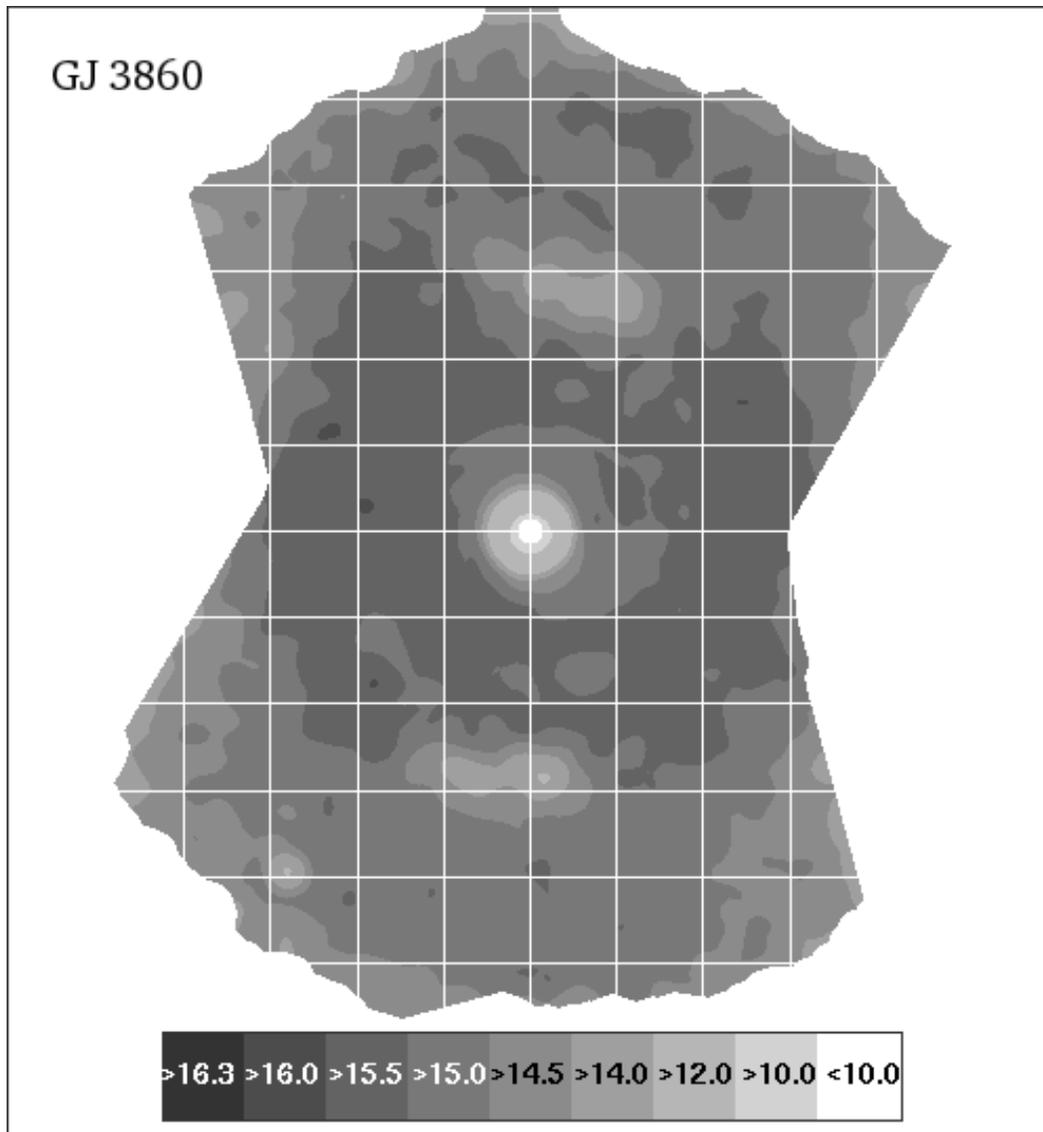


Figure 3.48 Sensitivity Contours for our  $L'$  images of GJ 3860. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.5$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.0$  but not as faint as  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

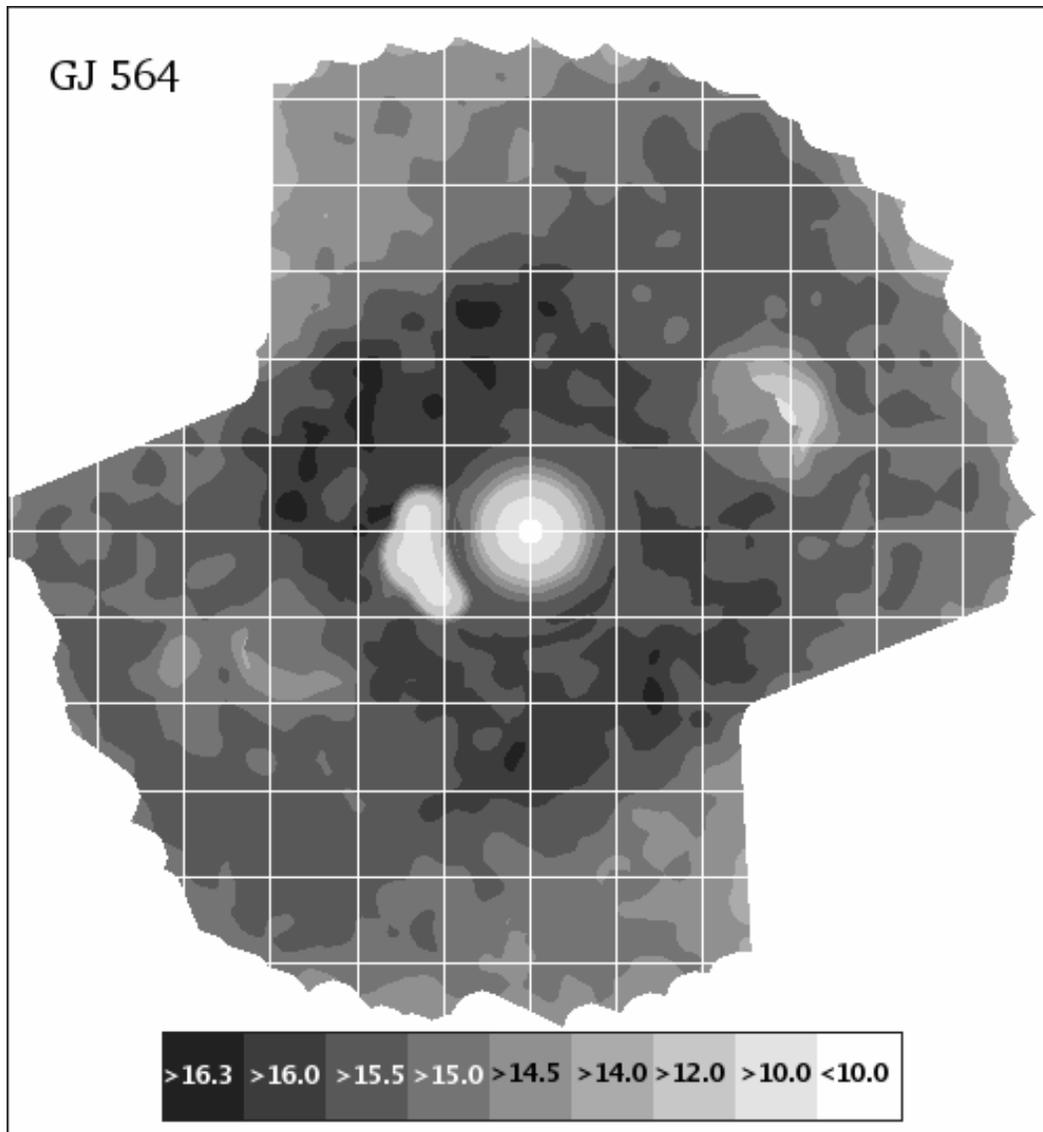


Figure 3.49 Sensitivity Contours for our  $L'$  images of GJ 564. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.3$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

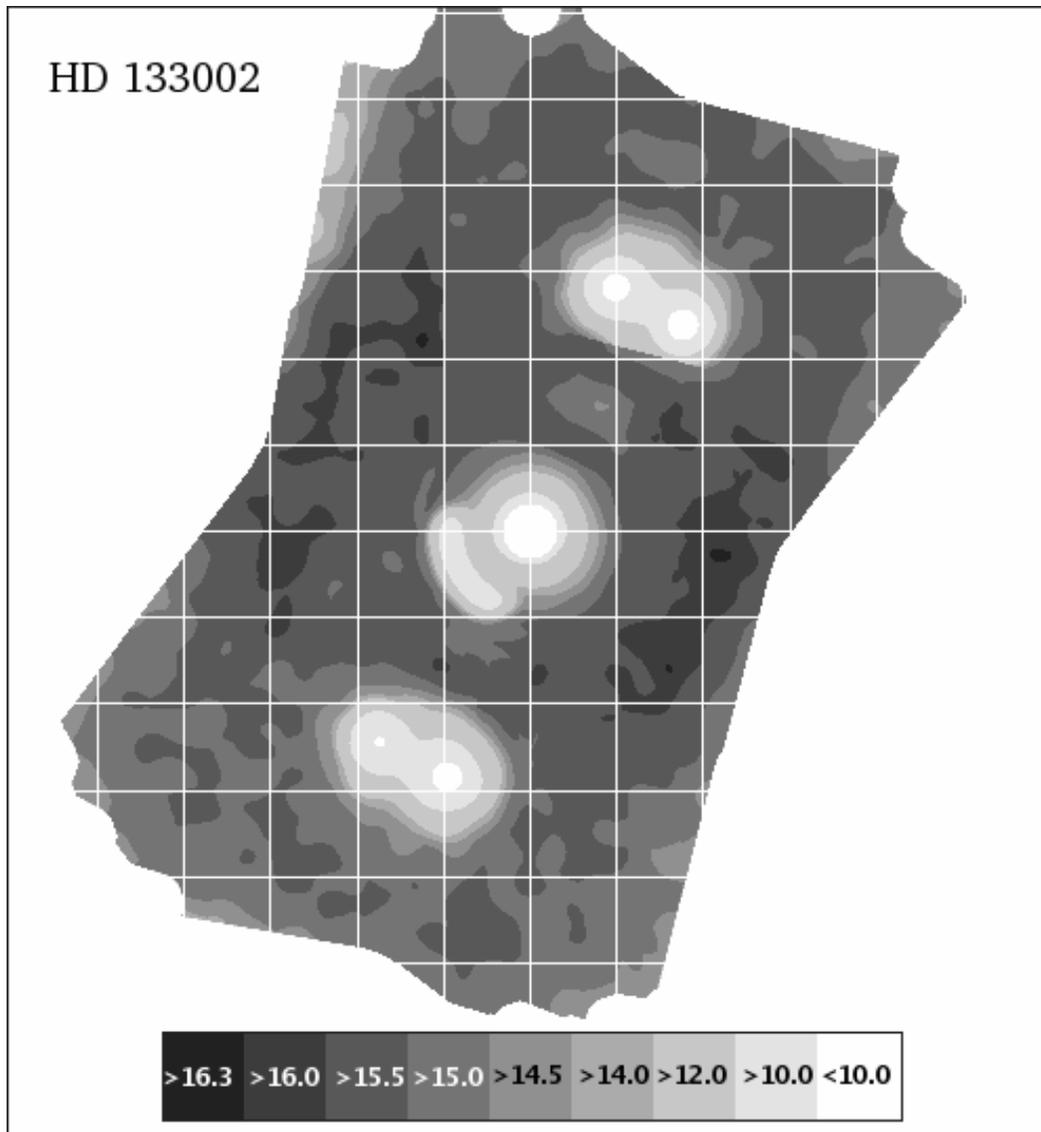


Figure 3.50 Sensitivity Contours for our  $L'$  images of HD 133002. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.3$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 16.0$  but not as faint as  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

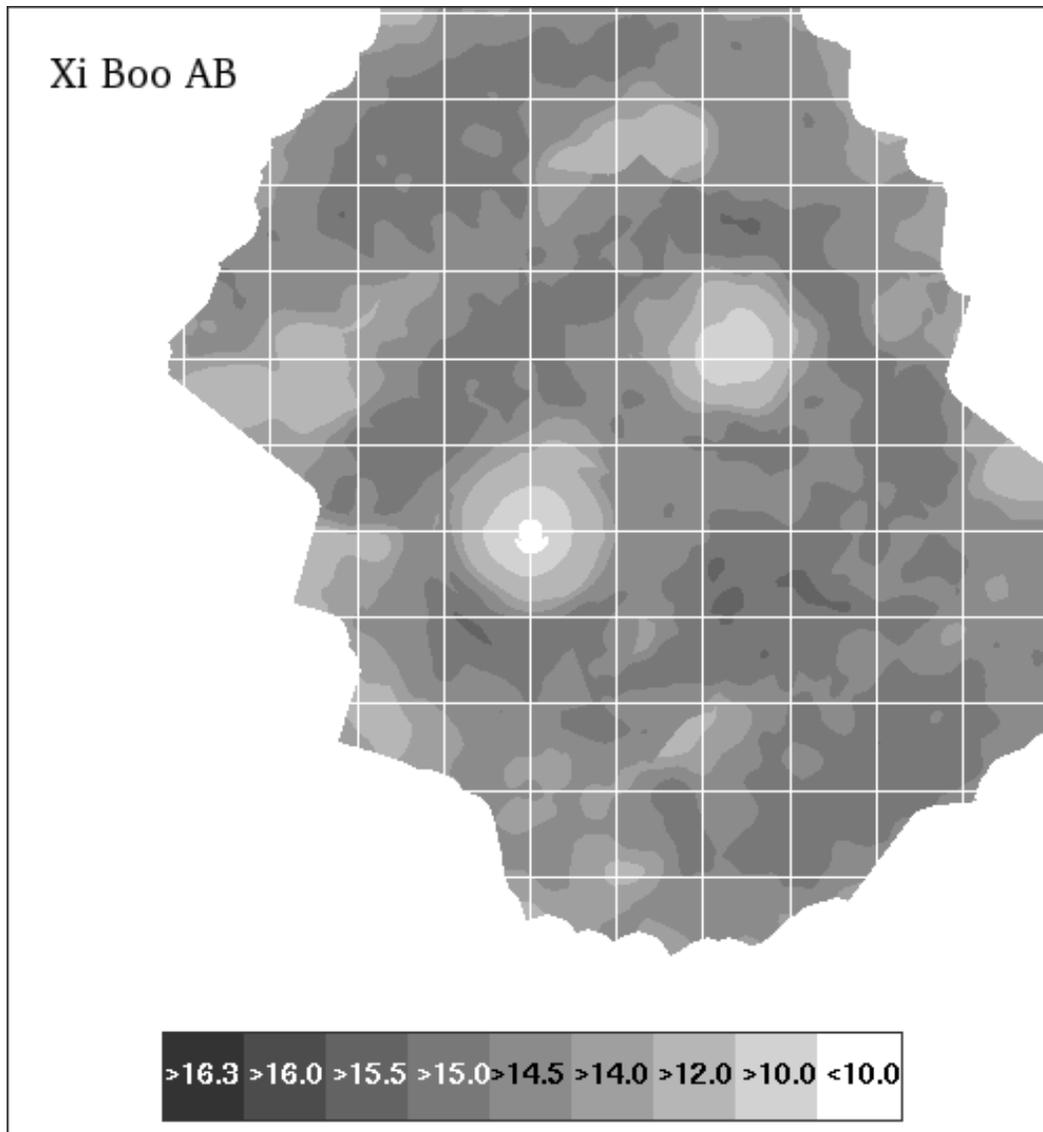


Figure 3.51 Sensitivity Contours for our  $L'$  images of  $\xi$  Boo AB. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.5$  in small areas; wider areas have sensitivities to objects fainter than  $L' = 15.0$  but not as faint as  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

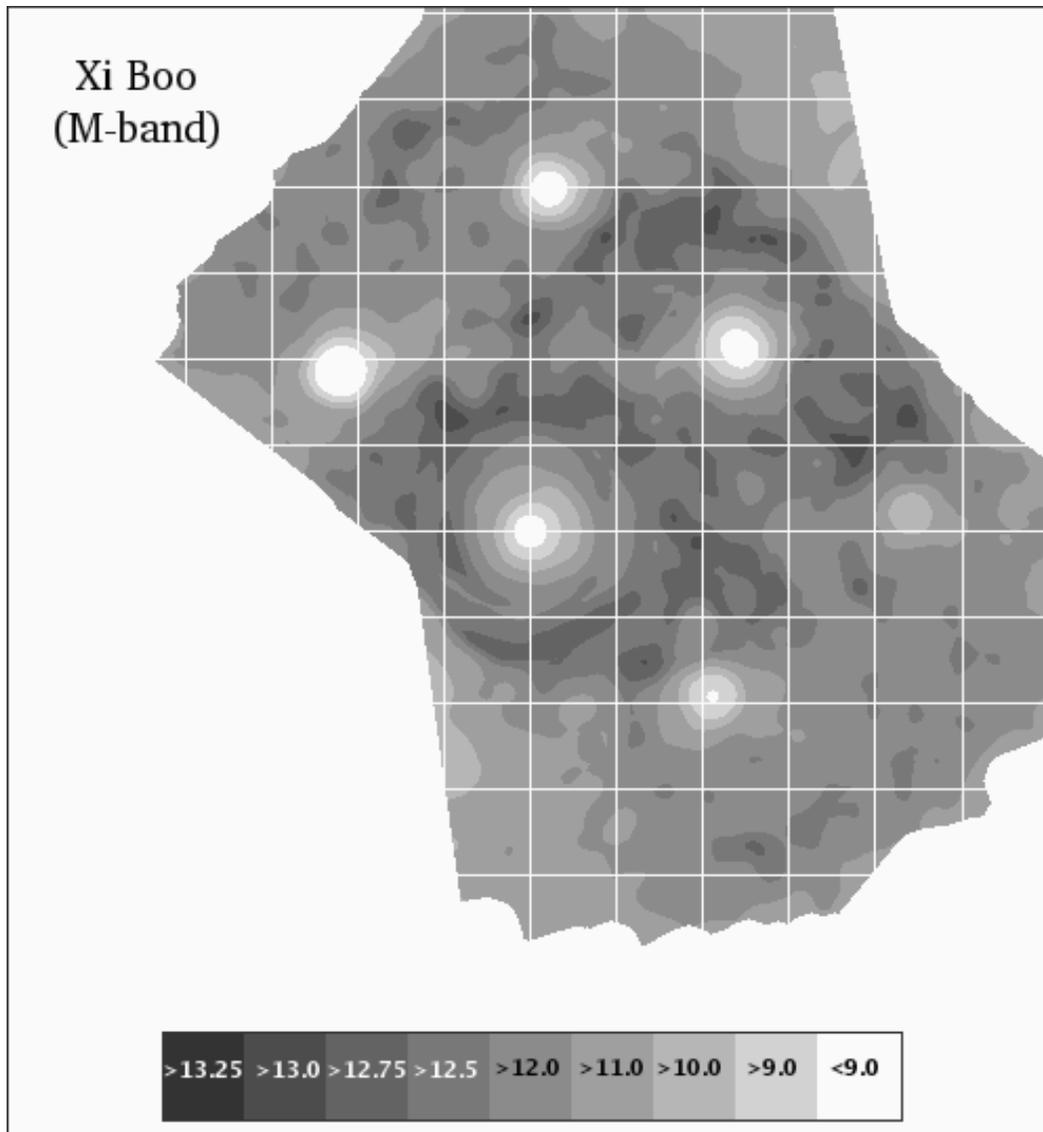


Figure 3.52 Sensitivity Contours for our  $M$  band images of  $\xi$  Boo AB. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $M = 13.25$  in small areas; wider areas have sensitivities to objects fainter than  $M = 13.0$  but not as faint as  $M = 13.25$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

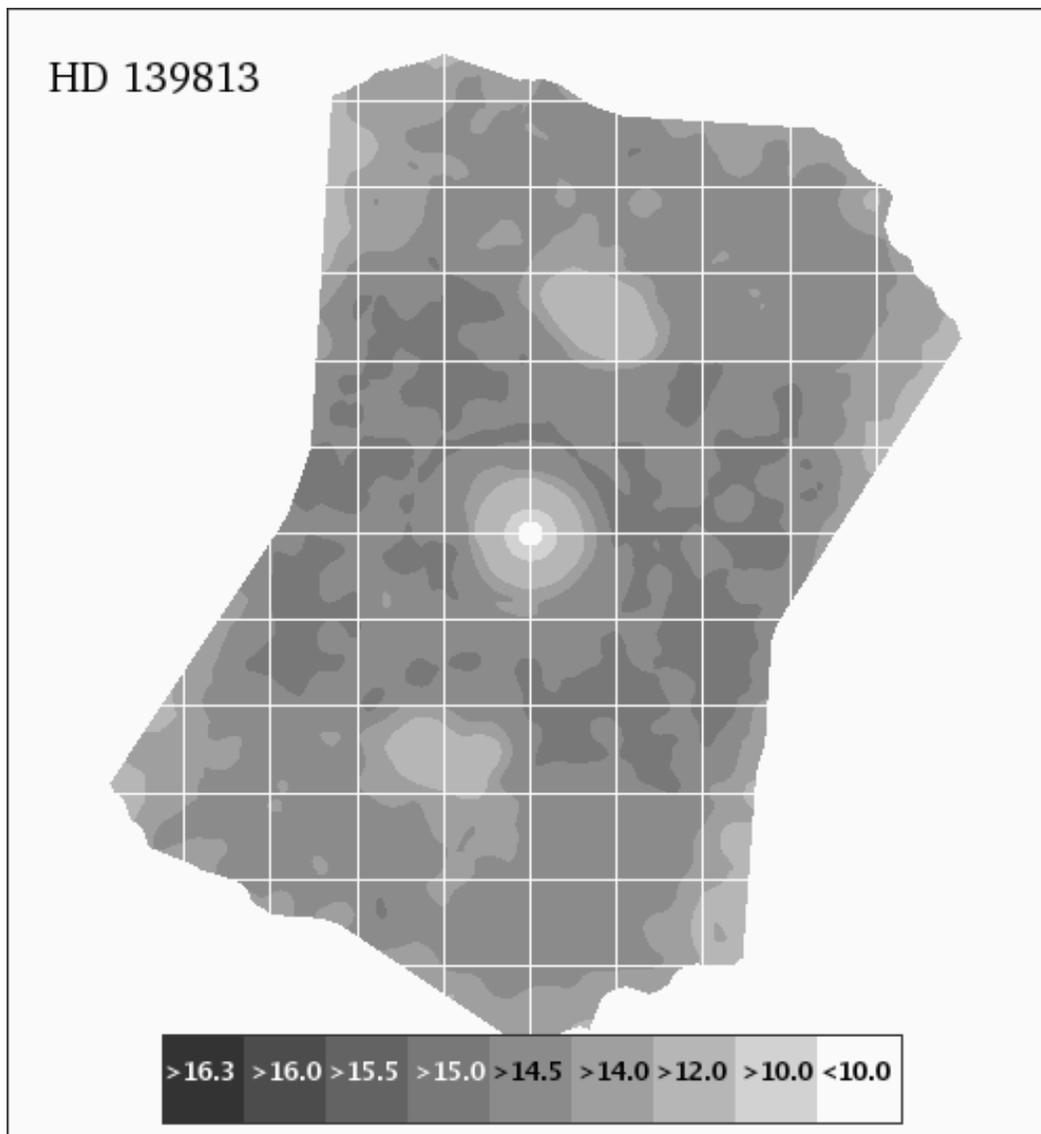


Figure 3.53 Sensitivity Contours for our  $L'$  images of HD 139813. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

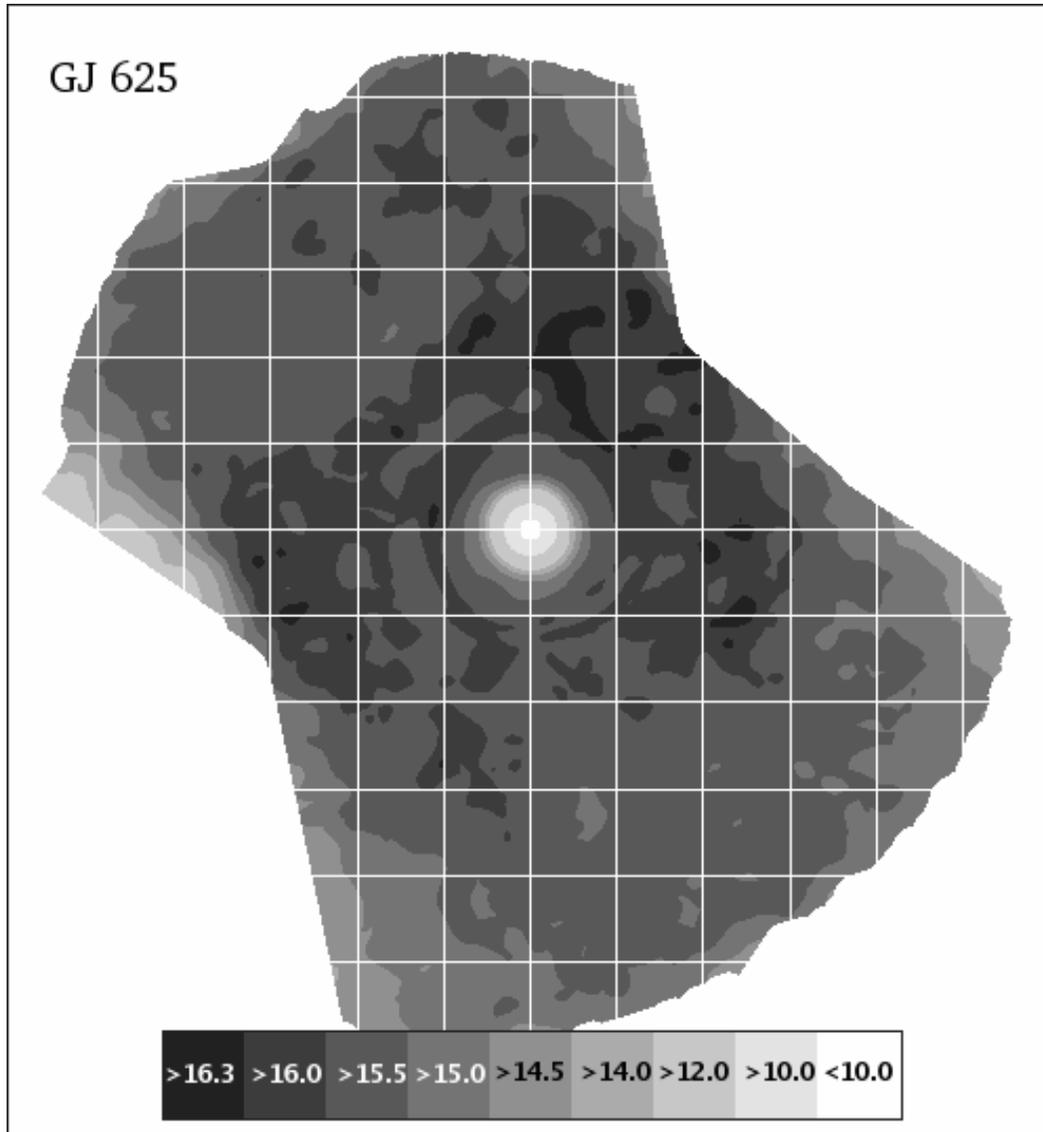


Figure 3.54 Sensitivity Contours for our  $L'$  images of GJ 625. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 16.3$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.



Figure 3.55 Sensitivity Contours for our  $L'$  images of GJ 659A. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 15.0$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 14.5$  but not as faint as  $L' = 15.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

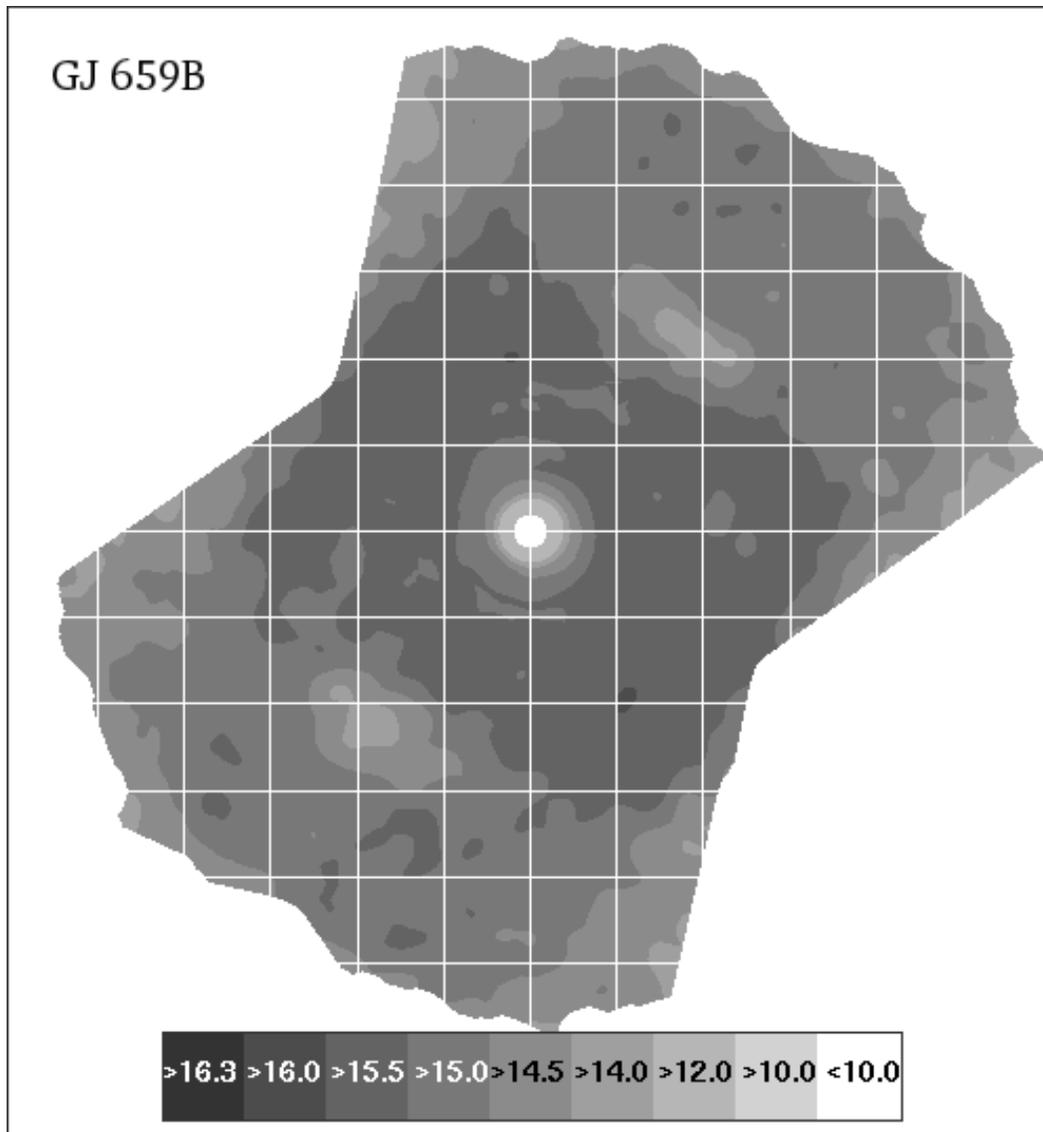


Figure 3.56 Sensitivity Contours for our  $L'$  images of GJ 659B. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.0$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 15.5$  but not as faint as  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

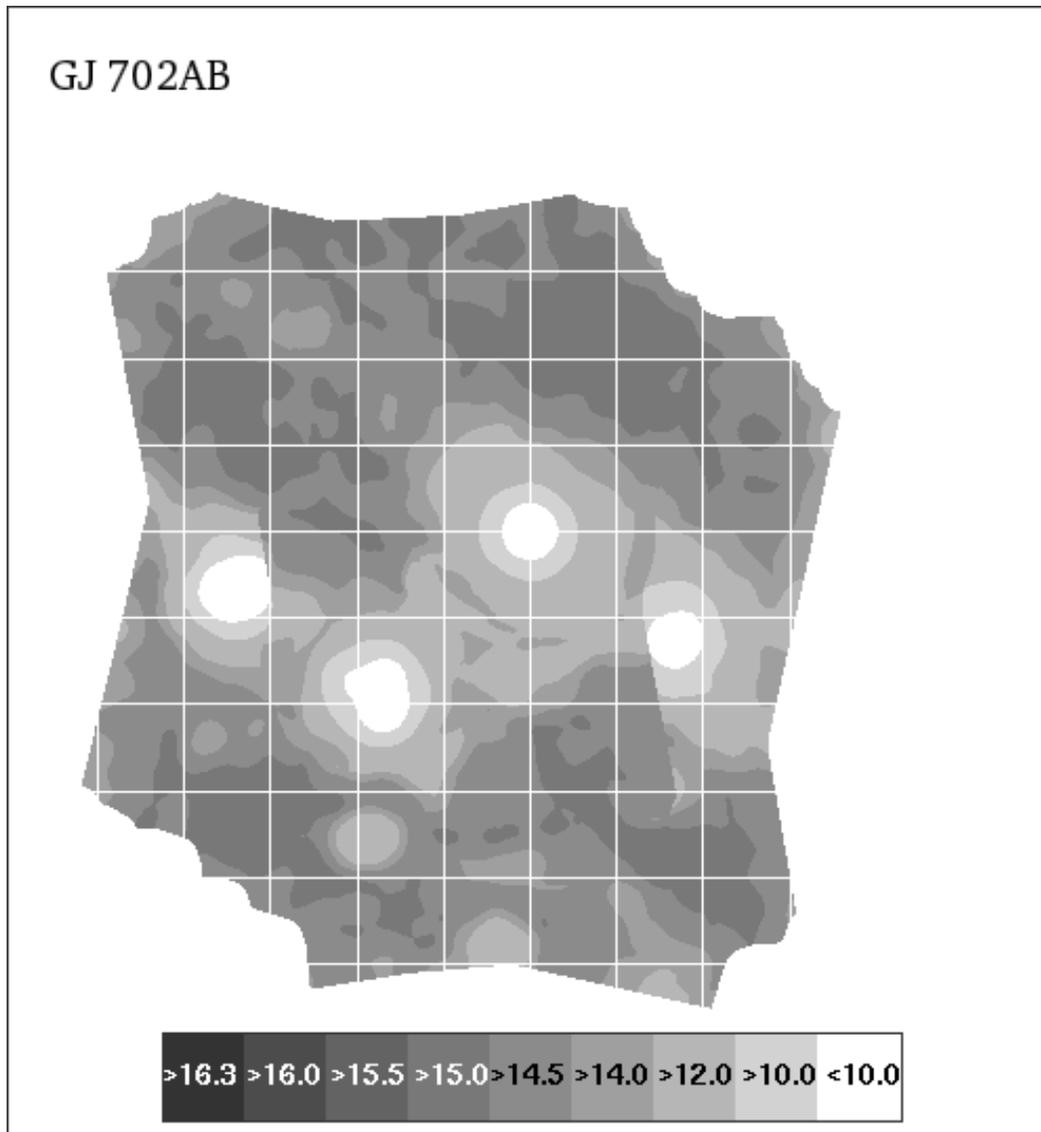


Figure 3.57 Sensitivity Contours for our  $L'$  images of GJ 702AB. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

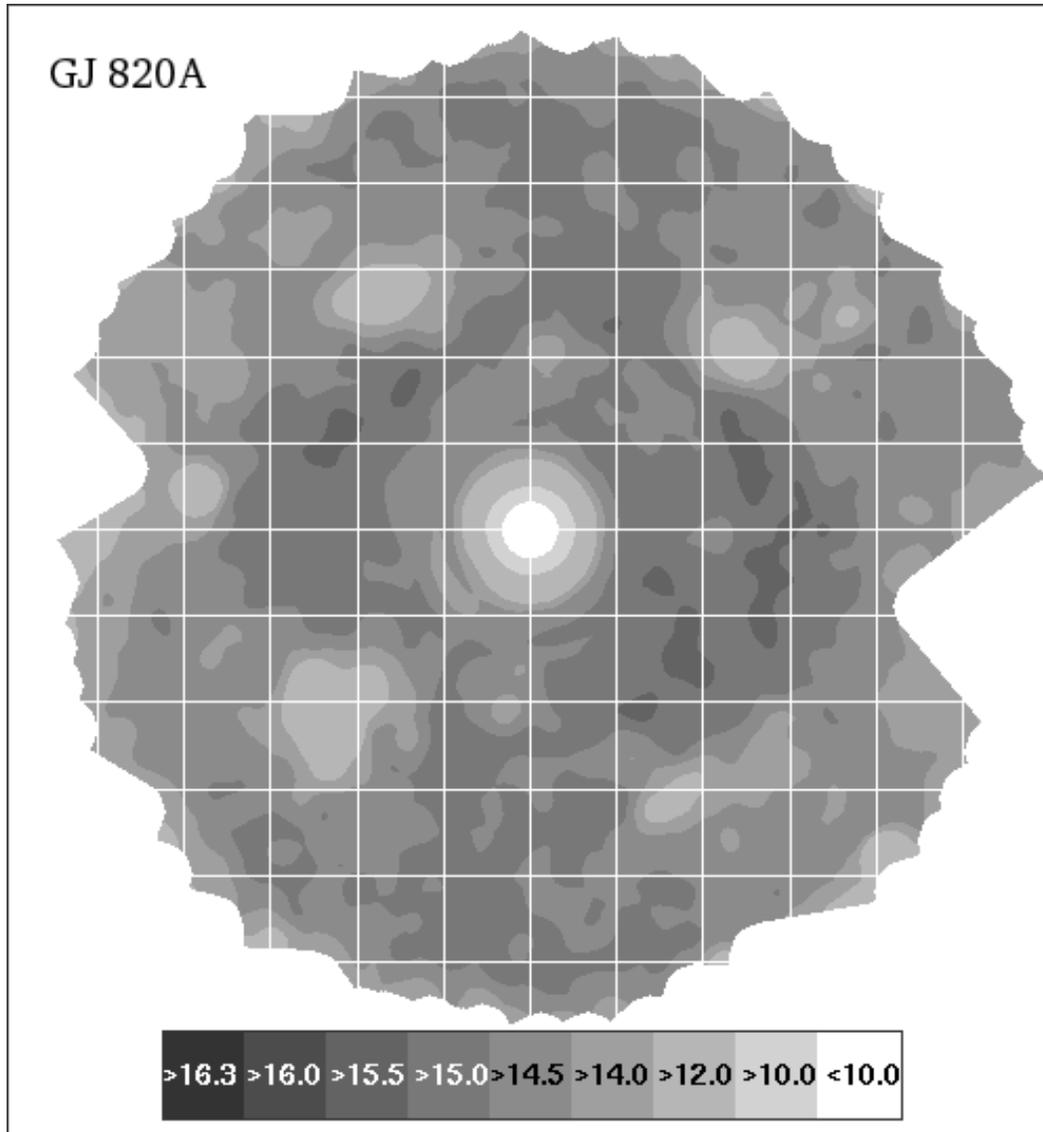


Figure 3.58 Sensitivity Contours for our  $L'$  images of 61 Cyg A. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

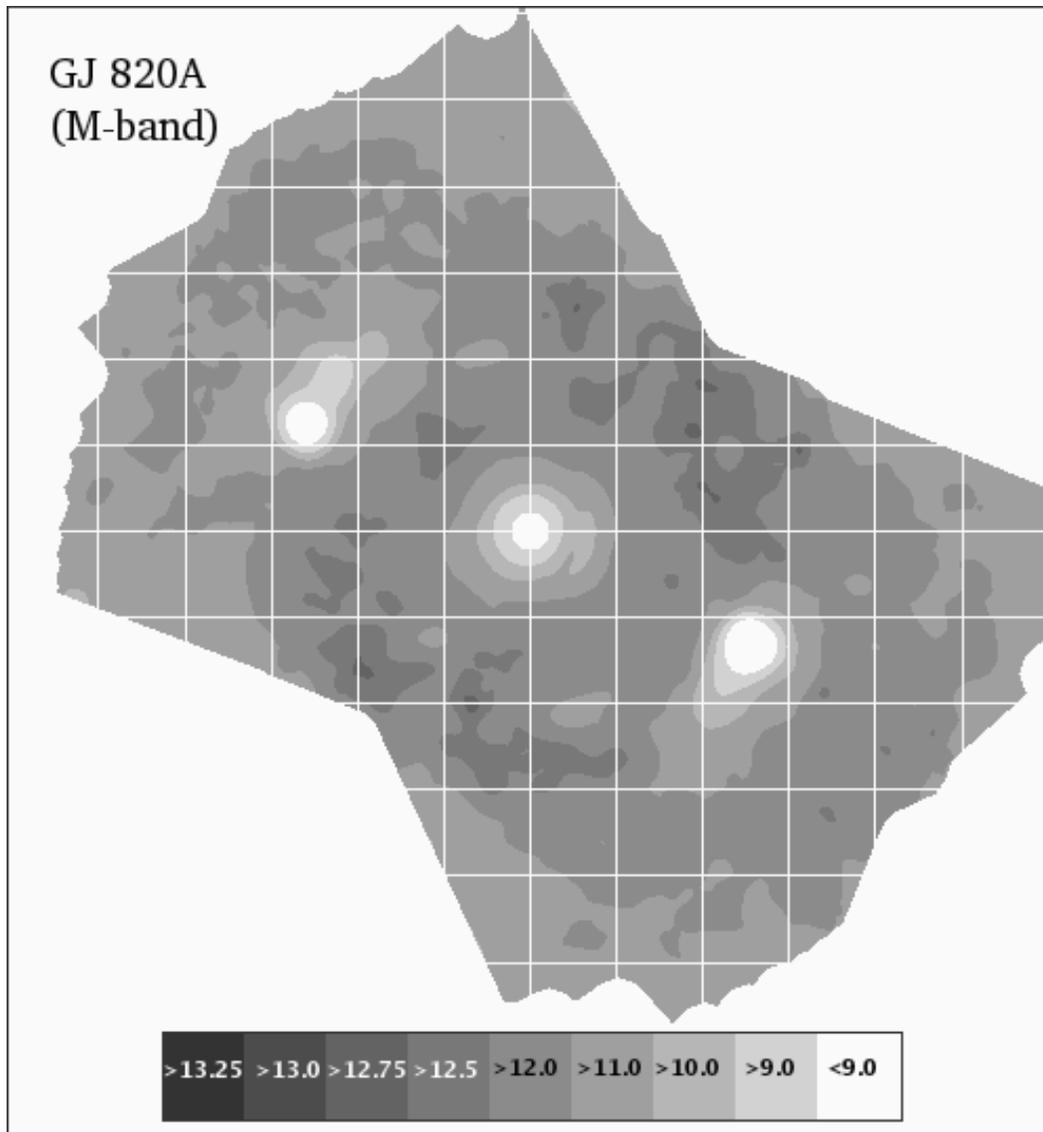


Figure 3.59 Sensitivity Contours for our  $M$  band images of 61 Cyg A. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $m = 12.75$  in a tiny area; wider areas have sensitivities to objects fainter than  $M = 12.5$  but not as faint as  $M = 12.75$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

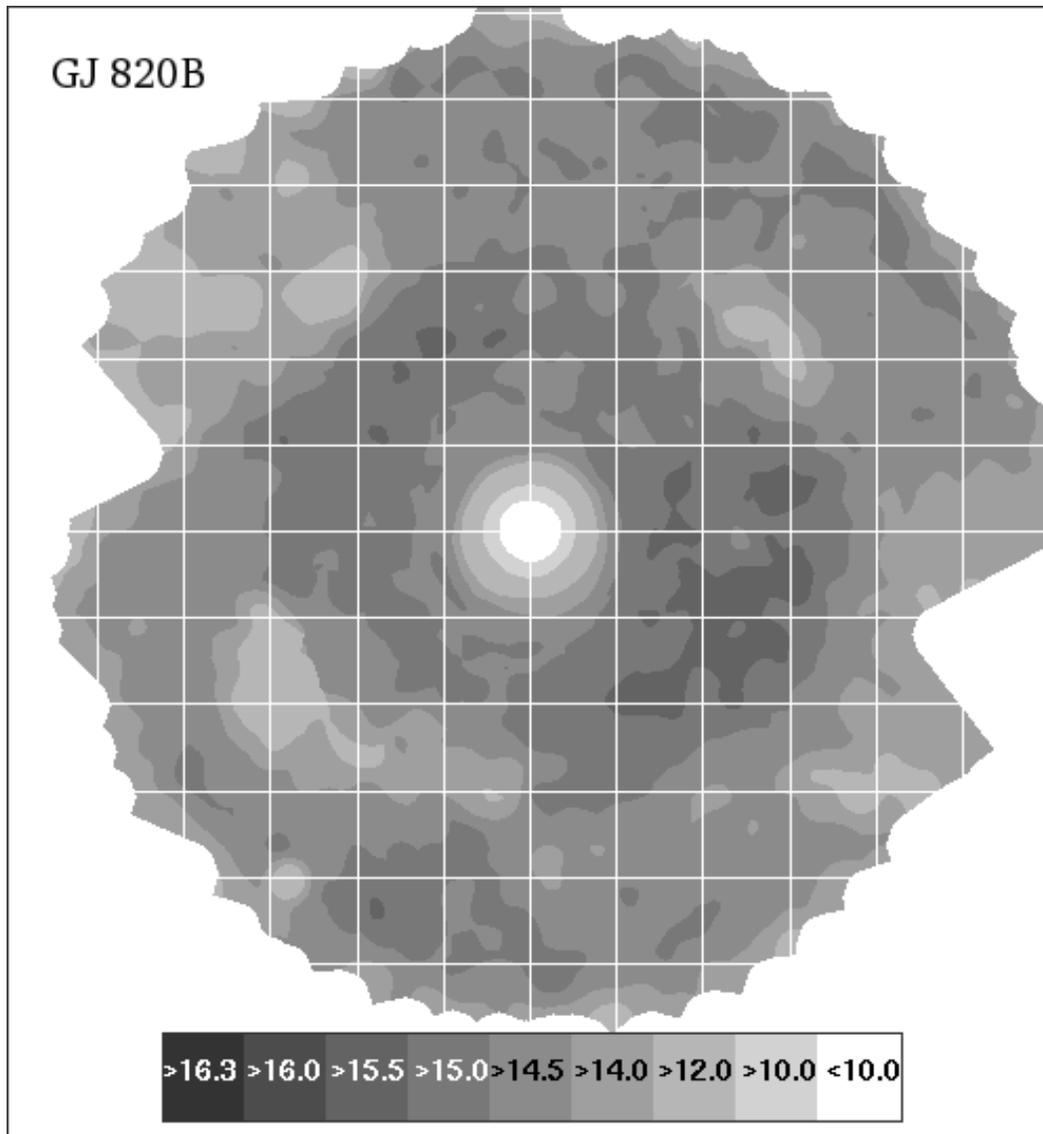


Figure 3.60 Sensitivity Contours for our  $L'$  images of 61 Cyg B. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 15.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

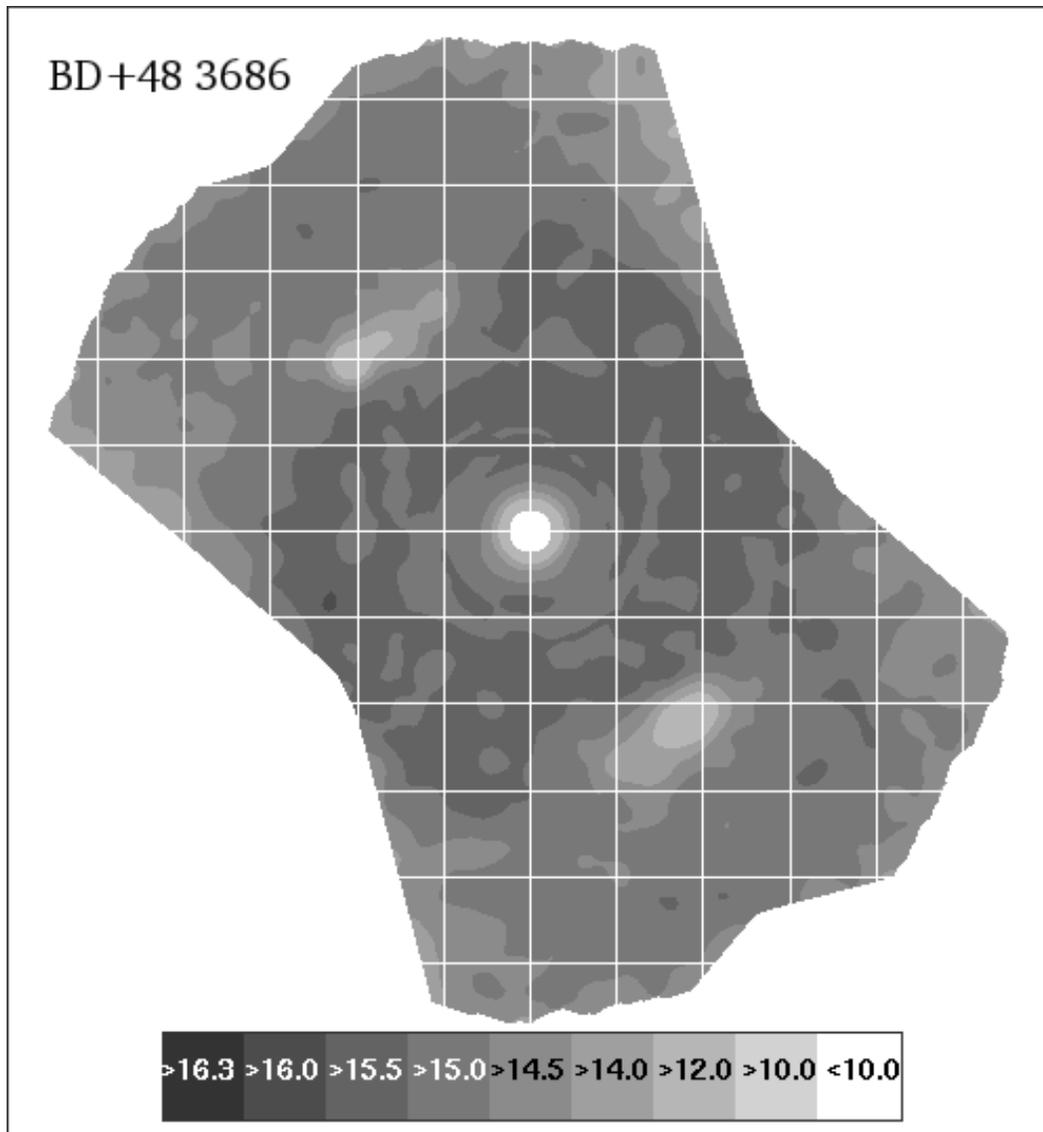


Figure 3.61 Sensitivity Contours for our  $L'$  images of BD+48 3686. The colorbar gives the sensitivities for each contour in magnitudes. This image attains a peak sensitivity better than  $L' = 16.0$  in a tiny area; wider areas have sensitivities to objects fainter than  $L' = 15.5$  but not as faint as  $L' = 16.0$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

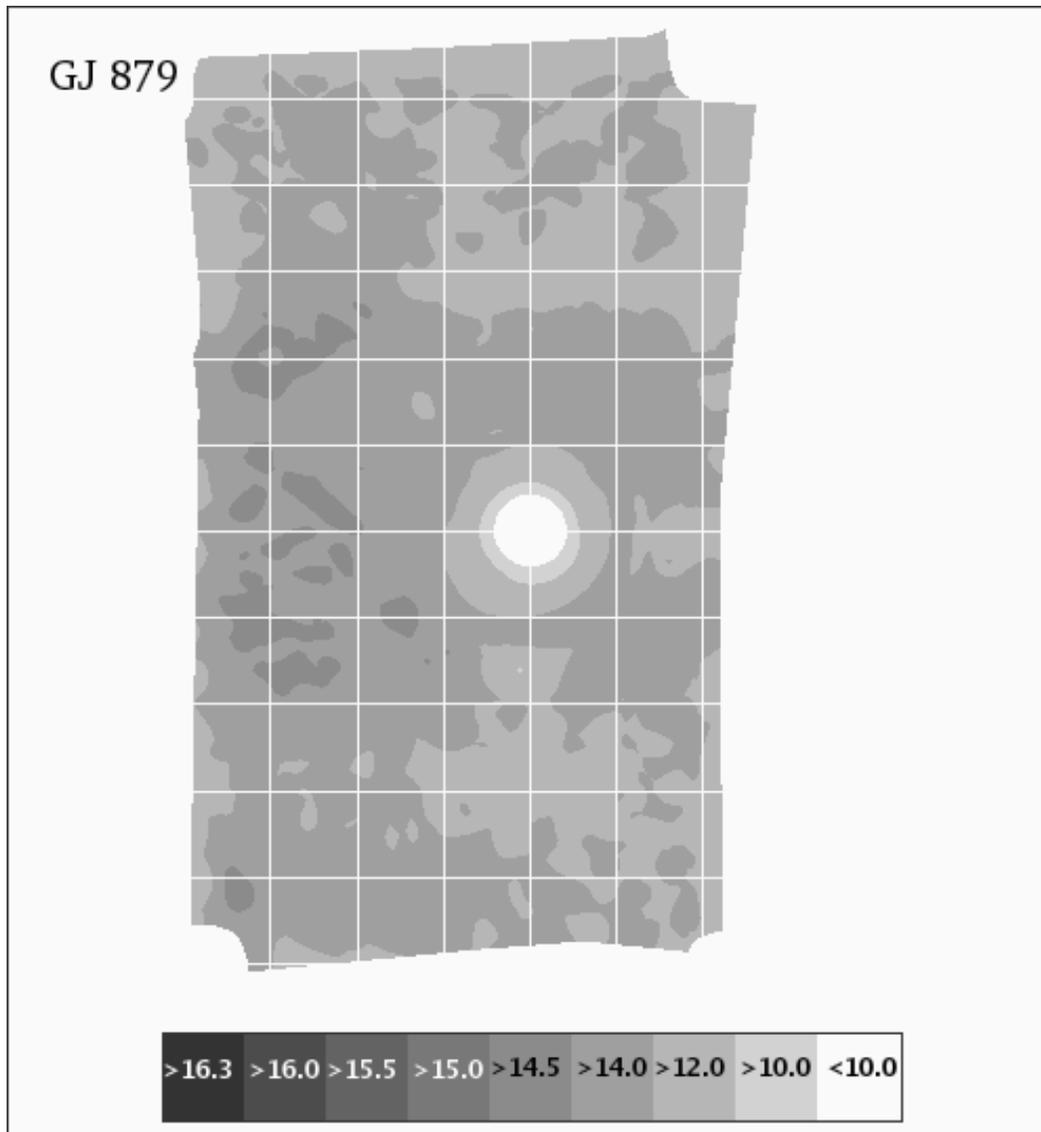


Figure 3.62 Sensitivity Contours for our  $L'$  images of GJ 879. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' = 14.5$ . These are nominal  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

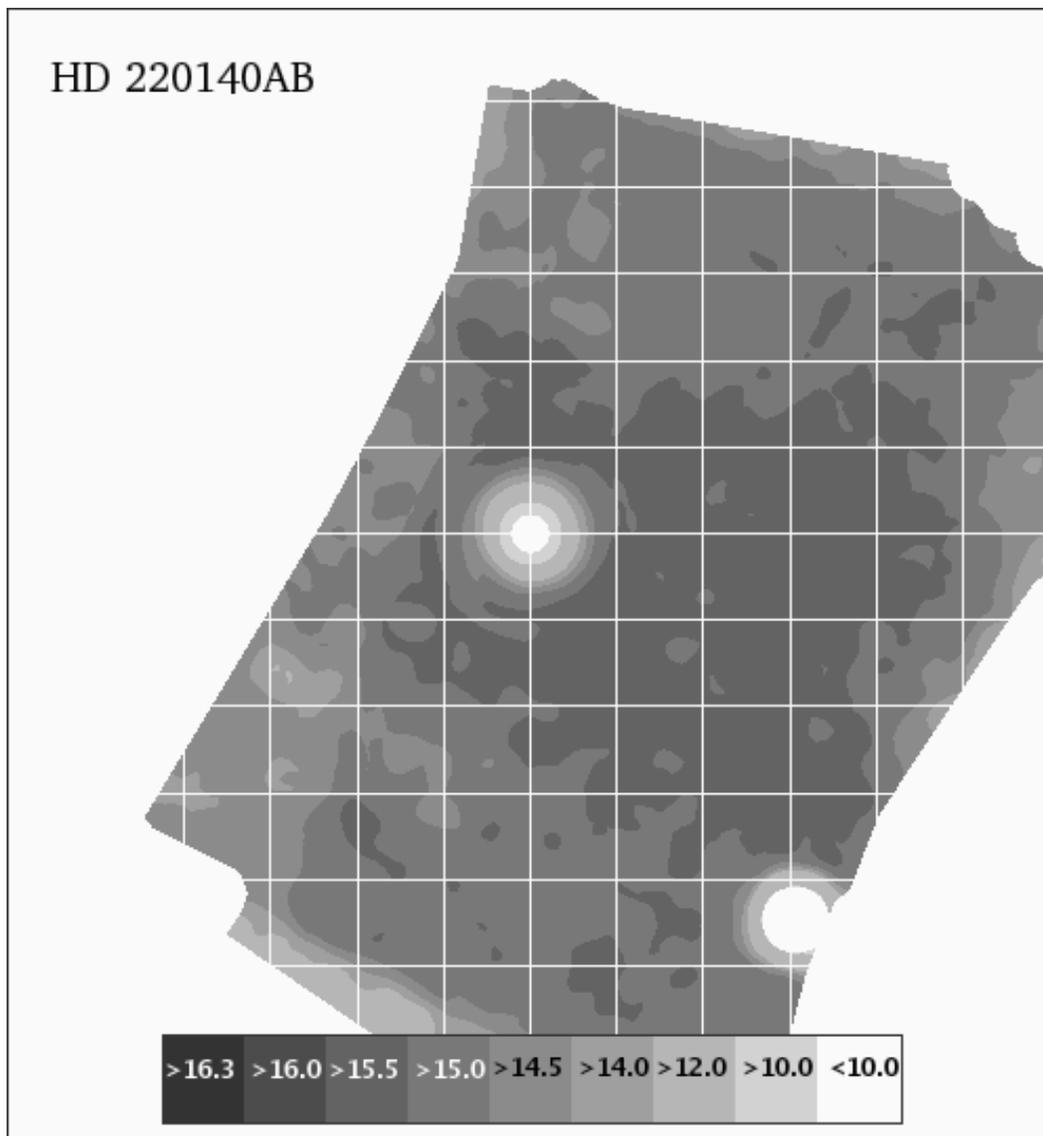


Figure 3.63 Sensitivity Contours for our  $L'$  images of HD 220140AB. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' =$  These are  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

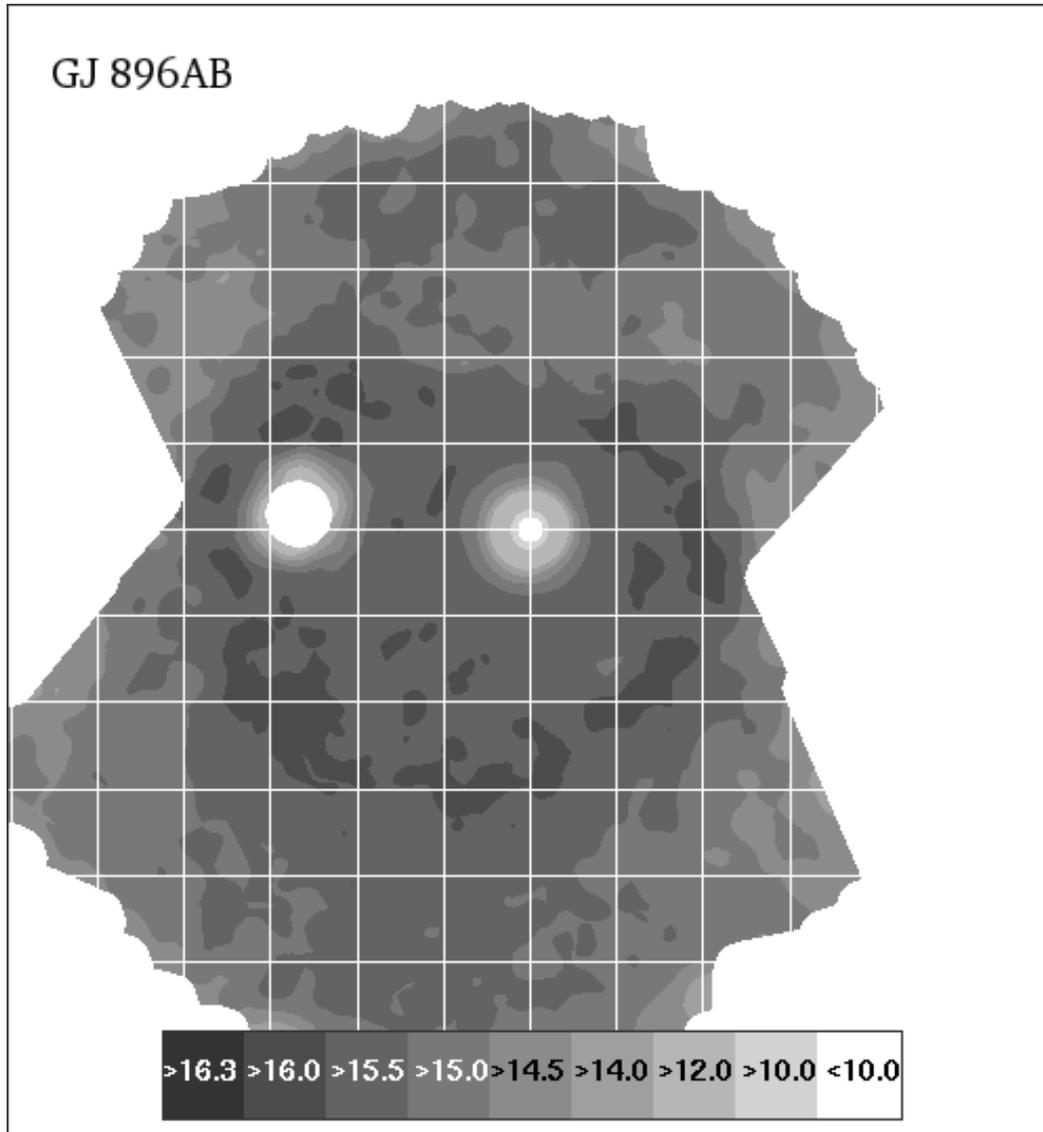


Figure 3.64 Sensitivity Contours for our  $L'$  images of GJ 896AB. The colorbar gives the sensitivities for each contour in magnitudes. The darkest regions in this image have sensitivity to objects fainter than  $L' =$  These are  $10\sigma$  sensitivities. North is up and east left; the grid squares are  $2 \times 2$  arcsec.

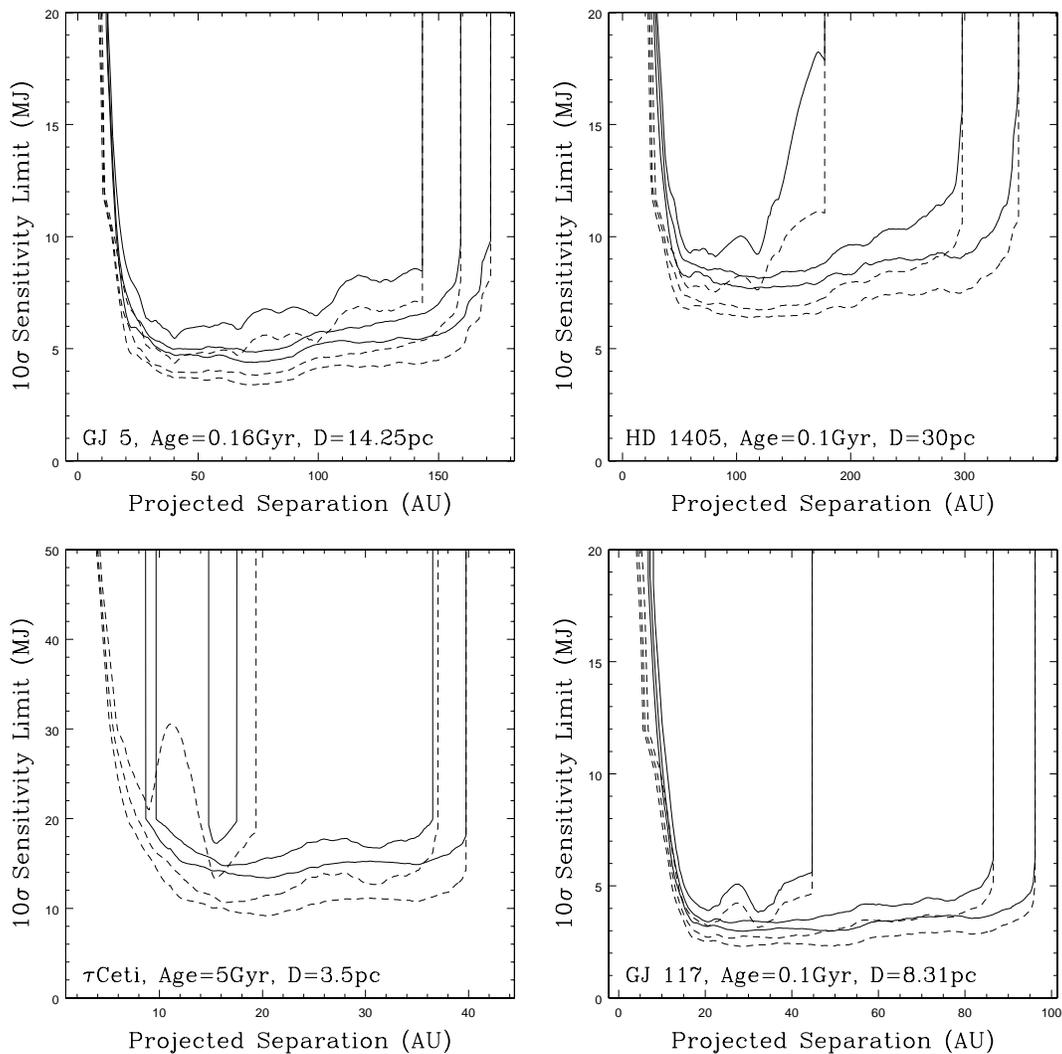


Figure 3.65 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003).

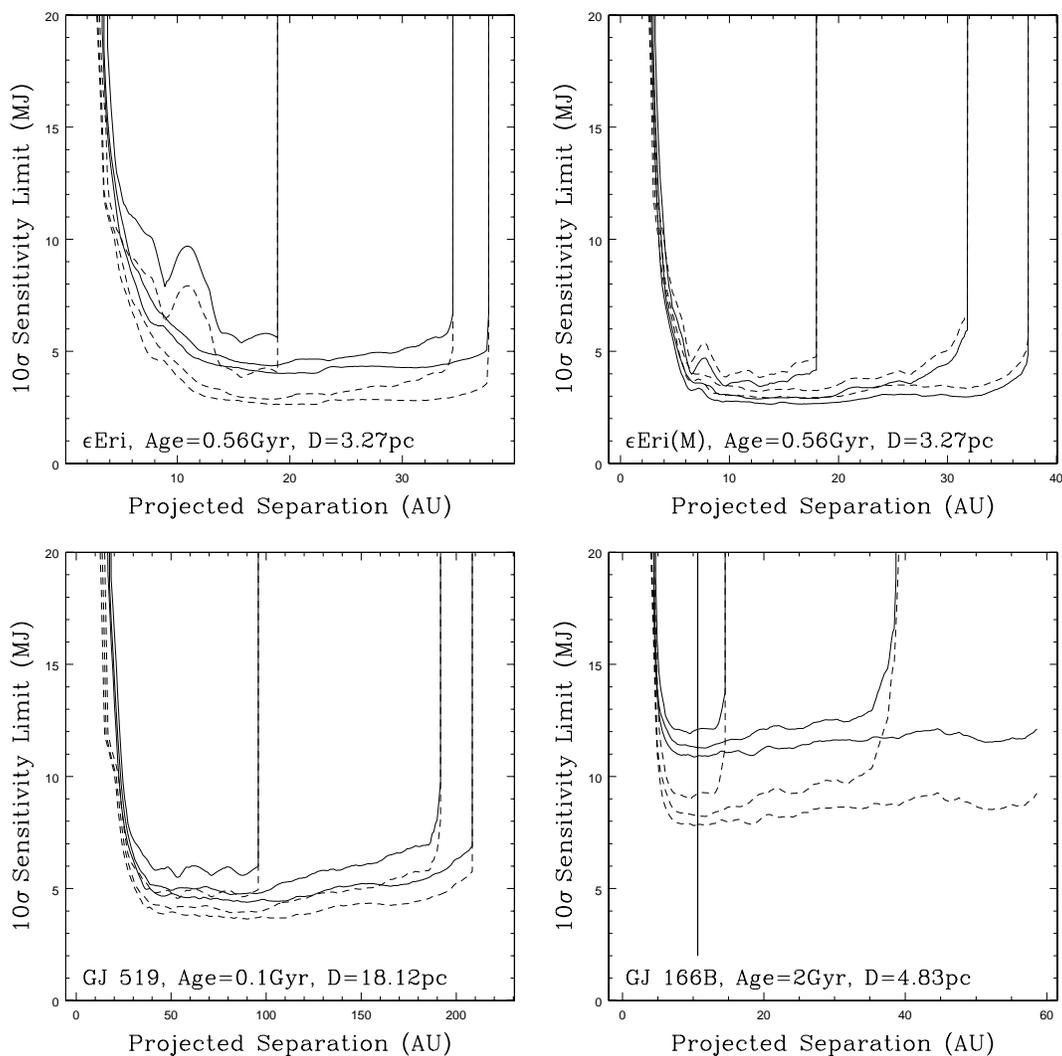


Figure 3.66 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits around GJ 166B are unstable at larger radii than the heavy vertical line, due to the companion star GJ 166C.

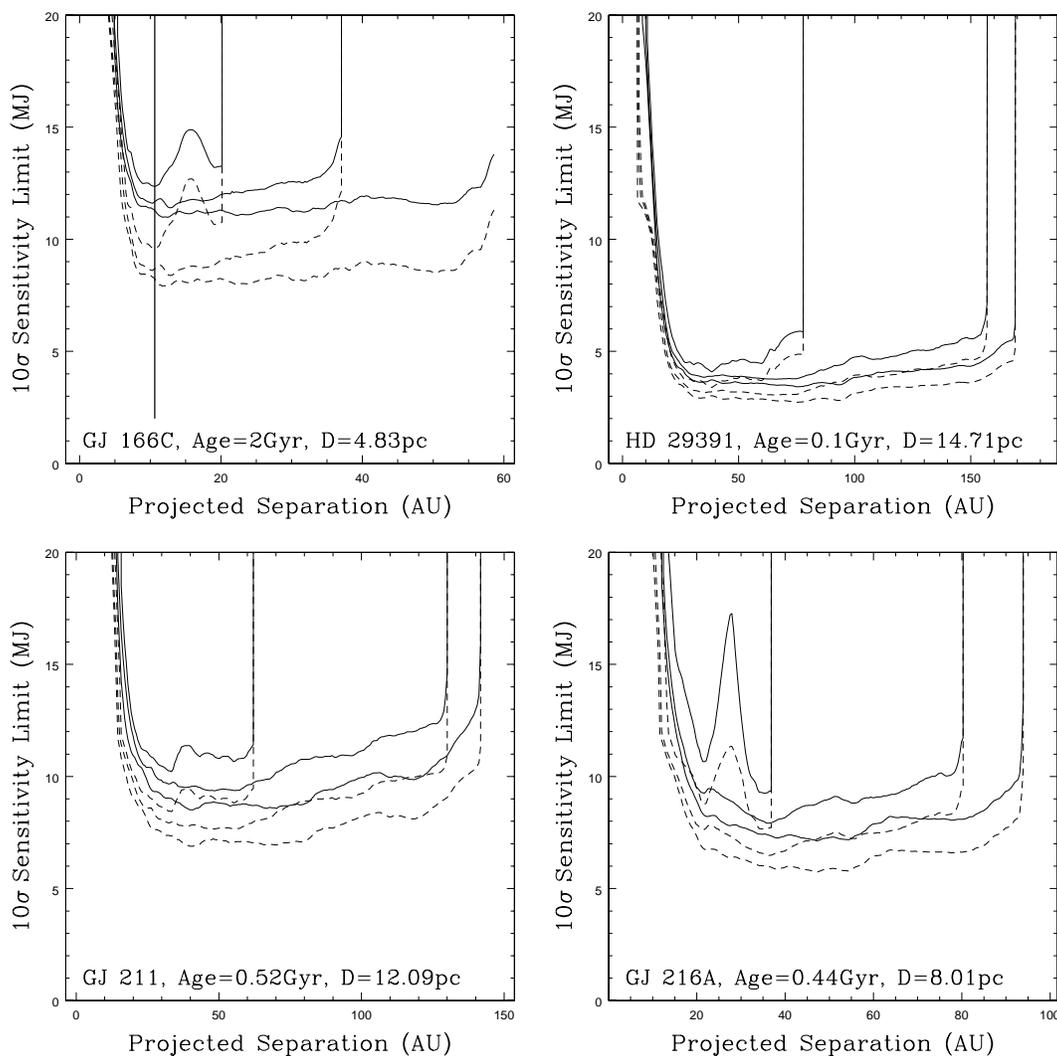


Figure 3.67 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits around GJ 166C are unstable at larger radii than the heavy vertical line, due to the companion star GJ 166B

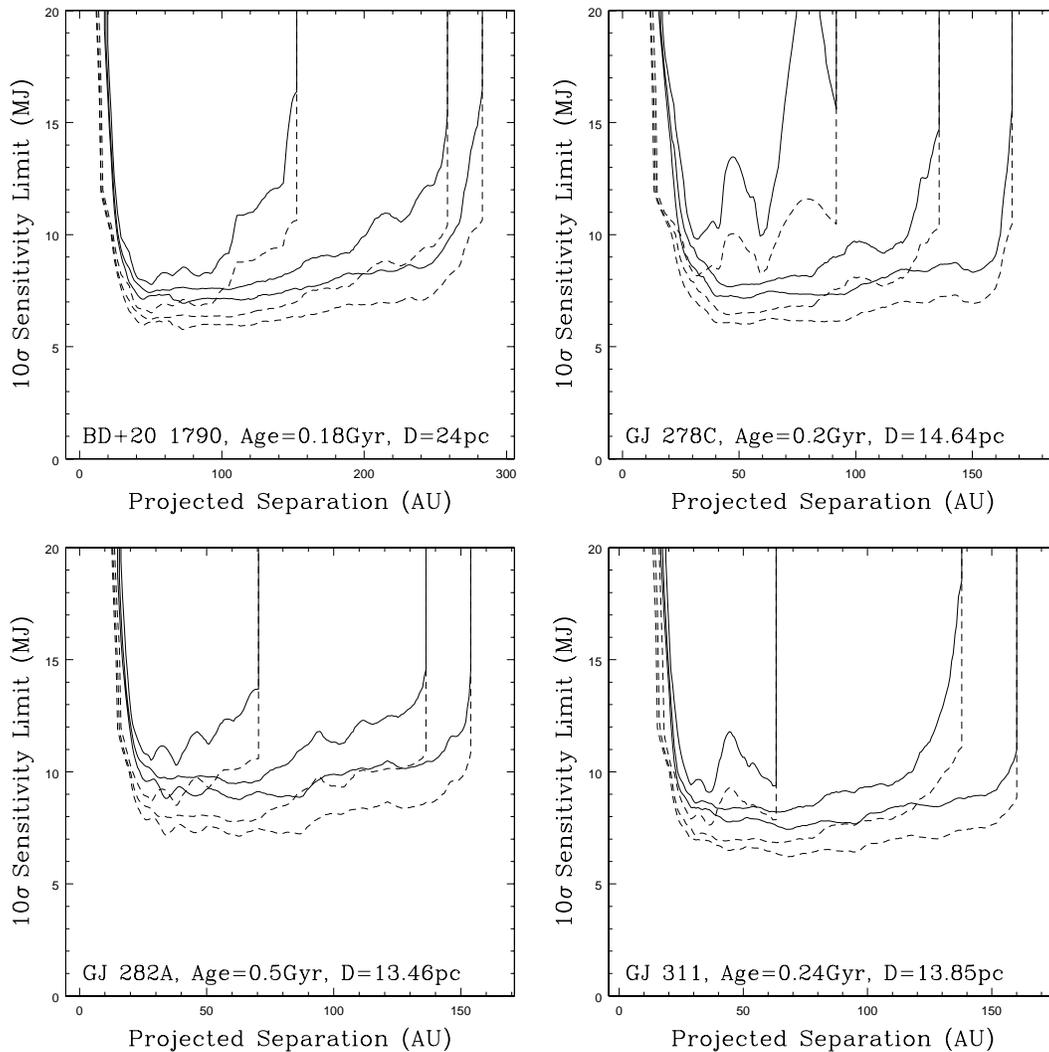


Figure 3.68 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003).

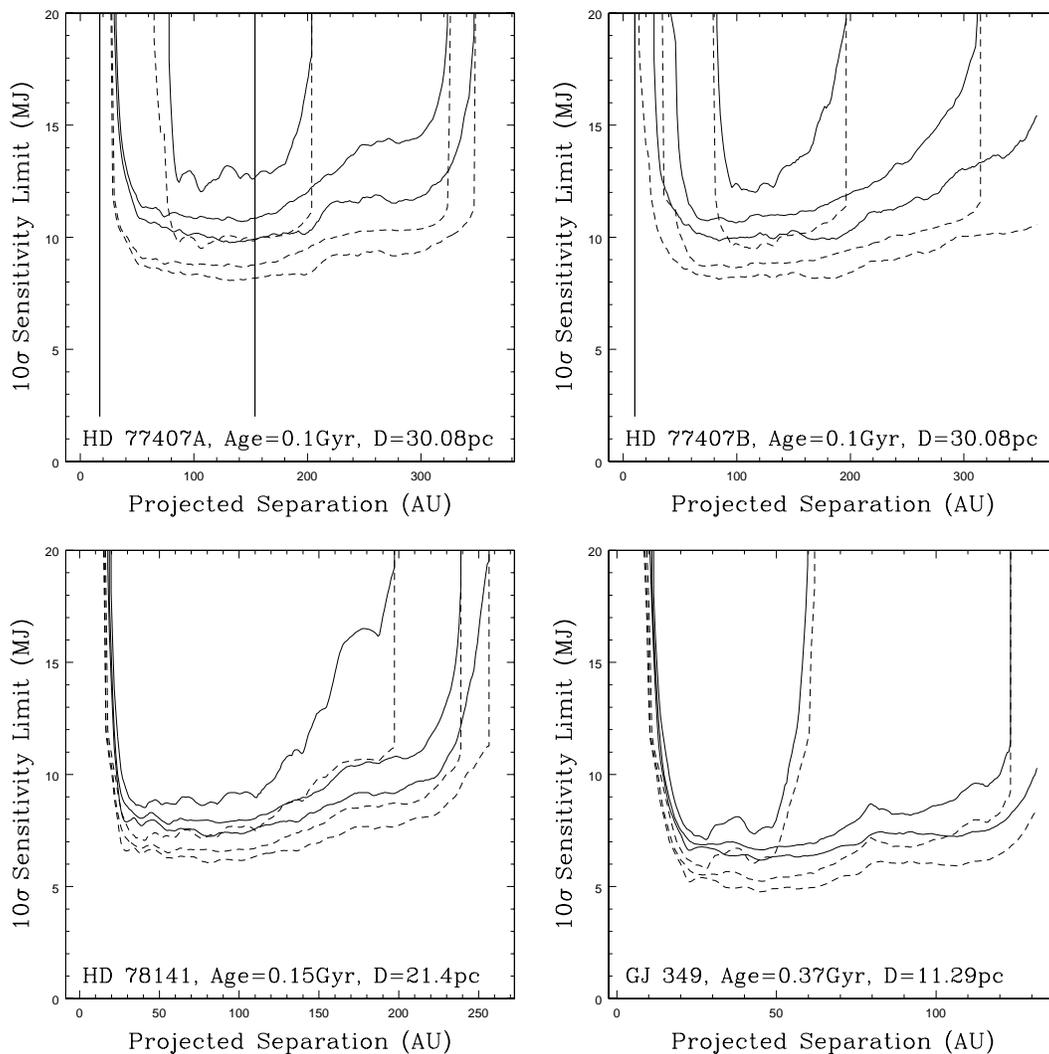


Figure 3.69 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits with radii between the heavy vertical lines on the HD 77407A plot, or outside the line on the HD 77407B plot, are unstable in this binary system.

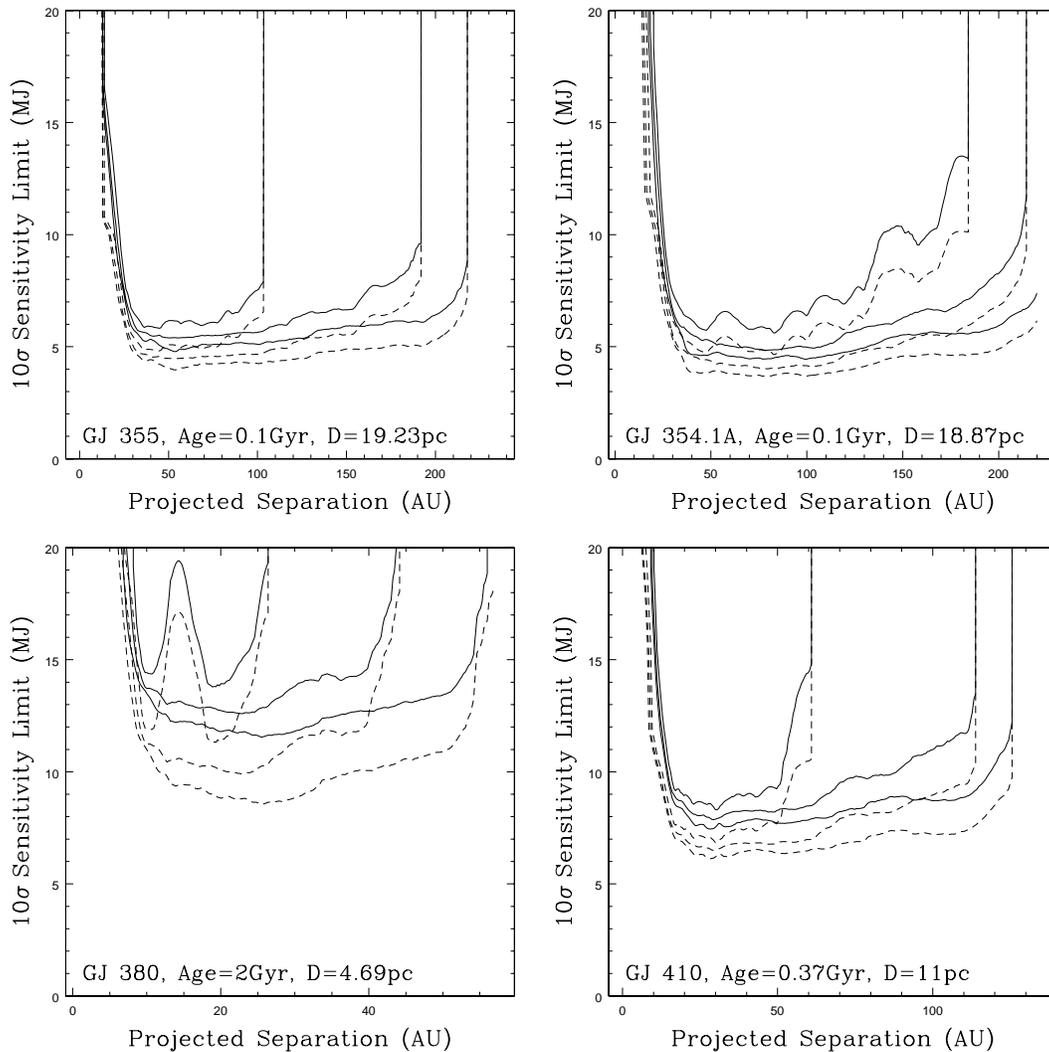


Figure 3.70 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003).

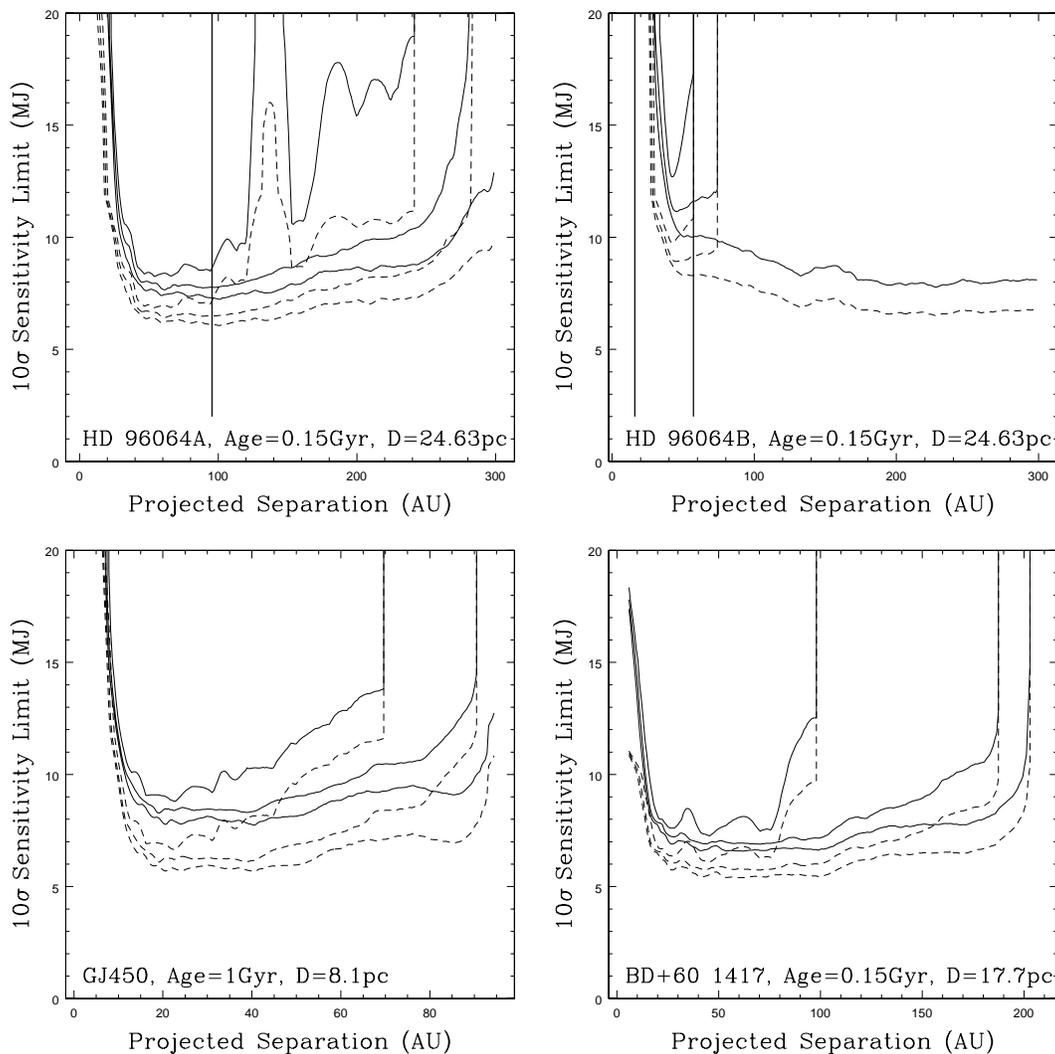


Figure 3.71 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits are stable in the HD 96064 triple system only inward of the heavy vertical line on the HD 96064A plot, or between the lines on the HD 96064B plot.

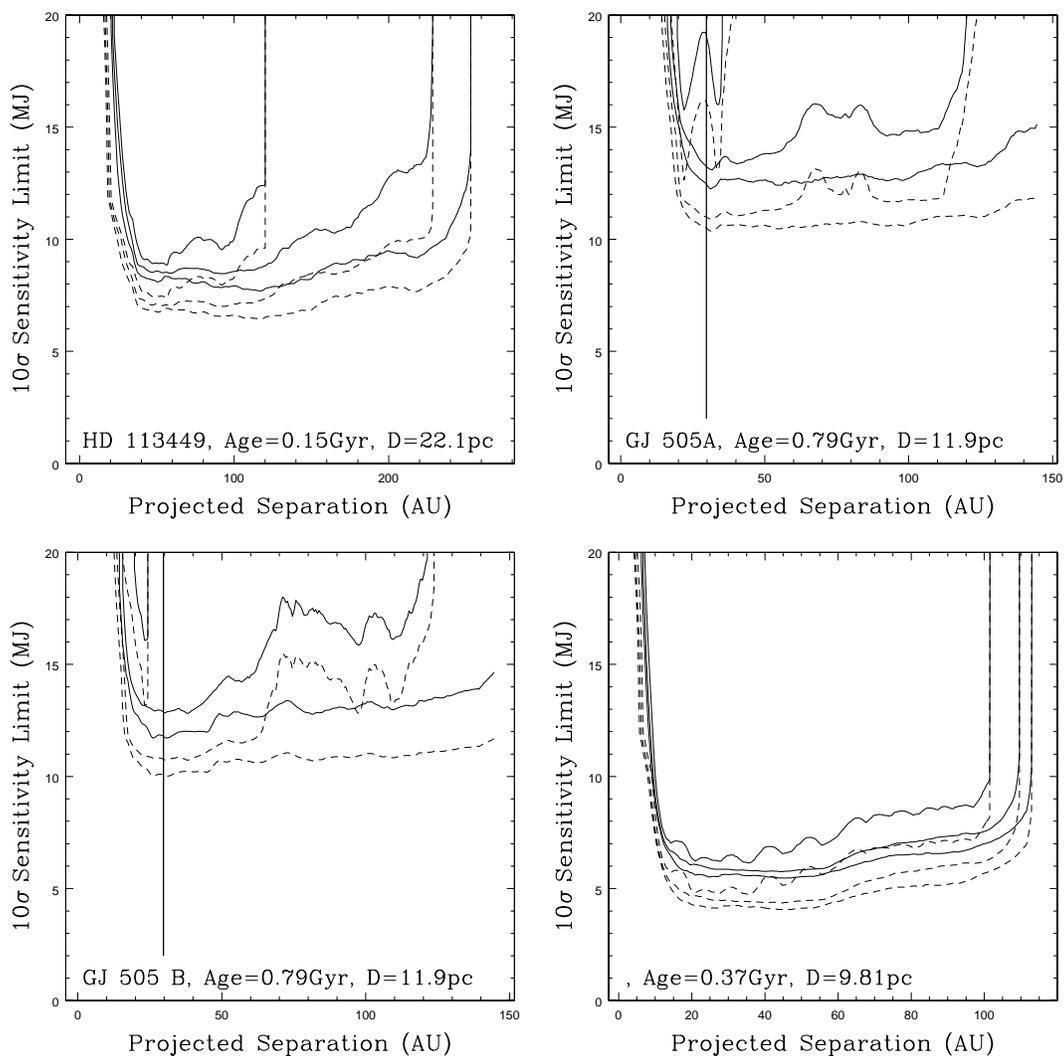


Figure 3.72 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits with radii beyond the heavy vertical lines on the GJ 505A and B plots are unstable in this binary system.

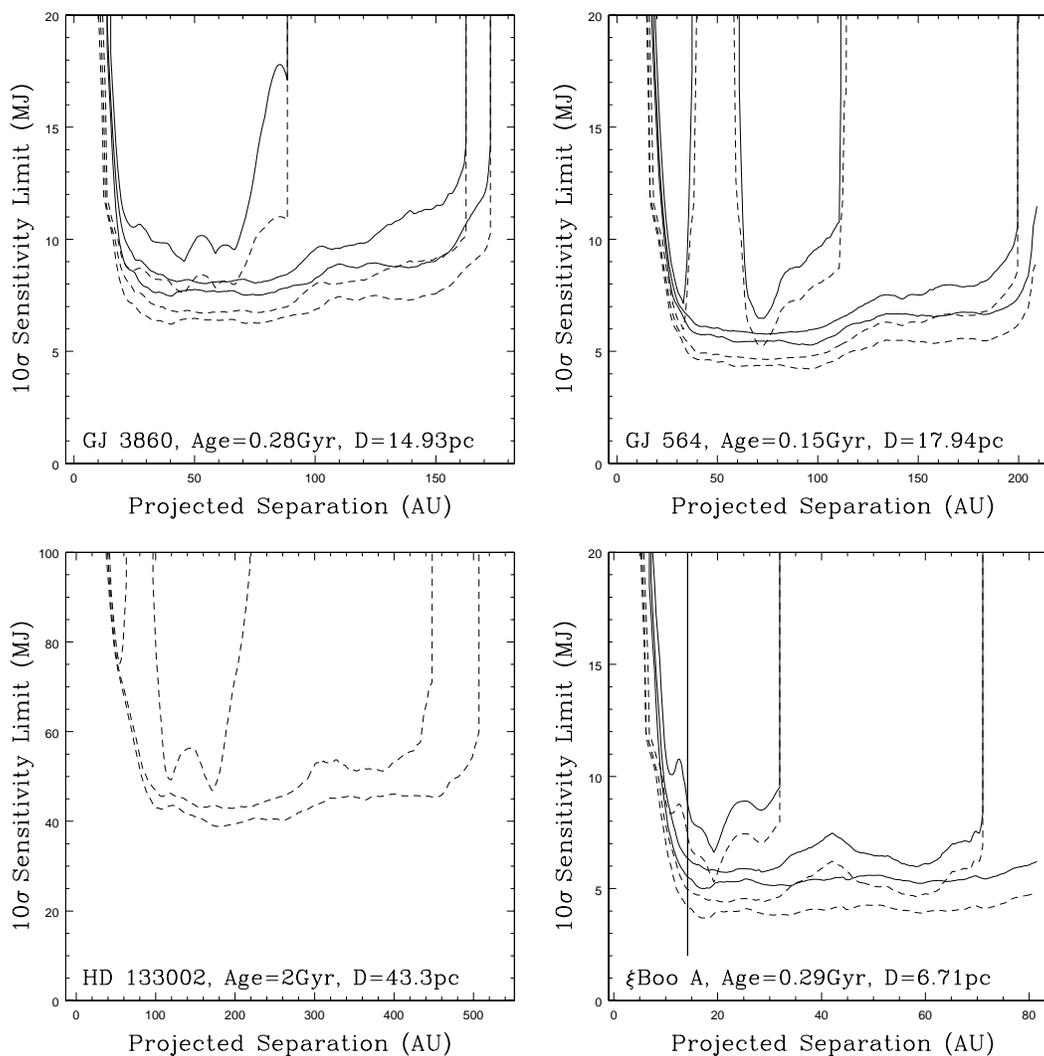


Figure 3.73 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Due to the binary companion, planetary orbits around  $\xi$  Boo A are stable only inward of the heavy vertical line.

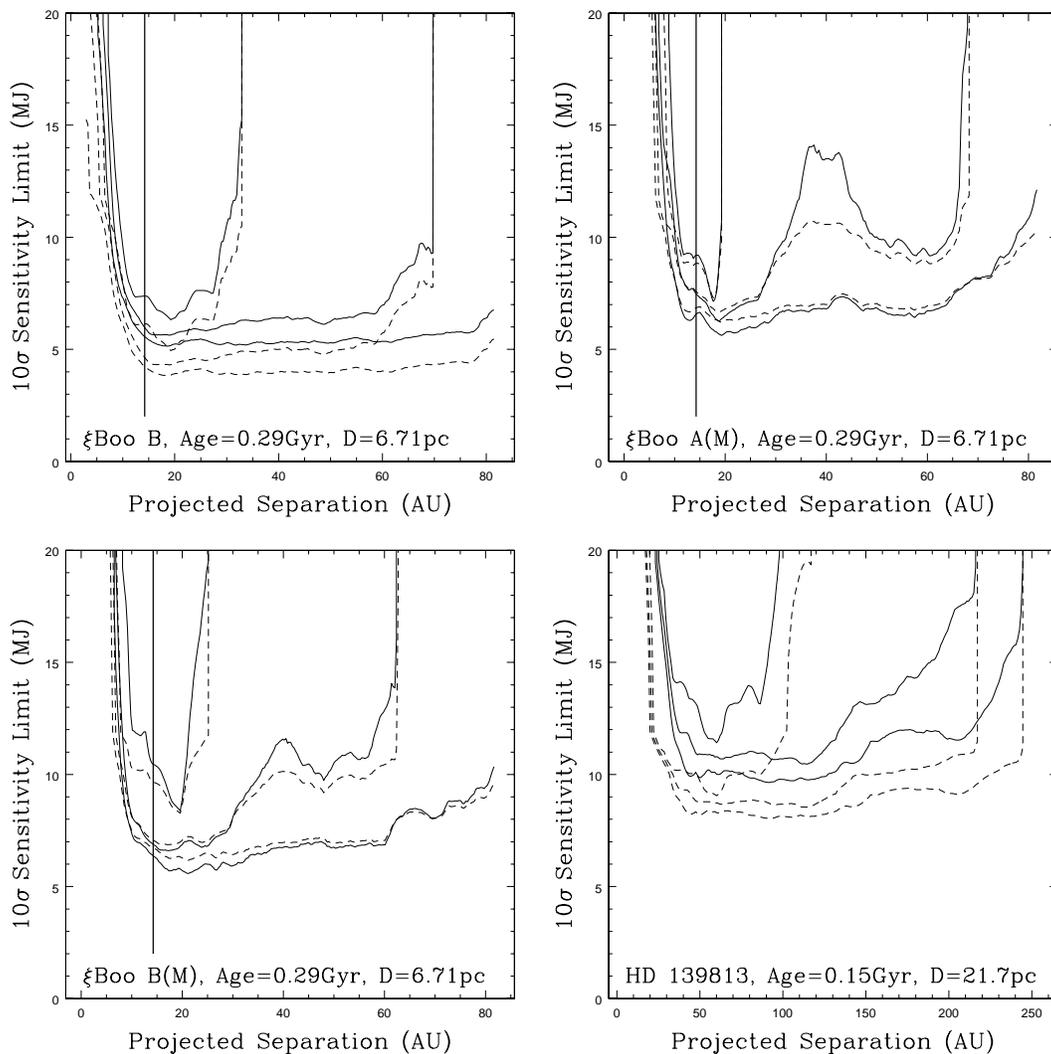


Figure 3.74 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits around  $\xi$  Boo A and B with radii outward of the heavy vertical lines are unstable in this binary system.

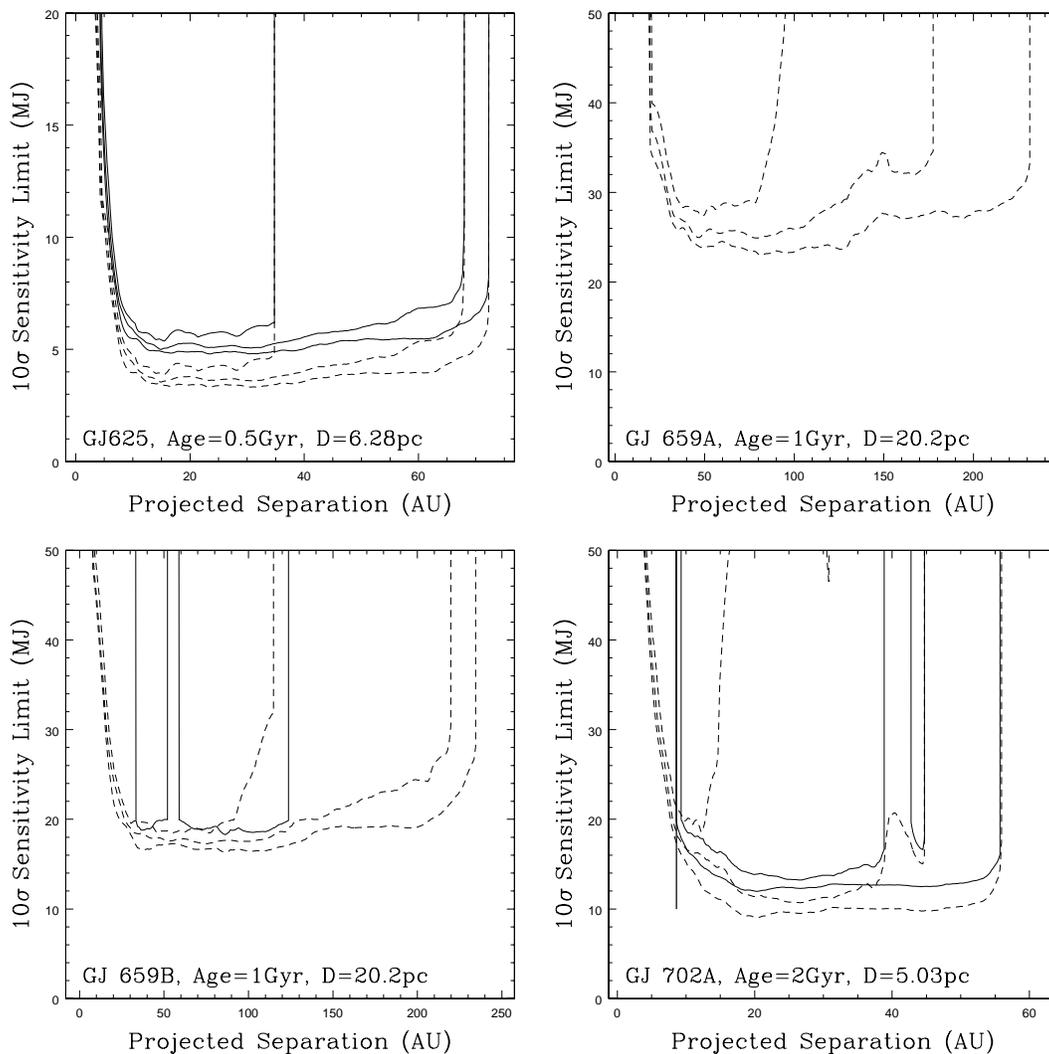


Figure 3.75 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Due to the binary companion, planetary orbits around GJ 702A are stable only inward of the heavy vertical line.

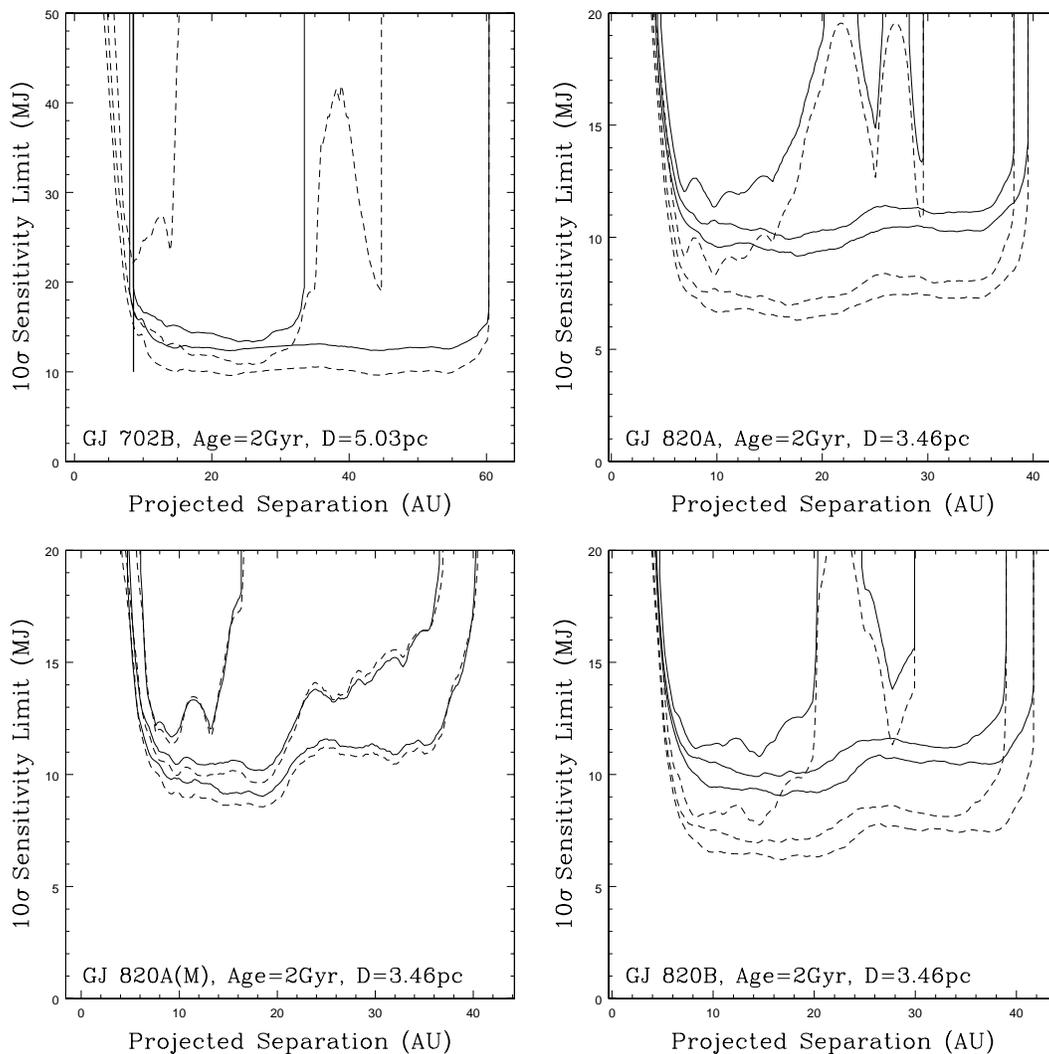


Figure 3.76 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Due to the binary companion, planetary orbits around GJ 702B are stable only inward of the heavy vertical line.

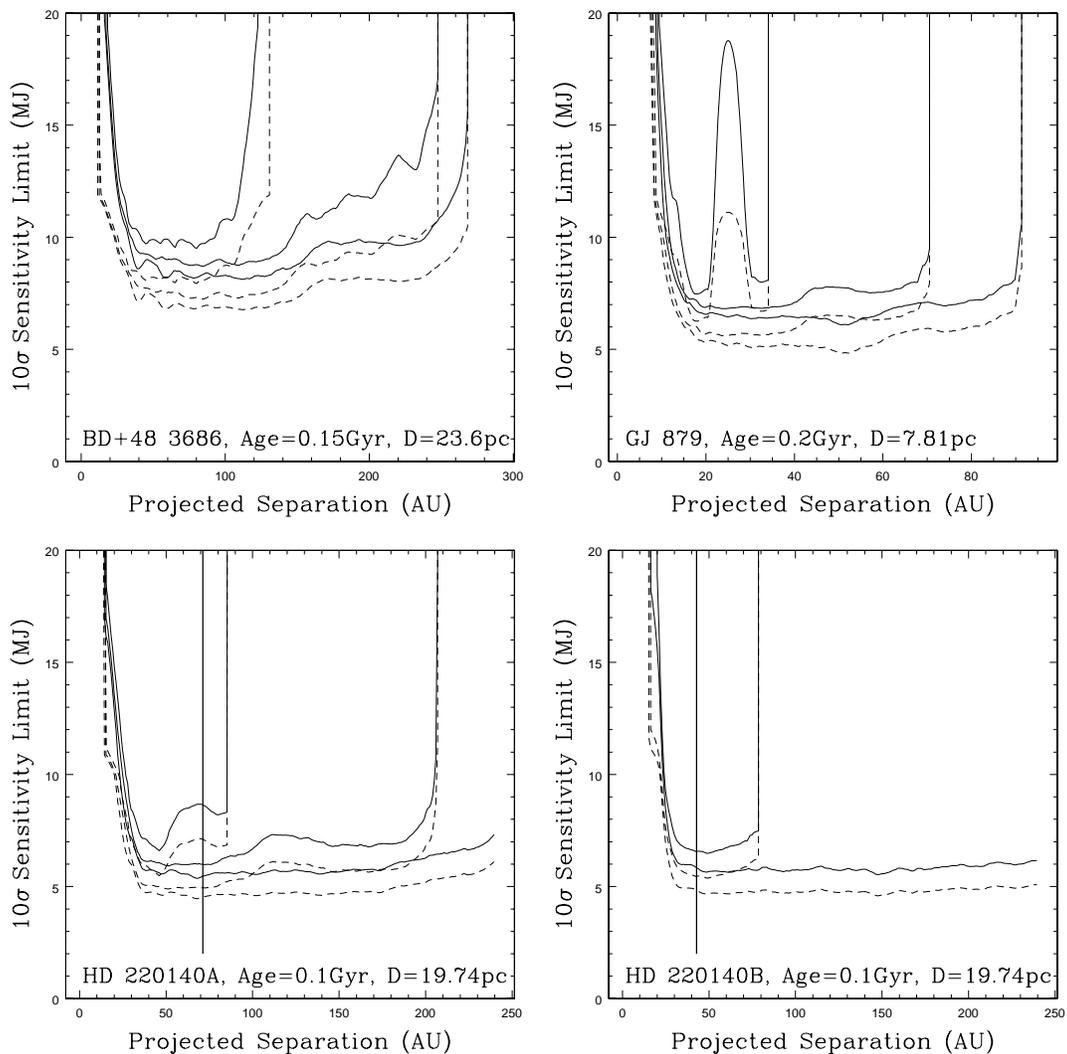


Figure 3.77 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits with radii beyond the heavy vertical lines on the HD 220140 A and B plots are unstable in this binary system.

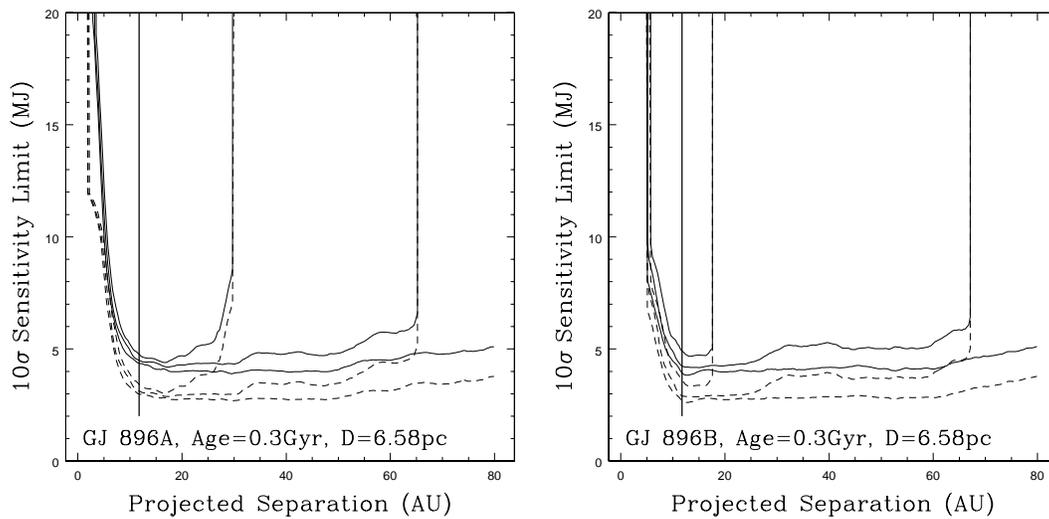


Figure 3.78 Minimum detectable planet mass vs projected separation in AU for stars surveyed. The 0th, 50th, and 90th percentile sensitivities vs separation are plotted. By definition these always nest, and so can easily be identified. For the solid lines, magnitudes have been converted to planet masses using the models of Burrows et al. (2003); for the dashed lines we have used Baraffe et al. (2003). Planetary orbits with radii beyond the heavy vertical lines are unstable in this binary system.

### 3.9 Sensitivities Compared to Other Surveys

The bulk of this section will be devoted to comparing our surveys to shorter wavelength ones. The  $L'$  survey of Kasper et al. (2007) will not be discussed in detail. However, we comment that Kasper et al. (2007) reach an average background-limited point-source  $5\sigma$  sensitivity of  $L' = 16 - 16.5$  in a 15 minute integration using counter-chopping, an advanced sky subtraction method for which we did not have the required hardware. In dithering operation, an observing strategy more comparable to our nodding, they obtain  $L' = 15.5 - 16.0$  instead. If we assume that the sensitivity of an observation goes as  $t^{-1/2}$ , where  $t$  is the total integration time, we would expect a  $5\sigma$  limits obtained in a 15 minute integration to be the same as a  $10\sigma$  limit obtained in a 1 hour integration. Thus, if our system is comparable to NACO we should have mean  $10\sigma$  background-limits of  $L' = 15.5 - 16.0$  in a one-hour exposure. However, NACO operates on the 8.2 meter VLT, whereas Clio uses the 6.5 meter MMT. For diffraction limited, background limited observations the sensitivity is expected to scale as the telescope diameter squared. This should give the VLT an advantage of 0.5 magnitudes, so that we might expect to attain a  $10\sigma$  sensitivity of  $L' = 15.0 - 15.5$  in a one hour exposure. In fact, we do somewhat better than this, with one-hour scaled background limited sensitivities peaking above  $L' = 16.0$  in the best data sets. This is probably due to the cleaner sky background delivered by the MMT adaptive secondary AO system.

We turn now to the comparison of our sensitivities with those obtained at shorter wavelengths. In Sections 3.2 and 3.3 we discussed the potential advantages of the  $L'$  and  $M$  bands over  $H$  for plane imaging surveys. We commented that planets are far brighter at the longer wavelengths, and in particular are brighter relative to their parent stars. As an example, according to the tables in Baraffe et al. (2003) a 5MJ planet orbiting a star of 0.5 Gyr age 10 pc from Earth

has an  $L'$  magnitude of 16.32. We could detect this planet easily in our best data sets provided it were far enough from the star. At  $H$  band the same planet has a magnitude of 20.11. The difference is 3.79 magnitudes, or a factor of 33. Even a rather red star, say spectral type K5, is just 0.22 magnitudes brighter at  $L'$  than at  $H$ . For such a star at 10 pc, the planet/star flux ratio for an 0.5 Gyr-old, 5MJ planet is still 25 times better at  $L'$  than at  $H$ .

The advantage is even more for a closer star, where the sensitivities at both  $L'$  and  $H$  band extend to cooler, fainter objects with even redder  $H - L'$  colors. For example, an  $L' = 16.3$  planet of age 2 Gyr orbiting a star 3.5 parsecs away has a mass of 7.6 MJ according to the Burrows et al. (2003) models, or 5.1 MJ according to Baraffe et al. (2003). At  $H$ -band the two different model sets predict the object will be 4.1 or 5.9 magnitudes fainter, respectively. If this planet is orbiting a K5 star the planet/star flux ratio at the  $L'$  band is 36 times or 190 times better than  $H$  band, using the Burrows et al. (2003) and Baraffe et al. (2003) model sets respectively. We note that the planet masses used here were the minimum masses detectable in a background limited region of an image from a good data set in our survey. They do not necessarily correspond to planets detectable sufficiently close to a star that  $H$  band images would be speckle-noise dominated and the planet/star contrast ratio would be critical.

As discussed in 3.2, there are balancing considerations to the large planet/star flux ratio advantages delivered by the  $L'$  and  $M$  bands. Here we will investigate in detail how our sensitivities compare with those of shorter-wavelength surveys.

First, we present Figures 3.79 and 3.80, showing the maximum detectable  $\Delta$ -magnitude for our survey and those of Biller et al. (2007) and Lafrenière et al. (2007b), around the relatively faint, distant star HD 96064 ( $H = 5.90$ , distance 24.63 pc) and the bright, nearby star  $\epsilon$  Eri ( $H = 1.88$ , distance 3.27 pc). We have

assumed that the  $5\sigma$  sensitivities from Biller et al. (2007) and Lafrenière et al. (2007b) are equivalent to our  $7\sigma$  sensitivities, as discussed in Subsection 3.7.6 above. Based on these assumptions, all sensitivities have been adjusted to approximate  $10\sigma$  values.

It is clear that the Biller et al. (2007) observations, with their powerful SDI speckle suppression method, are well tuned to deliver excellent  $\Delta$ -magnitude sensitivity close to the star. The Lafrenière et al. (2007b) observations are saturated close in and then deliver extremely good sensitivity at larger radii. The  $\Delta$ -magnitude sensitivity delivered by our own survey is substantially less good than either Biller et al. (2007) or Lafrenière et al. (2007b) at all distances from each star except a small range around  $\epsilon$  Eri that lies beyond the edge-of-field of Biller et al. (2007) and within the saturation radius of Lafrenière et al. (2007b).

When converting these  $\Delta$ -magnitude sensitivities to planet masses (see Chapter 4 for a detailed description of how we do this), the considerable planet/star flux ratio advantage delivered by the  $L'$  band (and the  $M$  band in the case of  $\epsilon$  Eri) will come into play. However, the  $H$  band  $\Delta$ -magnitude advantages are so large that it is not obvious the  $L'$  and  $M$  band observations will prove more sensitive. This is especially true for HD 96064, because it is a distant star and the faintest objects detectable will not be so red in  $H - L'$  color as for  $\epsilon$  Eri.

In Figures 3.81 and 3.82 we present minimum detectable planet mass versus separation from the star HD 96064 from Lafrenière et al. (2007b), Biller et al. (2007), Masciadri et al. (2005), and our own  $L'$  observation. Lafrenière et al. (2007b) actually observed this star; for Biller et al. (2007) we have used sensitivities from a similar brightness star (UY Pic). For the  $H$  and  $K_s$  band Masciadri et al. (2005) results we have used sensitivities from similar-brightness stars (HD 221503 and HIP 2729, respectively). For Lafrenière et al. (2007b) and Biller et al.

(2007), as mentioned above, we have assumed the quoted  $5\sigma$  limits correspond to our  $7\sigma$  limits, and converted the sensitivities to approximate  $10\sigma$  limits based on this assumption. Furthermore, we have applied corrections to the sensitivities based on the claims in each paper that for giant planet detection the narrowband filters used were superior to  $H$  band. Lafrenière et al. (2007b) claim planets are 1.5-2.5 times brighter in their band than in  $H$ ; accordingly we have adjusted their limits by 0.75 magnitudes (that is, a factor of 2.0) in the sense of improved sensitivity to get approximate equivalent  $H$  band magnitude limits. We have similarly adjusted the Biller et al. (2007) sensitivities by 0.84 magnitudes based on the suggestion in their paper. The Masciadri et al. (2005) data was for the standard  $H$  and  $K_s$  filters, and accordingly we have applied no correction except the conversion of their quoted  $5\sigma$  limits to  $10\sigma$  limits to match our own. Finally, we have converted the magnitudes to planet masses based on the age and distance of HD 96064, using interpolations of the planet models of Burrows et al. (2003) to produce Figure 3.81 and those of Baraffe et al. (2003) for Figure 3.82.

As is not surprising, for this rather distant star the shorter wavelength surveys do substantially better than our  $L'$  observation.  $L'$  observations of such systems diversify the investment of planet imaging efforts, hedging the astronomical community against possible large differences between the spectra of real giant planets and the current models, but if the models are largely correct the shorter wavelengths are substantially better.

It is interesting to note that there are significant differences between the prediction of the Burrows et al. (2003) and Baraffe et al. (2003) models. The Baraffe et al. (2003) models show a considerably smaller gap between our sensitivities and those of the Lafrenière et al. (2007b) survey. Also, the Burrows et al. (2003) models show the Masciadri et al. (2005)  $K_s$  band observation attaining considerable bet-

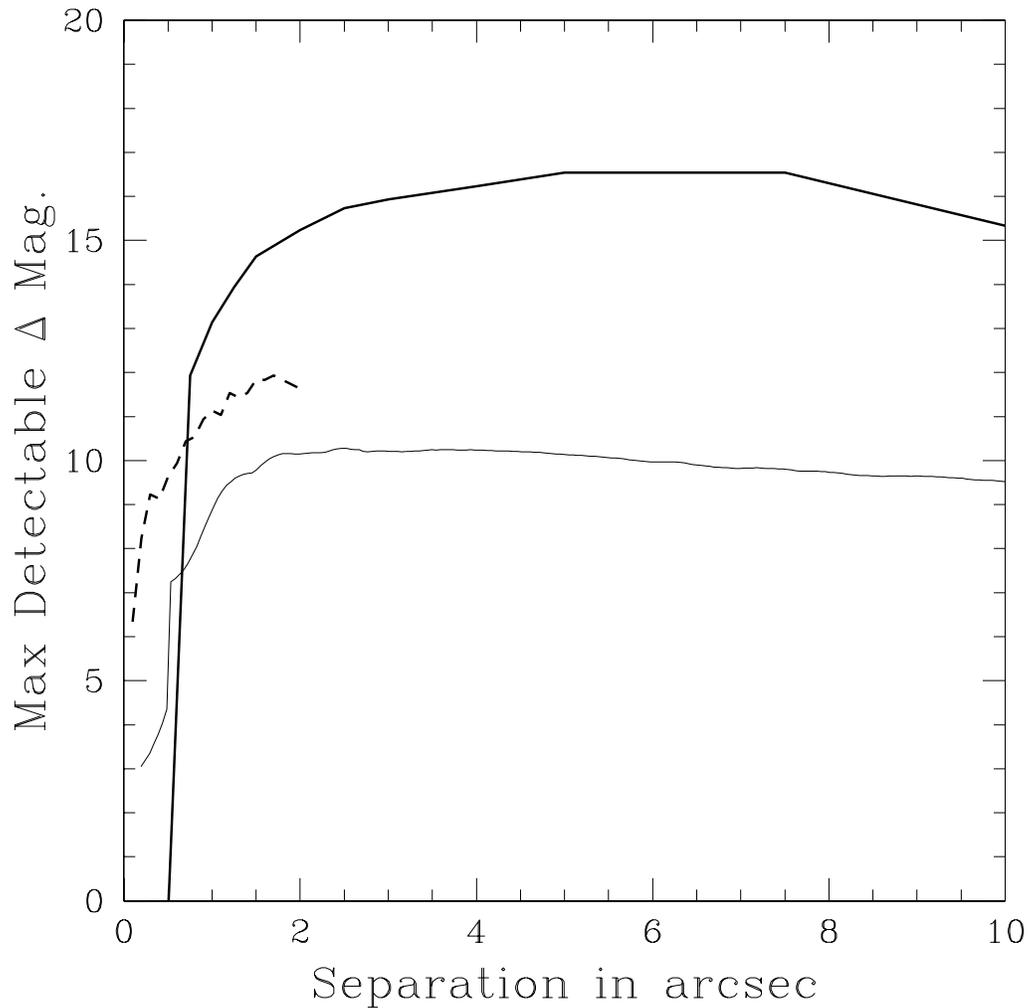


Figure 3.79  $\Delta$ -magnitude vs separation for HD 96064 ( $H = 5.90$ ;  $L' = 5.75$ ; distance 24.63 pc, age 0.15 Gyr) in various surveys. The heavy continuous line is from Lafrenière et al. (2007b), the heavy dashed line from Biller et al. (2007), and the light continuous line is from our  $L'$  observation.

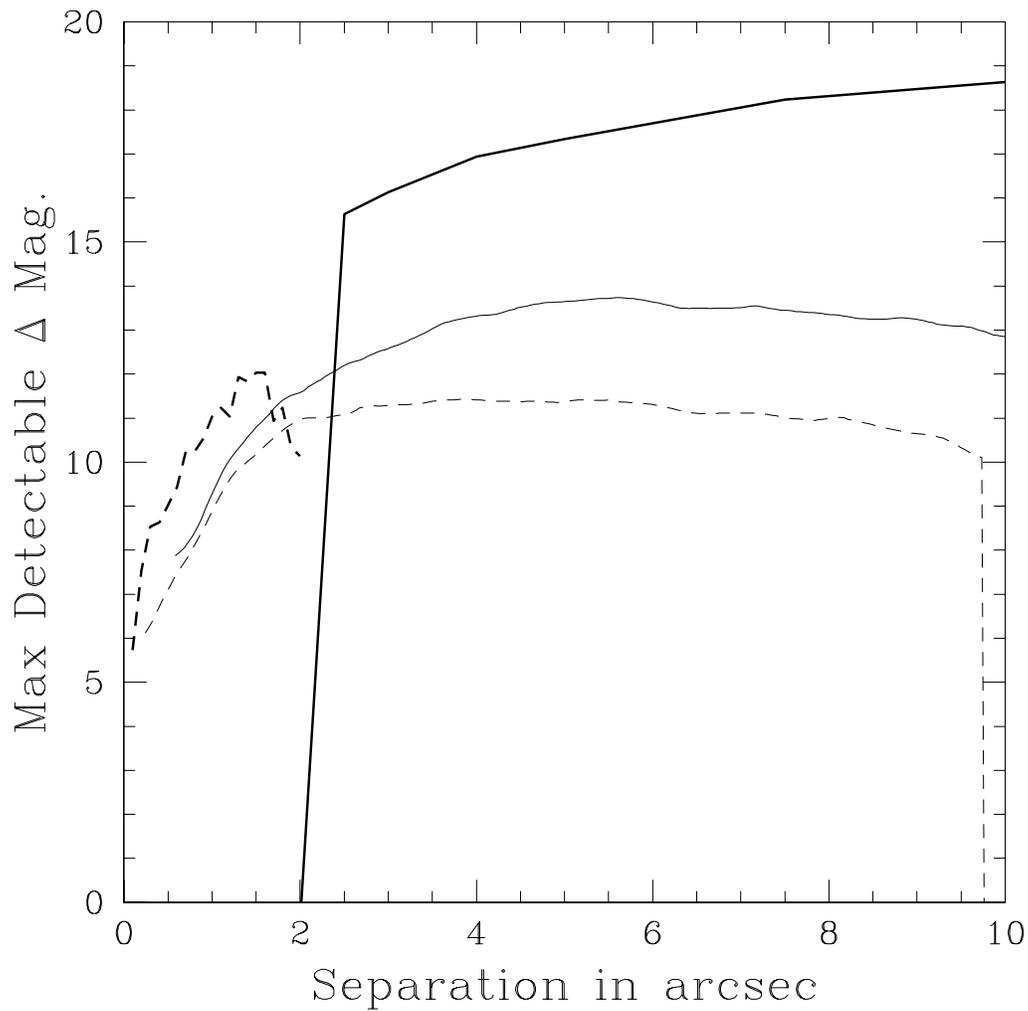


Figure 3.80  $\Delta$ -magnitude vs separation for  $\epsilon$  Eri ( $H = 1.88$ ;  $L' = 1.72$ ; distance 3.27 pc, age 0.56 Gyr) in various surveys. The heavy continuous line is from Lafrenière et al. (2007b), the heavy dashed line from Biller et al. (2007), the light continuous line from our  $L'$  observation, and the light dashed line from our  $M$  band observation.

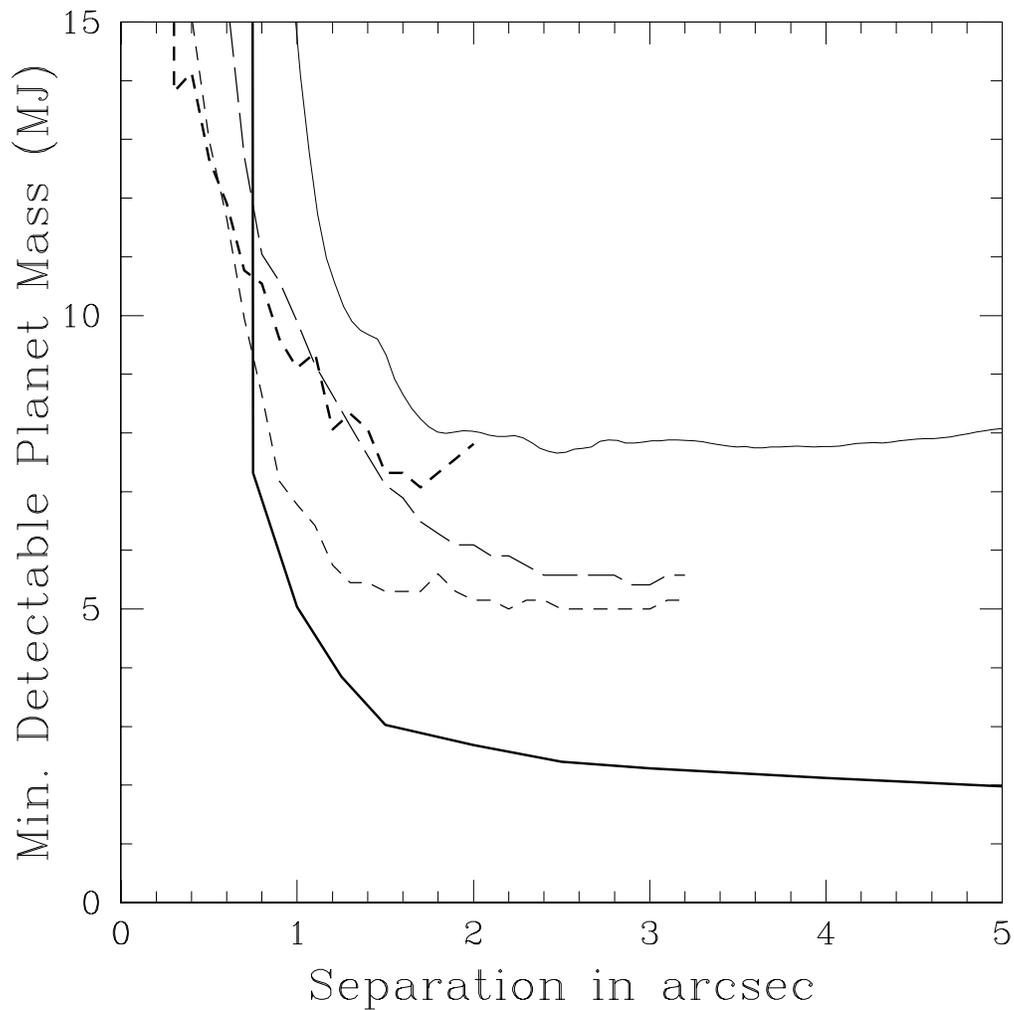


Figure 3.81 Minimum detectable planet mass vs separation for HD 96064 ( $H = 5.90$ ;  $L' = 5.75$ ; distance 24.63 pc, age 0.15 Gyr), based on the Burrows et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observations, the heavy short-dashed line for Biller et al. (2007), the light long-dashed and short-dashed lines, respectively, for  $H$  and  $Ks$  band data from Masciadri et al. (2005), and the light continuous line for our own  $L'$  observation.

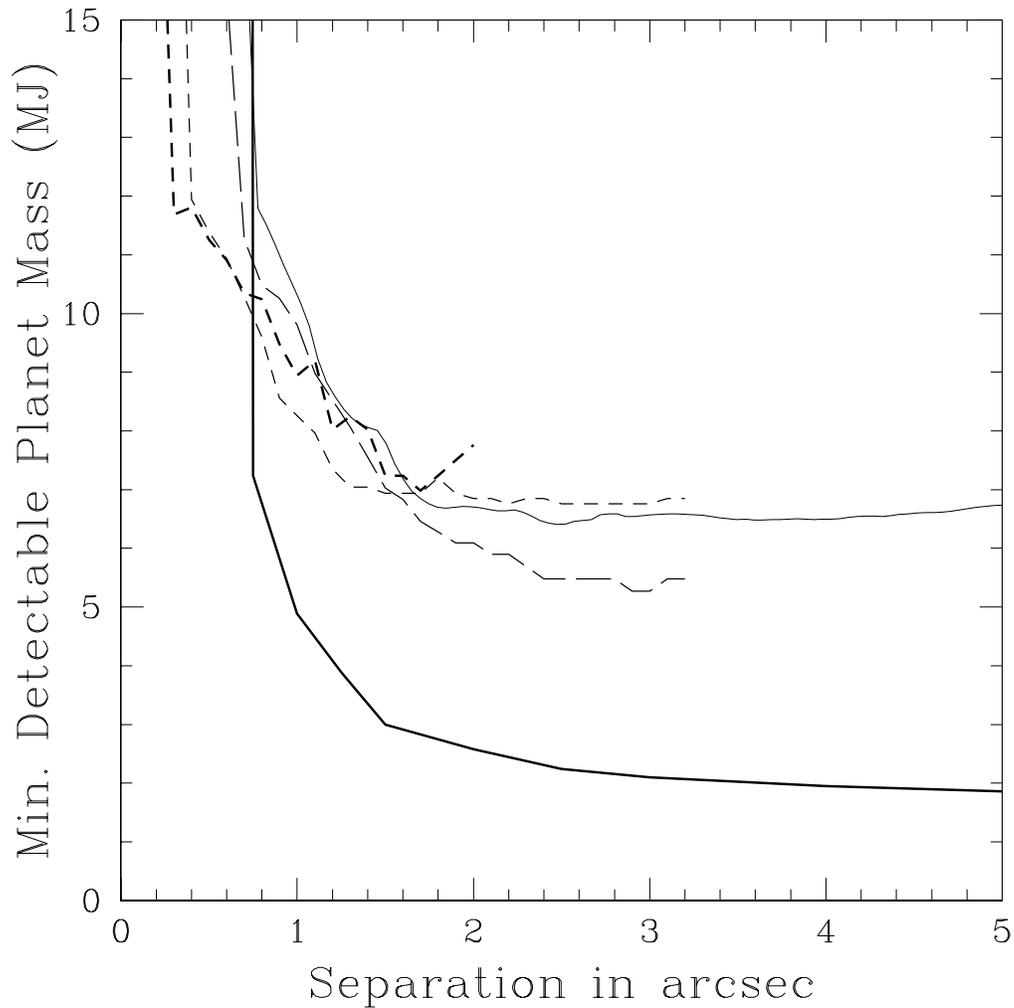


Figure 3.82 Minimum detectable planet mass vs separation for HD 96064 ( $H = 5.90$ ;  $L' = 5.75$ ; distance 24.63 pc, age 0.15 Gyr), based on the Baraffe et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observations, the heavy short-dashed line for Biller et al. (2007), the light long-dashed and short-dashed lines, respectively, for  $H$  and  $Ks$  band data from Masciadri et al. (2005), and the light continuous line for our own  $L'$  observation.

ter sensitivity than  $H$ , but for the Baraffe et al. (2003) models the case is reversed. There is a difference in the filter sets between the Burrows et al. (2003) and Baraffe et al. (2003) models: we have integrated the Burrows et al. (2003) model spectra using specifications for the MKO standard filters used in Clio, while the magnitudes given in Baraffe et al. (2003) are calculated using the CIT filter set ( $H$  and  $K$ ), the Johnson-Glass  $L'$  band, and the Johnson  $M$  band. As we have discussed in Chapter 2, the different filter sets may be responsible for some of the model disagreements, but inherent differences in the models exist and probably represent a more important effect.

In Figures 3.83 and 3.84 we present mass sensitivity curves for  $\epsilon$  Eri from Lafrenière et al. (2007b), Biller et al. (2007), and our own  $L'$  and  $M$  band observations. The Masciadri et al. (2005) survey observed only stars much fainter than  $\epsilon$  Eri, so no Masciadri et al. (2005) curve is shown. All sensitivity corrections and mass conversions were performed exactly as for HD 96064.

For this very bright star, it is clear that our observations attained better overall sensitivity than any others, exceeding even the Lafrenière et al. (2007b) sensitivities except at extremely large radii. Note that even the powerful SDI technique of Biller et al. (2007) is not able to overcome the much more unfavorable  $H$  band planet/star contrast ratio to deliver better sensitivity than our observations close to the star. Note also that the  $M$  band appears to deliver results comparable to those of the  $L'$  band for this very nearby star. Again, the results from the two model sets differ significantly, enough that they would predict different optimal filters for observations of this star. The Burrows et al. (2003) models indicate that the  $M$  band delivers the best sensitivity until beyond 7 arcsec, a very large separation. They also indicate  $L'$  is substantially less good than the Lafrenière et al. (2007b) result at large separations. The Baraffe et al. (2003) models show our  $L'$

observation much better relative to both  $M$  band and the Lafrenière et al. (2007b) result. The  $M$  band dominates  $L'$  by much less, and only within about 4 arc-sec. At large radii the  $L'$  sensitivity is very comparable to that of Lafrenière et al. (2007b), in contrast to the Burrows et al. (2003) result which showed it considerably inferior.

We have compared the sensitivities of  $H$  band regime surveys with our  $L'$  results for a young, distant, relatively faint star, HD 96064 ( $H = 5.90$ ;  $L' = 5.75$ ; distance 24.63 pc, age 0.15 Gyr), and with our  $L'$  and  $M$  band results for an older, very nearby, bright star,  $\epsilon$  Eri ( $H = 1.88$ ;  $L' = 1.72$ ; distance 3.27 pc, age 0.56 Gyr). In Figures 3.85 and 3.86, we present the comparison for a star intermediate between the previous extreme cases. GJ 117, or HD 17925, was observed by Lafrenière et al. (2007b), Biller et al. (2007), and our own survey. It is moderately bright ( $H = 4.23$ ), young enough to be very interesting (0.1 Gyr), and at 8.31 pc is quite nearby, though not nearly as close as  $\epsilon$  Eri.

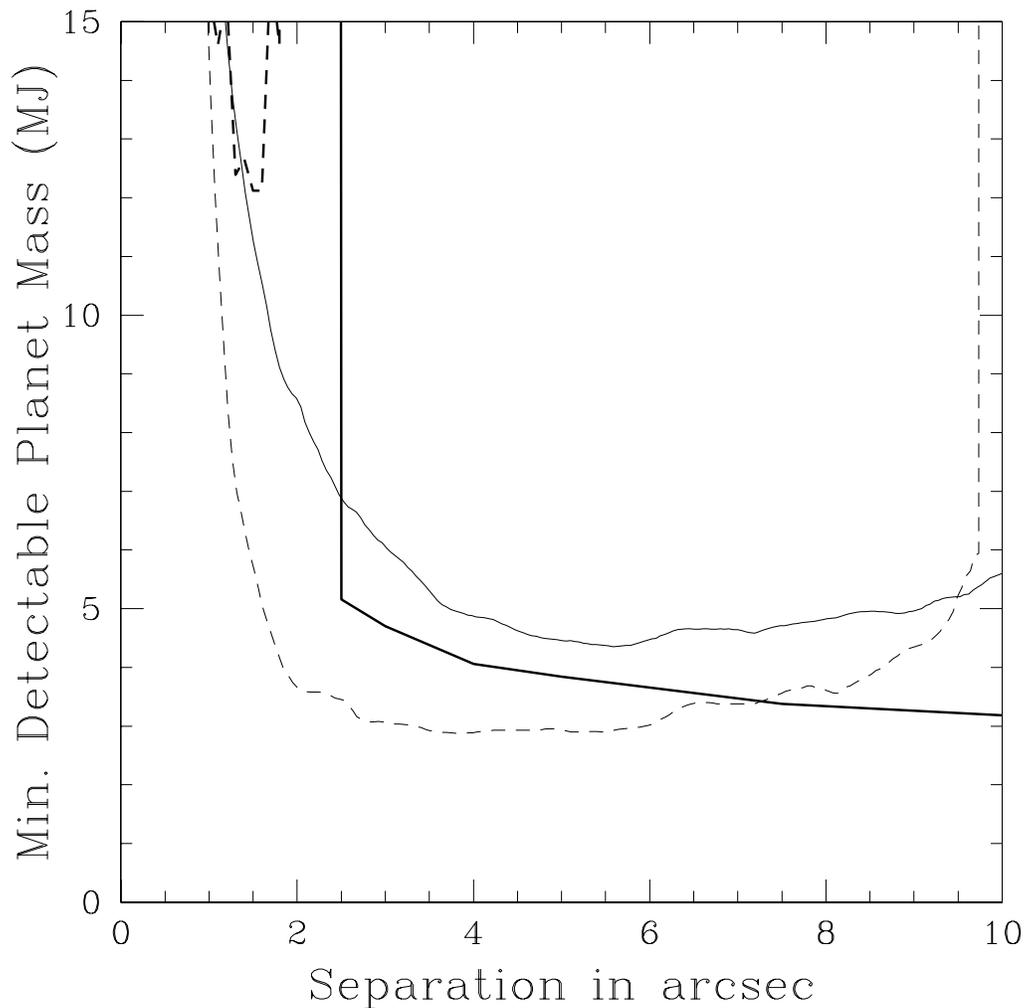


Figure 3.83 Min detectable planet mass vs separation for  $\epsilon$  Eri ( $H = 1.88$ ;  $L' = 1.72$ ; distance 3.27 pc, age 0.56 Gyr), based on the Burrows et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observation, the heavy dashed line for Biller et al. (2007) and the light continuous and dashed lines for our  $L'$  and  $M$  band observations, respectively.

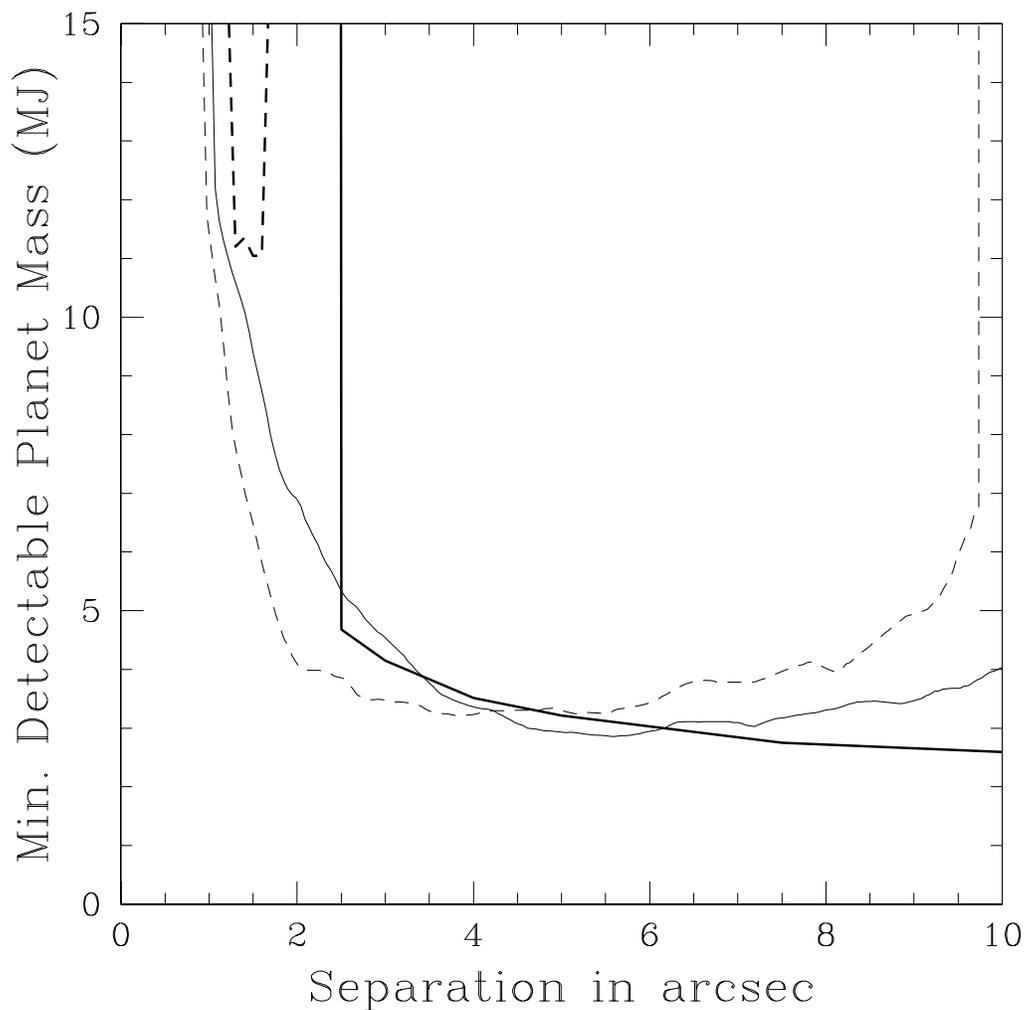


Figure 3.84 Min detectable planet mass vs separation for  $\epsilon$  Eri ( $H = 1.88$ ;  $L' = 1.72$ ; distance 3.27 pc, age 0.56 Gyr), based on the Baraffe et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observation, the heavy dashed line for Biller et al. (2007) and the light continuous and dashed lines for our  $L'$  and  $M$  band observations, respectively.

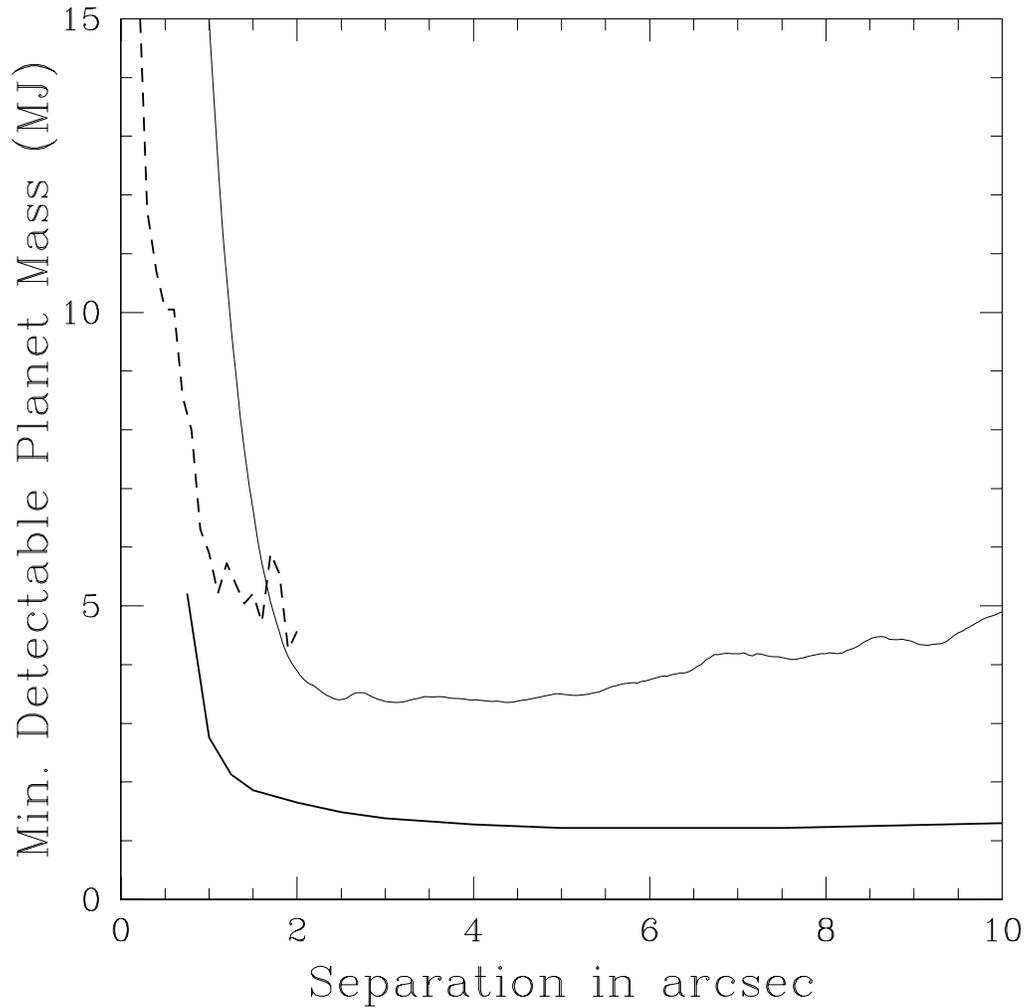


Figure 3.85 Min detectable planet mass vs separation for GJ 117 ( $H = 4.23$ ;  $L' = 4.11$ ; distance 8.31 pc, age 0.1 Gyr), based on the Baraffe et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observation, the heavy dashed line for Biller et al. (2007) and the light continuous line for our  $L'$  observation.

Even for this fairly nearby star, for which we obtained sensitivity to quite low mass planets, the Biller et al. (2007) results are better than ours at small separations, and the Lafrenière et al. (2007b) results are better further out. Note that for this fainter star the Lafrenière et al. (2007b) sensitivity is better than that of Biller et al. (2007) except for the very smallest separations. Note also that, once again, the Burrows et al. (2003) models show the gap between our sensitivity and that of Lafrenière et al. (2007b) a good deal larger than the Baraffe et al. (2003) models.

It is clear at this point that if the Burrows et al. (2003) and Baraffe et al. (2003) models are broadly correct our survey obtained better sensitivity than the Lafrenière et al. (2007b) observations only for the very brightest stars, considerably brighter than GJ 117. We have compared sensitivities for the bright, nearby, fairly young star  $\epsilon$  Eri. What would the result be for an older star with similar brightness and distance?

We have just such a star. 61 Cyg A, or GJ 820A, has an  $H$  band magnitude of 2.47, only slightly fainter than  $\epsilon$  Eri. At 3.46 pc it is only a little more distant, but at our adopted age of 2 Gyr it is much older. Lafrenière et al. (2007b) did not observe this star, or any star of comparable brightness. Scaling their  $\epsilon$  Eri results to a fainter star would not be a fair comparison. In order to avoid shortchanging the Lafrenière et al. (2007b) sensitivities, we have elected to scale their results from a star somewhat fainter than 61 Cyg A,  $\eta$  Crv, at  $H = 3.37$ . The result of the comparison is shown in Figures 3.87 and 3.88.

Here we see that although at very small separations our sensitivity is comparable to that of Lafrenière et al. (2007b), at large separations the plot based on the Burrows et al. (2003) models shows our observations delivering much poorer sensitivity than Lafrenière et al. (2007b). The gap is considerably larger than for the  $\epsilon$  Eri data. As usual, the Baraffe et al. (2003) plot shows us doing much bet-

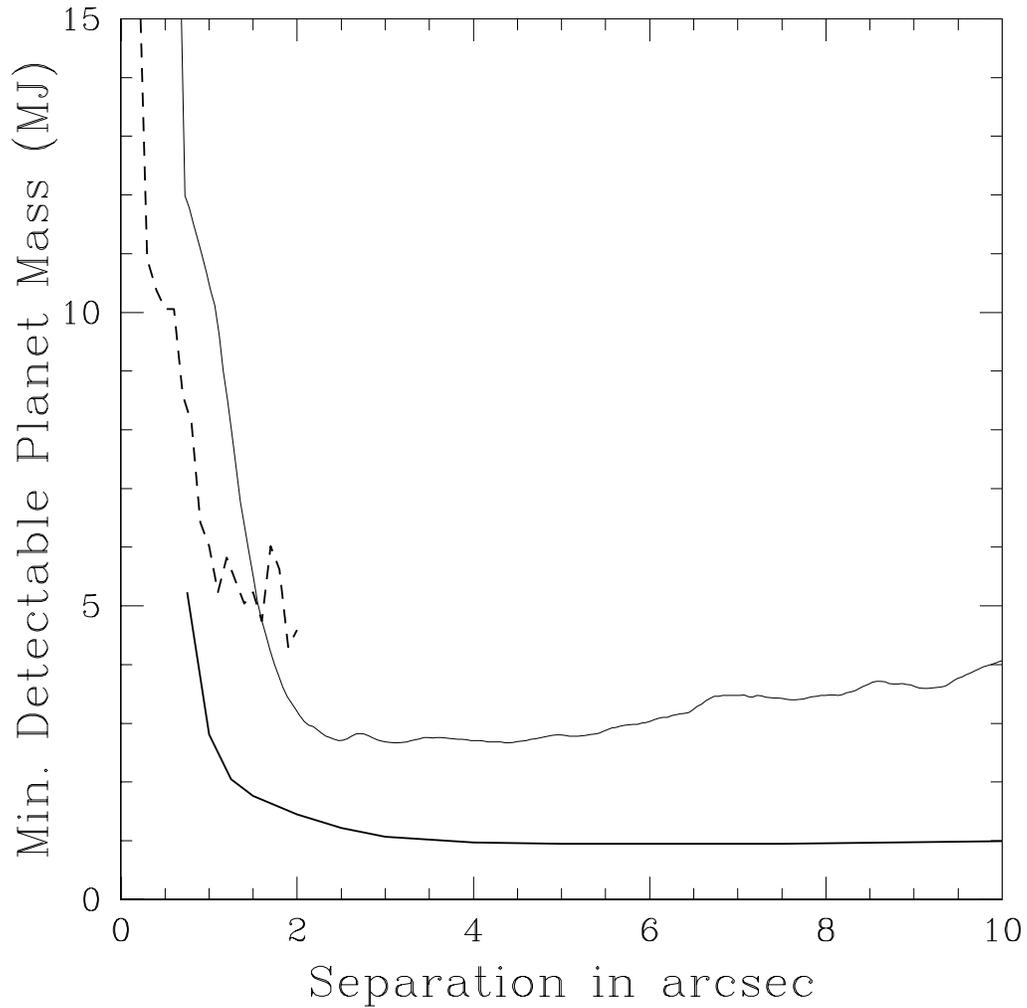


Figure 3.86 Min detectable planet mass vs separation for GJ 117 ( $H = 4.23$ ;  $L' = 4.11$ ; distance 8.31 pc, age 0.1 Gyr), based on the Baraffe et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observation, the heavy dashed line for Biller et al. (2007) and the light continuous line for our  $L'$  observation.

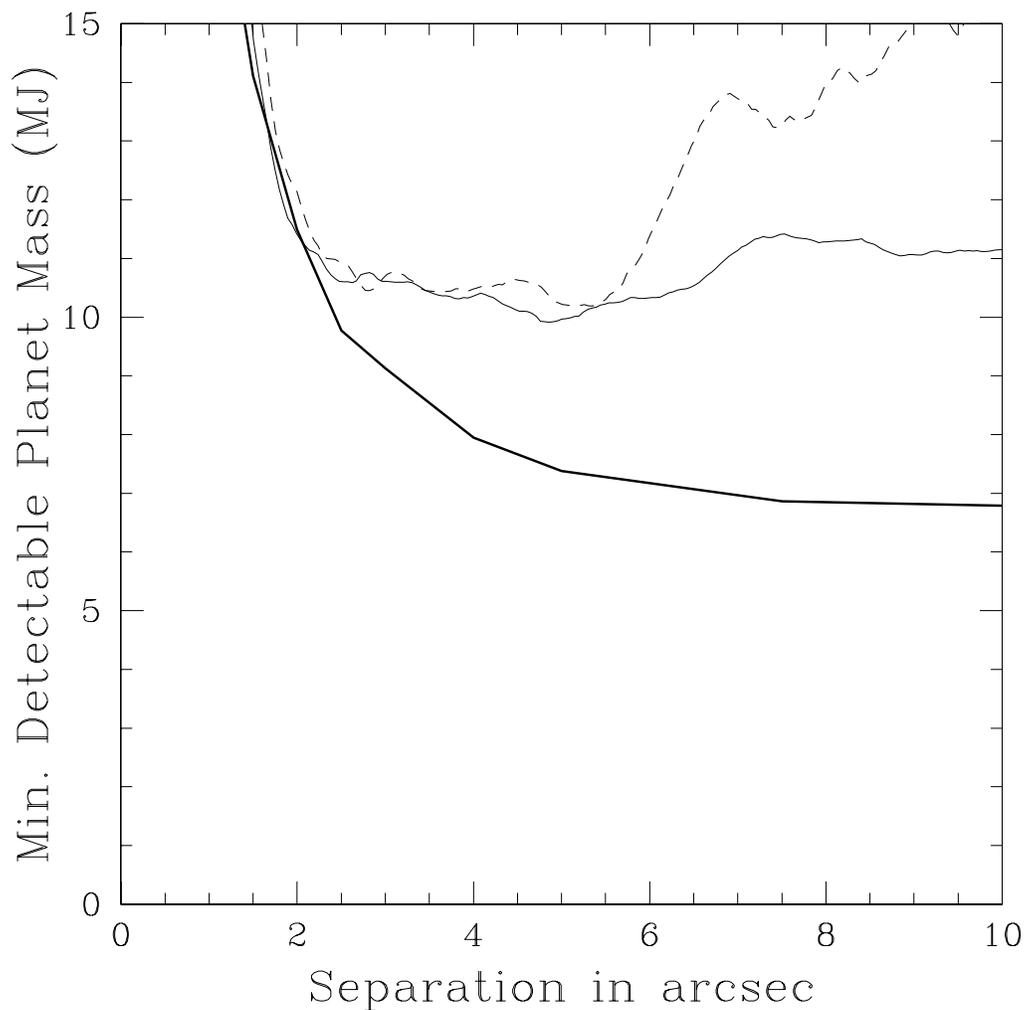


Figure 3.87 Min detectable planet mass vs separation for 61 Cyg A ( $H = 2.47$ ;  $L' = 2.25$ ; distance 3.46 pc, age 2 Gyr), based on the Burrows et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observation, and the light continuous and dashed lines are for our  $L'$  and  $M$  band observations, respectively.

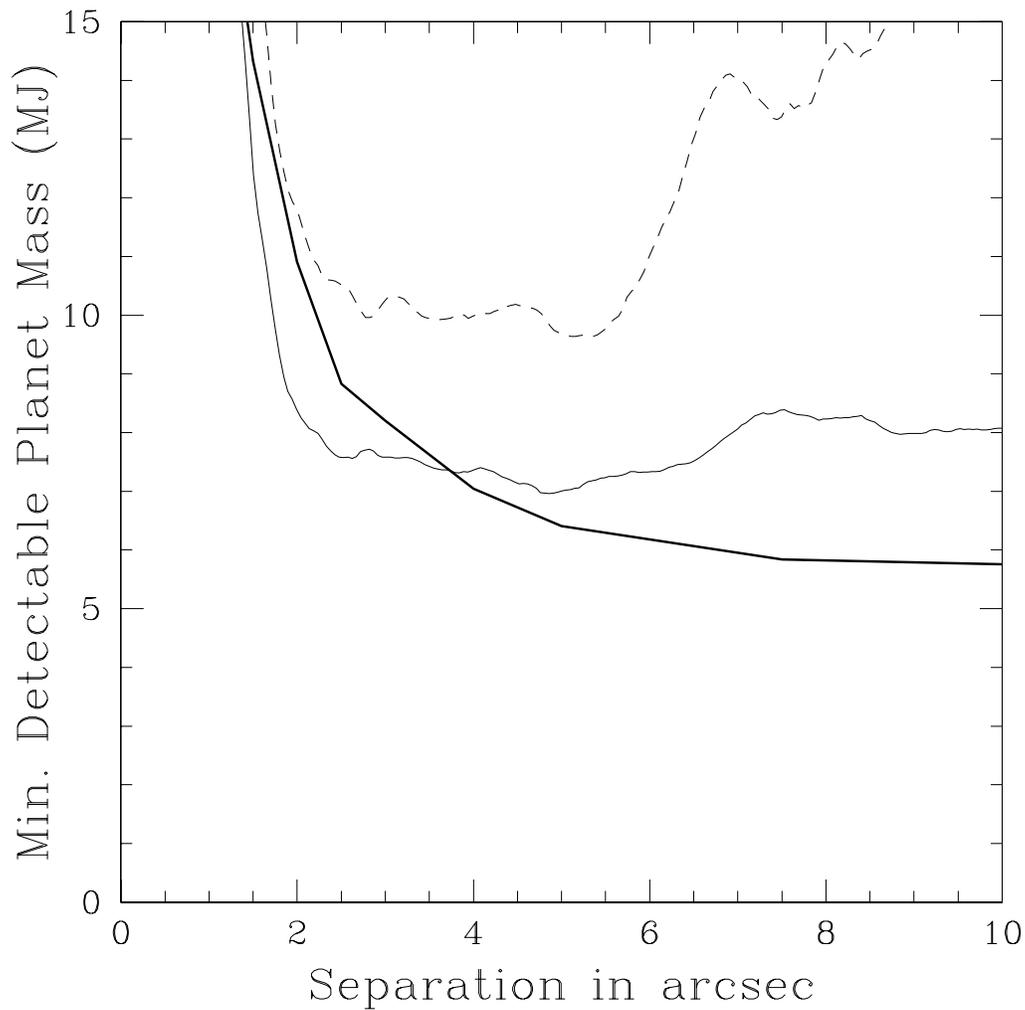


Figure 3.88 Min detectable planet mass vs separation for 61 Cyg A ( $H = 2.47$ ;  $L' = 2.25$ ; distance 3.46 pc, age 2 Gyr), based on the Baraffe et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observation, and the light continuous and dashed lines are for our  $L'$  and  $M$  band observations, respectively.

ter relative to Lafrenière et al. (2007b) than the Burrows et al. (2003) plot, with somewhat better sensitivity at  $L'$  out to about 4 arcsec.

This result seems inconsistent with the  $\epsilon$  Eri result, which showed our sensitivities generally at least as good as those of Lafrenière et al. (2007b). Does the utility of  $L'$  and  $M$  relative to  $H$  decrease for old bright nearby stars? Figures 3.87 and 3.88 are not a very good test, because our data sets on 61 Cyg A were considerably below average. Accordingly, we have calculated our planet mass sensitivity curves assuming we had attained the same magnitude sensitivity around 61 Cyg A that we actually attained around  $\epsilon$  Eri. The result is shown in Figures 3.89 and 3.90. We note here that our  $L'$  data for  $\epsilon$  Eri are only average, and our  $M$  band data, though above average, are not exceptional. Since 61 Cyg A is a fainter star than  $\epsilon$  Eri, it would be easy to acquire data of this quality or better around it using Clio.

In this comparison, unlike the one using our real 61 Cyg A data, our  $M$  band observations show sensitivity to considerably lower mass planets than the Lafrenière et al. (2007b) results for a wide range of separations from the star. As expected, since they are taken from a brighter star, the  $L'$  sensitivity is not as good close-in. It would be much better for a real, above average quality  $L'$  data set on 61 Cyg A.

For a final sensitivity comparison we selected the star around which we obtained sensitivity to the faintest magnitude objects. The star is GJ 450, and the excellent sensitivity is due to an unusually long 1.5 hour exposure of a relatively faint ( $H = 5.83$ ) star under very good conditions of seeing and transparency. The star is relatively nearby, at 8.1pc, and old, at 1 Gyr. Since it is about the same brightness as HD 96064, we have used the same comparison stars from the Biller et al. (2007) and Masciadri et al. (2005) surveys that we used for the HD 96064

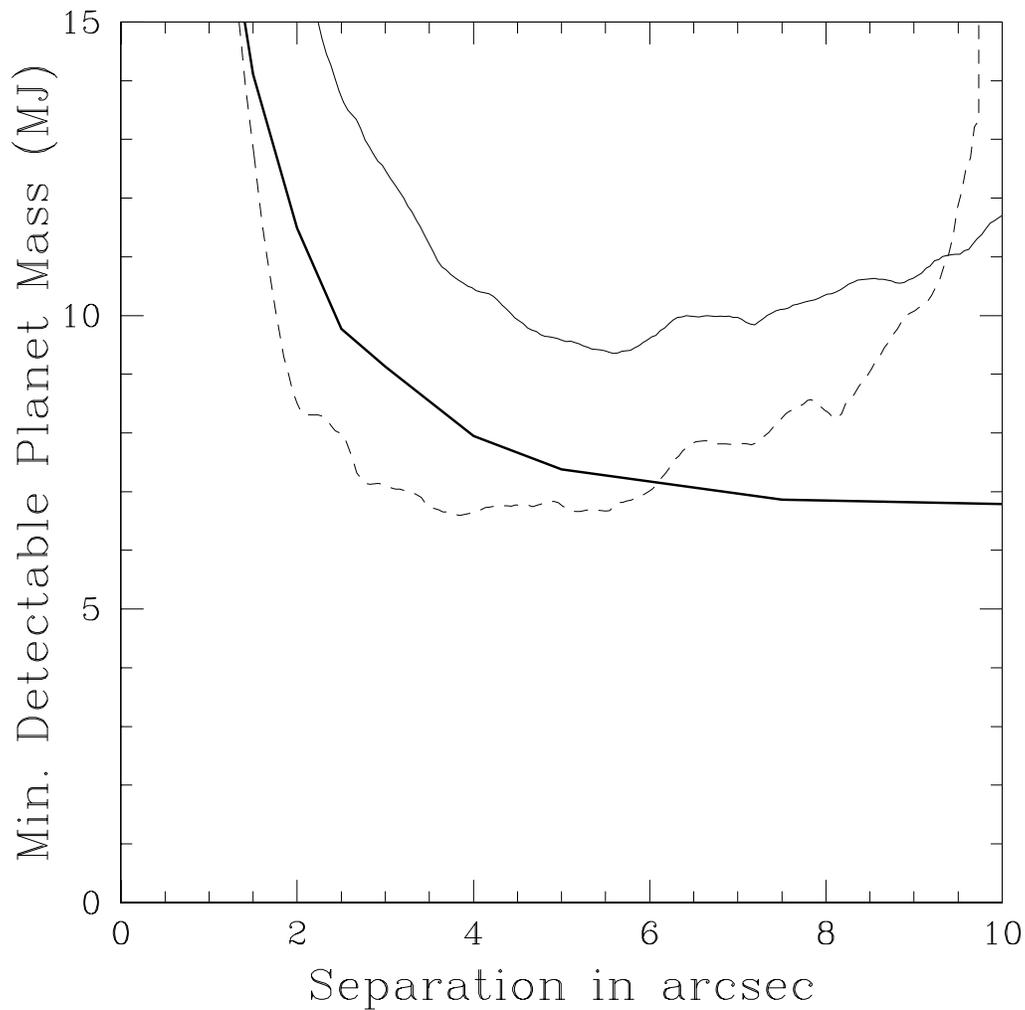


Figure 3.89 Min detectable planet mass vs separation for 61 Cyg A ( $H = 2.47$ ;  $L' = 2.25$ ; distance 3.46 pc, age 2 Gyr), based on the Burrows et al. (2003) models, assuming we had attained the same magnitude sensitivity around this star that we in fact obtained around  $\epsilon$  Eri. The heavy continuous line is for the Lafrenière et al. (2007b) observation, and the light continuous and dashed lines are for our  $L'$  and  $M$  band observations, respectively.

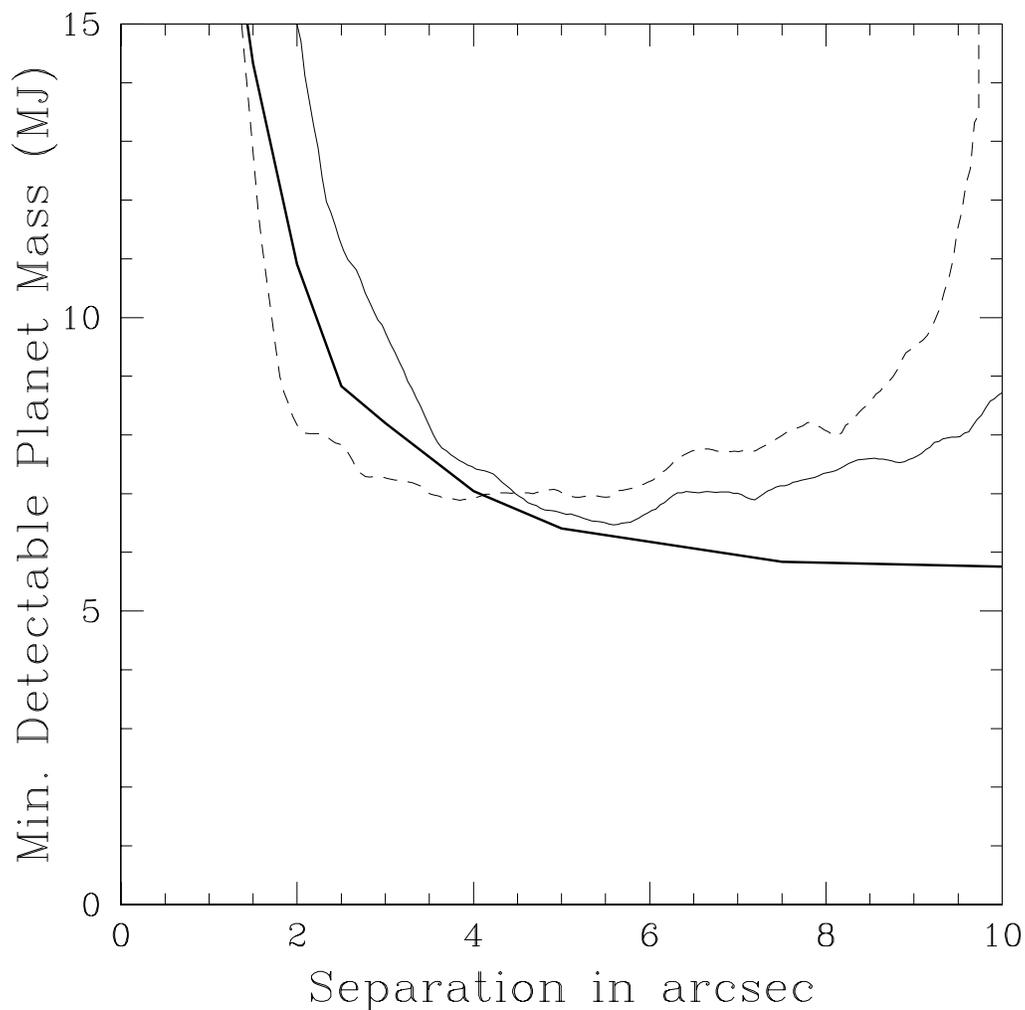


Figure 3.90 Min detectable planet mass vs separation for 61 Cyg A ( $H = 2.47$ ;  $L' = 2.25$ ; distance 3.46 pc, age 2 Gyr), based on the Baraffe et al. (2003) models, assuming we had attained the same magnitude sensitivity around this star that we in fact obtained around  $\epsilon$  Eri. The heavy continuous line is for the Lafrenière et al. (2007b) observation, and the light continuous and dashed lines are for our  $L'$  and  $M$  band observations, respectively.

plots. For Lafrenière et al. (2007b) we used HD 96064 itself, since they did not observe GJ 450.

We note that even though GJ 450 is quite faint, our very deep observations reached a background-limited sensitivity to planets only 1.75 times more massive than the Lafrenière et al. (2007b) limit from the Burrows et al. (2003) models, or 1.5 times more massive for the Baraffe et al. (2003) models. This is in contrast to factors nearer 3 or higher for HD 96064 and GJ 117. Our observation attains better sensitivity than the Masciadri et al. (2005) results for a wide range of separations, and better sensitivity than Biller et al. (2007) beyond about 1.4 arcsec in the Burrows et al. (2003) models or 1.1 arcsec in the Baraffe et al. (2003) models. Note how the Masciadri et al. (2005)  $H$  band observation is sensitive to considerably lower mass planets than their  $Ks$  result under either model set for this old, nearby star, in sharp contrast to the HD 96064 result. This is a manifestation of the closing of the  $Ks$  band and the anomalous brightening of the  $H$  band predicted by the models as planets age and cool.

The question of when the  $L'$  or  $M$  band background limited sensitivity actually becomes better than  $H$  is an interesting one. We explore it Table 3.9, taking a typical background  $H$  band sensitivity to be  $H = 23.0$  from Lafrenière et al. (2007b), typical  $L'$  sensitivity to be  $L' = 16.5$  (more typical of the Kasper et al. (2007) survey with the 8 meter VLT than our 6.5 meter work), and the typical  $M$  band sensitivity to be  $M = 13.5$  (again, scaling up our results somewhat to compare with an 8-meter telescope.)

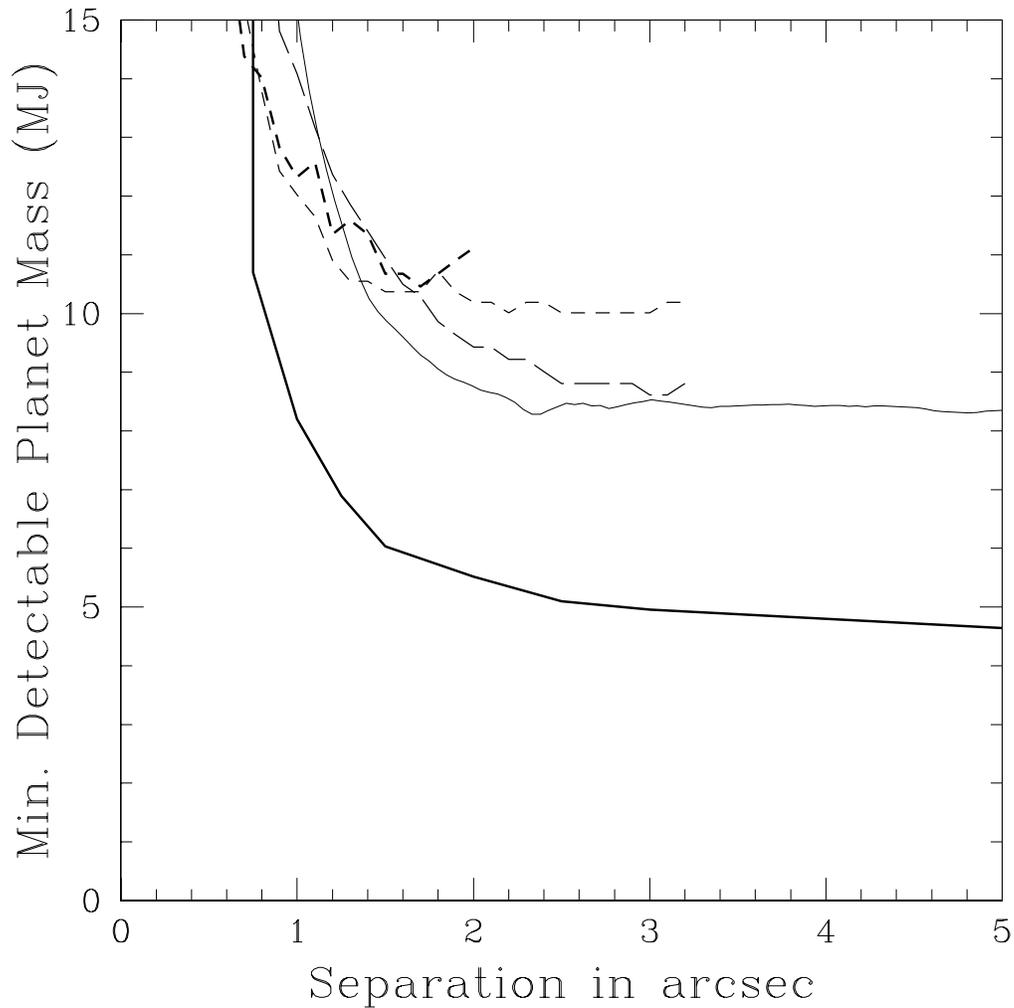


Figure 3.91 Minimum detectable planet mass vs separation for GJ 450 ( $H = 5.83$ ;  $L' = 5.40$ ; distance 8.1 pc, age 1 Gyr), based on the Burrows et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observations, the heavy short-dashed line for Biller et al. (2007), the light long-dashed and short-dashed lines, respectively, for  $H$  and  $Ks$  band data from Masciadri et al. (2005), and the light continuous line for our own  $L'$  observation.

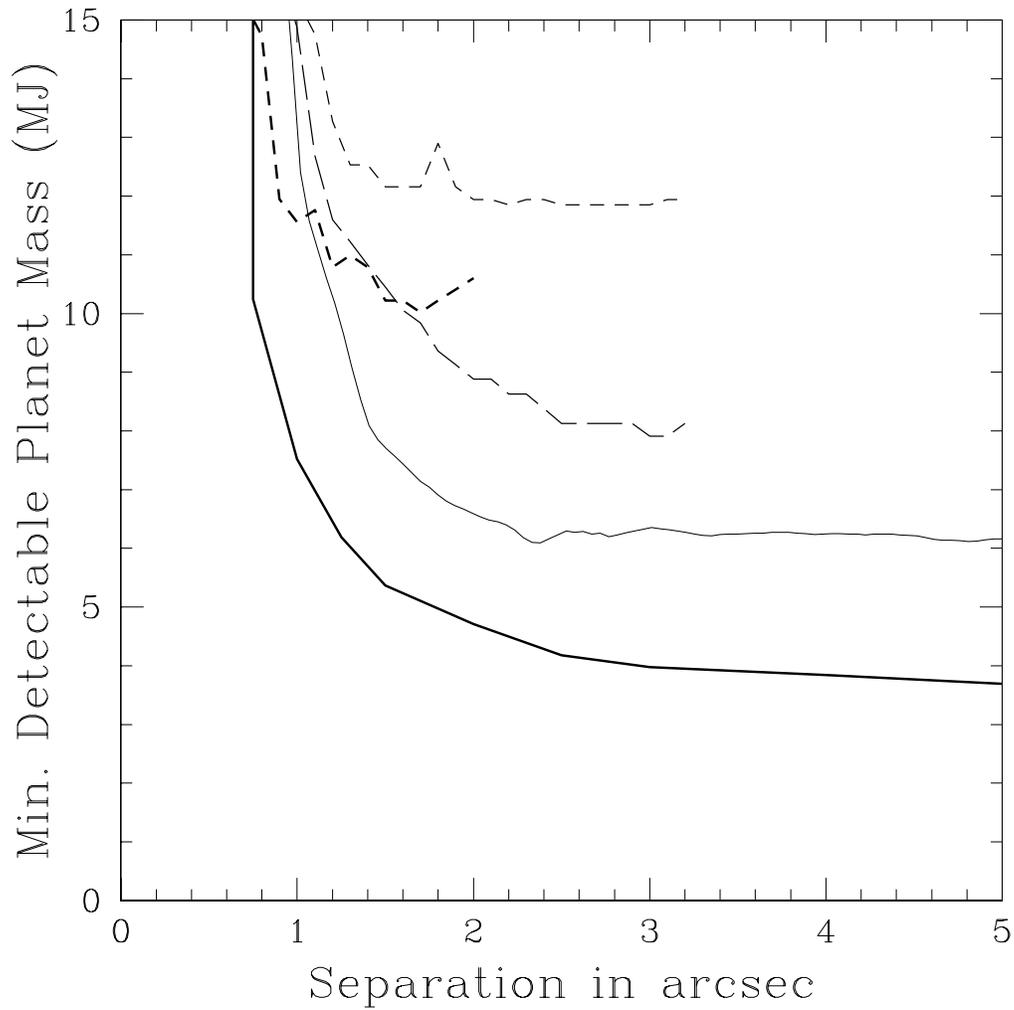


Figure 3.92 Minimum detectable planet mass vs separation for GJ 450 ( $H = 5.83$ ;  $L' = 5.40$ ; distance 8.1 pc, age 1 Gyr), based on the Baraffe et al. (2003) models. The heavy continuous line is for the Lafrenière et al. (2007b) observations, the heavy short-dashed line for Biller et al. (2007), the light long-dashed and short-dashed lines, respectively, for  $H$  and  $Ks$  band data from Masciadri et al. (2005), and the light continuous line for our own  $L'$  observation.

Table 3.16. Background-limited Planet Mass Sensitivities in Different Bands

Distance	Age	Model Set	<i>H</i> band Mass Limit	<i>L'</i> band Mass Limit	<i>M</i> band Mass Limit
10 pc	1.0 Gyr	Burrows	4.93	9.26	13.8
8 pc	1.0 Gyr	Burrows	4.67	7.88	10.7
6 pc	1.0 Gyr	Burrows	4.34	6.48	6.79
5 pc	1.0 Gyr	Burrows	4.13	5.80	5.42
4 pc	1.0 Gyr	Burrows	3.92	5.13	4.38
3 pc	1.0 Gyr	Burrows	3.71	4.51	3.49
2 pc	1.0 Gyr	Burrows	3.41	3.80	2.63
1.5 pc	1.0 Gyr	Burrows	3.20	3.51	2.15
1 pc	1.0 Gyr	Burrows	2.93	2.91	1.68
5 pc	5.0 Gyr	Burrows	10.0	13.8	14.9
4 pc	5.0 Gyr	Burrows	9.32	12.3	12.0
3 pc	5.0 Gyr	Burrows	8.58	10.8	9.30
2 pc	5.0 Gyr	Burrows	7.82	9.00	6.87
1.5 pc	5.0 Gyr	Burrows	7.47	8.06	5.98
1 pc	5.0 Gyr	Burrows	6.99	7.07	5.03
10 pc	1.0 Gyr	Barrafe	3.96	7.11	13.4
8 pc	1.0 Gyr	Barrafe	3.72	5.78	10.5
6 pc	1.0 Gyr	Barrafe	3.42	4.57	7.58
5 pc	1.0 Gyr	Barrafe	3.23	3.90	6.02

For a given pair of wavelengths, we expect there will be a break point in distance, inside of which the longer wavelength will provide background-limited sensitivity to lower mass objects. From Table 3.9, we see that according to the Burrows et al. (2003) models for objects of 1 Gyr age  $L'$  does better than  $H$  band inward of a point between 1.5 and 1.0 pc. For the 5 Gyr objects using the same models the breakpoint is within 1 pc. As both components of the nearest star system ( $\alpha$  Cen and its very long-period companion Proxima Cen) are both more than 1.0pc distant, the Burrows et al. (2003) models suggest that with current telescopes  $L'$  observations do not have better background-limited sensitivity than  $H$  band observations at any interesting distance. For the Baraffe et al. (2003) models, however, the  $L'$  vs  $H$  breakpoint occurs between 3 and 2 pc for objects of 1 Gyr age and between 4 and 3pc for objects of 5 Gyr age. According the the Burrows et al. (2003) models, the  $M$  vs  $L'$  breakpoint occurs between 6 and 5 pc for objects of 1 Gyr age, and between 5 and 4 pc for objects of 5 Gyr age. The Baraffe et al. (2003) models put this breakpoint between 4 and 3 pc for 1 Gyr-old planets, but show it never arriving at all for 5 Gyr planets. Finally, the Burrows et al. (2003) models put the  $M$  vs  $H$  breakpoint between 4 and 3 pc for 1 Gyr planets and between 3 and 2 pc for 5 Gyr objects. The Baraffe et al. (2003) models put this breakpoint between 1.5 and 1.0 pc for 1 Gyr planets and between 2 and 1.5 c for 5 Gyr objects. Note that in general the breakpoint is predicted closer in for older objects, but that this last example of  $M$  vs  $H$  under the Baraffe et al. (2003) models is an exception.

**The conclusion of this table is that an  $L'$  or  $M$  band survey of even very faint stars within 3 or 4 pc would attain better sensitivity limits, even at large separations from the stars, than a similar survey conducted in the  $H$  band regime. Note the interesting disagreement between the model sets: according to Burrows**

Table 3.16—Continued

Distance	Age	Model Set	$H$ band Mass Limit	$L'$ band Mass Limit	$M$ band Mass Limit
4 pc	1.0 Gyr	Barrafe	3.00	3.38	4.65
3 pc	1.0 Gyr	Barrafe	2.78	3.80	3.59
2 pc	1.0 Gyr	Barrafe	2.46	2.14	2.74
1.5 pc	1.0 Gyr	Barrafe	2.24	1.80	2.25
1 pc	1.0 Gyr	Barrafe	1.95	1.41	1.72
5 pc	5.0 Gyr	Barrafe	8.35	9.65	14.7
4 pc	5.0 Gyr	Barrafe	7.88	8.33	11.1
3 pc	5.0 Gyr	Barrafe	7.34	7.09	8.55
2 pc	5.0 Gyr	Barrafe	6.63	5.69	6.81
1.5 pc	5.0 Gyr	Barrafe	6.18	4.83	5.81
1 pc	5.0 Gyr	Barrafe	5.56	3.77	4.55

Note. — We have assumed background-limited sensitivities of  $H = 23.0$ ,  $L' = 16.5$ , and  $M = 13.5$ . The point of the table is to illustrate how as we move to smaller distances a breakpoint eventually occurs inward of which a longer wavelength is superior to a shorter one. The breakpoint occurs at different distances for different wavelength pairs, ages, and model sets.

et al. (2003), such a survey should be done in the  $M$  band, while according to Baraffe et al. (2003) the  $L'$  band is considerably better. Sensitivity to planetary mass objects would be attained even around 5 Gyr-old stars. We note, however, that there are only 22 known stars within 4 pc (Cox, 2000), and that a large fraction of them are  $M$  stars, which are thought to be less likely to harbor massive planets (Johnson et al., 2007).

Our final conclusion is that **based on current theoretical models planet imaging observations in the  $H$  band region are considerably more sensitive than  $L'$  or  $M$  band observations except for very nearby stars.** In this context,  $L'$  and  $M$  band surveys are useful mainly in the context of diversifying the investment of planet imaging effort in case unexpected atmospheric chemistry, clouds, or formation/evolution scenarios cause planets to be fainter in the  $H$  band region than expected.

As a counterpoint to this, we note that **our observations have been performed solely with a 6.5 meter telescope, while the bulk of observations for the other surveys have used 8 meter telescopes.** Even so, **our  $L'$  and  $M$  band observations of  $\epsilon$  Eri attain higher sensitivity than those of any other method, even SDI.** If we had obtained an average or above-average data set on 61 Cyg A, we could say the same of it also, and there are a handful of other bright, very nearby stars for which  $L'$  and  $M$  band planet search observations would attain considerably better sensitivity than  $H$  band regime observations.

As larger telescopes are built and the sensitivity in every wavelength regime increases, we will be able to detect much lower mass, redder objects in very nearby star systems, and the  $L'$  and  $M$  bands will become more and more useful. The breakpoints in Table 3.9 will move outward. Although they appear to be superior only for a very small sample of stars at present, **the bands we have**

**pioneered will become increasingly important in the future, and it may well at these wavelengths that the next generation of giant telescopes is used to image mature planetary systems around the nearest stars.** It is also possible that this will be accomplished first from space. JWST will have far better background-limited sensitivity at  $L'$ ,  $M$  band, and longer wavelengths than even planned 20-30 meter ground based telescopes such as the GMT. How good JWST will do in the contrast-limited regime close to bright stars remains to be seen, but it will certainly not be competitive with GMT for close-in planets. Space-based coronagraphs may also be built that will be capable of imaging mature planets by reflected starlight in the optical, but this will require exceptional coronagraphic performance.

### 3.10 Confirmed and Suspected Sources

#### 3.10.1 Our Source Detection Methods

We use both manual and automatic methods for detecting candidate sources in our data. Our automatic methods for source detection are exactly analagous to our sensitivity estimation methods.

We use the same aperture photometry and PSF amplitude maps that are used to make Method 2 and Method 3 sensitivity images, as described above. To make the sensitivity maps we calculated the RMS of these images within disks and arcs and fed the results into the sensitivity image. For source detection, we have used the exact same regions, but have punched roughly 5-pixel radius holes in them so that the RMS values will not be biased if a source happens to be located on the pixel under consideration.

Thus, for every pixel in the image that is sufficiently far from the star, we calculate the RMS within an annulus centered on that pixel with an inner radius of 6

pixels and an outer radius of 8 pixels, inclusive, on both the aperture photometry and the PSF amplitude images. If the central pixel is above the annular mean by more than a threshold times the RMS, a source is reported to be detected at that pixel. For pixels within a 60 pixel radius of the central star we calculate the RMS not within an annular region centered on the pixel, but along a 45-pixel arc at constant radius from the central star. The arc is centered on the pixel under consideration, but a region a bit wider than one PSF is ‘punched out’ of the arc at its center so that if there is a real source on the pixel, it will not bias the RMS. If the pixel value is above the arc mean by more than a threshold times the arc RMS, a detection is reported at that pixel.

Thus, we detect sources on every image based on both the aperture photometry map, (ie the second-to-last step image in the Method 2 sensitivity calculation), and the PSF amplitude map (ie the second-to-last step image in the Method 3 sensitivity calculation). Sometimes a source is detected using only one method, and sometimes it is detected using both. Bright real sources are always detected using both.

While in our sensitivity calculations we have set conservative  $10\sigma$  limits, in our automatic source detection we set much lower thresholds so that we would detect the faintest real sources possible. We developed a powerful technique to prevent our low thresholds from resulting in too many false positives: **We stacked each data set in four different ways, and required that any source must be automatically detected at a consistent location on all four images to be flagged as a candidate.** The first stacking method was the 20% creeping mean stack of all the data that we have discussed above in the data processing section. In addition to this **we made 20% creeping mean rejection stacks of the first and second halves of the data.** Essentially all ghosts and other artifacts rotated between the

first and second data halves, so they were extremely useful for sorting spurious detection from real. Often, an apparent faint source in the master stack would appear as a brilliant ghost in one half image and as empty sky in the other. The fourth image we made was a 50% rejection creeping mean stack of the entire data set. This image was very clean of artifacts, but had higher background noise than the 20% rejection master stack because so much data had been thrown away. We set detection thresholds of  $4.5\sigma$  for the master stack,  $3.0\sigma$  for each half-stack image, and  $3.5\sigma$  for the 50% rejection full stack. To be reported as a candidate by the automatic source detection algorithm, a source had to be detected above the appropriate threshold on each of the four images, at a location consistent within 2 pixels.

We made the four types of images for every processing method used on every star. We ran our automatic source detection algorithm independently on each set of four images, creating a separate list of automatically detected candidate sources for the images processed with each method. Thus, in a typical data set where the 'a', 'b', 'd', 'e', 'x', and 'y' processing methods had all been used, there would be six separate candidate lists. Each list would include all sources detected on the given image set using the aperture photometry method, and all sources detected using the PSF amplitude method. For clean data sets with good rotation on a faint star, one or fewer candidates might be detected on the images from each processing method. For bright-star data with insufficient rotation, especially binary-star data, there might be a dozen spurious detections on the image set from each method — one hundred or more for all methods combined. Many of these would be multiple detections of the same sources on images processed by different methods, or sources detected using both the aperture photometry and psf amplitude images from a given processing method. Bright real sources

or extended artifacts could result in multiple detections on a single data set using a single method, although we forbade the automatic algorithm to detect two sources separated by less than 2 pixels. In such cases it was required to choose the brighter one and delete the other from its reporting list.

We painstakingly checked out all of these automatic detections manually and classified them as either confirmed or unconfirmed. We brought a number of techniques to bear in carrying out this classification. To be confirmed as real, a source must look like a PSF, not an extended ghost that fooled the automatic algorithm. It must appear on the images processed with more than one method, unless there is a clear reason why it does not. It must appear with a consistent location, morphology, and brightness, within reasonable expectations set by the noise level, on the two half-stacks of each data set. In a number of cases where a given source was interesting but hard to confirm, we created additional 'custom' stacks of subsets of the data to obtain artifact-free sky around the candidate object, or to test a theory of its spurious origin.

For unconfirmed sources, the vast majority of which are certainly false positives, we assigned a reality likelihood ranging from 'certainly unreal' (ie, an obvious ghost or other type of artifact) and 'almost certainly unreal' up through 'possibly real,' 'provocative,' and 'very provocative.' We made careful log files from the investigation of each data set. These include the raw output lists of the automatic detection code, comments on each detection from the manual investigation, and then combined listings of all the detections of a given source or artifact, together with the final verdict on its nature and an explanation of why this conclusion was reached.

Since our automatic source detection algorithm had to detect objects on both halves of the data, it could not report even obvious sources for which there was

valid data only in the first or second half of the data set. Because of parallax rotation, this scenario was not unlikely. It was therefore necessary to carry out a purely manual examination of each data set in addition to running the automatic source detection and manually following up its candidates. This would have been desirable in any case, because of the incredible pattern-recognizing power of the human eye and brain. A human can learn to recognize a PSF much more easily than a computer can. In fact, in the blind sensitivity tests there were either zero or very few cases where the automatic source detection flagged a real source that was not also noticed during the purely manual investigation. The human observer also investigated far fewer spurious sources. The purely manual phase of each investigation, like the manual followup to the automatic phase, was carefully logged.

The reader may well imagine that with 24 images (four: the master stack, first half, second half, and 50% rejection full stack, from each of typically six processing methods, 'a', 'b', 'd', 'e', 'x', and 'y') to be examined carefully by multiple methods for each data set, with the investigation of some sources requiring the construction of additional stacked images, the source detection was a time-consuming process. Indeed, at the beginning it averaged something like one human-day per star system. Image sets from the blind sensitivity tests, with their larger number of sources to investigate, took more than a day. Toward the end of the source detection analysis, with increasing experience and time pressure, the author sometimes managed to analyze two or even three stars in a single day. The latter feat was possible only if one or more of the stars was unusually clear from spurious sources. Such clean images sets were the exception rather than the rule. The very time consuming nature of the source detection is the reason more blind sensitivity tests were not performed. The blind sensitivity tests did illustrate that

the large amount of time spent analyzing each data set was not wasted, as more real sources would continue to be flagged and confirmed throughout the process.

### 3.10.2 Confirmed and Suspected Sources

Eleven sources were confirmed as real. Of these, one is newly discovered low mass stellar companion to HD 133002, one is the previously known binary brown dwarf companion to GJ 564 (Potter et al., 2003), and the other 9 are background stars. This has been confirmed in every case. No planets, and no previously unknown brown dwarfs, were confirmed. Note that far fewer background stars are detected in an  $L'$  survey such as ours than at shorter wavelengths. Stars have IR colors close to zero, whereas planets have very red IR colors; thus a survey such as ours can obtain good sensitivity to planets while remaining blind to all but the brightest stars. We found 9 background stars; Lafrenière et al. (2007b), operating in the  $H$  band regime, found more than 300. More followup is thus needed for an  $H$ -regime survey. We note that because planets have very red  $K_s - L'$  and  $K_s - M$  colors, while stars have colors near zero, we can distinguish between a background star and a planet with a  $K_s$  band followup image taken immediately after the object is noticed in the  $L'$  data. There is no need to wait for a proper motion confirmation.

In Table 3.10.2 we list our detected sources, and in Figures 3.93 through 3.101 we present images showing each detected source.

The brighter of the two sources near BD+20 1790 and the sources near HD 96064A, BD+60 1417, and GJ 3860 were confirmed to be background stars based on their blue  $K_s - L'$  color in  $K_s$  followup observations. We note that the sources around HD 96064A and BD+60 1417 were independently detected in the Lafrenière et al. (2007b) survey, and confirmed to be background objects based on proper motions. The HD 96064A source looks double in our data, and was confirmed to

be so by Lafrenière et al. (2007b).

The source near GJ 354.1A is confirmed as a background star based on a previous image of it from Lowrance et al. (2005) showing it at a position consistent with a background object, not a proper motion companion. The fainter of the two sources near BD+20 1790 was found to be a background star based on a detection of it in an archival HST image showing it at a position consistent with a background object, not a proper motion companion. The two sources near 61 Cyg A and the one near 61 Cyg B were confirmed to be background stars based on detections on POSS plates from 1991 which showed them at positions consistent with background objects, not proper motion companions (see Figure 3.102). We note that proper motion had carried the 61 Cyg system almost an arcminute in the 15 years between the POSS observation and ours, which is how the companions we detected close to the stars were far enough out to clear the overexposed haloes and be detected on the POSS images.

The very bright source near GJ 564 is a close binary brown dwarf discovered by Potter et al. (2003). The similarly bright object near HD 133002 is a low-mass stellar companion to the star. We have confirmed common proper motion by followup observations several months after the initial detection. It appears to be the lone discovery of our survey. It does not seem to be of great interest since it is certainly not substellar.

Very many detections were suspected and found unconfirmable. The vast majority of these are certainly spurious. However, our blind sensitivity tests indicated that some real  $7\sigma$  sources and many real  $5\sigma$  sources would receive just such a 'suspected but not confirmed' classification from our source detection methods. Accordingly, we have provided a list of suspected sources in Table 3.10.2. The sources in this table represent all those that we classified with a reality likelihood

Table 3.17. Confirmed Sources in Our Survey

Star Name	Det. Sig.	$L'$ Mag	Sep (asec)	PA	Date Obs
BD+20 1790	$31.51\sigma$	$14.41 \pm 0.2$	8.73	$74.1^\circ$	07/01/04
BD+20 1790	...	$15.16 \pm 0.2$	6.42	$336.4^\circ$	07/01/04
GJ 354.1A	$4.93\sigma$	$16.37 \pm 0.2$	4.93	$187.3^\circ$	06/04/12
HD 96064A	$43.18\sigma$	$13.72 \pm 0.2$	5.57	$212.8^\circ$	07/01/04
BD+60 1417	$11.91\sigma$	$15.70 \pm 0.2$	1.93	$301.4^\circ$	06/06/10
GJ 3860	$19.21\sigma$	$14.53 \pm 0.2$	9.68	$144.4^\circ$	06/06/09
GJ 564	$175.68\sigma$	$10.80 \pm 0.2$	2.60	$103.0^\circ$	06/04/13
HD 133002	$246.38\sigma$	$10.92 \pm 0.2$	1.86	$118.2^\circ$	06/04/13
61 Cyg A	...	$12.43 \pm 0.2$	11.24	$227.5^\circ$	06/06/09
61 Cyg A	$32.82\sigma$	$13.05 \pm 0.2$	7.78	$83.2^\circ$	06/06/09
61 Cyg B	...	$14.04 \pm 0.2$	9.85	$145.4^\circ$	06/06/10

Note. — The detection significance column gives the highest significance with which the source was automatically detected on any image with any method. Blanks in this column imply sources that were detected only manually. The astrometry should be accurate to about 0.05 arcsec or better. The photometry of the GJ 564 companion is probably biased downward because of its binarity.

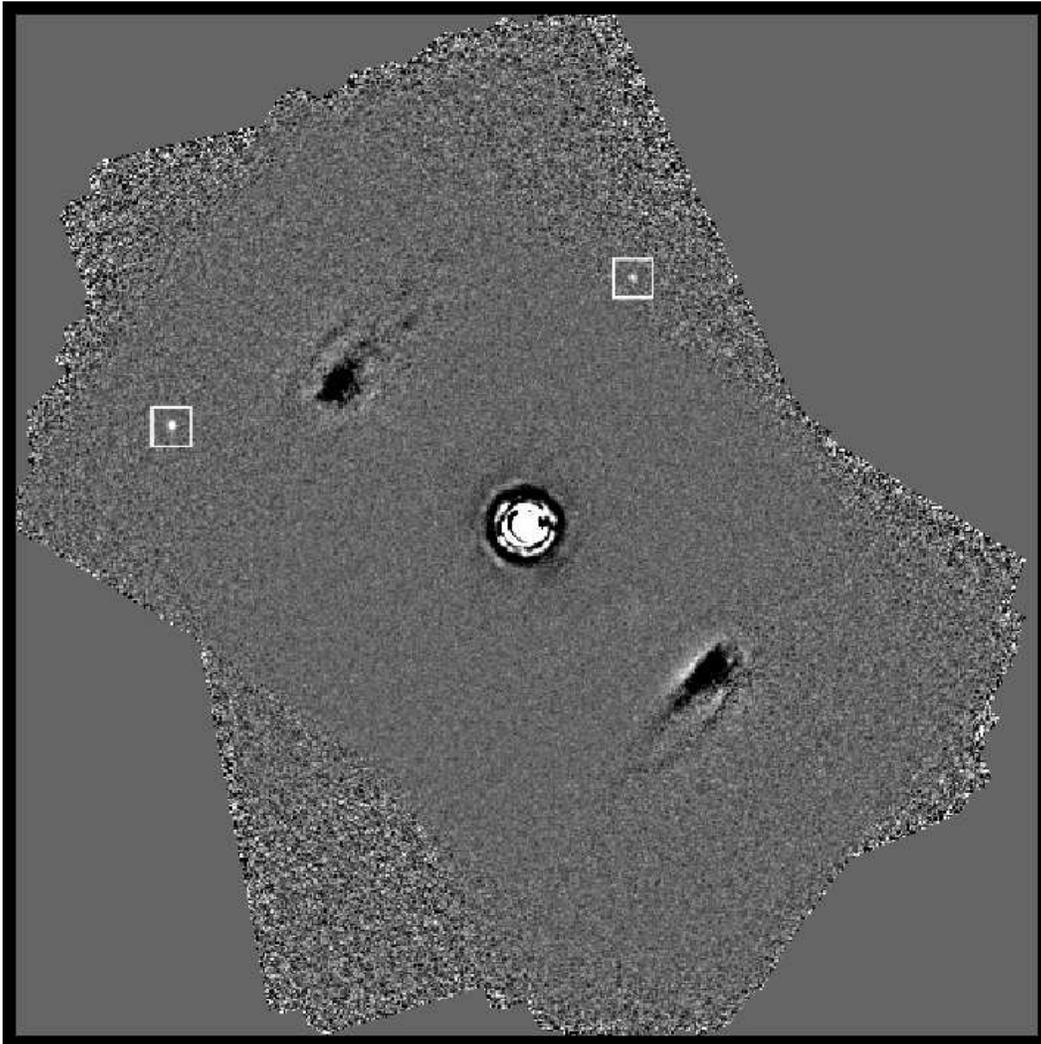


Figure 3.93 Our  $L'$  image of BD+20 1790. The two apparent companions we detected around this star are highlighted with white boxes. They turned out to be background stars.

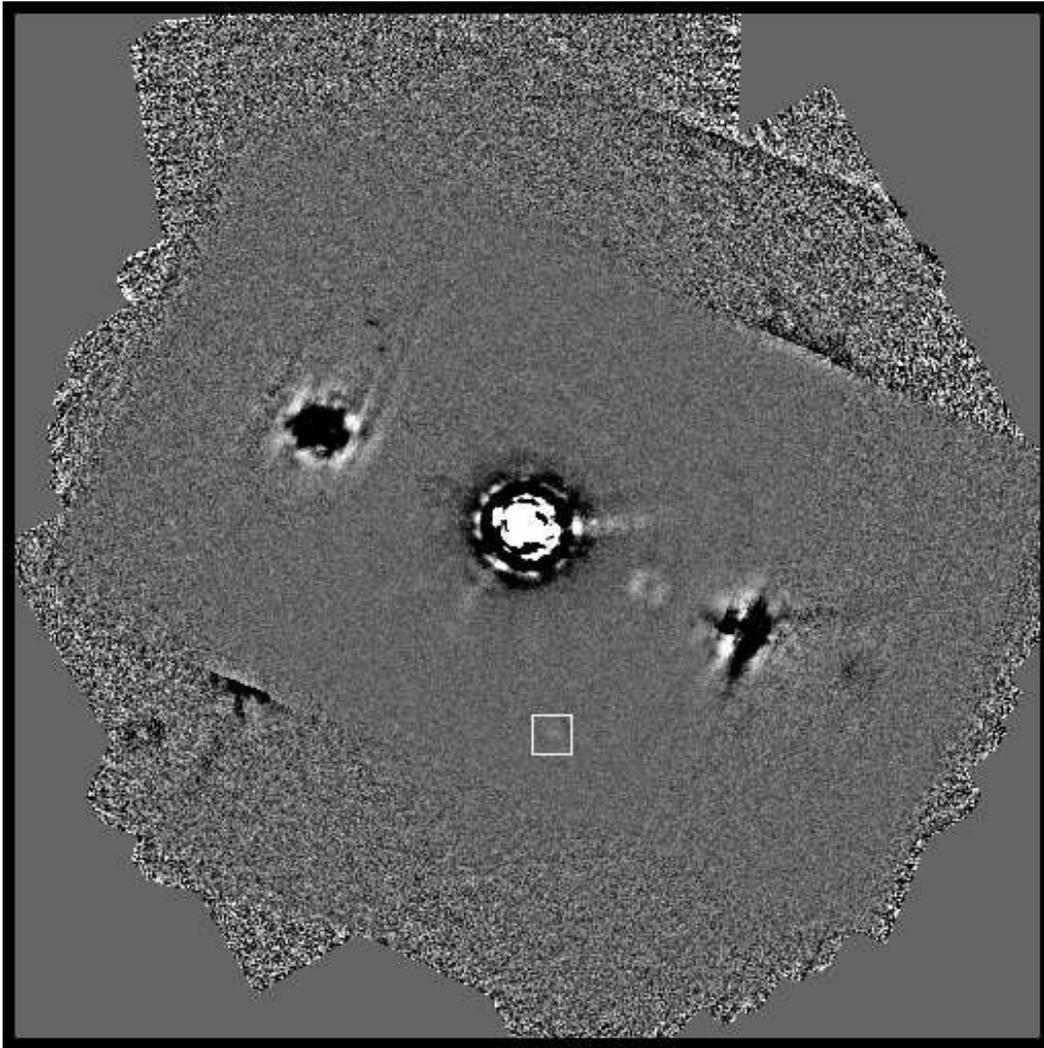


Figure 3.94 Our  $L'$  image of GJ 354.1A. The apparent companions we detected around this star is highlighted with a white box. It turned out to be a background star.

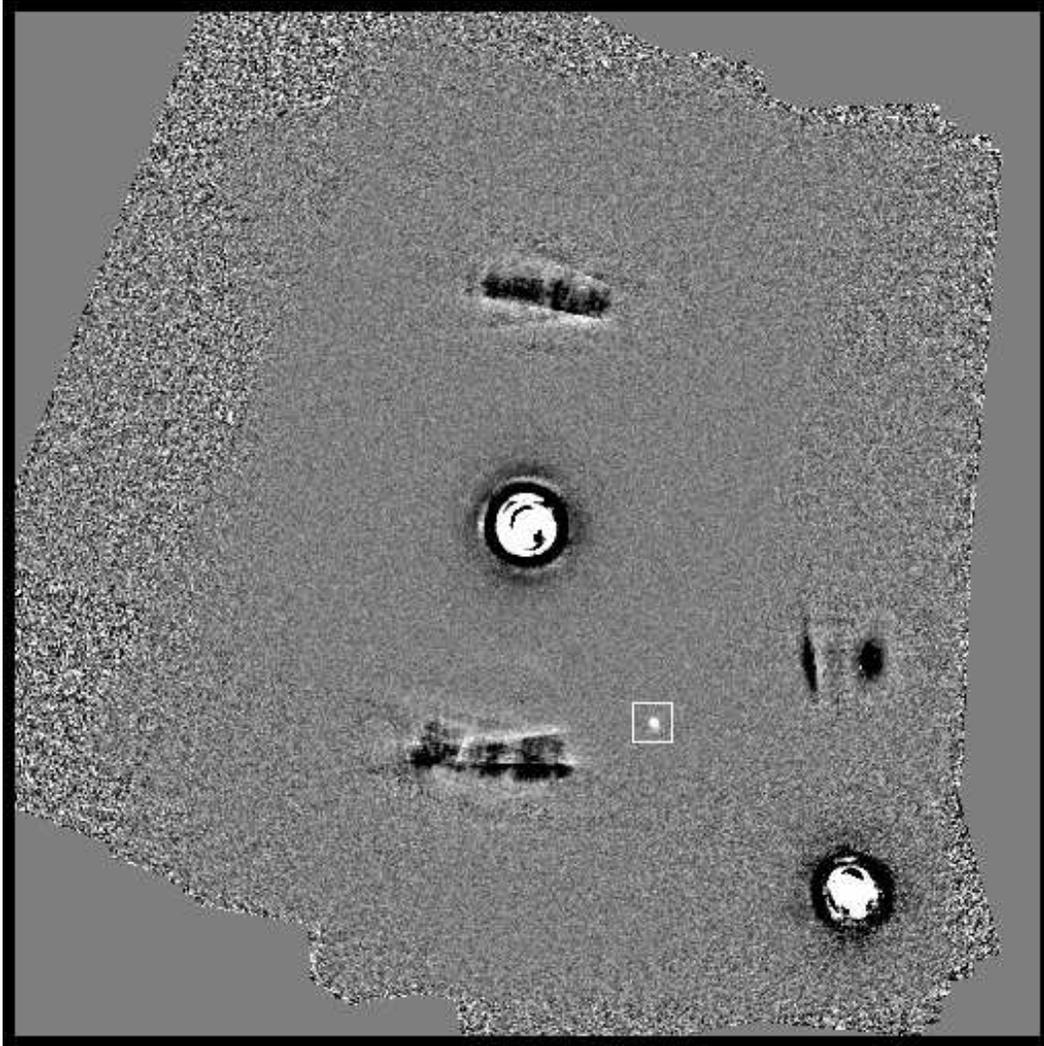


Figure 3.95 Our  $L'$  image of HD 96064. The apparent companion we detected around this star is highlighted with a white box. It turned out to be a background star.

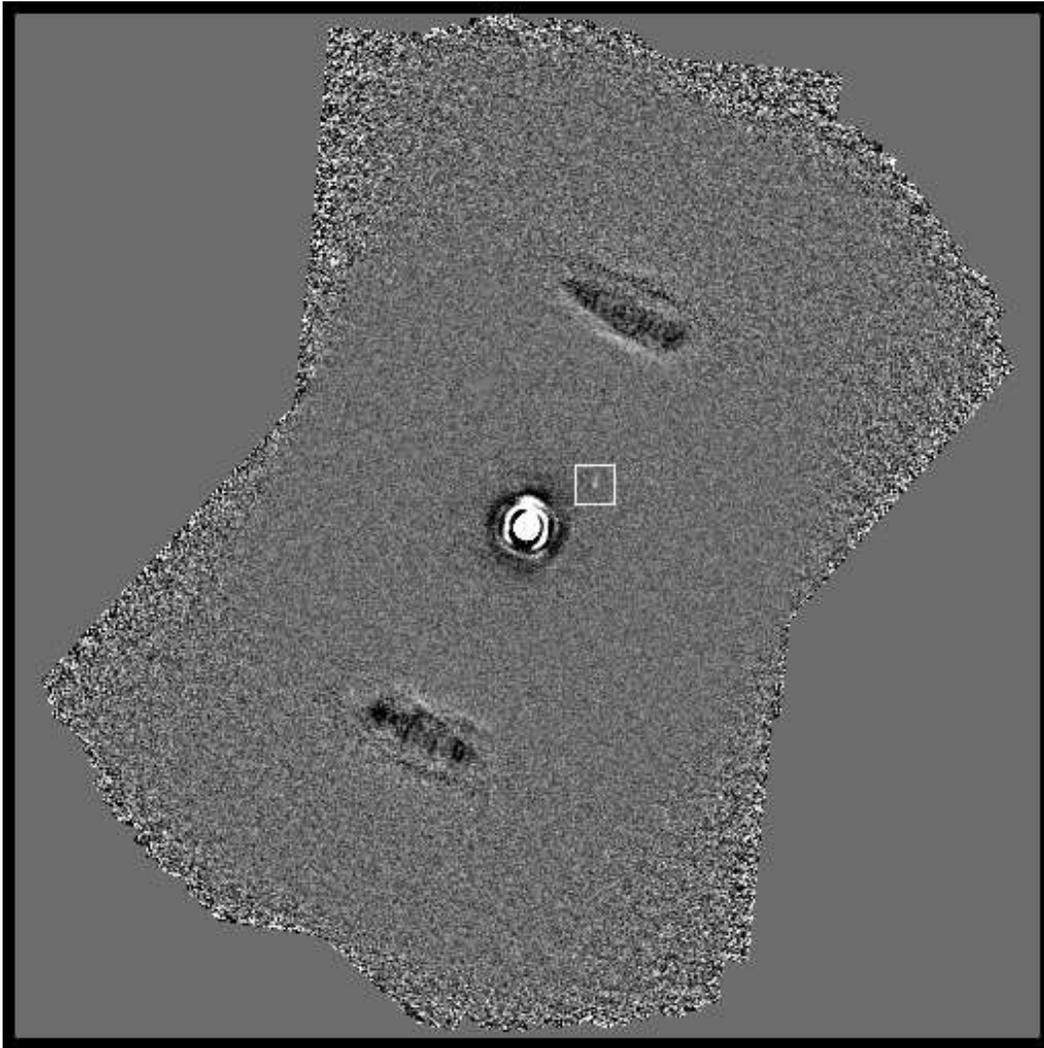


Figure 3.96 Our  $L'$  image of BD+60 1417. The apparent companion we detected around this star is highlighted with a white box. It turned out to be a background star.

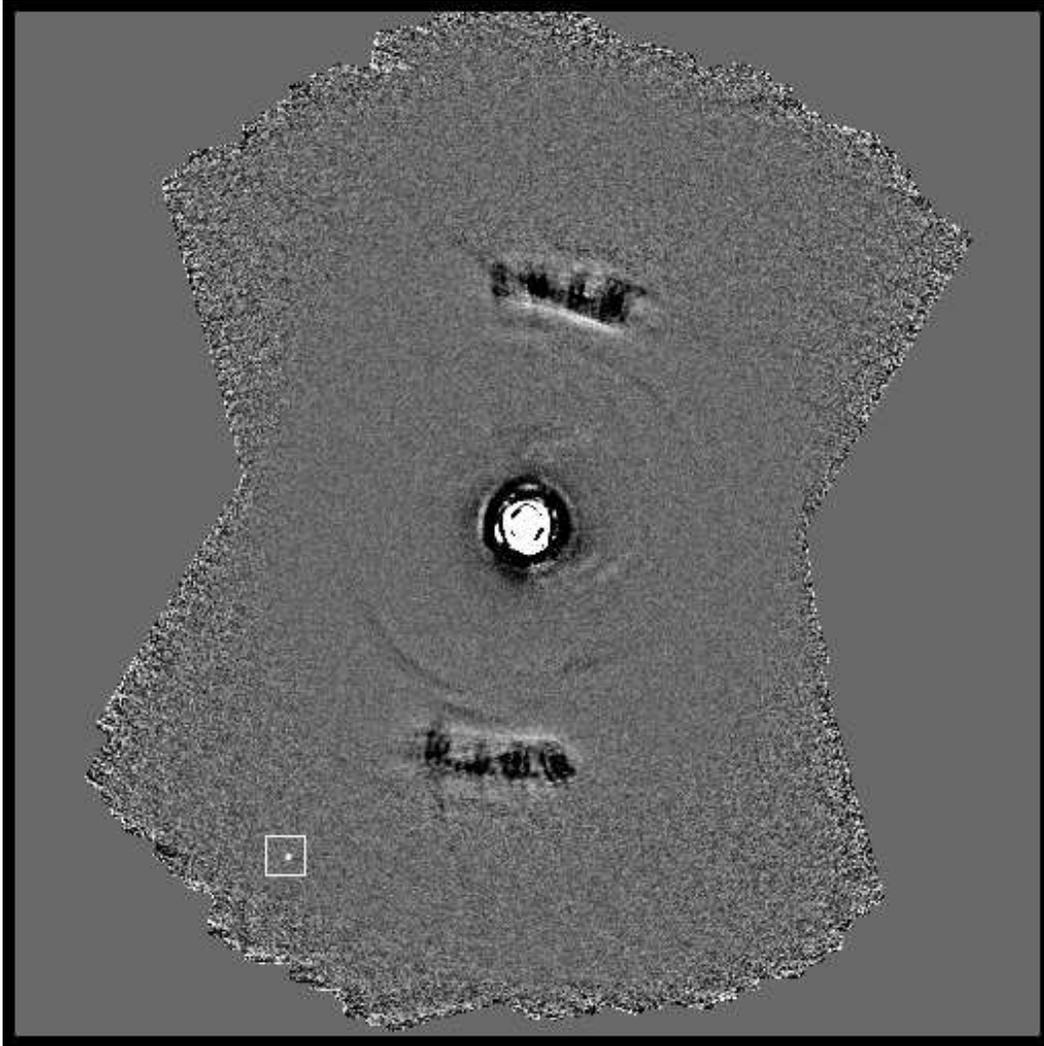


Figure 3.97 Our  $L'$  image of GJ 3860. The apparent companion we detected around this star is highlighted with a white box. It turned out to be a background star.

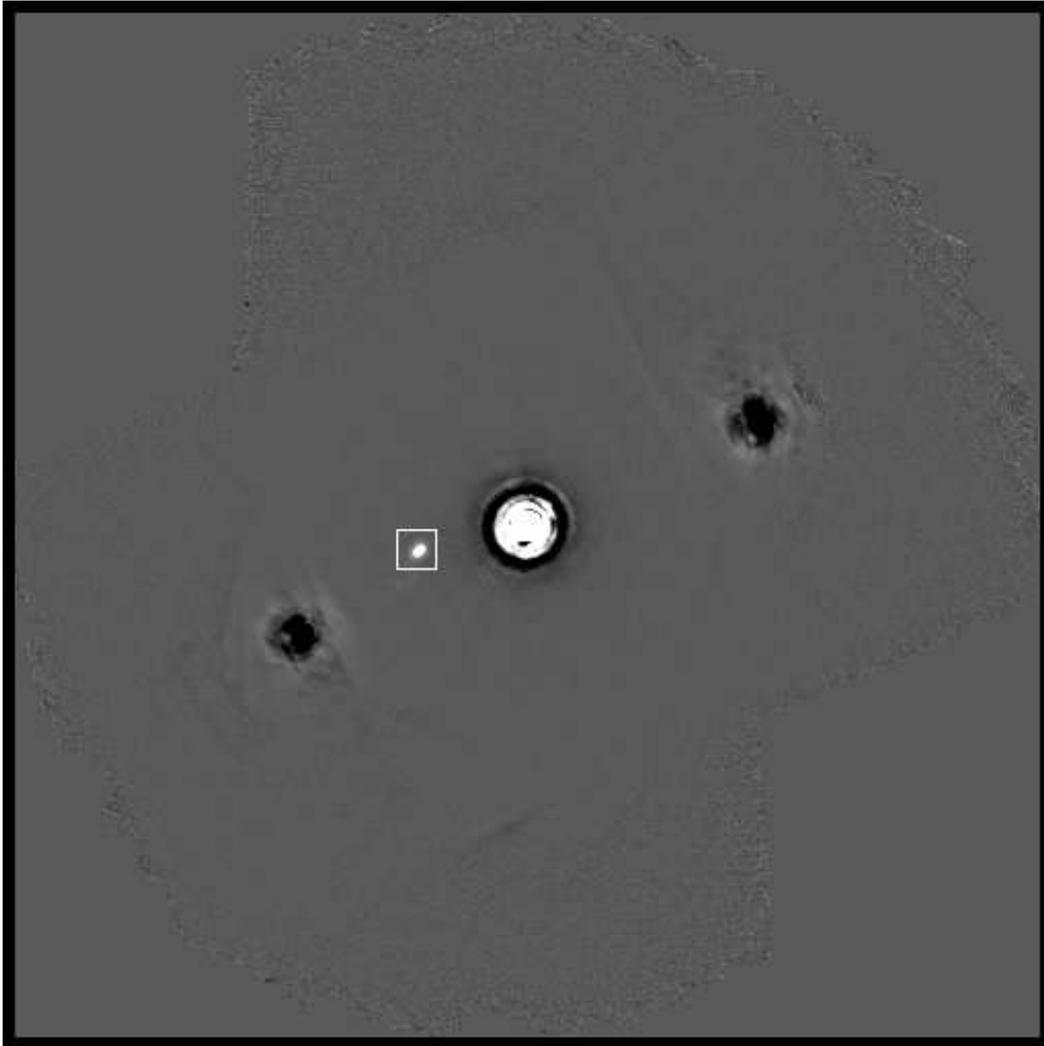


Figure 3.98 Our  $L'$  image of GJ 564. The apparent companion we detected around this star is highlighted with a white box. It is a binary brown dwarf previously discovered by Potter et al. (2003). The clear elongation in this image shows the binary is partially resolved.

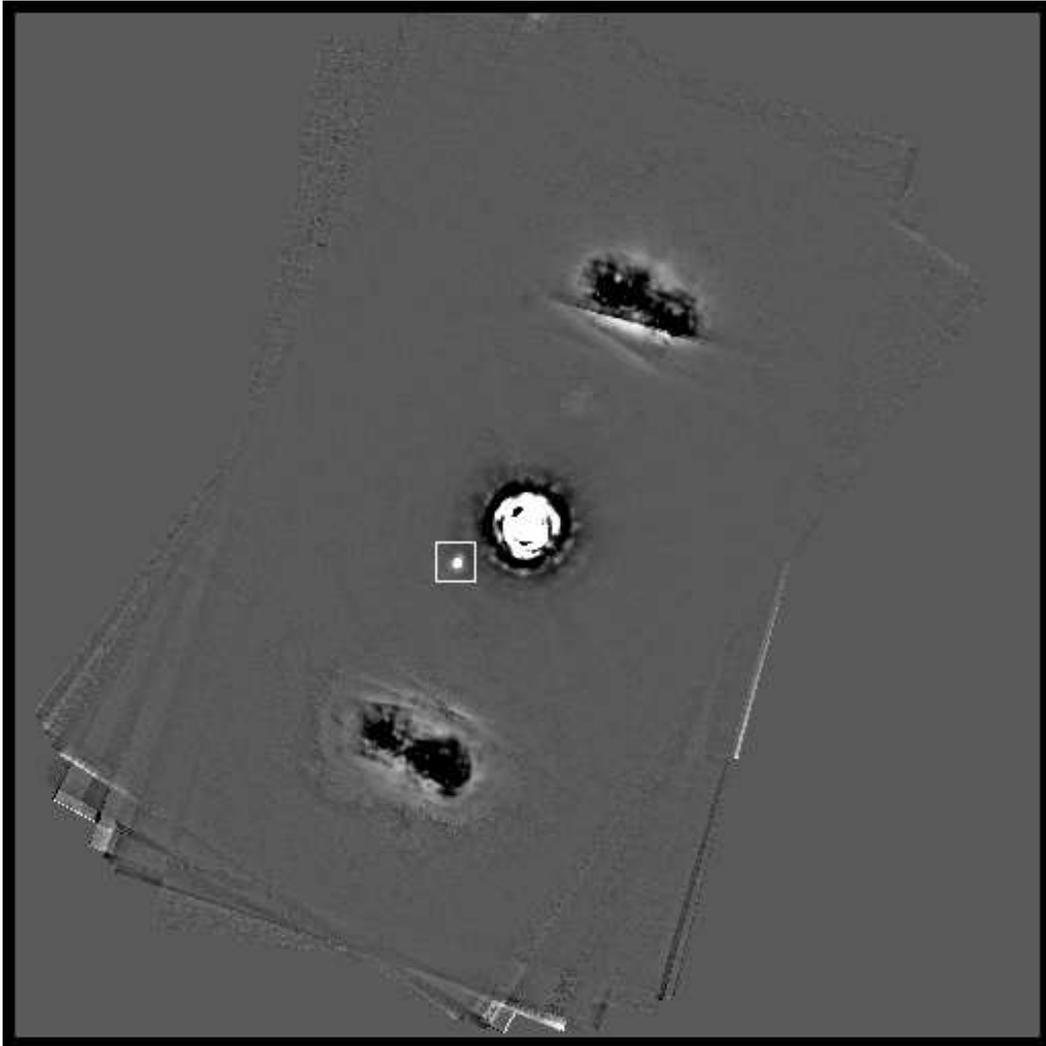


Figure 3.99 Our  $L'$  image of HD 133002. The apparent companion we detected around this star is highlighted with a white box. It is a low mass star that is a common proper motion companion of the primary, and is a discovery of our survey.

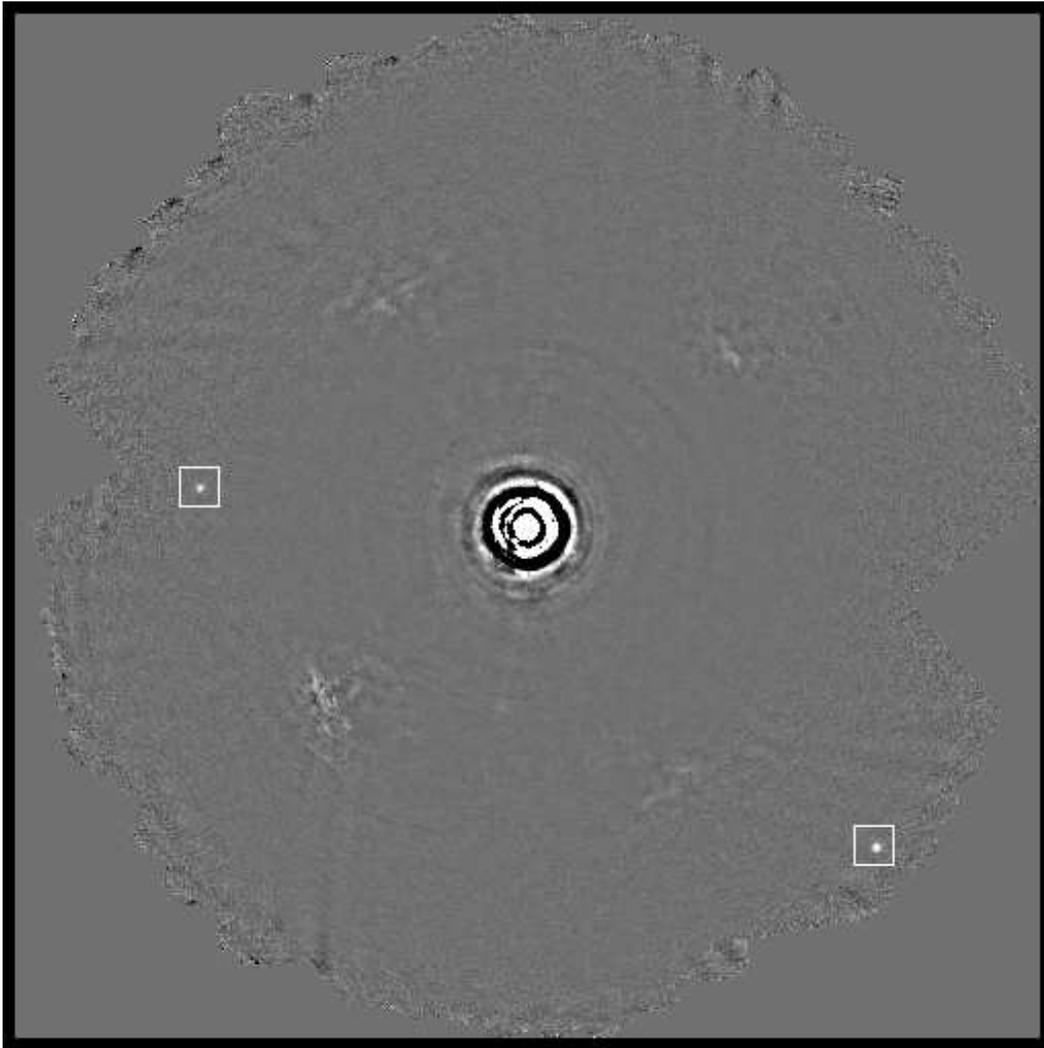


Figure 3.100 Our  $L'$  image of 61 Cyg A. The two apparent companions we detected around this star are highlighted with white boxes. They turned out to be background stars.

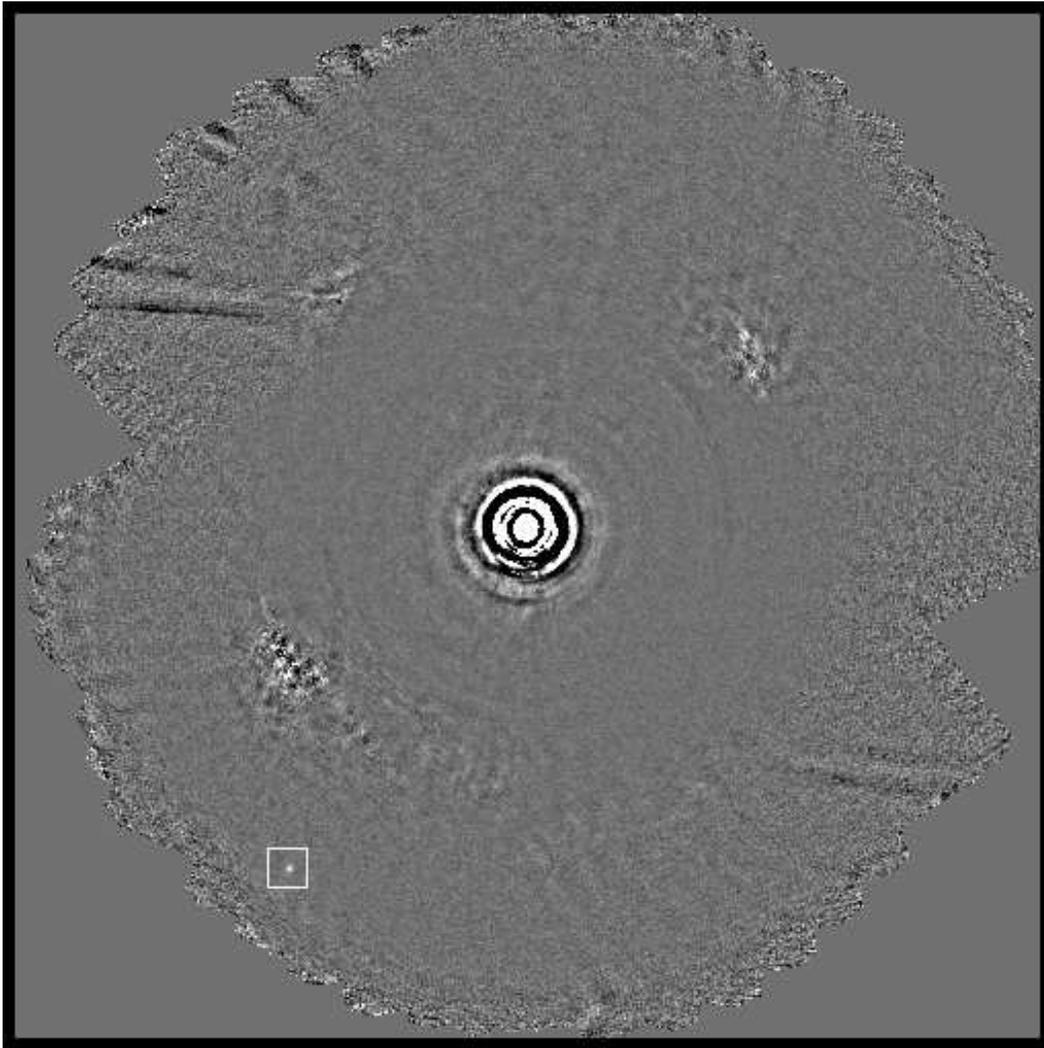


Figure 3.101 Our  $L'$  image of 61 Cyg B. The apparent companion we detected around this star is highlighted with a white box. It turned out to be a background star.

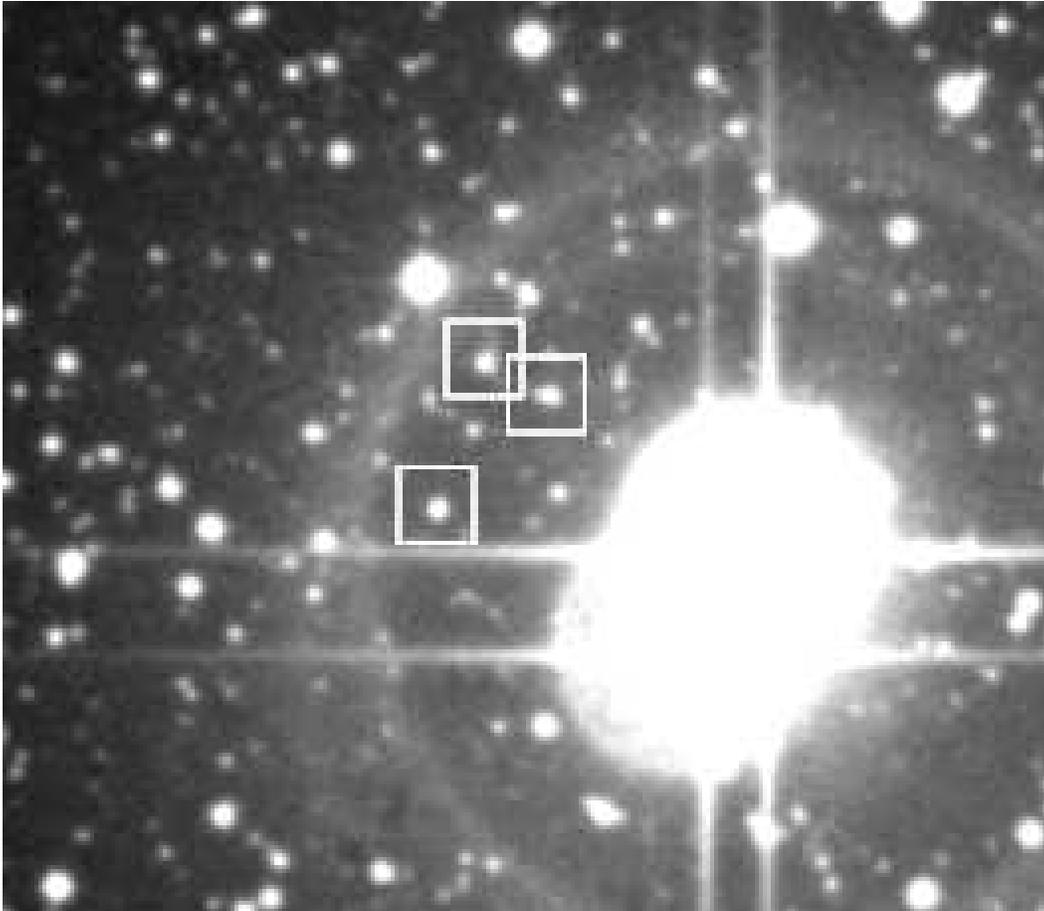


Figure 3.102 A 1991 image of the 61 Cyg system from the POSS. The 3 boxed stars are the apparent companions we detected 15 years later much closer to the star; the good match shows they are background objects.

higher than 'almost certainly unreal'. We have assigned to each a numerical reality likelihood, with 1 meaning probably unreal, 2 meaning likely unreal, 3 meaning provocative, 4, very provocative, and 5, confirmed real. Sources classified as 1 or 2 are certainly not worthy of followup, and are provide solely in case they turn out to correspond to a secure detection in some future, more sensitive survey. Sources with reality likelihood 3 or 4 are possible candidates for followup, depending on one's philosophy about the optimal use of telescope time. They are certainly good targets if, as not infrequently happens, other Clio/MMTAO observers find themselves at the telescope under usable conditions that do not (for example, because of wind direction) allow them to observe their intended program targets.

The '# Dets' column in the table is meaningful only for automatically detected sources, where it signifies the total times a given source was flagged. For example, a source detected by both the aperture photometry and the PSF amplitude methods on the images from just one processing method is said to have been detected twice. A source detected by both methods on all 6 image sets from a typical star gets 12 detections. Sources that were detected both manually and automatically are listed twice. The separation and PA values are probably good to about 1 pixel, or roughly 0.05 arcsec. They are very unlikely to be off by more than 0.1 arcsec. The magnitudes, and corresponding masses from the Burrows et al. (2003) and Baraffe et al. (2003) models, are very uncertain indeed. If a given candidate is real, its brightness probably lies within 0.5 mag of the value given. It is extremely unlikely that it is more than 1.0 mag brighter than this value. The photometry for manual detections is even more uncertain than that for automatic ones, because it is based on a subjectively assigned significance level in  $\sigma$ . Finally, we remind the reader once again that it is certain the vast majority of sources in

this table are unreal.

Table 3.18. List of Suspected Sources

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 5	4	4	9.84	80.63	Auto.	16.35	4.34	3.35
GJ 5	1	7	6.9	46.42	Man.	16.22	4.59	3.58
GJ 5	1	2	7.07	59.91	Auto.	16.45	4.16	3.18
GJ 5	1	1	8.8	272.53	Auto.	16.71	3.74	2.8
HD 1405	3	12	6.7	7.07	Man.	15.85	7.81	6.49
HD 1405	3	5	8.5	103.81	Auto.	16.21	6.73	5.61
HD 1405	1	4	0.68	68.96	Man.	12.82	>20	21.32
HD 1405	1	1	0.86	16.39	Auto.	15.42	9.41	7.79
$\tau$ Ceti	3	8	10.4	40.46	Man.	15.63	14.15	10
$\tau$ Ceti	1	1	1.31	90	Auto.	12.96	>20	28.24
$\tau$ Ceti	1	3	1.62	352.52	Auto.	13.9	>20	20.04
$\tau$ Ceti	1	1	2.82	125.88	Auto.	15.7	13.9	9.74
$\tau$ Ceti	1	1	2.98	59.68	Auto.	15.46	14.8	10.68
$\tau$ Ceti	1	2	4.84	226.22	Man.	15.73	13.81	9.65
$\tau$ Ceti	1	4	5.05	97.74	Auto.	16.39	11.82	7.94
$\tau$ Ceti	1	6	8.89	156.51	Man.	15.71	13.88	9.72
$\tau$ Ceti	1	1	10.49	1.33	Auto.	16.1	12.63	8.57
GJ 117	3	12	2.91	255.5	Man.	15.53	2.79	2.01
GJ 117	2	1	1.75	180	Auto.	15.24	3.08	2.41
GJ 117	2	1	7.24	66.7	Auto.	16	2.36	1.68
GJ 117	1	4	0.64	81.25	Man.	10.91	19.91	11.43
GJ 117	1	4	1.19	281.77	Man.	12.83	8.62	7.09
GJ 117	1	1	1.2	68.63	Auto.	13.89	5.51	4.57
GJ 117	1	8	1.2	133.36	Man.	13.12	7.65	6.36

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 117	1	4	1.24	135.8	Auto.	13.16	7.51	6.25
GJ 117	1	3	1.41	94.6	Auto.	13.84	5.62	4.66
GJ 117	1	1	1.51	271.85	Auto.	14.15	4.9	4.07
GJ 117	1	2	2.43	126.87	Auto.	15.83	2.51	1.8
GJ 117	1	2	2.5	296.57	Auto.	15.71	2.62	1.89
GJ 117	1	2	5.18	74.79	Auto.	15.85	2.49	1.79
GJ 117	1	1	7.04	321.16	Auto.	15.06	3.34	2.65
GJ 117	1	1	8.04	187.98	Auto.	15.25	3.07	2.4
GJ 117	1	2	8.57	4.55	Auto.	15.97	2.39	1.71
GJ 117	1	2	9.63	176.24	Auto.	16	2.36	1.68
GJ 117	1	4	10.78	188.55	Auto.	15.71	2.62	1.88
$\epsilon$ Eri	2	4	0.85	103.24	Man.	11.25	18.91	13.2
$\epsilon$ Eri	2	8	6.15	238.57	Man.	15.74	4.01	2.62
$\epsilon$ Eri	2	2	7.86	217.72	Auto.	15.94	3.84	2.45
$\epsilon$ Eri	1	4	1.65	137.39	Man.	13.53	7.88	6.3
$\epsilon$ Eri	1	7	1.96	330.26	Man.	13.45	8.1	6.5
$\epsilon$ Eri	1	8	2.19	36.87	Man.	13.3	8.55	6.88
$\epsilon$ Eri	1	1	2.25	37.12	Auto.	13.51	7.93	6.35
$\epsilon$ Eri	1	2	3.65	86.95	Man.	14.69	5.41	3.85
$\epsilon$ Eri	1	4	3.73	8.24	Man.	15.61	4.17	2.73
$\epsilon$ Eri	1	1	4.9	107.9	Auto.	16.25	3.58	2.19
$\epsilon$ Eri	1	4	6.57	329.83	Man.	15.5	4.31	2.83
$\epsilon$ Eri	1	4	7.59	342.49	Man.	15.91	3.86	2.48
$\epsilon$ Eri	1	7	7.87	132.75	Man.	15.92	3.85	2.47

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
$\epsilon$ Eri	1	2	8.08	318.17	Auto.	15.24	4.62	3.07
$\epsilon$ Eri	1	6	8.08	318.17	Man.	15.37	4.46	2.93
$\epsilon$ Eri	1	2	8.73	337.43	Auto.	15.86	3.91	2.52
$\epsilon$ Eri ( $M$ )	1	4	1.42	139.18	Man.	12.58	3.84	4.36
$\epsilon$ Eri ( $M$ )	1	1	3.07	64.65	Auto.	13.93	1.99	2.4
$\epsilon$ Eri ( $M$ )	1	8	9.85	157.62	Man.	13.67	2.28	2.63
GJ 159	1	4	1.7	143.13	Man.	15.75	5.1	4.25
GJ 159	1	5	3.24	330.98	Auto.	16.53	3.66	2.95
GJ 159	1	2	3.53	301.5	Auto.	16.55	3.63	2.93
GJ 166B	2	4	0.49	51.91	Man.	14.01	16.86	15.26
GJ 166B	1	4	0.42	315.51	Man.	13.43	>20	18.57
GJ 166B	1	4	0.48	186.35	Man.	14.17	15.89	14.16
GJ 166B	1	4	3.32	105.11	Man.	16.08	9.69	6.76
GJ 166B	1	4	4.65	17.59	Auto.	16.46	8.74	5.93
GJ 166B	1	4	11.86	314.19	Man.	15.59	10.96	7.9
GJ 166C	2	12	1.57	201.8	Man.	15.63	10.86	7.81
GJ 166C	1	12	3.12	225.63	Man.	15.64	10.85	7.79
GJ 166C	1	4	4.63	274.81	Man.	15.59	10.96	7.9
GJ 166C	1	4	6.32	123.08	Auto.	16.46	8.74	5.93
GJ 166C	1	4	11.25	140.95	Man.	16.08	9.69	6.76
HD 29391	1	5	0.45	319.4	Man.	11.69	>20	14.24
HD 29391	1	2	1.54	161.57	Auto.	16.45	3.12	2.45
HD 29391	1	8	2.09	357.34	Man.	16.22	3.46	2.77
HD 29391	1	4	2.85	47.76	Auto.	16.64	2.9	2.19

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
HD 29391	1	4	5.49	183.04	Man.	15.77	4.15	3.45
HD 29391	1	6	5.93	123.3	Man.	16.61	2.93	2.23
HD 29391	1	8	6.8	9.87	Man.	15.52	4.63	3.84
GJ 211	1	4	0.45	220.6	Man.	10.63	>20	57.56
GJ 211	1	4	0.53	33.69	Man.	10.92	>20	49.58
GJ 211	1	8	8.92	296.15	Man.	16.03	8.53	6.9
GJ 216A	2	8	9.86	333.06	Man.	15.01	8.07	6.61
GJ 216A	1	4	0.7	213.69	Man.	9.75	>20	52.38
GJ 216A	1	5	2.87	113.96	Man.	15.33	7.24	5.85
GJ 216A	1	2	3.99	124.08	Auto.	15.94	5.91	4.47
BD+20 1790	5	12	6.41	336.35	Man.	14.83	14.7	10.29
BD+20 1790	5	12	8.73	74.19	Man.	14.35	18.21	11
BD+20 1790	5	12	8.73	74.19	Auto.	14.63	16.14	10.58
BD+20 1790	2	4	0.65	116.57	Man.	14.38	18	10.96
BD+20 1790	1	4	7.72	263.86	Man.	16.64	6.59	5.37
BD+20 1790	1	4	8.34	5.68	Man.	16.4	7.19	6.02
BD+20 1790	1	2	9.08	76.71	Auto.	16.69	6.46	5.22
GJ 278C	1	4	0.5	11.31	Man.	12.01	>20	23.6
GJ 278C	1	1	0.79	317.49	Auto.	14.3	11.86	9.29
GJ 278C	1	1	1.5	35.75	Auto.	15.77	6.47	5.2
GJ 278C	1	2	2.84	56.85	Auto.	16.06	5.84	4.61
GJ 278C	1	2	9.22	228.42	Auto.	16.19	5.58	4.37
GJ 282A	3	8	10.53	205.97	Man.	16.09	8.86	7.23
GJ 282A	1	4	0.46	212.01	Man.	11.4	>20	43.87

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 282A	1	1	1.04	280.78	Auto.	14.26	18.53	11.81
GJ 282A	1	4	1.33	349.51	Auto.	15.91	9.47	7.77
GJ 282A	1	1	6.55	22.2	Auto.	16.31	8.2	6.65
GJ 282A	1	1	8.11	232.79	Auto.	16.23	8.43	6.84
GJ 282A	1	8	8.12	233.27	Man.	15.75	10.02	8.27
GJ 311	1	5	0.86	242.57	Auto.	12.38	>20	20.98
GJ 311	1	6	0.94	21.25	Man.	12.28	>20	21.99
GJ 311	1	14	4.37	40.95	Man.	15.7	6.98	5.79
GJ 311	1	1	4.41	40.53	Auto.	15.9	6.52	5.22
HD 77407A	3	1	3.67	83.16	Auto.	15.75	8.18	6.76
HD 77407A	1	7	0.65	26.57	Man.	12.04	>20	33.92
HD 77407A	1	2	1.07	357.4	Auto.	13.91	18.36	11.15
HD 77407A	1	4	3.31	304.85	Man.	15.51	9.04	7.48
HD 77407A	1	8	4.89	20.96	Man.	15.86	7.81	6.49
HD 77407A	1	4	6.61	130.83	Man.	15.07	10.98	8.85
HD 77407A	1	3	6.71	130.6	Auto.	15.34	9.76	8.07
HD 77407A	1	8	6.97	108.69	Man.	15.2	10.36	8.46
HD 77407B	3	1	3.95	108.51	Auto.	15.75	8.18	6.76
HD 77407B	1	2	0.63	174.67	Auto.	11.97	>20	35.55
HD 77407B	1	4	2.62	274.41	Man.	15.68	8.44	6.93
HD 77407B	1	8	3.42	32.85	Man.	15.86	7.81	6.49
HD 77407B	1	8	7.77	120.33	Man.	15.2	10.36	8.46
HD 77407B	1	4	7.89	139.66	Man.	15.07	10.98	8.85
HD 77407B	1	3	7.99	139.35	Auto.	15.34	9.76	8.07

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
HD 78141	3	4	2.96	332.18	Auto.	16.67	5.32	4.25
HD 78141	2	1	3.49	218.78	Auto.	16.39	5.93	4.77
HD 78141	2	3	7.84	94.26	Auto.	16	6.89	5.73
HD 78141	1	4	0.5	78.69	Man.	11.94	>20	29.67
HD 78141	1	1	0.59	55.01	Man.	13.78	18.58	11.09
HD 78141	1	1	1.6	1.74	Auto.	16.39	5.93	4.78
HD 78141	1	1	3.18	58.74	Auto.	16.72	5.23	4.17
HD 78141	1	8	3.66	349.29	Man.	16.34	6.06	4.88
HD 78141	1	1	3.68	143.58	Auto.	16.76	5.13	4.08
HD 78141	1	1	5.05	181.1	Auto.	16.52	5.65	4.53
HD 78141	1	2	5.63	55.9	Auto.	16.43	5.84	4.7
HD 78141	1	1	5.73	211.13	Auto.	16.56	5.58	4.47
HD 78141	1	3	7.05	266.84	Auto.	16.09	6.67	5.5
HD 78141	1	4	9.34	338.97	Auto.	16.05	6.78	5.61
HD 78141	1	8	9.85	21.38	Man.	15.68	7.78	6.5
HD 78141	1	8	10.22	20.59	Man.	15.77	7.52	6.3
GJ 349	2	1	4.18	9.35	Auto.	16.51	5.73	4.34
GJ 349	1	8	0.69	98.13	Man.	13	>20	16.63
GJ 349	1	1	1.72	49.57	Auto.	16.98	4.91	3.51
GJ 349	1	1	2.02	117.18	Auto.	16.7	5.39	3.98
GJ 349	1	8	2.75	222.14	Man.	16.16	6.42	4.97
GJ 349	1	8	2.96	209.51	Man.	16	6.77	5.4
GJ 349	1	3	2.97	209.78	Auto.	16.3	6.12	4.72
GJ 349	1	8	3.8	335.07	Man.	16.21	6.32	4.89

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 349	1	1	4.37	268.09	Auto.	16.56	5.64	4.24
GJ 355	1	8	6.89	338.5	Man.	15.85	5.17	4.3
GJ 354.1A	5	12	4.93	187.93	Man.	16.2	4.38	3.63
GJ 354.1A	5	4	4.94	187.33	Auto.	16.6	3.7	2.99
GJ 354.1A	2	1	3.24	289.25	Auto.	16.92	3.23	2.55
GJ 354.1A	2	4	5.06	111.39	Man.	16.97	3.16	2.48
GJ 354.1A	2	4	6.16	67.28	Man.	16.13	4.5	3.73
GJ 354.1A	2	8	8.49	171.45	Man.	15.24	6.62	5.51
GJ 354.1A	1	3	0.48	350.22	Man.	11.23	>20	29.6
GJ 354.1A	1	3	1.55	187.79	Auto.	16.55	3.76	3.06
GJ 354.1A	1	3	2.8	117.9	Auto.	16.92	3.22	2.54
GJ 354.1A	1	1	3.78	146.51	Auto.	16.67	3.59	2.89
GJ 354.1A	1	1	4.86	296.05	Auto.	16.55	3.76	3.06
GJ 354.1A	1	2	8.19	69.52	Auto.	16.78	3.44	2.74
GJ 354.1A	1	2	8.4	89.34	Auto.	16.01	4.74	3.93
GJ 354.1A	1	1	8.56	254.55	Auto.	15.57	5.74	4.75
GJ 354.1A	1	2	8.59	247.06	Auto.	15.85	5.07	4.23
GJ 380	3	6	0.88	167.24	Auto.	11.61	>20	35.39
GJ 380	1	4	1.03	199.29	Man.	12.77	>20	23.18
GJ 410	2	8	4.24	35.69	Man.	16.22	6.18	4.77
GJ 410	1	1	2.49	110.56	Auto.	16.76	5.18	3.78
GJ 410	1	1	3.23	132.56	Auto.	15.97	6.72	5.34
GJ 410	1	6	9.91	343.78	Man.	15.76	7.2	5.89
HD 96064A	5	12	5.59	212.59	Man.	14.58	15.02	10.38

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
HD 96064A	5	12	5.59	212.59	Auto.	14.08	18.63	11.1
HD 96064A	2	1	5.35	92.6	Auto.	16.7	5.92	4.77
HD 96064A	2	1	5.83	230.07	Auto.	16.72	5.87	4.73
HD 96064A	2	6	5.98	277.93	Man.	16.35	6.77	5.6
HD 96064A	1	4	0.46	71.57	Man.	12.6	>20	25.54
HD 96064A	1	1	0.83	20.56	Auto.	15.25	10.67	8.66
HD 96064A	1	2	1.48	242.59	Auto.	16.35	6.77	5.6
HD 96064A	1	2	5.73	219.49	Auto.	16.94	5.41	4.33
HD 96064A	1	6	7.43	70.14	Man.	15.01	12.01	9.4
HD 96064A	1	6	9.99	315.59	Man.	16.08	7.51	6.29
HD 96064B	5	12	6.17	49.78	Man.	14.58	15.02	10.38
HD 96064B	5	12	6.17	49.78	Auto.	14.08	18.63	11.1
HD 96064B	2	1	5.92	33.29	Auto.	16.72	5.87	4.73
HD 96064B	2	6	9.68	10.68	Man.	16.35	6.77	5.6
HD 96064B	2	1	15.55	57.1	Auto.	16.7	5.92	4.77
HD 96064B	1	2	5.91	43.66	Auto.	16.94	5.41	4.33
HD 96064B	1	6	15.84	2.63	Man.	16.08	7.51	6.29
HD 96064B	1	6	18.5	52.68	Man.	15.01	12.01	9.4
GJ 450	1	4	1.94	318.54	Man.	16.61	7.69	5.62
GJ 450	1	1	1.97	285.75	Auto.	17.14	6.49	4.59
GJ 450	1	1	2.02	99.69	Auto.	16.86	7.06	5.13
GJ 450	1	1	3.87	340.2	Auto.	17.33	6.13	4.23
GJ 450	1	4	4.07	274.45	Man.	16.15	8.95	6.79
GJ 450	1	12	6.14	60.9	Man.	16.84	7.1	5.17

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 450	1	8	6.18	164.98	Man.	16.34	8.4	6.21
BD+60 1417	5	8	1.9	302.47	Auto.	15.82	6.31	5.1
BD+60 1417	5	8	1.94	301.7	Man.	15.6	6.85	5.69
HD 113449	3	2	1.62	228.65	Auto.	16.21	6.55	5.36
HD 113449	3	1	4.06	294.73	Auto.	16.24	6.47	5.28
HD 113449	2	2	7.55	52.84	Auto.	16.69	5.44	4.35
HD 113449	1	4	0.76	39.81	Man.	13.54	>20	11.54
HD 113449	1	2	1.91	125.71	Auto.	16.33	6.26	5.04
HD 113449	1	2	2.05	211.43	Auto.	15.83	7.55	6.32
HD 113449	1	8	10.14	33.08	Man.	15.61	8.23	6.85
GJ 505A	1	8	0.87	270	Man.	13	>20	27.22
GJ 505A	1	1	2.76	100.12	Auto.	15.62	11.87	10.12
GJ 505A	1	1	6.1	205.95	Auto.	16.14	10.25	8.33
GJ 505A	1	8	9.91	151.3	Man.	15.38	12.73	10.68
GJ 505B	1	1	4.76	287.69	Auto.	15.09	13.9	11.34
GJ 505B	1	8	7.21	200.29	Man.	15.38	12.73	10.68
GJ 505B	1	1	10.54	250.3	Auto.	16.14	10.25	8.33
GJ 519	3	2	9.34	154.1	Auto.	16.64	4.96	3.58
GJ 519	1	8	10.68	346.72	Man.	15.54	7.12	5.81
GJ 3860	5	8	9.67	144.4	Auto.	14.68	12.58	9.75
GJ 3860	3	4	5.51	75.72	Auto.	16.69	5.71	4.39
GJ 3860	1	4	2	221.05	Man.	16.22	6.7	5.4
GJ 3860	1	4	2.3	42.44	Man.	15.98	7.25	6.03
GJ 3860	1	3	2.42	22.44	Auto.	16.6	5.88	4.56

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 3860	1	4	5.55	76.33	Man.	16.52	6.04	4.71
GJ 564	5	11	2.59	102.99	Auto.	11.03	>20	39.44
GJ 564	5	4	2.64	102.76	Man.	11.55	>20	29.8
GJ 564	1	2	0.98	243.43	Auto.	14.85	9.33	7.8
GJ 564	1	1	1.54	151.82	Auto.	15.86	6.29	5.08
GJ 564	1	4	4.22	84.06	Auto.	16.86	4.22	3.25
GJ 564	1	8	4.67	226.68	Man.	16.76	4.4	3.42
GJ 564	1	3	4.68	226.54	Auto.	16.82	4.29	3.32
GJ 564	1	8	9.04	159.23	Man.	16.22	5.47	4.38
HD 133002	5	4	1.87	117.9	Man.	11.5	>20	>100
HD 133002	5	12	1.87	117.9	Auto.	10.99	>20	>100
HD 133002	2	4	0.72	42.27	Man.	11.82	>20	>100
HD 133002	2	8	4.54	232.39	Man.	16.51	>20	34.49
HD 133002	1	8	1.46	248.55	Man.	14.57	>20	62.73
HD 133002	1	1	2.77	267.99	Auto.	16.9	>20	29.68
HD 133002	1	16	3.57	281.77	Man.	15.85	>20	43
HD 133002	1	2	5.23	48.01	Auto.	16.66	>20	32.61
HD 133002	1	4	5.46	226.8	Man.	16.49	>20	34.79
HD 133002	1	4	8.4	0.99	Man.	16.13	>20	39.23
$\xi$ Boo A	2	2	8.06	271.73	Man.	15.08	5.59	4.26
$\xi$ Boo A	2	2	11.62	253.99	Auto.	15.73	4.45	3.09
$\xi$ Boo A	1	4	0.49	84.29	Auto.	10.51	>20	26.51
$\xi$ Boo B	2	2	5.25	219.91	Man.	15.08	5.59	4.26
$\xi$ Boo B	2	2	9.89	220.9	Auto.	15.73	4.45	3.09

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
$\xi$ Boo A ( $M$ )	1	4	9.03	29.74	Man.	13.11	4.86	5.61
$\xi$ Boo B ( $M$ )	1	4	0.49	57.22	Man.	10.8	>20	22.38
HD 139813	2	2	5.62	108.12	Auto.	16.14	6.61	5.43
GJ 625	2	12	7.24	127.1	Man.	15.75	5.68	4.17
GJ 625	1	4	0.78	291.8	Man.	14.25	9.48	7.78
GJ 625	1	1	1.61	205.02	Auto.	16.19	4.94	3.47
GJ 625	1	8	1.63	206.57	Man.	15.96	5.32	3.82
GJ 625	1	2	2.14	291.32	Auto.	16.98	3.94	2.59
GJ 625	1	1	2.55	124.9	Auto.	16.44	4.62	3.1
GJ 625	1	2	2.9	148.71	Auto.	16.99	3.94	2.58
GJ 625	1	2	5.05	172.26	Auto.	16.32	4.77	3.28
GJ 625	1	2	7.45	70.97	Auto.	16.87	4.07	2.69
GJ 659A	2	2	7.22	118.98	Man.	15.1	>20	21.41
GJ 659A	1	1	0.74	156.8	Auto.	12.23	>20	61.54
GJ 659A	1	2	1.28	351.25	Auto.	14.64	>20	25.74
GJ 659A	1	1	2.55	34.9	Auto.	15.77	19.32	16.86
GJ 659A	1	2	2.69	310.6	Auto.	15.49	>20	18.53
GJ 659A	1	1	2.73	106.5	Auto.	15.7	>20	17.3
GJ 659A	1	4	5.84	70.06	Auto.	16.43	14.39	11.94
GJ 659A	1	2	7.35	167.8	Auto.	16.08	16.41	15.05
GJ 659B	2	4	0.45	45	Man.	13.66	>20	36.82
GJ 659B	1	2	7.5	352.18	Auto.	16.23	15.43	13.67
GJ 702A	1	9	1.2	46.64	Man.	12.32	>20	28.81
GJ 702A	1	9	1.42	67.83	Man.	12.51	>20	27.07

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 702A	1	1	1.87	233.43	Auto.	13.35	>20	19.51
GJ 702A	1	1	2.03	213.31	Auto.	14.75	13.74	11.26
GJ 702A	1	8	2.51	131.08	Man.	13.97	17.78	16
GJ 702A	1	1	2.61	331.05	Auto.	15.78	10.7	7.66
GJ 702A	1	8	3.09	343.56	Man.	15.22	12.23	9.4
GJ 702A	1	2	5.13	10.35	Man.	15.63	11.09	8.01
GJ 702B	1	6	2.66	320.2	Man.	13.94	17.97	16.15
GJ 702B	1	9	4.81	331.6	Man.	11.56	>20	37.93
GJ 702B	1	1	5.12	292.96	Auto.	13.23	>20	20.3
GJ 702B	1	8	5.28	328.94	Man.	11.54	>20	38.23
GJ 702B	1	1	5.72	296.82	Auto.	12.63	>20	25.93
GJ 702B	1	1	7.71	320.89	Auto.	15.49	11.49	8.48
GJ 702B	1	8	8.02	326.11	Man.	15.22	12.23	9.4
GJ 702B	1	2	9.15	342.99	Man.	15.63	11.09	8.01
61 Cyg A	5	8	7.77	83.19	Auto.	13.23	17.25	15.58
61 Cyg A	5	8	11.27	227.44	Man.	12.38	>20	20.65
61 Cyg A	2	2	1.28	240.52	Auto.	14.26	12.7	10.05
61 Cyg A	1	2	6.73	135	Auto.	14.19	12.89	10.3
61 Cyg A	1	2	9.36	14.11	Auto.	15.98	8.14	5.54
61 Cyg A( $M$ )	2	8	8.6	59.09	Man.	12.52	>20	19.6
61 Cyg A( $M$ )	1	1	1.35	115.64	Auto.	12.51	>20	19.65
61 Cyg A( $M$ )	1	1	2.42	202.44	Auto.	13.39	16.25	14.61
61 Cyg A( $M$ )	1	8	7.58	260.78	Man.	12.65	>20	18.89
61 Cyg B	5	8	9.84	145.45	Man.	14.01	13.5	10.99

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
61 Cyg B	5	16	9.84	145.25	Auto.	14.46	12.09	9.21
61 Cyg B	2	4	3.31	259.85	Man.	16.05	7.98	5.43
61 Cyg B	2	8	3.43	116.93	Man.	15.57	9.15	6.28
61 Cyg B	2	8	10.42	296.57	Man.	15.48	9.39	6.5
61 Cyg B	1	9	1.04	50.92	Man.	12.49	>20	19.76
61 Cyg B	1	1	2.33	357.61	Auto.	16.05	7.98	5.43
61 Cyg B	1	1	3.87	51.12	Auto.	16.41	7.31	4.87
61 Cyg B	1	8	8.94	351.25	Man.	15.65	8.93	6.09
61 Cyg B	1	8	9.23	261.53	Man.	14.96	10.72	7.68
61 Cyg B	1	2	9.68	319.68	Auto.	16.01	8.09	5.5
BD+48 3686	1	8	11.3	199.44	Man.	15.63	8.65	7.19
GJ 879	3	4	5.6	6.97	Auto.	14.7	5.83	4.6
GJ 879	2	10	2.96	137.66	Man.	14.91	5.41	4.22
GJ 879	1	1	0.65	243.43	Auto.	10.43	>20	25.89
GJ 879	1	1	6.29	171.12	Auto.	14.93	5.36	4.17
GJ 879	1	4	7.71	187.59	Auto.	15.01	5.21	4.03

Table 3.18—Continued

Star Name	Conf. Rating	# Dets	Sep (arcsec)	PA (deg)	Det Method	Mag ( $L'$ or $M$ )	Mass (Bur)	Mass (Bar)
GJ 896A	1	1	6.38	34.78	Auto.	16.7	3.18	2.07
GJ 896B	1	1	5.18	340.7	Auto.	16.7	3.18	2.07
HD 220140A	1	2	1.94	107.53	Auto.	15.7	5.66	4.68
HD 220140A	1	8	6.5	119.05	Man.	14.66	8.82	7.27
HD 220140B	1	8	13.15	63.96	Man.	14.66	8.82	7.27

Note. — For the Conf. Rating (confidence rating) column, 1 means probably unreal, 2 means likely unreal, 3 means provocative, 4 means very provocative, and 5 means confirmed. Sources were listed twice if they were detected with both manual and automatic methods. The # detections column is not very relevant for manual detections; for automatic detections more detections means the source is more likely to be real. All observations are  $L'$  except those noted as  $M$  in the Star Name column. Masses are in MJ. Magnitudes for manually detected sources have little meaning. For automatic detections the magnitudes are probably accurate to  $\pm 0.7$ . The objects are most likely fainter than the quoted magnitudes.

### 3.11 Conclusion

We have obtained sensitive  $L'$  images of 42 star systems containing 50 stars around which we obtained meaningful sensitivity to substellar objects. We also observed three of these star systems in the  $M$  band. We have processed the data intensively using a sophisticated pipeline.

We have not detected any planets, nor any very promising candidates. To make our null result secure and meaningful, we have analyzed our sensitivity more rigorously than any prior planet imaging survey. We find that our pixel-to-pixel RMS (Method 1) sensitivity estimator, which may be analogous to the estimators used in a number of other planet imaging surveys (but not Lafrenière et al. (2007b) and Biller et al. (2007)) overestimates the sensitivity by up to 1 magnitude in regions where the noise is strongly spatially correlated. Since regions without spatially correlated noise are the exception rather than the rule, this is a potentially serious problem. If some of the surveys did indeed use a Method 1-type estimator, the sensitivities probably need revision to considerably more conservative values. The papers in question did not give sufficient details of their sensitivity calculations for us to be sure whether the estimator used was Method 1 or something more similar to the Biller et al. (2007) estimator, which does not have the same problems. Future papers should devote more space to the important issue of the sensitivity estimator. Also, all future surveys should follow our example of carrying out blind sensitivity tests with simulated planets inserted at random positions. This is by far the surest way to quantify the completeness vs significance level relation for a given survey. Our results showed lower completeness at each significance level than we would naively have expected, although after consideration the statistical reasons for this became clear. This may be true of other surveys as well.

We find that for the great majority of the stars in our survey, observations using the SDI method of Biller et al. (2007) and the narrowband  $H$  regime imaging of Lafrenière et al. (2007b) would have obtained sensitivity to lower mass planets than we did at small and large separations from the host star, respectively. It appears, therefore, that the main value of our survey is not that we obtained sensitivity to planets no other method could have obtained, but that we used an independent method. Our survey is similar to that of Kasper et al. (2007) in this regard, though the target samples are different (see Section 3.3), and they carried out observations only at  $L'$ . Our survey (and that of Kasper et al. (2007)) have helped diversify the investment of the astronomical community in planet imaging surveys, hedging the community against the possibility that planets may turn out to be far fainter at  $H$  regime wavelengths than current models indicate. This is a real possibility, since the chemistry, cloud structures, and typical evolutionary histories of giant extrasolar planets are almost entirely without observational constraint. We note also that the remarkable usefulness of the  $H$  band depends on strongly super-blackbody emission, which is expected to occur in this wavelength range because of strong absorption elsewhere. This situation may be very unstable to errors in the real atmospheric chemistry, especially for the coldest detectable planets.

We note that although we would expect the  $L'$  and  $M$  band brightness of planets to be more robust, since the band is located nearer to their natural blackbody peak, the longer wavelengths have uncertainties too. Leggett et al. (2007) and Reid & Cruz (2002) find that brown dwarfs with spectral types from T2 to T8 have  $M$  band fluxes 0.2-0.7 magnitudes fainter than models predict. This range of spectral types corresponds to effective temperature ( $T_{\text{eff}}$ ) from 1300-700 K (Leggett et al., 2007). Reid & Cruz (2002) and Leggett et al. (2007) suggest that

the suppression of  $M$  band flux in these objects could be due to CO in their atmospheres, which absorbs in the  $M$  band. For low temperature objects such as these, local thermodynamic equilibrium (LTE) chemistry increasingly favors  $\text{CH}_4$  over CO at the photosphere. However, vertical mixing could bring up CO from deeper, higher-temperature regions, resulting in above-LTE concentrations of CO, which would explain the lower  $M$  band fluxes (Leggett et al., 2007). No extrasolar planets or brown dwarfs with  $T_{\text{eff}}$  below 700 K have yet been detected, and it is not known whether they would also show suppressed  $M$  band flux Leggett et al. (2007). We would expect that as  $T_{\text{eff}}$  decreases and LTE favors CO over  $\text{CH}_4$  only at increasingly great depths below the photosphere, the effect would eventually vanish. Most of our  $M$  band observations are sensitive to objects with  $T_{\text{eff}}$  well below 700 K (close to 300 K in the best cases), which may be below the coolest  $T_{\text{eff}}$  where CO absorption is important. The majority of our observations are conducted in the  $L'$  band, where above-LTE CO abundance would increase the flux, or at most leave it unchanged.

No self-luminous objects with temperatures between that of Jupiter (roughly 100 K) and the coolest T dwarfs ( $\geq 700$  K) have yet been detected. Models of the flux of such objects are therefore uncertain at all wavelengths. There are reasons to consider the  $H$  band flux more uncertain, with potential large sensitivity to unknown parameters such as clouds, but the longer wavelengths have their own uncertainties. What is clear is that observations using a range of wavelengths and techniques are desirable to increase the likelihood of discoveries and the confidence with which we can set limits based on null results.

Although there was considerable overlap between our survey and that of Lafrenière et al. (2007b), we probed older stars on average. We have thus placed constraints on the distribution of planets around older stars than any other planet-

imaging survey. This is of potential importance since the orbits of giant planets are expected to continue evolving due to planet-planet interactions for at least 0.1 Gyr after the planets form.

There was one star system for which the better performance of  $H$  regime methods described above strikingly did not apply. This was the very interesting  $\epsilon$  Eri system. We obtained better sensitivity than Biller et al. (2007) through to whole range of separations of which the latter had data. For the 2-3 arcsecond separation range we set good limits where no other deep planet-imaging survey had data at all. Beyond 3 arcsec, our results remained better than or comparable to those of Lafrenière et al. (2007b) until at least 8 arcseconds out. We note that the Burrows et al. (2003) models indicated that our  $M$  band observation attained significantly higher sensitivity than our  $L'$  image, but the Baraffe et al. (2003) models showed them as more comparable.

There are two reasons we did so well on  $\epsilon$  Eri. First, it is a bright star, so the large planet/star flux advantage delivered by the  $L'$  and  $M$  bands relative to  $H$  comes strongly into play. Second, and probably even more significant, it is a nearby star. For nearby stars the sensitivity in any wavelength regime extends to intrinsically fainter, cooler objects. This is regardless of the age of the star, though of course the mass of objects at a given temperature depends strongly on the age. The planet models robustly predict that the  $H - L'$  and  $H - M$  colors of giant planets redden with decreasing temperature. Thus **for the nearest stars, around which the lowest temperature objects can be detected, the longer wavelengths will give the greatest advantage.** As the  $\epsilon$  Eri data show, this is an extremely important effect.

There are three other stars in our sample that share a similar brightness and distance with  $\epsilon$  Eri. These are 61 Cyg A and B, and  $\tau$  Ceti. Unfortunately, the data

for all three are a bit below average, so it is not clear that we did strongly surpass the sensitivity Lafrenière et al. (2007b) type observations could have obtained around each star. We note, however, that 1.5 hour integrations on each of these stars under good conditions would produce data at least as good as our  $\epsilon$  Eri data. Even the data sets we have acquired are sensitive to planetary mass objects despite the greater age of these systems relative to  $\epsilon$  Eri. **The bands we have helped pioneer will yield considerably better sensitivity than any  $H$  regime technique for stars closer than 4 pc and brighter than  $H = 2-3$ .**

In closing, we note that just as observing a more nearby star allows us to see fainter, redder objects for which the  $L'$  and  $M$  bands are the best choice, so would using a bigger telescope. The  $L'$  and  $M$  bands will be superior to  $H$  out to larger distances for the next generation of giant ground based telescopes with AO. We have therefore helped develop a wavelength regime that will become increasingly important as larger telescopes are built and the science of extrasolar planets advances.

## CHAPTER 4

## CONSTRAINTS ON PLANET ABUNDANCE FROM THE CLIO SURVEY OF 50 STARS

## 4.1 Introduction

The vast majority of the roughly 250 known extrasolar planets were discovered by the radial velocity (RV) method. In a radial velocity detection, the reflex motion of the star in response to an orbiting planet's gravity is detected in the form of a periodic variation in the star's line-of-sight (or radial) velocity. RV measurements are now extremely precise thanks to several ingenious developments that have fueled the boom in extrasolar planet detections, most notably the iodine cell (Marcy & Butler, 1992). Thanks to the great success of RV observations, enough extrasolar planets are now known that meaningful estimates can be made of how their masses and orbital parameters are distributed statistically.

Despite its great success, the RV method has its limitations. The amplitude of the RV signature of a planet is proportional to the planet's orbital velocity, which goes as the inverse square root of the orbital radius, so the sensitivity declines somewhat at larger orbital radii. Far more importantly, a well-constrained RV detection requires good observations spanning at least one full orbital period of the planet, and the orbital period goes up as the  $3/2$  power of the orbital radius. Jupiter, with an orbital radius of about 5AU, already has a period of 11 years, and Saturn at about 10AU takes 30 years to complete an orbit. Since precise RV observations have only been possible since the early 1990's, planets in Saturn-like orbits can not yet be detected, while Jupiter-like orbits are accessible only around stars that have good observations stretching back to near the first availability of the required technology. Planets in orbits much more distant than Saturn's could take centuries to find. New methods of planet detection are required to discover

extrasolar planets in Saturn-like and larger orbits.

Direct imaging of extrasolar planets has been possible in principle for several years thanks to advances in adaptive optics (AO) and astronomical infrared detectors. AO reverses the blurring effect of turbulence in the Earth's atmosphere and dramatically sharpens the images from large telescopes, while new infrared detectors take clear, sensitive images of celestial objects. The most promising method for imaging extrasolar planets is to try to see them not by the starlight they reflect (as human eyes see the planets in our own solar system), but by their own infrared emission. Giant planets emit infrared light because the gravitational potential energy converted to heat in their formation and subsequent slow contraction is slowly radiated into space. The planets radiatively cool over time from an initial very hot state. As they get colder, of course, their infrared glow slowly fades. According to theoretical models (ie Burrows et al. (2003) and Baraffe et al. (2003)), giant planets are bright enough to be detected around nearby stars using large telescopes equipped with AO and current-generation IR cameras, provided the planets are not too old (and thus too cool and faint).

The first attempts to image extrasolar planets were done at the astronomical  $H$  and  $K$  (or, more accurately, the better optimized  $K_s$ ) bands (see Table 4.1). The detector technology for these shorter wavelengths advanced more quickly than for the longer wavelength  $L'$  and  $M$  bands. Also, the thermal background from telescope optics and sky is very high at the  $L'$  and  $M$  bands and is aggravated by the additional optical surfaces needed in most AO systems.

Planet-imaging surveys have moved from using both the  $H$  and  $K_s$  bands (ie Masciadri et al. (2005)) to using narrowband filters within the  $H$  band tuned to predicted flux peaks in extrasolar planet models (Biller et al. (2007), Lafrenière et al. (2007b)). Theoretical models predict that these optimized filters, unlike the  $K_s$

band, give good sensitivity even to fairly old planets. around nearby stars. This is important, as the young systems more easily imaged at  $K_s$  are rare and therefore tend to be distant. The advances in  $H$  regime imaging strategies brought closer star systems in range, around which planets in physically smaller orbits more comparable to those found in RV surveys should be detectable. Earlier surveys that focused on distant, young stars were sensitive only to planets in wide orbits.

However, the planet/star flux ratio in the  $H$  band wavelength regime is not nearly as favorable as for longer wavelengths such as the  $L'$  and  $M$  bands. The difference is especially pronounced when we consider observations of very nearby stars, around which intrinsically faint, cool objects with very red  $H - L'$  and  $H - M$  colors can be detected. For detecting close-in planets against the bright stellar haloes of nearby stars, the  $L'$  and  $M$  bands have a large advantage over  $H$  regime wavelengths. It would appear that for the nearest bright stars  $L'$  and  $M$  band observations may have the best potential to close the gap between the very close-in planets detectable in RV surveys, and the more distant planets to which  $H$  band regime observations are sensitive.

The MMT telescope has the only fully operational adaptive secondary mirror in the world. This remarkable innovation removed a key obstacle to the use of the longer wavelength  $L'$  and  $M$  bands for AO planet searches: the thermal glow from additional optical surfaces in AO systems. The MMT AO system has only two non-cryogenic surfaces — the primary and (adaptive) secondary mirrors. It delivers sharp, AO-corrected images with no higher thermal background than a telescope without AO. The Clio instrument was designed and built with the specific purpose of taking advantage of this unique AO system to search for extrasolar planets at the  $L'$  and  $M$  bands. Therefore with Clio and the MMTAO

system we can target very nearby bright stars, and take advantage of the favorable planet/star flux ratio to see planets in close-in orbits since the glare from the bright parent star is less intense. Stars at moderate ages are much more common than very young stars, and thus more of them can be found nearby. Clio's ability to target nearby, moderate-age stars with sensitivity to planets at small angular separations means a Clio survey can detect planets orbiting at substantially smaller physical radii than shorter wavelength surveys can, at least for the very brightest nearby stars (see Chapter 2). The  $L'$  and  $M$  band capability that Clio delivers when combined with the MMT AO system provides a way to fill in the gap between RV surveys which are sensitive to planets at small orbital radii, and shorter wavelength IR imaging surveys, which for very bright, nearby stars are sensitive mainly to planets orbiting at large radii.

We have used the MMT's unique capability to deliver AO corrected images with low thermal background to carry out a planet imaging survey of nearby middle aged and moderately young stars. We designed this survey with the goal of constraining the statistical distribution of extrasolar planets at larger orbital radii. In particular, we planned to observe 50 stars, obtaining sufficient sensitivity around each that a null result for the entire survey would rule out simple extrapolations of power law fits to the distributions of planet mass and orbital radius for the known RV planets.

This survey is now complete. Amusingly, despite drawing from a target sample of considerably more than 50 stars as weather and scheduling allowed, and despite concerns at various times that weather would limit us to substantially fewer than 50 stars, the final number we have analyzed and around which we obtained meaningful sensitivity to substellar objects, is exactly 50. As expected, we have obtained sensitivity to planets in orbits ranging inward to overlap the

outer edge of the RV sensitivity region. We have found no confirmed planets, and in the final analysis no promising planet candidates. Our survey null result is thus secure.

In practice, we find that according to the theoretical models of Burrows et al. (2003) and Baraffe et al. (2003) our  $L'$  and  $M$  band observations attain markedly better sensitivity than the clever, optimized  $H$  regime observations of Biller et al. (2007) and Lafrenière et al. (2007b) only for very nearby, bright stars such as  $\epsilon$  Eri. For the majority of the stars in our sample,  $H$  regime observations could have attained better sensitivity. Given a power law for the planet mass  $m$  of  $dN/dm \propto m^{-1.44}$ , our survey null result rules out a power law for the semimajor axis  $a$  of  $dN/da \propto a^{-0.2}$  extending beyond 155 AU, or  $dN/da$  constant extending beyond 70 AU, at the 95% confidence level.

The surveys of Biller et al. (2007) and Lafrenière et al. (2007b) used observations in the  $H$  band wavelength regime to set stronger constraints on the planet distributions than we did. However, setting constraints and making sensitive observations at a wide range of wavelengths is very desirable in the present situation where models of giant extrasolar planets are not observationally well constrained. Observations such as ours (and those of Kasper et al. (2007)) increase the confidence with which we as a community can constrain the populations of extrasolar planets. Reasonable constraints from our survey and that of Kasper et al. (2007) will still apply even if we find that unexpected atmospheric chemistry, clouds, or other effects make giant planets much fainter in the  $H$  band than current models indicate. The models are, of course, uncertain at the longer wavelengths too (see Reid & Cruz (2002) and Leggett et al. (2007); also the discussion thereof in Chapter 3); the point is not that only the  $H$  band is uncertain but that observing at a wide range of wavelengths increases the confidence we can have

in the results. We note, also, that the  $L'$  and  $M$  bands will become increasingly important when the next generation of giant telescopes begins operations. This is because larger telescopes will be sensitive to objects in nearby star systems that are so faint, cool, and red in  $H - L'$  and  $H - M$  color that even background limited  $L'$  and  $M$  band observations will be much better than those in the  $H$  band regime. These wavelengths will be important from space too: JWST will deliver background-limited sensitivity far superior to any ground-based telescope at  $L'$ ,  $M$  band, and longer wavelengths, though its performance in the contrast-limited regime close to bright stars is uncertain. Twenty or thirty meter-class ground based telescopes such as the GMT will certainly deliver better sensitivity very close to bright stars than the GMT.

The completeness of RV planet samples has grown rapidly over the years between our survey's initial conception and final completion, and the power law fits from known RV planets have generally evolved toward predicting fewer massive planets and fewer planets at large orbital radii. Because of this, our adopted power laws based on a largely complete samples of RV planets, extrapolated out to a reasonable truncation radius of 40 AU, are not ruled out by our our null result. They do predict that the survey had roughly a 50% chance of detecting a planet, and thus it was not a long shot and more extended similar surveys in the near future may well turn up planets. Power law fits with more optimistic slopes still consistent with the RV detections, and larger truncation radii for the distribution of planet orbits, are in fact ruled out by our survey. We consider it a satisfying, though not spectacular, success. It has increased our understanding of how extrasolar planets are distributed. It also represents a significant advance in the field of  $L'$  and  $M$  band planet searches, which have great promise for the future.

We present the nature of the survey and discuss the resulting constraints in this chapter.

In Section 4.2 we discuss how we chose our sample and list the stars observed with the ages we have adopted for each one.

Section 4.3 details the radial velocity sample from which we extracted power law fits and normalizations, and explains how these were constructed mathematically.

In Section 4.4 we discuss the theoretical models we used to estimate the brightness of planets with different masses and ages, and the methods we used to interpolate across the model grids.

In Section 4.5 we introduce the main results of this chapter: Monte Carlo simulations of our survey in which we calculate the likelihood of obtaining the observed null result, and the expected properties of planets to which the survey was sensitive.

Section 4.6 gives detailed results from a Monte Carlo simulation in which the power law fits to the RV planets were extrapolated out to a maximum orbital radius of 40 AU, beyond which there were assumed to be no planets. Fifty thousand realizations of the survey were carried out, in which planets were randomly drawn from the truncated power law distribution, placed around the stars in our survey, and then classified as detected or undetected according to whether the magnitude of the planet from theoretical models was brighter or fainter than our measured sensitivity at the appropriate projected distance from the star under consideration. We find that the probability of detecting no planets given the 40 AU truncation radius and the adopted power laws is about 50%, so that although our survey had a good chance of finding a planet, the null result does not rule out the adopted power laws from the RV planets truncated at 40 AU. The char-

acteristics of detected planets in the Monte Carlo simulation are quite interesting both because they clarify what types of planets the survey was sensitive to (and therefore what types it proves are not extremely common), and also because they provide guidelines for the types of planets likely to be detected by future surveys using the same promising wavelength regime. We therefore conclude the section by analyzing the distributions of these simulated planets in detail, and commenting on the implication for future surveys.

Section 4.7 carries on from Section 4.6 to consider what distributions of planets may be confidently ruled out by our survey. We present the key result in Figures 4.40 through 4.43; these are contour plots showing the probability of a null result for our survey as a function of the orbital semimajor axis power law slope and the truncation radius for the planet distribution. These plots illustrate, i.e., what distributions may be ruled out at the 95% confidence level by our survey: they are the ones giving less than a 5% chance of detecting no planets.

Finally, in 4.8 we present our conclusions on what the survey tells us about the distribution of planets, and discuss briefly the place similar surveys have in the broader picture of planet imaging searches over the next decade.

## 4.2 The Survey Sample

The key criterion for our sample was to choose the nearest stars around which, based on the theoretical models of Burrows et al. (2003) and Baraffe et al. (2003), we could detect planets down to 10 Jupiter Masses (MJ) or below. This practically meant that very nearby stars were potential targets up to ages of several Gyr, while at larger distances we would consider only fairly young stars. We set out initially to investigate only FGK stars within 25pc of the sun, in order to make our sample comparable in spectral type to the samples of the RV surveys and to

Table 4.1. Astronomical Infrared Filter Bands

Band Name	Central Wavelength ( $\mu\text{m}$ )	Band FWHM( $\mu\text{m}$ )
<i>J</i>	1.215	0.26
<i>H</i>	1.654	0.29
<i>K<sub>s</sub></i>	2.157	0.32
<i>K</i>	2.179	0.41
<i>L'</i>	3.761	0.65
<i>M</i>	4.769	0.45

Note. — The values in this table are taken directly from Table 7.5 in Cox (2000).

focus on the nearest stars at which the  $L'$  and  $M$  bands offer the largest advantage over shorter wavelengths. In the end we included a few M stars and a few stars slightly beyond 25pc. We did this because it was difficult to come up with the desired 50 targets using the more stringent constraints, and because some M stars or stars slightly beyond 25pc were very interesting. Our sample is discussed in detail in Chapter 3.

#### 4.3 Statistical Distributions from RV Planets

Enough planets have now been detected by the RV method that meaningful estimates may be made of their statistical distribution. As is traditional in astronomy, power laws were the first distributions tried, and they fit fairly well. Both the mass and the orbital semimajor axes of RV planets seem to be distributed approximately according to power laws. In the case of the semimajor axis distribution, the 'hot Jupiters', a distinct population of planets with extremely small orbital radii, must be excluded before the power law fit is tried.

Several papers have presented power law fits to the population of extrasolar planets, but since we required a normalization (number of stars that have planets within a given interval in mass and orbital semimajor axis) as well as power law slopes, we decided to construct our analysis from a known, reliable set of planets. We chose the Carnegie Planet Sample, as described in Fischer & Valenti (2005). The electronic version of this paper gives two very helpful tables. Table 1 lists 850 stars that have been thoroughly investigated with RV. The paper states that all planets with a mass of at least 1 MJ and an orbital period less than 4 years have been detected around these stars, with completeness down to even lower masses for shorter orbital periods. Forty-seven of these stars are marked in the table as having RV planets. Table 2 of Fischer & Valenti (2005) gives the measured

properties of 124 RV planets, including those orbiting 45 of the 47 stars listed as planet-bearing in Table 1. The stars left out are HD 18445 and and HD 225261. We cannot find any record of these stars having planets, and therefore as far as we can tell they are typos in Table 1. The 45 remaining stars are available to us for use in normalizing the power law distributions.

To do this, we had to choose intervals on mass and orbital semimajor axis over which the RV surveys were guaranteed to be complete, calculate how many stars had planets in these intervals, and divide that by 850, the total number of stars surveyed. For the lower mass limit, we selected 1.0 MJ. This is a bit below the lowest-mass planet our survey could have detected, and is the lowest mass Fischer & Valenti (2005) guarantee would have been detected even in a 4-year orbit. For the upper limit we select 13.0 MJ. This is approximately the lowest mass at which objects experience appreciable amounts of core deuterium fusion, and is therefore often quoted as the planet/brown dwarf boundary mass. We have adopted this definition for expediency. It has little else to recommend it since deuterium fusion is not very important in the evolution of these objects, but more physically motivated definitions of planets (based on, ie, the way they formed) are simply impossible to apply at present. For the inner limit on orbital semimajor axes we chose 0.3 AU to exclude the 'hot Jupiters', and for the outer limit we chose 2.5 AU, which for a sunlike star corresponds to the maximum 4-year orbital period at which the survey is guaranteed complete down to 1 MJ.

Twenty-eight stars, or 3.29 % of the 850 in the Fischer & Valenti (2005) list, have planets of between 1 and 13 MJ orbiting between 0.3 and 2.5 AU. Some of these stars have more than one planet matching these criteria, but of course these were still counted only once since we want the normalized probability of a star having at least one such planet, not the odds of a given planet in the universe

falling in our range of parameter space.

We could have attempted to compare power laws to the planets orbiting this set of 28 stars, but in order to get better statistics for our fits we decided to draw from the larger set of planets in Table 2 of Fischer & Valenti (2005). Fifty-one out of the 124 planets in this table have masses between 1 and 13 MJ and orbital semimajor axes between 0.3 and 2.5 AU, and it is to this set of planets that we fit power laws.

A concern, of course, is that this larger set of planets does not come with the same guarantees of completeness as the set of 45 that we used to get the normalization. However, 1.0 MJ is a fairly large mass by RV standards, and most surveys have been going on for longer than 4 years, so the effects of incompleteness in the subsample we chose should be small. We determined the 90% and 95% confidence intervals on the power law slopes using the Kolmogorov-Smirnov (KS) test, as implemented in the routines `ksone` and `probks` from Press et al. (1992). Briefly, as described in Press et al. (1992) this implementation of the KS test takes as input a mathematically defined model for the probability distribution of a single quantity, and a set of data points simply consisting of measured values of that quantity. The test involves constructing the cumulative distributions for both the model distribution and the measured data points, and calculating the likelihood that the data distribution would deviate from the model distribution as much as it in fact does, under the assumption that the input model distribution does in fact describe the data. If the observed deviation is very unlikely, the model probably does not describe the data.

Table 4.2 gives the 90% and 95% confidence intervals on the mass and semimajor axis power laws, based on the KS test. We adopt a slope of -1.44 for the mass power law, and -0.35 for the semimajor axis power law; these values lie

near the center of the confidence intervals we have determined. The truncation radius of the distribution, which is presumably related physically to the truncation radii of the circumstellar disks from which the planets formed, is therefore a very important parameter.

In Figure 4.1, we show a histogram of the masses of the 51 planets used in our KS power law analysis, with our adopted power law for the mass distribution overlaid as a smooth curve. Figure 4.2 shows the same thing for semimajor axis distribution. That the data match the power law curve less well can immediately be seen, but it seems this could well be due to small-number statistics and not to a systematic departure from the power law form. In Figure 4.3 we show the eccentricity distribution of these 51 planets, with a smooth curve overlaid that is the eccentricity distribution from Juric & Tremaine (2007):  $P(\epsilon) = \epsilon e^{-\epsilon^2/(2\sigma^2)}$ , where  $\epsilon$  is the eccentricity,  $e$  is the root of the natural logarithm, and  $\sigma = 0.3$ . The significance of this distribution is that we have used it to produce the random orbital eccentricities for planets in the Monte Carlo simulations discussed in Sections 4.5 through 4.7. The figure indicates that this distribution does indeed describe the eccentricities of RV planets well, as Juric & Tremaine (2007) claim.

In conclusion, we find that about 3.3% of stars have planets with masses between 1.0 and 13.0 MJ, in orbits with semimajor axes between 0.3 and 2.5 AU. In this interval, the distributions of both planet masses and planet orbital semimajor axes are consistent with power law fits with a range of slopes; we have set confidence intervals for these slopes using the KS test (see Table 4.2). We have adopted a slope of -1.44 for the mass power law and -0.35 for the semimajor axis.

Table 4.2. Power Law Fits to the Masses and Semimajor Axes of RV Planets

Variable	Statistic	Power Law Slopes
Planet Mass	90% Conf. Interv.	-1.803 to -1.156
Planet Mass	95% Conf. Interv.	-1.881 to -1.098
Semimajor Axis	90% Conf. Interv.	-0.494 to -0.196
Semimajor Axis	95% Conf. Interv.	-0.569 to -0.117

Note. — The KS test allows us to determine confidence intervals for the slopes of power laws describing the distributions of planet masses and orbital semimajor axes. This has the advantage over histogram fits that it does not depend on a bin-width parameter.

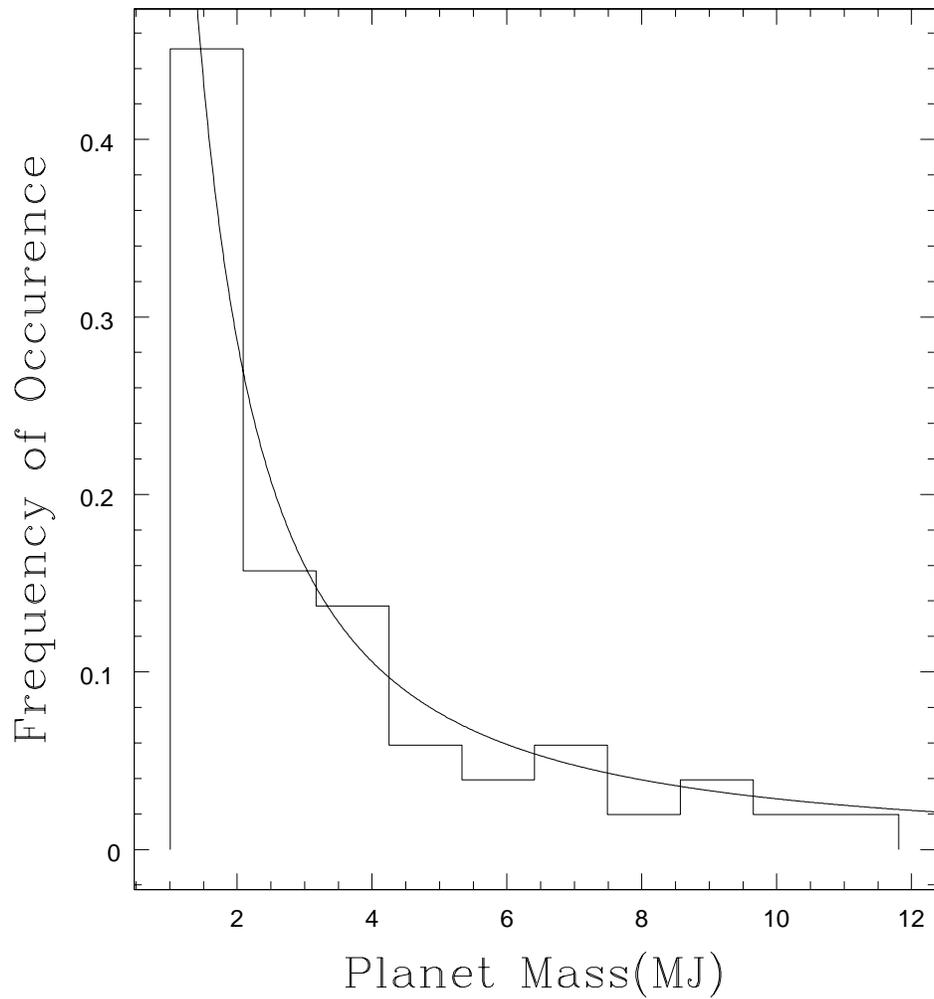


Figure 4.1 Histogram of planet masses for the 51 RV planets used in our power law fitting, with our adopted law (slope -1.44) overlaid as a continuous curve. As can easily be seen, the power law matches the data well.

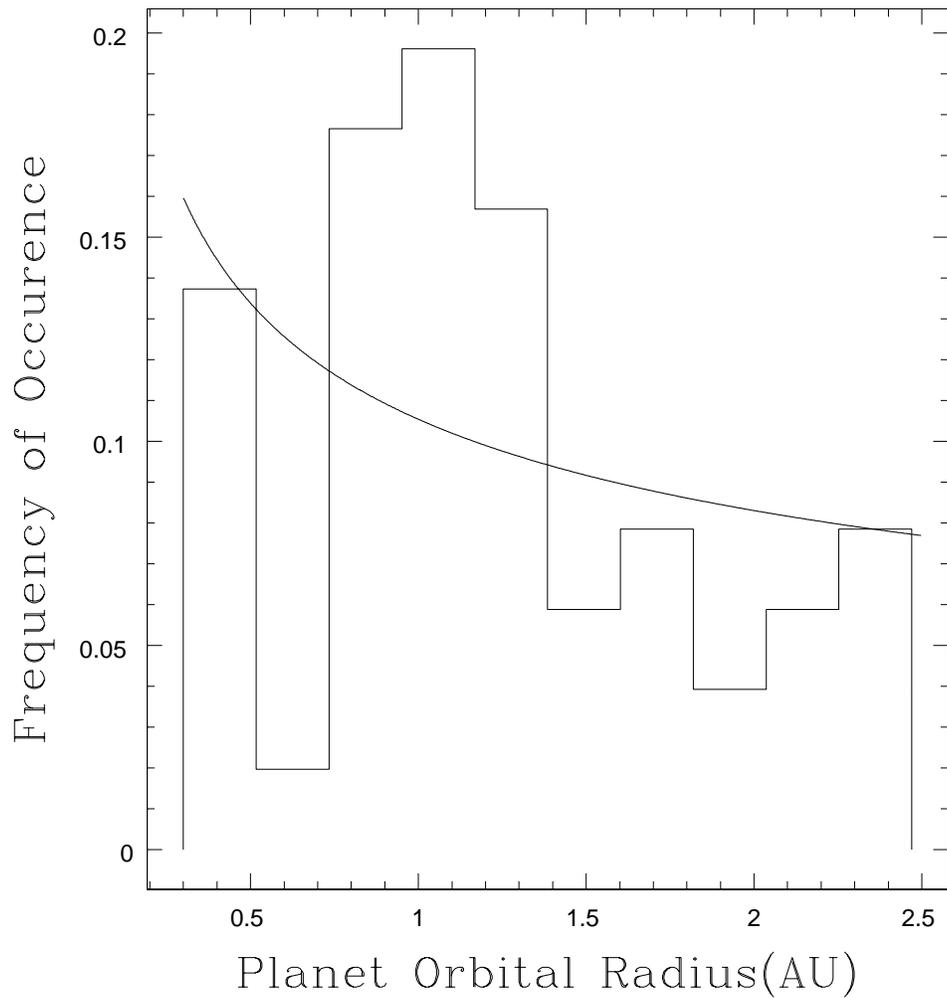


Figure 4.2 Histogram of orbital semimajor axes for the 51 RV planets used in our power law fitting, with our adopted law (slope -0.345) overlaid as a continuous curve. The match is poorer than for the mass distribution, but this could well be due to small-number statistics and not to a systematic departure from the power law form.

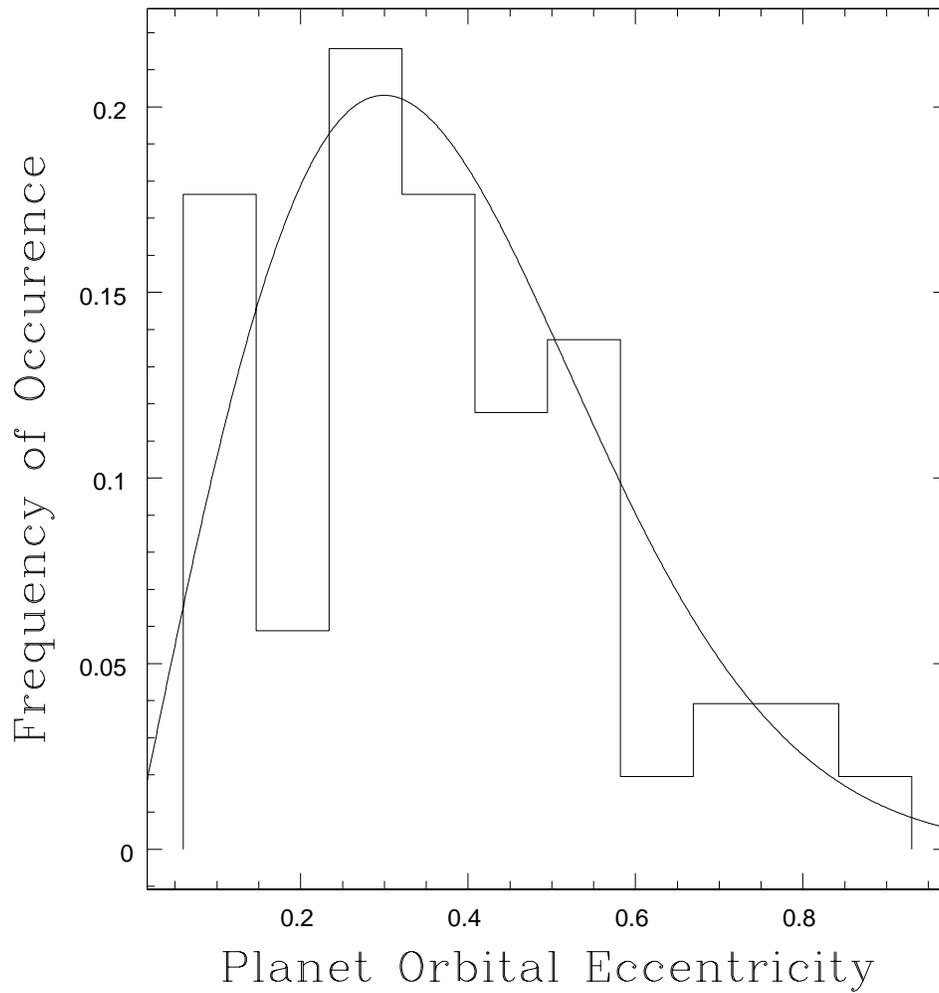


Figure 4.3 Histogram of orbital eccentricities for the 51 RV planets used in our power law fitting to mass and radius. The overplotted continuous curve is the distribution from Juric & Tremaine (2007), which we used to generate random eccentricities for planets in our Monte Carlo simulations. The fit appears very good.

## 4.4 Theoretical Models of Giant Planets

### 4.4.1 Two Sets of Models

There are two excellent 2003 papers presenting models of extrasolar giant planets from which the observables we need to use can easily be derived. We use non-irradiated models, of course, because in contrast to the RV-detected ‘hot Jupiters’ the stellar flux does not appreciably contribute to the internal energy or thermal infrared brightness of planets orbiting at large distances from their parent star. Burrows et al. (2003) present high resolution, flux-calibrated theoretical spectra of giant planets and brown dwarfs for various ages and masses. Baraffe et al. (2003) give theoretical absolute magnitudes at all of the bands typically used in AO imaging, including  $L'$  and  $M$ , for giant planets, brown dwarfs, and low-mass stars covering a similar range of age as the Burrows et al. (2003) models but with a wider range and finer spacing in mass. We have integrated the Burrows et al. (2003) spectra to give absolute magnitudes in the  $L'$  and  $M$ -bands, and have found that both sets of models have enough grid points in mass, age space to be reasonably interpolated to give the  $L'$  or  $M$  band magnitudes for all planets of interest for our survey.

### 4.4.2 Absolute Magnitudes from the Burrows et al. (2003) Models

The Burrows et al. (2003) models, which are available for download from <http://zenith.as.arizona.edu/~burrows/>, give the flux received from each object at a distance of 10pc, for 5000 finely spaced wavelengths ranging from 0.43 to 300  $\mu\text{m}$ , in units of milli-Janskys (1 milli-Janskys is  $10^{-29}\text{Wm}^{-2}\text{Hz}^{-1}$ ). To convert this into integrated absolute magnitude for the  $L'$  and  $M$  bands, we sum over the appropriate wavelength intervals, using the boxcar approximations to the two filters given in Table 7.5 of Cox (2000), assuming full transmission be-

tween the cut-on and cut-off wavelengths. As can be seen in Figure 4.4, these box-car approximations fit the measured curves of the Clio filters quite well, which is what we should expect since they are based on the same filter set as we use in Clio (the MKO filters; see Cohen et al. (1992), on which the Cox (2000) table is based). We note that this is not quite the same filter set as was used in Baraffe et al. (2003). We would expect this to cause only slight differences in the model predictions, as the filter sets are similar. See Section 4.4.4 for a fuller discussion.

For each spectral point in the Burrows et al. (2003) models that falls within the range of our filter, we convert the flux from milli-Janskys to  $\text{Wm}^{-2}\text{Hz}^{-1}$ . We form the product  $F_\nu \times \Delta\nu$  for each spectral point in the model, giving us a series of ‘monochromatic’ fluxes in units of  $\text{Wm}^{-2}$ . These could simply be summed over the bandpass of our filters to give the total flux in  $\text{Wm}^{-2}$ , but units of photons/ $\text{m}^2$ /sec are really more appropriate to model the response of a photoelectric detector. We divide each point by the energy per photon  $h\nu$  and then sum the points over the bandpasses of our filters to get the predicted flux from each model planet in units of photons/ $\text{m}^2$ /sec. This flux, of course, assumes the model planet is at a distance of 10pc.

To convert this photon flux to an absolute magnitude we again refer to Table 7.5 in Cox (2000), which gives the flux from Vega, a zero magnitude standard star, in commonly used astronomical infrared filters including  $L'$  and  $M$ , in units of photons/ $\text{m}^2$ /sec/ $\mu\text{m}$  (these values, again, are derived from Table 1 in Cohen et al. (1992)). This is easily converted to photons/ $\text{m}^2$ /sec by multiplying by the filter bandwidth in  $\mu\text{m}$ . Having calculated the photon flux in our filter from the model planet and from Vega, the magnitude of the planet is easily calculated using the formula  $\text{mag} = 2.5 \times \log_{10}(F_{\text{Vega}}/F_{\text{planet}})$ . Tables 4.3 and 4.4 give the absolute magnitudes we derive for all the Burrows et al. (2003) planet models, the the  $L'$

and  $M$  bands respectively.

It is important to note that what we have described above is simply a reduction of the Burrows et al. (2003) theoretical spectra to magnitudes, and has nothing to do with the photometric calibration of Clio. If we were making an attempt at photometric calibration from theoretical spectral models, our boxcar approximation to the filters, with no attempt made directly to include the effect of atmospheric absorption bands, could very reasonably be expected to cause problems! We do not at any point in this work try to predict from theory how many counts/sec Clio should receive in a given filter from a given object. All of our instrumental photometric calibrations were obtained empirically by observing standard stars of known brightness (see Chapter 3).

#### 4.4.3 Interpolating Between Model Points

Tables 4.3 through 4.6 give, respectively the  $L'$  absolute magnitudes of planets from Burrows et al. (2003), the  $M$  band values from Burrows et al. (2003), and then the corresponding  $L'$  and  $M$  band magnitudes from Baraffe et al. (2003). Both model sets extend from Jovian-mass planets up to brown dwarfs, and the models of Baraffe et al. (2003) go all the way up to low mass stars of 100 MJ, or 0.10 solar masses. For purposes of the planet-search simulations discussed in this chapter, objects above 13 MJ matter only to secure endpoints for interpolations, but in Chapter 3 we will present sensitivity results around each star that extend to the upper limits of the respective sample sets.

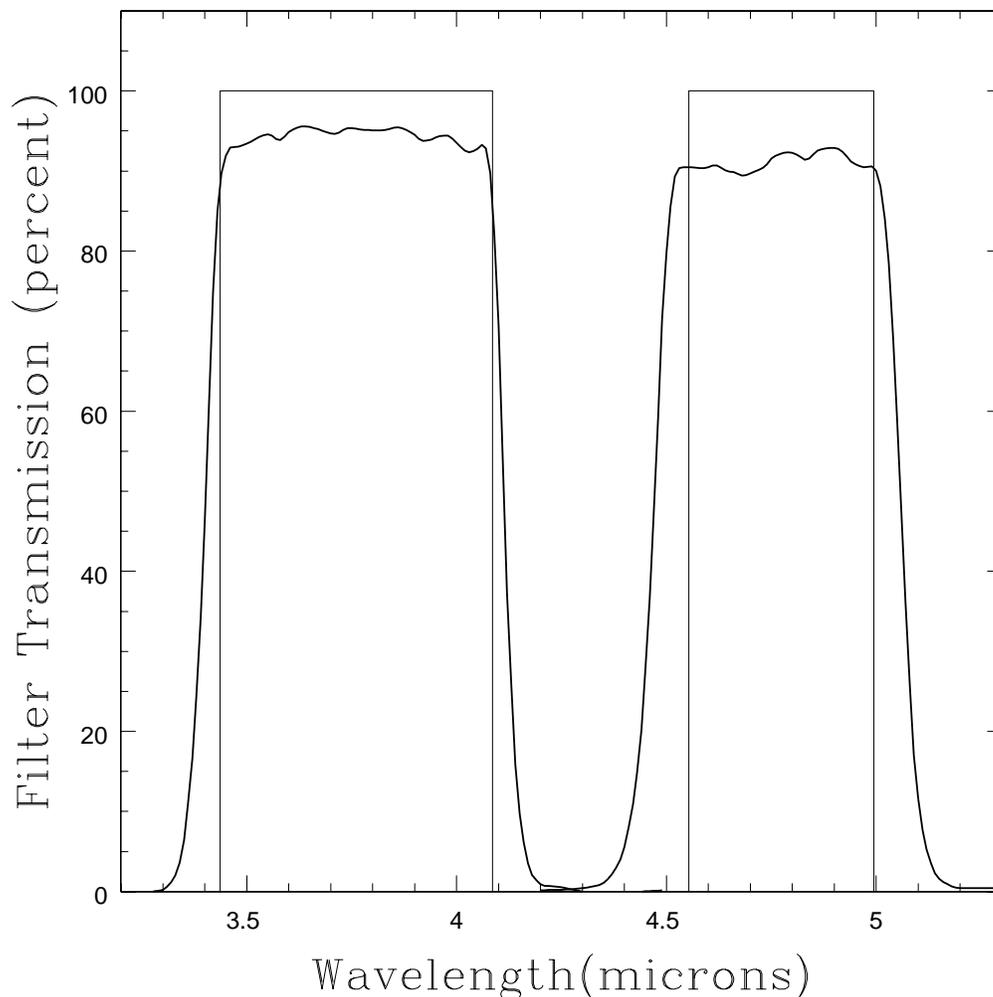


Figure 4.4 Clio L' (left) and M-band (right) filter transmission curves, with the boxcar approximations from Cox (2000) overplotted in thin lines. The boxcar approximations were used along with the magnitude-to-Jansky conversions from Cohen et al. (1992) to convert the theoretical spectra of Burrows et al. (2003) into band-averaged magnitudes like those provided in Baraffe et al. (2003). The height of the boxcar is of no importance since we were trying merely to match theoretical models, and our Clio observations were photometrically calibrated by other means.

Table 4.3. *L'* Band Absolute Mags from Burrows et al. (2003)

Planet Mass in MJ	Mag at 0.10 Gyr	Mag at 0.32 Gyr	Mag at 1.0 Gyr	Mag at 3.2 Gyr	Mag at 5.0 Gyr
1.0	19.074	23.010	27.870	33.50 <sup>a</sup>	35.50 <sup>a</sup>
2.0	16.793	19.351	23.737	28.398	29.479
5.0	14.500	16.397	18.588	22.437	24.407
7.0	13.727	15.390	17.336	20.131	21.574
10.0	12.888	14.437	16.246	18.480	19.466
15.0	12.00 <sup>b</sup>	13.61 <sup>b</sup>	14.773	16.816	17.691
20.0	11.30 <sup>b</sup>	12.98 <sup>b</sup>	14.190	15.967	16.766

Note. — The values in this table are calculated from spectral models which are discussed in Burrows et al. (2003) and are available online.

<sup>a</sup>No brightnesses for planets of these masses and ages appear in Burrows et al. (2003). Presumably the authors deemed such planets far too faint to be detected with any current or near-future telescope. This is almost certainly correct, but we have inserted ad hoc values to smooth the interpolations. They have zero or negligible effect on the interpolated brightnesses of planets we could actually detect.

<sup>b</sup>No brightnesses for planets of these masses and ages appear in Burrows et al. (2003), apparently because the surface temperatures were so high they fell out of the purview of the paper's subject (the coolest brown dwarfs). We have added values from Baraffe et al. (2003) and then fudged them slightly fainter to match the generally more pessimistic predictions of Burrows et al. (2003) and to insure smooth interpolations.

Table 4.4. *M* Band Absolute Mags from Burrows et al. (2003)

Planet Mass in MJ	Mag at 0.10 Gyr	Mag at 0.32 Gyr	Mag at 1.0 Gyr	Mag at 3.2 Gyr	Mag at 5.0 Gyr
1.0	14.974	16.995	19.987	25.0 <sup>a</sup>	26.0 <sup>a</sup>
2.0	14.023	15.313	17.807	21.295	22.163
5.0	13.014	14.017	15.153	17.167	18.537
7.0	12.618	13.561	14.558	16.126	16.909
10.0	12.189	13.096	14.093	15.315	15.951
15.0	11.55 <sup>b</sup>	12.60 <sup>b</sup>	13.370	14.512	14.990
20.0	11.29 <sup>b</sup>	12.21 <sup>b</sup>	13.069	14.122	14.580

Note. — The values in this table are calculated from spectral models which are discussed in Burrows et al. (2003) and are available online.

<sup>a</sup>No brightnesses for planets of these masses and ages appear in Burrows et al. (2003). Presumably the authors deemed such planets far too faint to be detected with any current or near-future telescope. This is almost certainly correct, but we have inserted ad hoc values to smooth the interpolations. They have zero or negligible effect on the interpolated brightnesses of planets we could actually detect.

<sup>b</sup>No brightnesses for planets of these masses and ages appear in Burrows et al. (2003), apparently because the surface temperatures were so high they fell out of the purvue of the paper’s subject (the coolest brown dwarfs). We have added values from Baraffe et al. (2003) and then fudged them slightly fainter to match the generally more pessimistic predictions of Burrows et al. (2003) and to insure smooth interpolations.

Table 4.5.  $L'$  Absolute Mags from Baraffe et al. (2003)

Planet Mass in MJ	Mag at 0.1 Gyr	Mag at 0.5 Gyr	Mag at 1.0 Gyr	Mag at 5.0 Gyr	Mag at 10.0 Gyr
0.5	19.57	23.09	24.15	29.0 <sup>a</sup>	31.5 <sup>a</sup>
1	17.41	20.93	22.4	27.0 <sup>a</sup>	30.0 <sup>a</sup>
2	15.94	18.66	20.18	23.33	25.0 <sup>a</sup>
3	15.21	17.52	18.84	22.22	23.36
4	14.59	16.85	17.91	21.28	22.59
5	14.06	16.32	17.38	20.48	21.92
6	13.67	15.92	16.87	19.78	21.29
7	13.26	15.52	16.54	19.16	20.71
8	12.97	15.19	16.18	18.64	20.18
9	12.63	14.88	15.93	18.18	19.7
10	12.34	14.59	15.65	17.91	19.26
12	10.9	13.51	14.88	17.41	18.48
15	10.83	13.41	14.56	16.96	17.93
20	10.53	12.78	13.72	16.18	17.21
30	9.82	11.7	12.67	15.04	16.04
40	9.39	10.95	11.88	14.17	15.03
50	9.04	10.44	11.25	13.38	14.23
60	8.78	10.1	10.77	12.73	13.5
70	8.55	9.81	10.32	11.6	12.27

Table 4.5—Continued

Planet Mass in MJ	Mag at 0.1 Gyr	Mag at 0.5 Gyr	Mag at 1.0 Gyr	Mag at 5.0 Gyr	Mag at 10.0 Gyr
72	8.51	9.75	10.23	11.16	11.55
75	8.44	9.67	10.08	10.54	10.56
80	8.36	9.53	9.84	9.97	9.97
90	8.19	9.25	9.4	9.42	9.41
100	8.05	9	9.07	9.07	9.07

Note. — The values in this table are taken from Tables 1-5 in Baraffe et al. (2003).

<sup>a</sup>No brightnesses for planets of these masses and ages appear in Baraffe et al. (2003). Presumably the authors deemed such planets far too faint to be detected with any current or near-future telescope. This is almost certainly correct, but we have inserted ad hoc values to smooth the interpolations. They have zero or negligible effect on the interpolated brightnesses of planets we could actually detect.

Table 4.6. *M* Band Absolute Mags from Baraffe et al. (2003)

Planet Mass in MJ	Mag at 0.1 Gyr	Mag at 0.5 Gyr	Mag at 1.0 Gyr	Mag at 5.0 Gyr	Mag at 10.0 Gyr
0.5	17.64	20.59	21.49	25.0 <sup>a</sup>	27.0 <sup>a</sup>
1	15.69	18.68	19.95	23.0 <sup>a</sup>	25.0 <sup>a</sup>
2	14.55	16.58	17.94	20.80	22.0 <sup>a</sup>
3	13.93	15.53	16.66	19.81	20.89
4	13.50	15.05	15.73	18.92	20.19
5	13.14	14.67	15.36	18.16	19.57
6	12.83	14.36	15.01	17.49	18.97
7	12.55	14.06	14.76	16.88	18.41
8	12.35	13.83	14.50	16.36	17.89
9	12.13	13.63	14.31	15.91	17.42
10	11.96	13.43	14.11	15.71	16.98
12	11.17	12.68	13.60	15.32	16.20
15	11.15	12.60	13.38	14.97	15.76
20	10.99	12.21	12.80	14.45	15.19
30	10.38	11.68	12.15	13.71	14.37
40	9.84	11.33	11.80	13.13	13.75
50	9.37	10.99	11.53	12.63	13.21
60	9.03	10.63	11.26	12.27	12.73
70	8.75	10.26	10.85	11.77	12.08

We needed accurate interpolation schemes that would give both the magnitude of a planet given its mass and age, and the mass of a planet given its age and magnitude. The former is needed for the Monte Carlo simulations discussed later in this Chapter, while the latter is required to translate measured sensitivities from our Clio observations into mass upper limits for planets orbiting the stars we observed.

For the Burrows et al. (2003) models, we handled the first case, the mass-to-magnitude conversion given age, by first performing a spline interpolation on magnitude vs.  $\log(\text{age})$  for each of the seven different masses in the Burrows et al. (2003) model grid. Thus we obtained magnitudes corresponding to each model mass for the specific input age. Then, we performed a spline interpolation on magnitude vs.  $\log(\text{mass})$  using the age-specific magnitude values, to output an interpolated magnitude given the input mass and age.

For the reverse process, obtaining mass given input age and magnitude, we began in the same way, by performing a spline interpolation on magnitude and age at each value of planet mass to get age-specific magnitudes for each mass. Then, we used spline interpolation on age-specific magnitude vs.  $\log(\text{mass})$  to get age-specific magnitudes for a more finely sampled set of masses (the set chosen was masses from 1 MJ to 20 MJ at 1 MJ intervals). The mass was obtained from the input magnitude and the finely sampled magnitude values by simple linear interpolation.

For the Baraffe et al. (2003) models our procedure was almost identical but slightly simpler. We began each case exactly as before by finding an age-specific magnitude corresponding to each model mass using spline interpolation on magnitude vs.  $\log(\text{age})$ . However, since the Baraffe et al. (2003) models had more finely spaced mass points, the interpolations on magnitude vs. mass were han-

Table 4.6—Continued

Planet Mass in MJ	Mag at 0.1 Gyr	Mag at 0.5 Gyr	Mag at 1.0 Gyr	Mag at 5.0 Gyr	Mag at 10.0 Gyr
72	8.70	10.19	10.75	11.54	11.75
75	8.63	10.08	10.58	11.06	11.08
80	8.53	9.89	10.27	10.43	10.43
90	8.35	9.53	9.72	9.73	9.73
100	8.19	9.22	9.31	9.31	9.30

Note. — The values in this table are taken from Tables 1-5 in Baraffe et al. (2003).

<sup>a</sup>No brightnesses for planets of these masses and ages appear in Baraffe et al. (2003). Presumably the authors deemed such planets far too faint to be detected with any current or near-future telescope. This is almost certainly correct, but we have inserted ad hoc values to smooth the interpolations. They have zero or negligible effect on the interpolated brightnesses of planets we could actually detect.

dled by simple linear interpolation for both the case of obtaining magnitude given mass and the case of obtaining mass given magnitude.

All of the spline interpolations on the model grids were done using the routines `spline` and `splint` from Press et al. (1992), with the edge conditions set for a 'natural' spline.

Figures 4.5 through 4.8 give the magnitude vs. age plots for the entire model grids of, respectively, the Burrows et al. (2003) models in the  $l'$  and  $M$  bands, and the citetbar models in the  $L'$  and  $M$  bands. Our interpolations are shown as solid lines through the points.

Open circles in Figures 4.5 through 4.8 are points that were not included in the model grids of the respective papers, which we inserted to get our interpolation codes to run straightforwardly. Burrows et al. (2003) do not give a model for a 1 MJ planets at ages 3.2 or 5.0 Gyr, or for 15 MJ or 20 MJ brown dwarfs at 0.1 and 0.32 Gyr. The former objects are presumably omitted because of their exceeding faintness; the latter objects because they were too hot to fall properly under the purvue of the Burrows et al. (2003) paper, which is about the coolest brown dwarfs and giant planets. We have inserted an ad hoc values for the 1 MJ planets at 3.2 and 5.0 Gyr, adjusting them to give the interpolations as natural an appearance as possible. This should have no practical effect on our Monte-Carlo simulations or other conclusions, since a 1 MJ planet even at age 3.2 Gyr is orders of magnitude too faint for us to detect even orbiting the nearest stars in our sample. For the young, hot brown dwarfs we have inserted values from Baraffe et al. (2003) or our interpolations thereof, and then fudged them slightly fainter to match the generally more pessimistic predictions of Burrows et al. (2003), and to make the interpolations look smooth and reasonable. Similarly, Baraffe et al. (2003) do not give magnitudes for an 0.5 or 1 MJ planet at age 5 or 10 Gyr, nor for

a 2 MJ planet at age 10 Gyr. We have, again, filled in the missing grid points with ad hoc values adjusted to give the spline interpolations a natural appearance. Our sensitivity never comes close to planets as old and low-mass as these, so the ad-hoc values will have no appreciable effect on our results.

Figures 4.9 through 4.12 give magnitude vs. mass plots for planets from 1 MJ to 20 MJ at various ages. As with the previous figures, the lines are our interpolations, solid points are from the theory papers, and open circles are values not supplied from theory, which we inserted to get the interpolations to work smoothly. These figures can be thought of as presenting isochrones; each line representing magnitude vs mass for a single freeze-frame in the evolution of an ensemble of planets. They illustrate the behavior of planet brightnesses with mass for the different models, and also verify that we implemented our interpolation schemes properly. As a final test of this, each figure also has a heavy line spanning most of the 1 Gyr isochrone. This is the result of the putting a range of magnitudes chosen to span this line into our interpolation function that takes magnitude and age as given and outputs mass (the rest of the lines, of course, were made by the reverse process of interpolating on mass and age to get magnitude). Our interpolation scheme for the Burrows et al. (2003) models was not precisely symmetric: getting magnitude from mass and age involved only spline interpolations, but getting mass from age and magnitude required a linear interpolation on a finely sampled vectors constructed by means of a spline. These figures indicate the linear interpolation tracks the line more than well enough for our purposes. Our interpolation scheme for the Baraffe et al. (2003) models was precisely symmetrical, so the fact that the heavy and fine lines for 1 Gyr overlay perfectly simply illustrates there were no errors in our implementation of the algorithm.

In conclusion, Figures 4.5 through 4.12 exist mainly to demonstrate that our

interpolations schemes work properly and produce reasonable curves. Having confirmed this, we can move on to more interesting subjects.

#### 4.4.4 Comparing the two Model Sets

Figures 4.13 and 4.14 present the differences between the Burrows et al. (2003) and Baraffe et al. (2003) models for the  $L'$  and  $M$  bands, respectively. The Baraffe et al. (2003) models predict substantially brighter planets at  $L'$  relative to the Burrows et al. (2003) models, but in the  $M$  band the models agree more closely, with the Baraffe et al. (2003) models sometimes fainter and sometimes brighter. Since most of our observations were done in the  $L'$  band, the Baraffe et al. (2003) models predict we should have detected more planets than the Burrows et al. (2003) models. The difference, however, is not extreme. We were not sensitive to the very low mass, old planets for which the models diverge most strongly. The odds of detecting planets using the different models sets are discussed in Sections 4.6 and 4.7.

We note that some of the difference between the Burrows et al. (2003) and Baraffe et al. (2003) models may be due to the different filter sets used. We integrated the Burrows et al. (2003) model spectra using the MKO filter set used in Clio, but the magnitudes provided in Baraffe et al. (2003) are based on the Johnson  $M$  band and the Johnson-Glass  $L'$  filters. Based on the transmission curve in Bessell & Brett (1988), it would appear that the Johnson  $M$  band is indistinguishable from the Clio MKO filter, and Reid & Cruz (2002) comment that only subtle differences exist between different  $L'$  filters. We expect, therefore, that much of the difference shown in Figures 4.13 and 4.14 represents inherent differences in the models. Filter differences probably could not explain the magnitude and behavior of the model disagreements, especially the extreme differences seen for cold, low-mass objects at advanced ages. We note that two model sets do differ

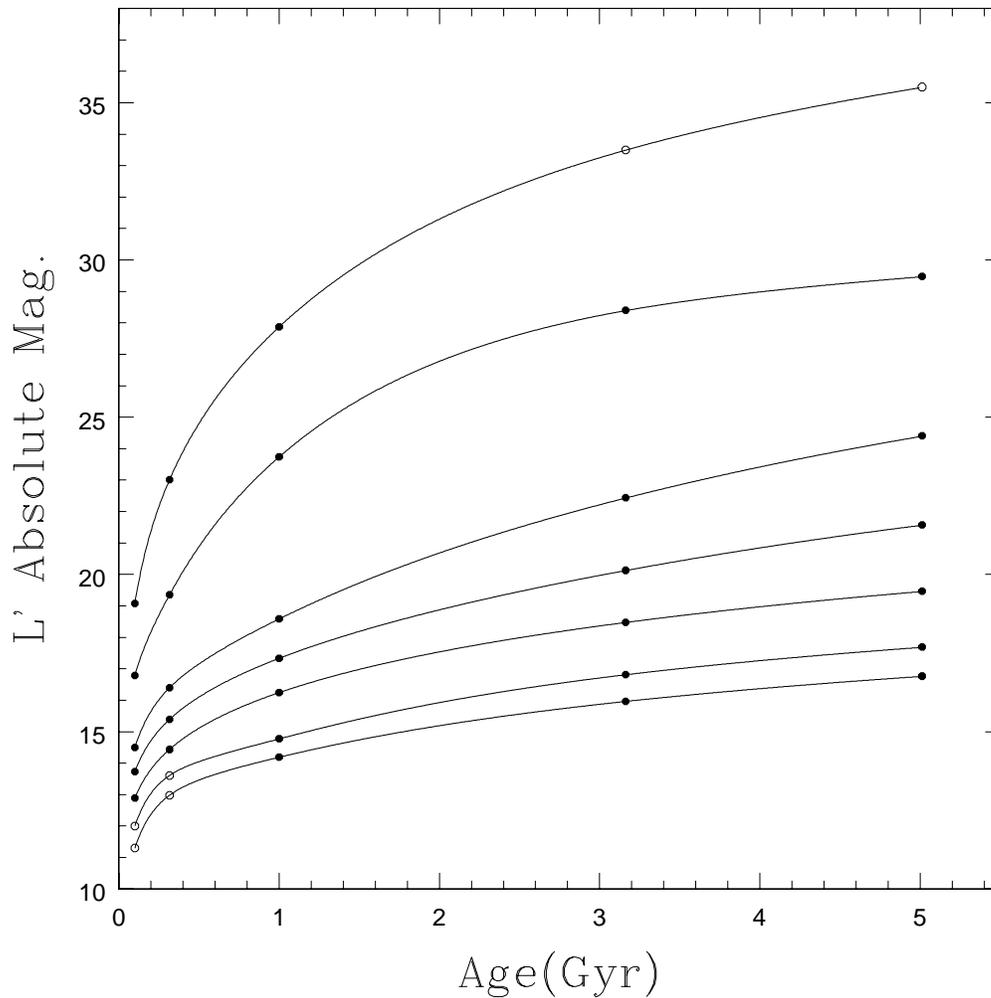


Figure 4.5  $L'$  Absolute Magnitude vs Age. The solid points are from our integrations of the planet models discussed in Burrows et al. (2003); the lines are our interpolation between the points. In order from top to bottom, the curves are for planets/brown dwarfs of 1, 2, 5, 7, 10, 15, and 20 MJ. The open circles are grid points for which no Burrows et al. (2003) model is available. For the 1 MJ planet at 5 Gyr we inserted an ad hoc value to smooth the interpolations. This has no practical effect since the object represented is far too faint to detect. For the two most massive planets at the youngest ages we inserted values from Baraffe et al. (2003) and then fudged them slightly fainter to be consistent with the generally more pessimistic predictions of Burrows et al. (2003) and to smooth the interpolations.

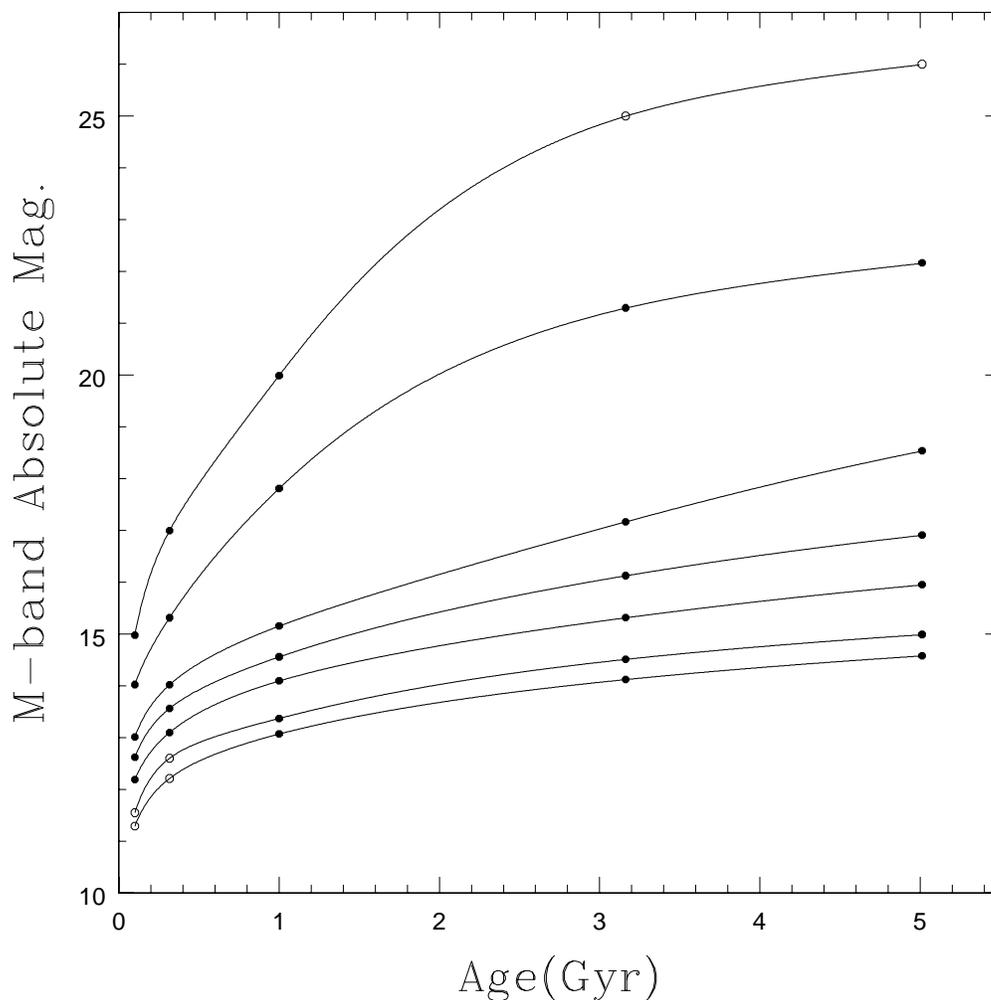


Figure 4.6 *M* band Absolute Magnitude vs Age. The solid points are from our integrations of the planet models discussed in Burrows et al. (2003); the lines are our interpolation between the points. In order from top to bottom, the curves are for planets/brown dwarfs of 1, 2, 5, 7, 10, 15, and 20 MJ. The open circles are grid points for which no Burrows et al. (2003) model is available. For the 1 MJ planet at 5 Gyr we inserted an ad hoc value to smooth the interpolations. This has no practical effect since the object represented is far too faint to detect. For the two most massive planets at the youngest ages we inserted values from Baraffe et al. (2003) and then fudged them slightly fainter to be consistent with the generally more pessimistic predictions of Burrows et al. (2003) and to smooth the interpolations.

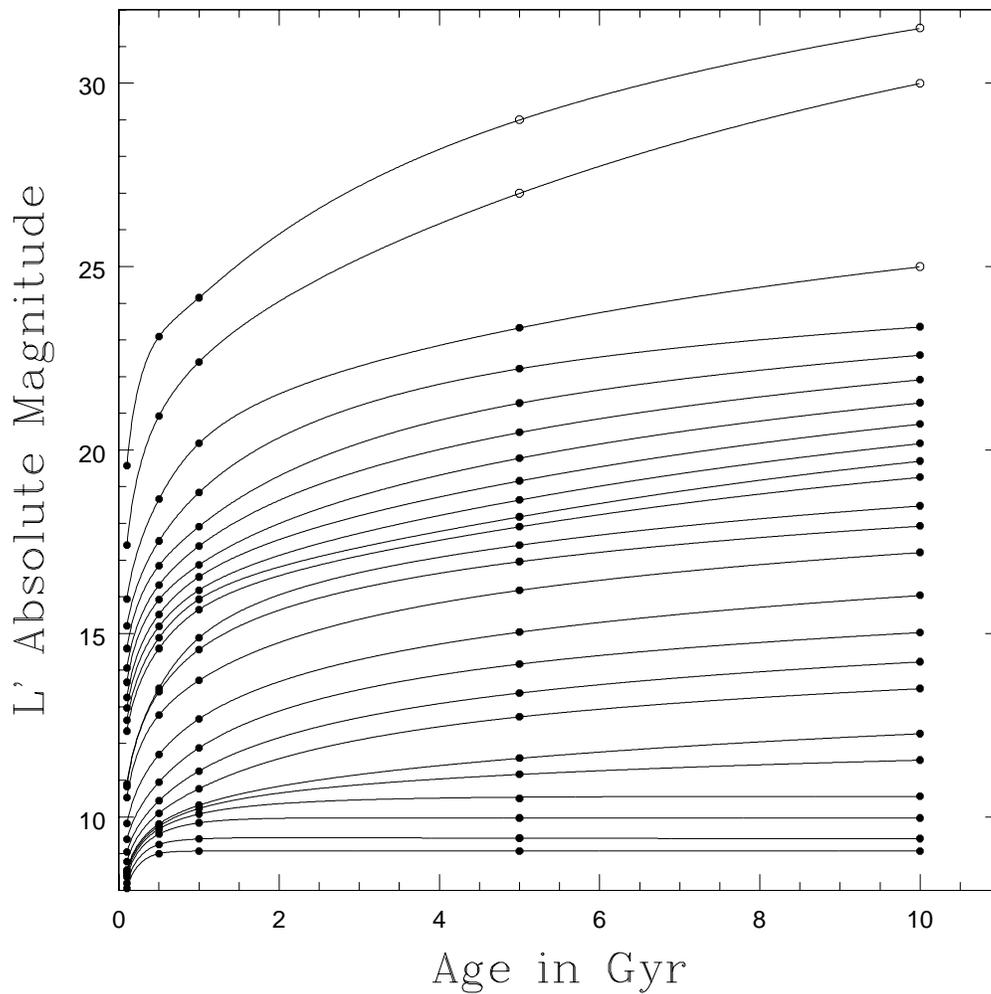


Figure 4.7  $L'$  Absolute Magnitude vs Age. The solid points are from Table 1-5 in Baraffe et al. (2003); the lines are our interpolation between the points. In order from top to bottom, the curves are for planets/brown dwarfs/low-mass stars of 0.5, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 15, 20, 30, 40, 50, 60, 70, 72, 75, 80, 90, and 100 MJ. The open circles at 5 and 10 Gyr ages for the 0.5, 1, and 2 MJ objects are not found in Baraffe et al. (2003) since the objects were so faint; ad hoc values have been inserted to make the interpolations look good. This has no practical effect since the objects represented are much too faint to detect.

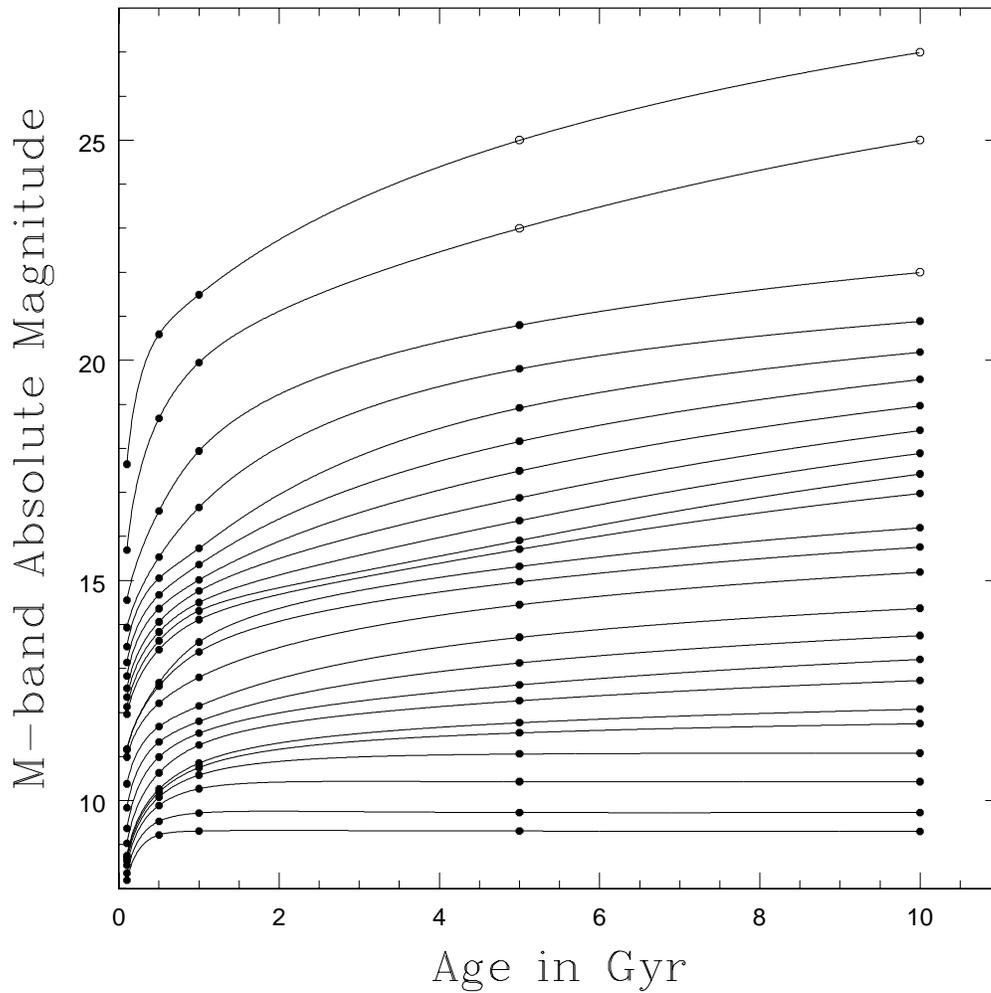


Figure 4.8 *M* band Absolute Magnitude vs Age. The solid points are from Table 1-5 in Baraffe et al. (2003); the lines are our interpolation between the points. In order from top to bottom, the curves are for planets/brown dwarfs/low-mass stars of 0.5, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 15, 20, 30, 40, 50, 60, 70, 72, 75, 80, 90, and 100 MJ. The open circles at 5 and 10 Gyr ages for the 0.5, 1, and 2 MJ objects are not found in Baraffe et al. (2003) since the objects were so faint; ad hoc values have been inserted to make the interpolations look good. This has no practical effect since the objects represented are much too faint to detect.

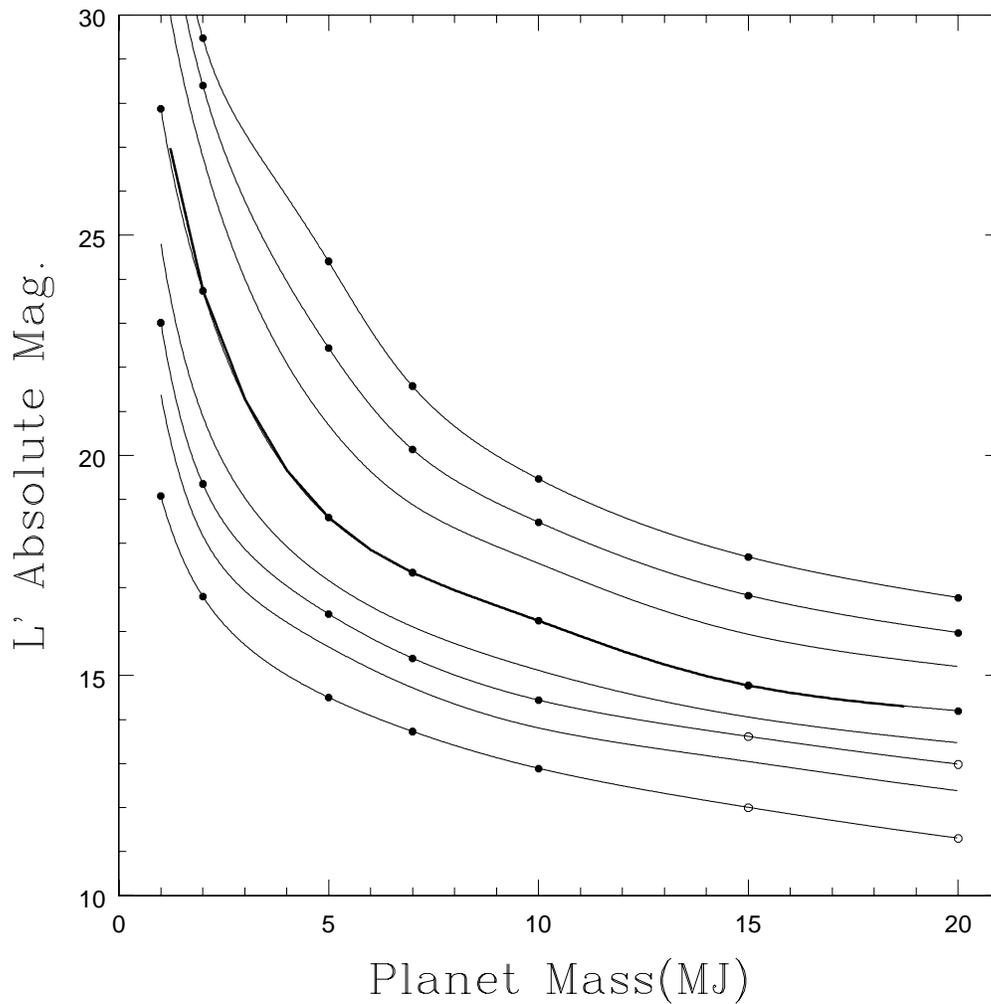


Figure 4.9  $L'$  Absolute Magnitude vs Mass. The solid points are from our integrations over the Burrows et al. (2003) theoretical spectra. The open circles are as discussed in the preceding figures. The thin lines are from our interpolation method with mass and age as input and magnitude as output. In order from top to bottom, the lines are for 5.0, 3.2, 2.0, 1.0, 0.5, 0.32, 0.2 and 0.1 Gyr ages. The thicker line along most of the 1.0 Gyr line is from our interpolation with magnitude and age as input and mass as output, testing its accuracy and consistency with the reverse case. It appears satisfactory.

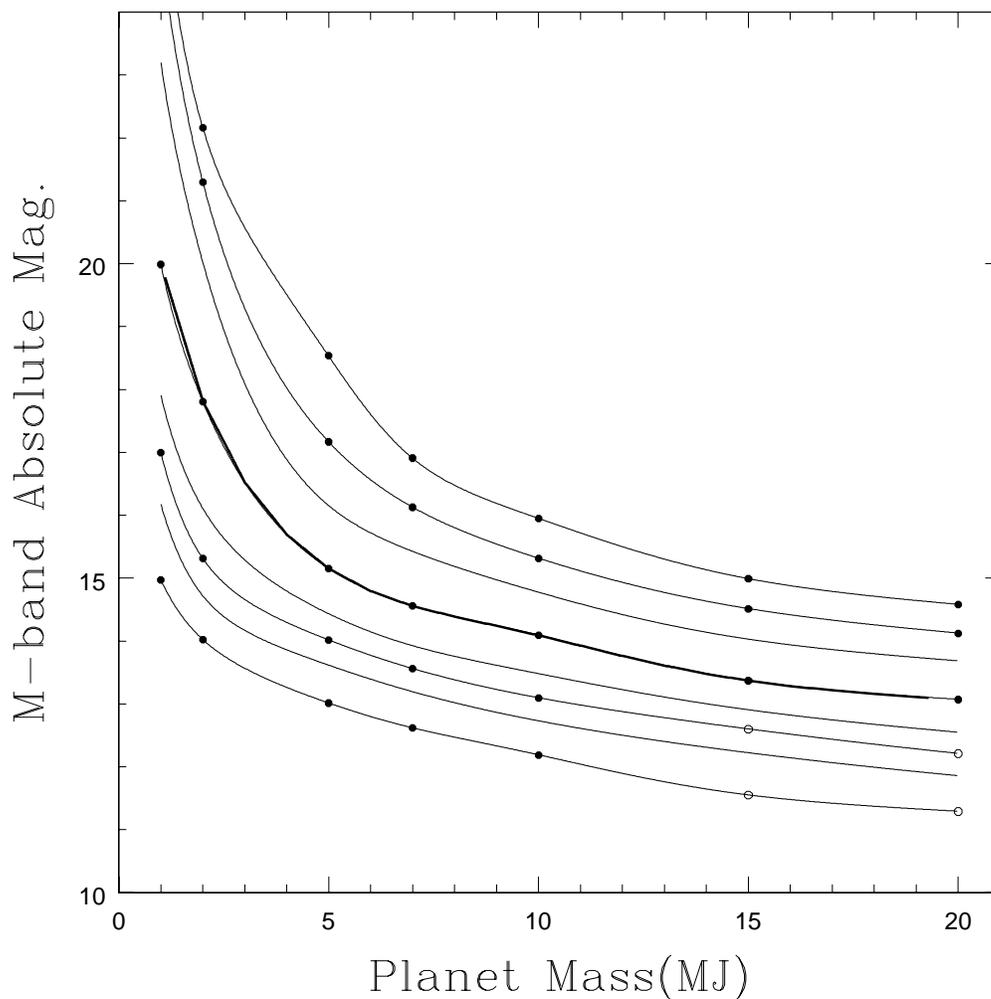


Figure 4.10 *M* Band Absolute Magnitude vs Mass. The solid points are from our integrations over the Burrows et al. (2003) theoretical spectra. The open circles are as discussed in the preceding figures. The thin lines are from our interpolation method with mass and age as input and magnitude as output. In order from top to bottom, the lines are for 5.0, 3.2, 2.0, 1.0, 0.5, 0.32, 0.2 and 0.1 Gyr ages. The thicker line along most of the 1.0 Gyr line is from our interpolation with magnitude and age as input and mass as output, testing its accuracy and consistency with the reverse case. It appears satisfactory.

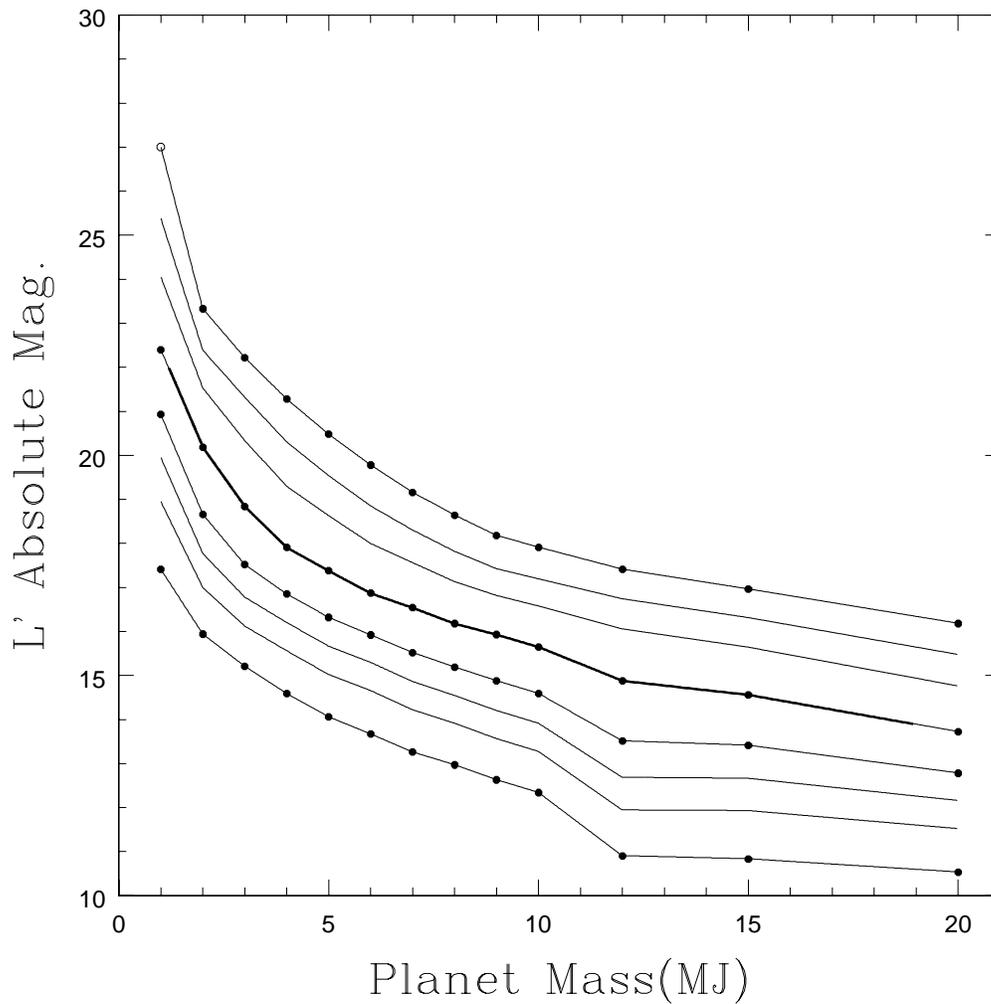


Figure 4.11  $L'$  Absolute Magnitude vs Mass. The solid points are from Tables 1-5 in Baraffe et al. (2003). The open circles are as discussed in the preceding figures. The thin lines are from our interpolation method with mass and age as input and magnitude as output. In order from top to bottom, the lines are for 5.0, 3.2, 2.0, 1.0, 0.5, 0.32, 0.2 and 0.1 Gyr ages. The thicker line along most of the 1.0 Gyr line is from our interpolation with magnitude and age as input and mass as output. Mathematical logic indicates the two interpolation cases used for the Baraffe et al. (2003) models should match exactly; that they do shows the algorithms were implemented correctly.

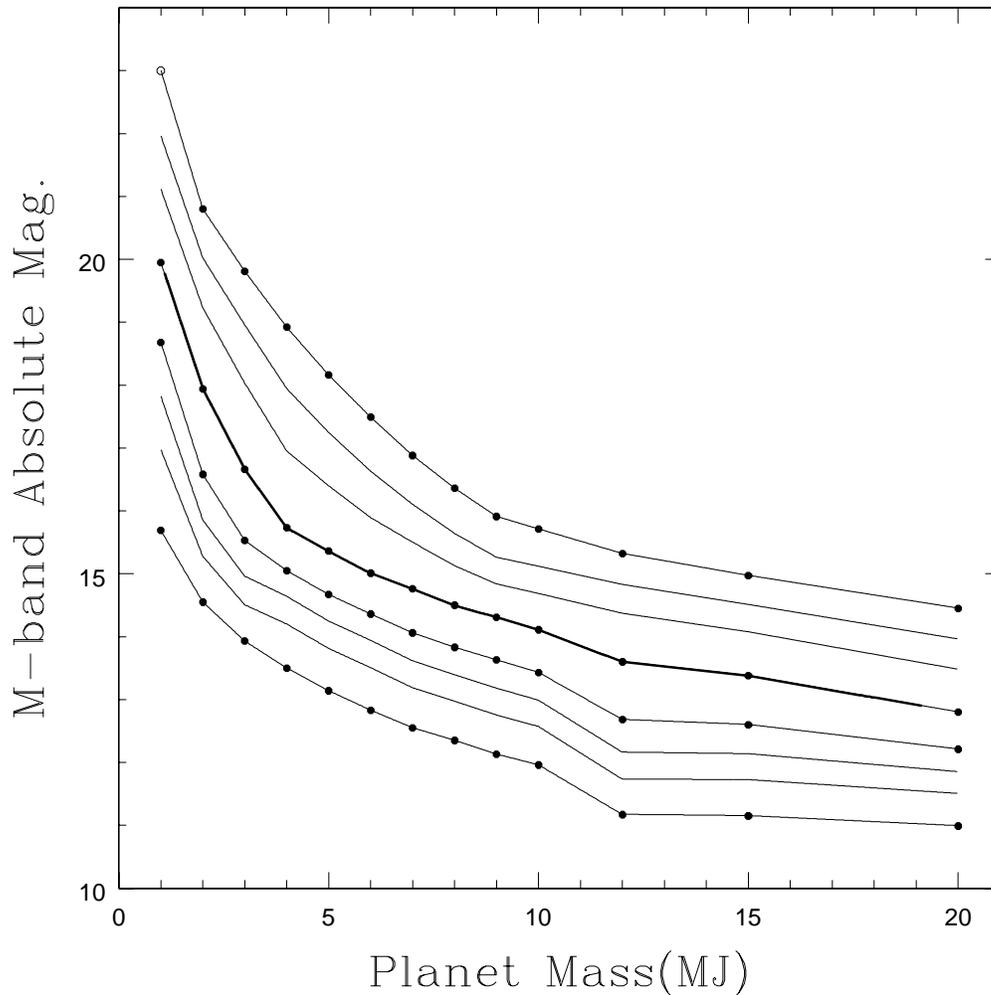


Figure 4.12 *M* Band Absolute Magnitude vs Mass. The solid points are from Tables 1-5 in Baraffe et al. (2003). The open circles are as discussed in the preceding figures. The thin lines are from our interpolation method with mass and age as input and magnitude as output. In order from top to bottom, the lines are for 5.0, 3.2, 2.0, 1.0, 0.5, 0.32, 0.2 and 0.1 Gyr ages. The thicker line along most of the 1.0 Gyr line is from our interpolation with magnitude and age as input and mass as output. Mathematical logic indicates the two interpolation cases used for the Baraffe et al. (2003) models should match exactly; that they do shows the algorithms were implemented correctly.

in the  $T_{\text{eff}}$  values they predict for a planet with a given mass and age.

#### 4.5 Introducing the Monte Carlo Simulations

Chapter 3 to this point has chiefly been leading up to the results we will now discuss. Since we did not find any planets in our survey, the key question of interest is how likely this result was. If the most reasonable distributions of planets predict that we had only a small chance of detecting one, it is not at all surprising that we found none. If, on the other hand, the expected planet distributions indicate it was overwhelmingly probable we would find at least one planet, then the fact that we detected none is surprising and may rule out distributions previously considered likely to apply.

We have chosen to address the question of how likely it was that our survey would find no planets given various input distributions by means of Monte Carlo simulations. We have constructed Monte Carlo simulations in which many realizations of our survey are run, and the fraction in which no planets were detected serves as an estimate of the likelihood of a null result given the input distribution. In each survey realization each real star we observed is considered in turn. It is randomly determined to have zero, one, or more planets, with probabilities set by the input power law slopes, the truncation radius for the distribution of planet orbits, and the normalization described in section 4.3. If the star is assigned planets, their properties are randomly drawn from the input distributions. The magnitude of each planet is then calculated based on the planet mass and the known age and distance to the star. This magnitude is then compared with our measured sensitivity at the appropriate projected separation from the star, and the planet is classified as detected or undetected, accordingly.

In Section 4.6 we present detailed results from Monte Carlo simulations using

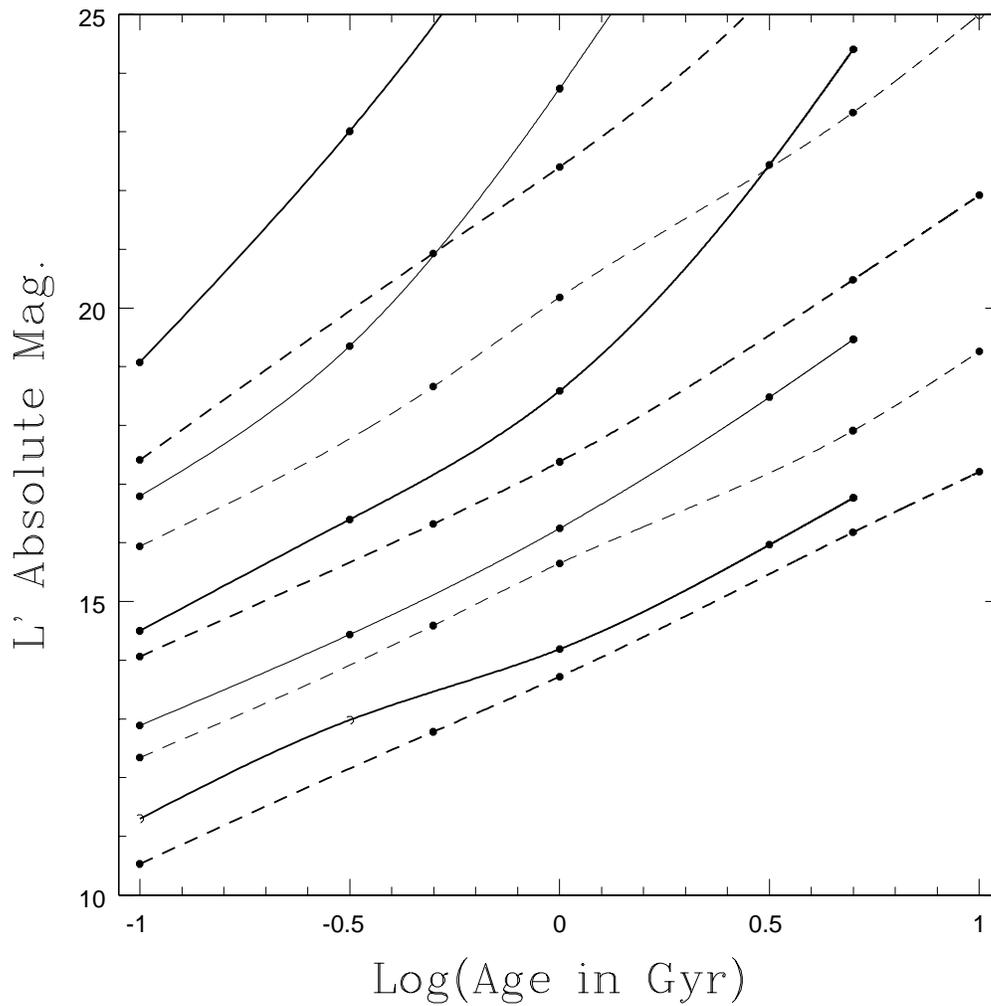


Figure 4.13  $L'$  magnitude vs  $\log(\text{age})$  for the Burrows et al. (2003) models (solid lines) and the Baraffe et al. (2003) models (dashed lines). From top to bottom the lines are for planets/brown dwarfs of 1, 2, 5, 10, and 20 MJ. The lines alternate between bold and normal from mass to mass to guide the eye in making comparisons: that is, the two curves for 1 MJ planets are bold, those for 2 MJ planets are thin, those for 5 MJ planets are bold, etc. The Baraffe et al. (2003) models predict substantially brighter planets in general.

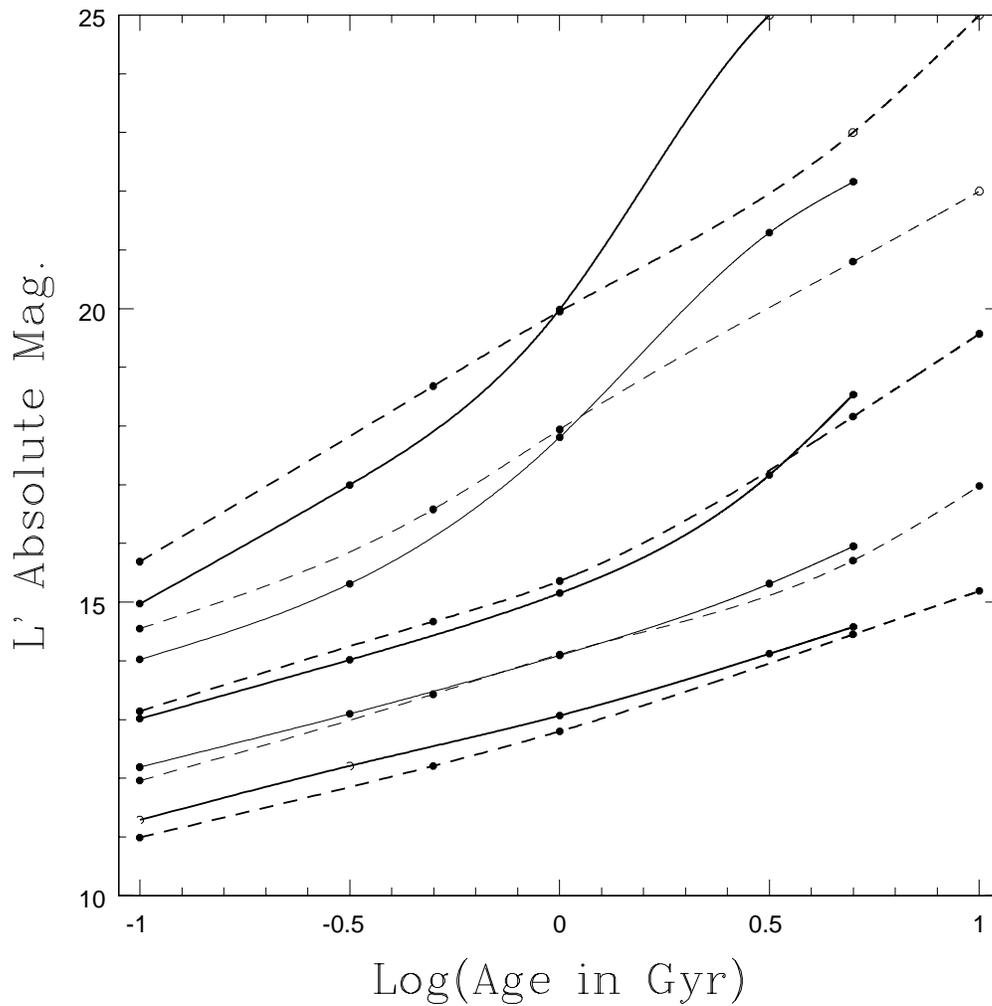


Figure 4.14  $M$  band magnitude vs  $\log(\text{age})$  for the Burrows et al. (2003) models (solid lines) and the Baraffe et al. (2003) models (dashed lines). From top to bottom the lines are for planets/brown dwarfs of 1, 2, 5, 10, and 20 MJ. The lines alternate between bold and normal from mass to mass to guide the eye in making comparisons: that is, the two curves for 1 MJ planets are bold, those for 2 MJ planets are thin, those for 5 MJ planets are bold, etc. The models agree fairly well, with the Baraffe et al. (2003) models sometimes predicting fainter planets than Burrows et al. (2003), in strong contrast to the  $L'$  results.

what seems to us the most reasonable guess at what the true extrasolar planet distribution will be: power laws in mass and radius with the adopted slopes recorded in 4.3, and a truncation radius of 40 AU, chosen to be just slightly beyond the orbit of our own solar system's outermost sizeable planet, Neptune. These simulations tell us not only how surprising it is that we found no planets, but also what kind of planets we might have expected to find — and what sort of planets a new, similar survey with improved sensitivity might actually detect. In Section 4.7 we consider a broad range of possible planet distributions, and show which ones are confidently ruled out by the fact that we found no planets.

## 4.6 Monte-Carlo Simulations: A Detailed Look

### 4.6.1 Constructing the Simulations

We have already described the basic outline of our Monte Carlo simulations in Section 4.5. Here a more detailed picture will be given, designed to allow the reader to perform identical simulations if desired.

The statistical distributions of planets input to our simulations had 4 parameters: the power law scaling for mass, the power law scaling for orbital semimajor axis, the truncation radius for the semimajor axis distribution, and the normalization. The normalization, defined as the fraction of stars bearing planets with masses between 1.0 and 13.0 MJ in orbits between 0.3 and 2.5 AU, was kept fixed at the value of 0.0329 determined in Section 4.3. For the simulations described in detail here, the power law slopes were set at their adopted values from Section 4.3, and the truncation radius was set to 40 AU, corresponding very roughly to Neptune's orbit as described in Section 4.5.

The simulation began with a calculation of the expected number of planets a given star would have. This was done using the following formula, where  $N_{RV}$  is

the RV normalization of 0.0329,  $\alpha$  is the power law slope of the mass distribution,  $\gamma$  is the power law slope of the semimajor axis distribution,  $M_{1RV}$  and  $M_{2RV}$  are the lower and upper limit masses for the RV normalization region (always 1.0 and 13.0 MJ),  $a_{1RV}$  and  $a_{2RV}$  are the inner and outer limit semimajor axes for the RV normalization region (0.3 and 2.5 AU),  $M_{1S}$  and  $M_{2S}$  are the lower and upper limit masses for the simulation, and  $a_{1S}$  and  $a_{2S}$  are the inner and outer limit semimajor axes for the simulation:

$$N_S = N_{RV} \times \frac{\int_{M_{1S}}^{M_{2S}} M^\alpha dM}{\int_{M_{1RV}}^{M_{2RV}} M^\alpha dM} \times \frac{\int_{a_{1S}}^{a_{2S}} a^\gamma da}{\int_{a_{1RV}}^{a_{2RV}} a^\gamma da}. \quad (4.1)$$

In general we have used the same upper and lower mass limit for the simulation as for the RV normalization:  $M_{1S} = M_{1RV} = 1.0\text{MJ}$  and  $M_{2S} = M_{2RV} = 13.0\text{MJ}$ , so the first ratio of integrals has been 1.0. We have also set  $a_{1S} = a_{1RV} = 0.3\text{AU}$ . However, the truncation radius  $a_{2S}$  has generally been much higher than the outer limit for the RV normalization,  $a_{2RV} = 2.5\text{AU}$ . For the simulations described in this Section,  $a_{2S} = 40\text{AU}$ , and the second integral ratio in the above equations is substantially greater than 1.0. Specifically, with the truncation radius  $a_{2S}$  set to 40 AU we get  $N_S = 0.258$ : the expected number of planets for each star is 0.258. This is higher than the 0.0329 from the RV normalization because the ratio

$$\frac{\int_{a_{1S}}^{a_{2S}} a^\gamma da}{\int_{a_{1RV}}^{a_{2RV}} a^\gamma da}. \quad (4.2)$$

is greater than 1.0; this quantifies mathematically the obvious expectation that more planets will be found between 0.3 and 40 AU than between 0.3 and 2.5 AU.

Once the expected number of planets per star has been determined, our simulation code loops through the stars, performing the desired number of realizations of an observation on each star. For the two simulations discussed in this Sec-

tion, one using the Burrows et al. (2003) models and the other using the Baraffe et al. (2003) models, we ran 50000 realizations of the survey.

A realization of survey observations of a single star begins with the assignment of a number of planets. This number is randomly drawn from a Poisson distribution with mean equal to the expected number of planets per star calculated using Equation 4.1. For the simulations in this Section, zero was the most likely number of planets, but stellar systems could have 1 or 2 planets with non-negligible likelihood.

If the star has planets, the mass and orbital radius of each planet are drawn from the input power law distributions. The eccentricity is drawn from the distribution given in Juric & Tremaine (2007), which we verify to be a good fit to the observed eccentricities of RV planets in Figure 4.3. From the eccentricity and the semimajor axis, the apastron and periastron distances can be calculated. If the planet is orbiting the primary star of a binary, we require that its apastron distance be no more than  $1/3$  the projected distance to the other star. If it is orbiting the secondary we require its apastron to be no more than  $1/4$  the projected distance to the primary if the secondary is only slightly less massive than the primary, and no more than  $1/5$  the distance if the secondary is much less massive. If the planet is orbiting both stars of a binary, its periastron distance is required to be at least three times the projected separation of the two stars. Planets in binary systems whose orbits do not satisfy these conditions are considered unstable and ejected from the simulation without being recorded.

For planets passing the orbital stability tests, a full set of parameters describing the orientation of the orbit and the planet's orbital phase are drawn from the appropriate distribution, and input into Matt Kenworthy's binary star code to calculate a projected separation for the planet. Besides those already listed, these

parameters are the orbital inclination,  $i$ , chosen from the distribution  $P(i) \propto \sin(i)$  for  $i \in (0^\circ, 90^\circ)$ , the longitude of periastron,  $\omega$ , distributed uniformly between  $0^\circ$  and  $360^\circ$  the PA of the ascending node,  $\Omega$ , distributed the same way, and the phase of the orbit, distributed uniformly between 0.0 and 1.0, with 0.0 or 1.0 corresponding to the periastron and 0.5 to apastron. Note that  $\Omega$  has no influence on the separation and is included only to make the orbit code run (and produce cool, realistic-looking orbits for Figures 4.6.2 through 4.6.2).

Once the separation of the planet has been obtained using the randomly selected orbital parameters and Matt Kenworthy's orbit code, the magnitude is easily calculated from the planet mass and the known age and distance to the star, using the models and interpolation methods described in Section 4.4. The magnitude and separation of the planet are compared against our observational sensitivity curves around that specific star to determine if the planet is detected.

Our sensitivity files give  $10\sigma$  sensitivity in magnitudes vs. separation from the star. Rays, ghosts, and other image artifacts described in more detail in Chapter 3 caused our sensitivity at a given radius to vary azimuthally. For this reason, we calculated 10 different sensitivity values for each separation. These values are the 0th, 10th, 20th, 30th, 40th, 50th, 60th, 70th, 80th, and 90th percentiles of the sensitivity for that radius as it varies azimuthally around the star. Figure 4.15 gives a plot of our full sensitivity file for the star GJ 117 as an example. Our files record  $10\sigma$  sensitivities, but we have determined by careful blind tests that we can recover  $7\sigma$  point sources with at least 50% completeness (see Chapter 3). To account for this, our Monte Carlo simulations randomly choose to compare the simulated planet's magnitude against either the observational  $10\sigma$  or  $7\sigma$  limits, with an even chance on either. The simulation then randomly selects a number between 1 and 10 to decide which percentile sensitivity will be used. Finally, the

planet is recorded as detected or undetected.

#### 4.6.2 Simulation Results

The most significant single number to come out of each Monte Carlo simulation is the probability of detecting zero planets. If this is low enough, the model planet distributions input to the simulation is ruled out by our survey. For this Section we carried out two identical Monte Carlo simulations, each involving 50000 realizations of our survey with planets drawn from the adopted power laws truncated at 40 AU. However, one simulation used the Burrows et al. (2003) planet models to convert mass and age to magnitude, and the other used the Baraffe et al. (2003) models. The simulation using the Burrows et al. (2003) models predicted a 63.4% chance of finding no planets, while the one using the more optimistic Baraffe et al. (2003) models gave the probability of a null result at 49.7%. Table 4.7 gives the probabilities of detecting one or more planets for the respective surveys.

Perhaps the next most interesting question is what were the properties of the detected planets. Table 4.8 gives the min, median, and max for the masses, orbital semimajor axes, separations, magnitudes, and detection levels for all the detected planets in each of the two simulations.

The most interesting thing about this table is the large median detection significance. Using either model set, the significance of more than half the detected planets is above  $16\sigma$ . This corresponds to a glaringly obvious detection. When an extrasolar planet is finally imaged, the observer's comment may well be, "Oh my goodness, what is THAT?" rather than, "There seems to be a subtle noise burst here."

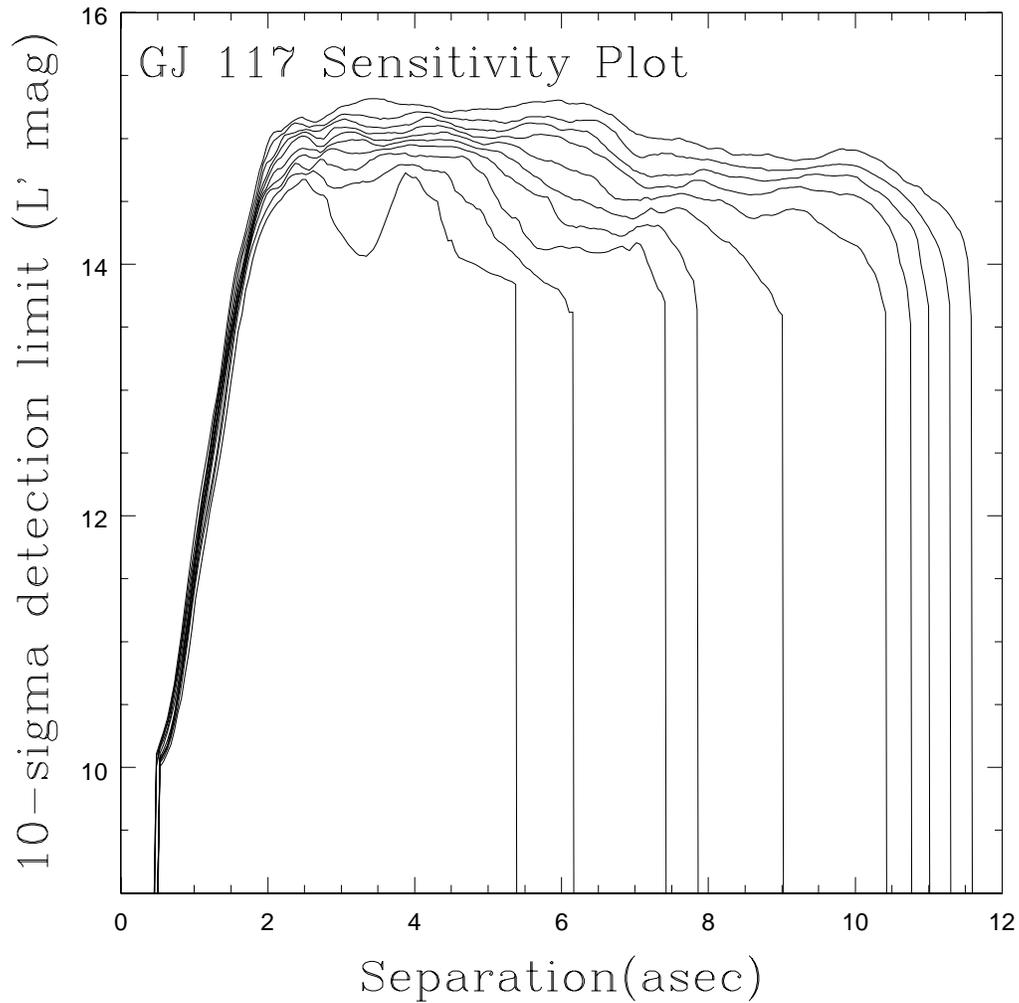


Figure 4.15  $10\sigma$  sensitivity obtained around the star GJ 117,  $L'$  magnitudes plotted versus separation in asec. From top to bottom, the lines are the 0th, 10th, 20th, 30th, 40th, 50th, 60th, 70th, 80th, and 90th percentiles in sensitivity for each separation. The sensitivity varies azimuthally about the star due to ghosts, rays, variable coverage, and other observational issues.

Table 4.7. Probabilities of Detecting  $N$  Planets

$N$	Prob. of Detecting $N$ Planets (Burrows Models)	Prob of Detecting $N$ Planets (Baraffe Models)
0	63.4%	49.7%
> 0	36.6%	50.3%
1	28.7%	34.6%
2	6.8%	12.3%
3	1.0%	2.8%
4	0.1%	0.5%
5	0.0%	0.1%

Table 4.8. Statistics of Detected Planets

Quantity	Simulation	Min	Median	Max
<b>Planet Mass</b>	<b>Burrows</b>	<b>2.24 MJ</b>	<b>8.92 MJ</b>	<b>13.00 MJ</b>
Planet Mass	Baraffe	1.93 MJ	8.51 MJ	13.00 MJ
<b>Semimaj. Axis</b>	<b>Burrows</b>	<b>2.75 AU</b>	<b>26.60 AU</b>	<b>40.00 AU</b>
Semimaj. Axis	Baraffe	1.73 AU	26.07 AU	40.00 AU
<b>Separation</b>	<b>Burrows</b>	<b>0.494 asec</b>	<b>2.257 asec</b>	<b>11.487 asec</b>
Separation	Baraffe	0.291 asec	2.164 asec	11.487 asec
<b>Det. Significance</b>	<b>Burrows</b>	<b>7.00 <math>\sigma</math></b>	<b>16.47 <math>\sigma</math></b>	<b>283.42 <math>\sigma</math></b>
Det. Significance	Baraffe	7.00 $\sigma$	19.23 $\sigma$	851.45 $\sigma$

The min and max values are not of great interest. The minimum values for the planet mass and semimajor axis may seem impressively small, and of course they do correspond to planets detectable in principle with  $L'$  and  $M$  band AO imaging, but it must not be forgotten that these are the minimum values drawn from 50,000 realizations of our survey. To say we would be lucky to find planets like this is an understatement.

The median values are the ones of real interest. Clearly, the planets our survey was most likely to find were very massive, nearly 9MJ. They were found mostly in orbits comparable to Neptune's in our own solar system. We were likely to find them at separations of around 2 asec from their parent star.

With 50,000 realizations, subtle differences between the surveys clearly appear. The Baraffe et al. (2003) models, with their more optimistic planet brightness values, predict we should detect planets at higher median significance than do the Burrows et al. (2003) models. They also predict we should find planets ranging to lower masses, and that we will be able to see them at slightly smaller distances from their parent stars. These observations fit in perfectly with the broad picture of the Baraffe et al. (2003) models being more optimistic: they predict planets will be brighter down to lower masses, so lower mass planets will be detectable. They predict planets will be brighter over all, so they will be detectable closer in against the bright parent stars.

Figures 4.16 through 4.23 present the details of the planet distributions from the simulation using the Burrows et al. (2003) models in histogram form. The histogram of all the input planets is shown as a thin line, while the histogram of detected planets is indicated in bold. Observational biases thus become clear. Part of the use of these figures is as a check that, in fact, there are no bugs in the Monte Carlo code and the input distributions and observational biases are what

they should be.

Figures 4.16 and 4.17, giving the histograms for planet masses and orbital semimajor axes, show the expected huge observational biases toward massive planets in distant orbits. The continuous curve on the semimajor axis histogram plot, which traces the input power law distribution, does not quite match the histogram of input planets. The reason for this is rather interesting: planets orbiting one component of a binary system that came too close to the other to be stable were not allowed in the simulation. Since we had a number of binary stars, this has affected the distribution in the sense that some planets with large orbital semimajor axes are missing.

Figure 4.18, the eccentricity histogram, shows a small observational bias toward high eccentricities. The effect is so small it could not be seen in a simulation with only 5000 survey realizations due to statistical noise. The reason for the bias seems clear: planets in highly eccentric orbits spend more time near apastron, where they are most distant from the star and are therefore easier to detect against the stellar glare.

The orbital inclination histogram (Figure 4.19) shows a similar but clearer bias: toward face-on ( $i = 0^\circ$ ) orbits in which the planet is always far from the star. Histograms for the orbital orientation parameters  $\omega$  and  $\Omega$  were made but are not included, since they simply showed uniform distributions with no observational bias, as they should. Figure 4.20 shows the magnitude distributions of input and detected planets. Of course only the very brightest ones are detected; not surprisingly the distribution of input magnitudes extends to planets so faint that even the most ambitious new telescopes will not be able to detect them. Figure 4.21 shows the separation distributions of input and detected planets, which do not exactly mimic the orbital semimajor axis distributions mainly because of the wide

range of distances to stars in our sample. While the majority of simulated planets are undetectably close to their parent stars, some (presumably those orbiting our nearest target stars in very large orbits) actually fell beyond our approximately 11 arcsecond field of view!

Figure 4.22, a histogram of the log of the detection significance for planets detected in the simulation, illustrates again the non-negligible likelihood of planets being detected at very high significance. There are many  $7\sigma$  detections to be sure, but it is more likely that the first extrasolar planet image will be bright and obvious. Finally, Figure 4.23 shows the histograms for the planet orbital phase. The input, of course, is a uniform distribution, and the observational bias, not surprisingly, is strongly concentrated toward apastron (phase 0.5), where the planets are at their maximum distance from their stars and can most easily be detected without stellar interference.

A concern with any planet imaging survey is how strongly the results hinge on the best few stars. A survey of 50 stars may have far less statistical power than the number would imply if the best two or three stars had the lion's share of the probability of hosting detectable planets. The survey could really be a survey of 3 stars with 47 more irrelevantly tacked on.

Table 4.9 gives the fraction of planets detected around each star in our sample, using both the Burrows et al. (2003) models and the Baraffe et al. (2003) models to predict the planet flux. The stars are ordered from most to least probability of hosting a detectable planet based on the Burrows et al. (2003) planet models. Table 4.10 presents the same information with the stars binned by detected planet fraction, so it can be seen how many stars had very low detection fractions and were thus basically irrelevant to the Monte Carlo survey simulations, and on what fraction of the stars the simulation chiefly depended.

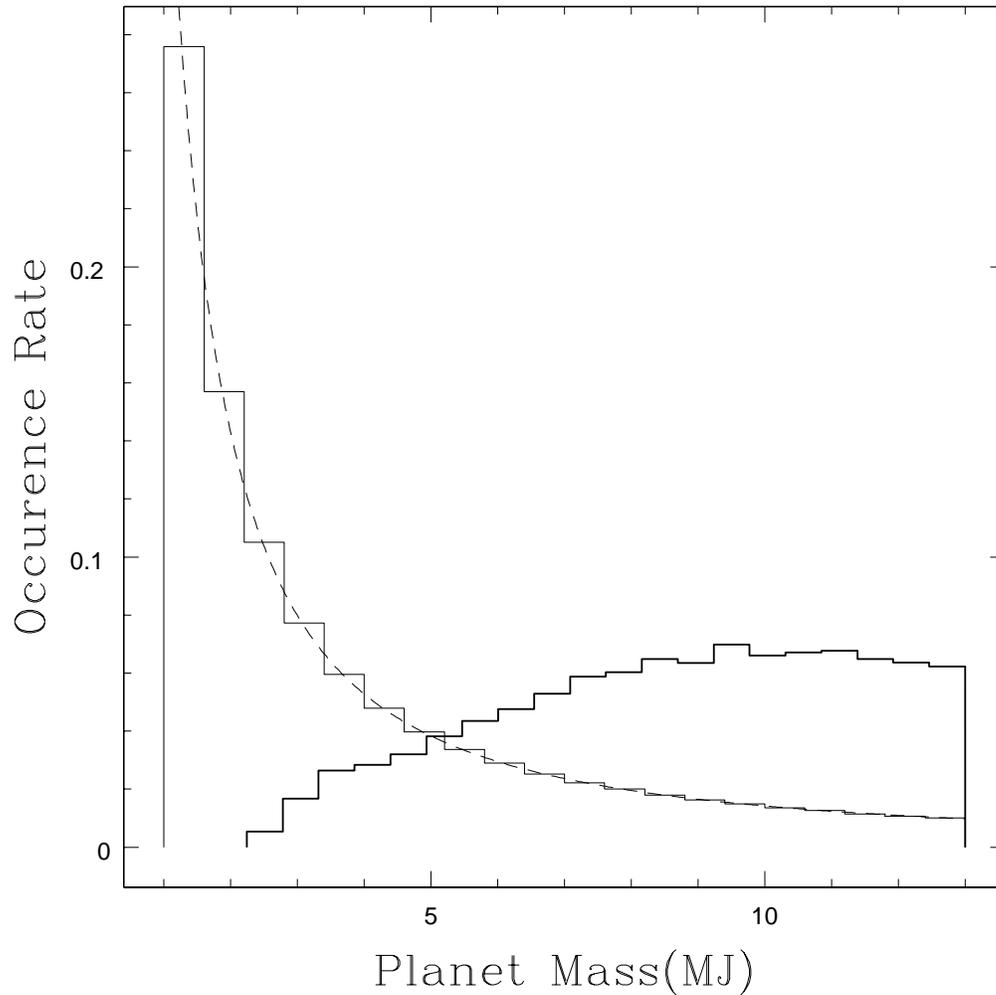


Figure 4.16 Planet mass histogram for the planets in the Monte Carlo simulation based on the Burrows et al. (2003) models. The thin line histogram is the mass distribution of all the planets in the simulation; the heavy line gives the distribution of those that were detected. The thin continuous curve is the input mass power law.

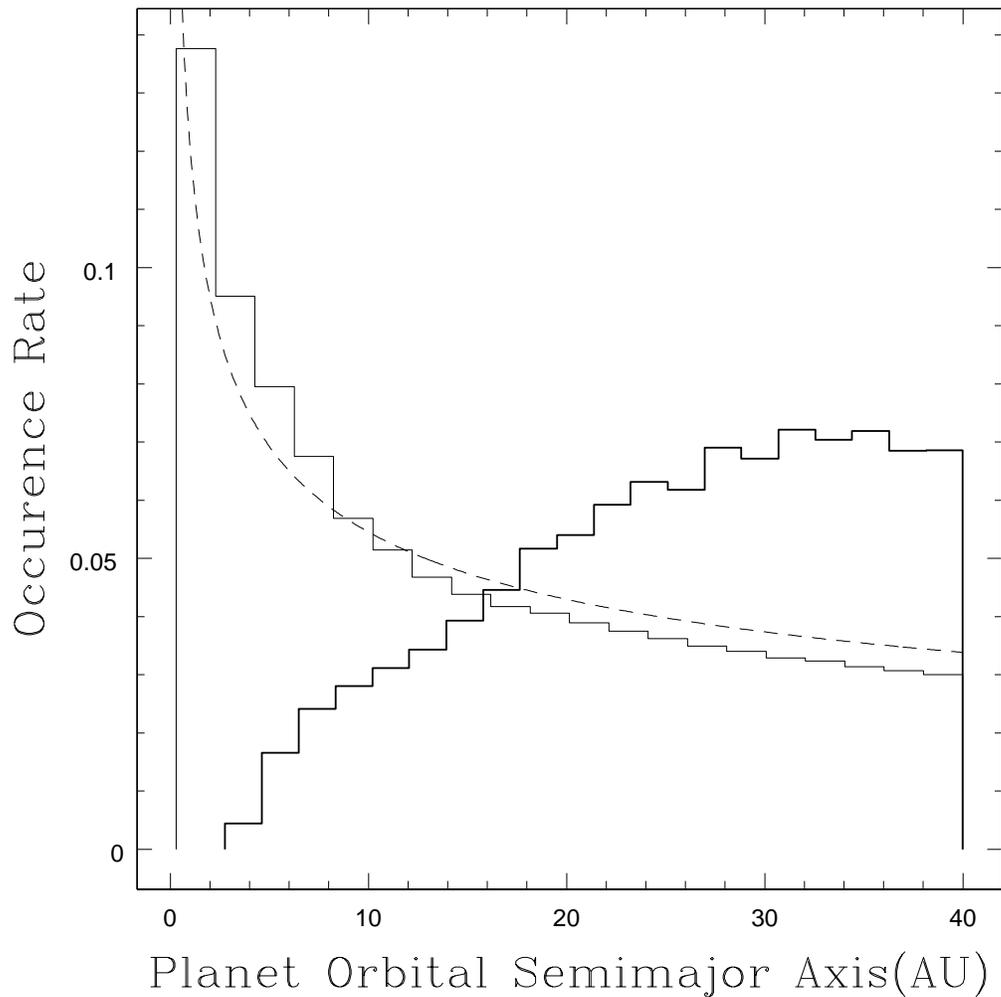


Figure 4.17 Planet orbital semimajor axis histogram for the planets in the Monte Carlo simulation based on the Burrows et al. (2003) models. The thin line histogram is the semimajor axis distribution of all the planets in the simulation; the heavy line gives the distribution of those that were detected. The thin continuous curve is the input semimajor axis power law.

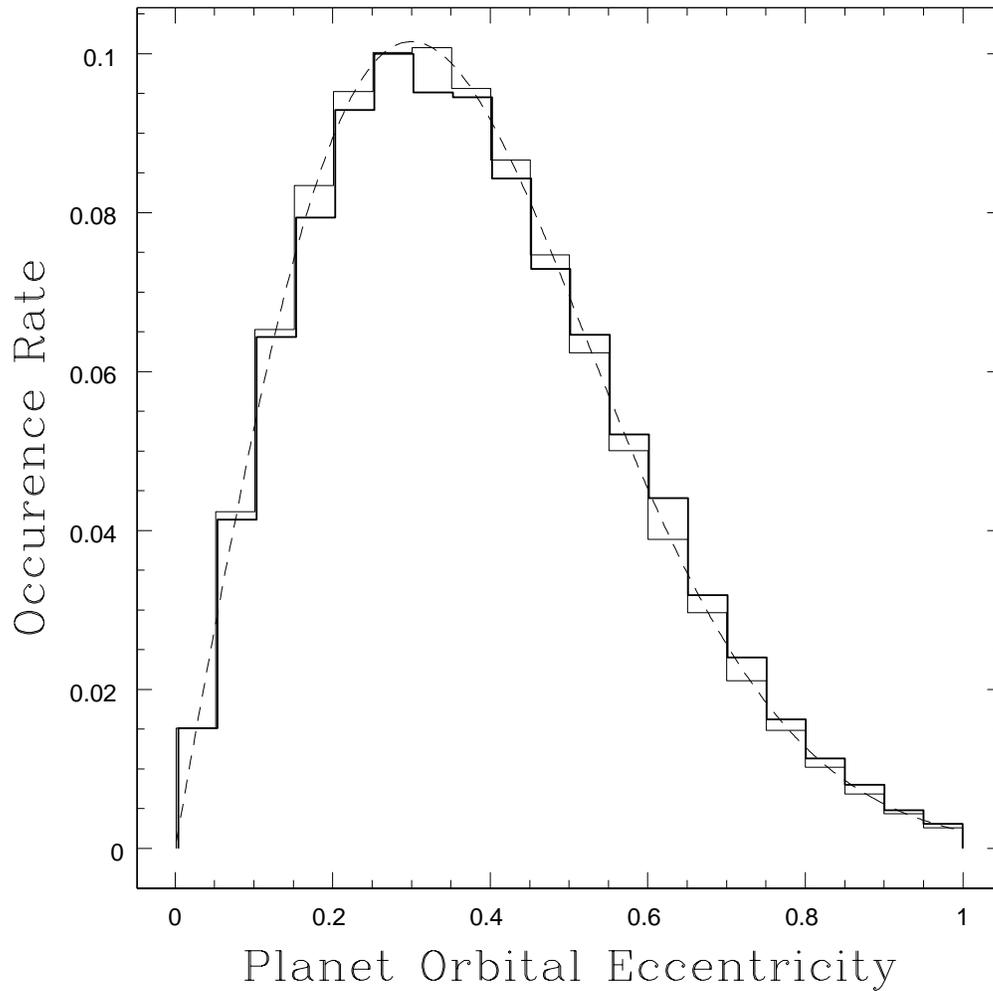


Figure 4.18 Planet orbital eccentricity histogram for the planets in the Monte Carlo simulation based on the Burrows et al. (2003) models. The thin line histogram is the semimajor axis distribution of all the planets in the simulation; the heavy line gives the distribution of those that were detected. The thin continuous curve is the Juric & Tremaine (2007) distribution.

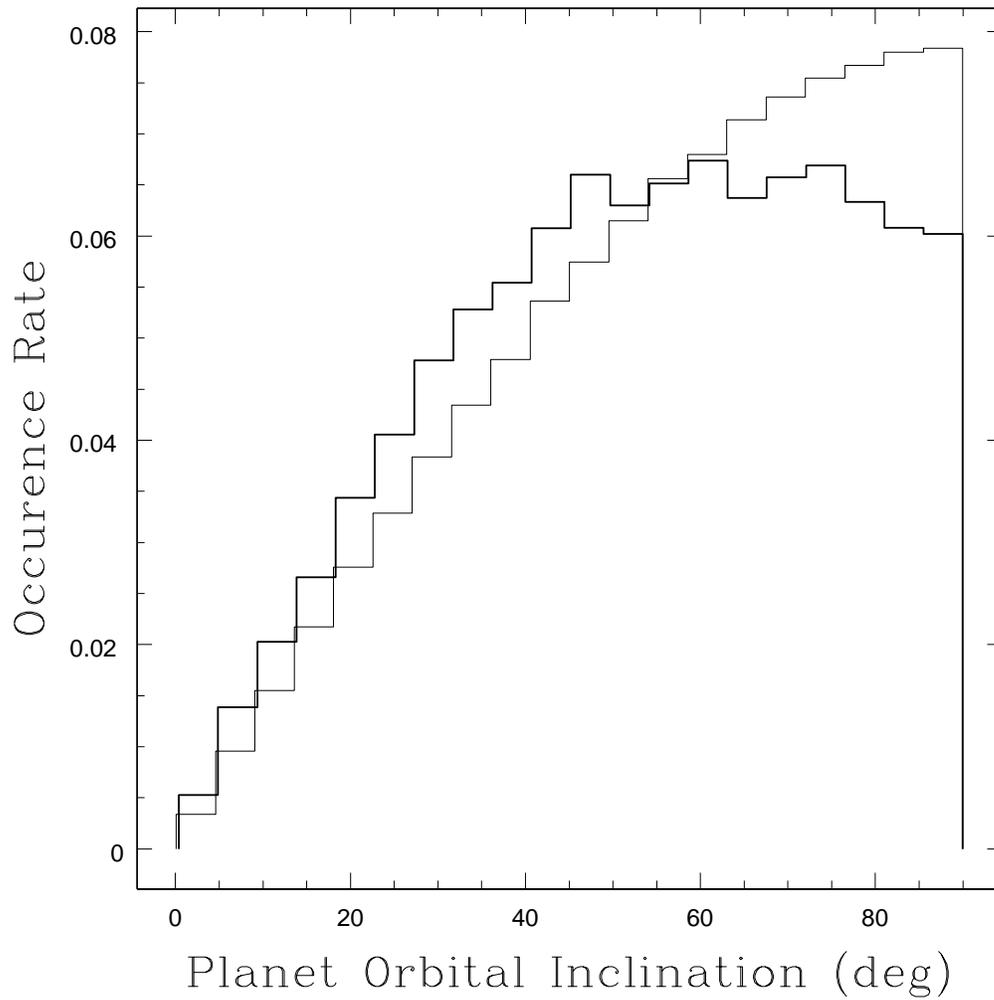


Figure 4.19 Planet orbital inclination histogram for the planets in the Monte Carlo simulation based on the Burrows et al. (2003) models. The thin line histogram is the inclination distribution of all the planets in the simulation; the heavy line gives the distribution of those that were detected.

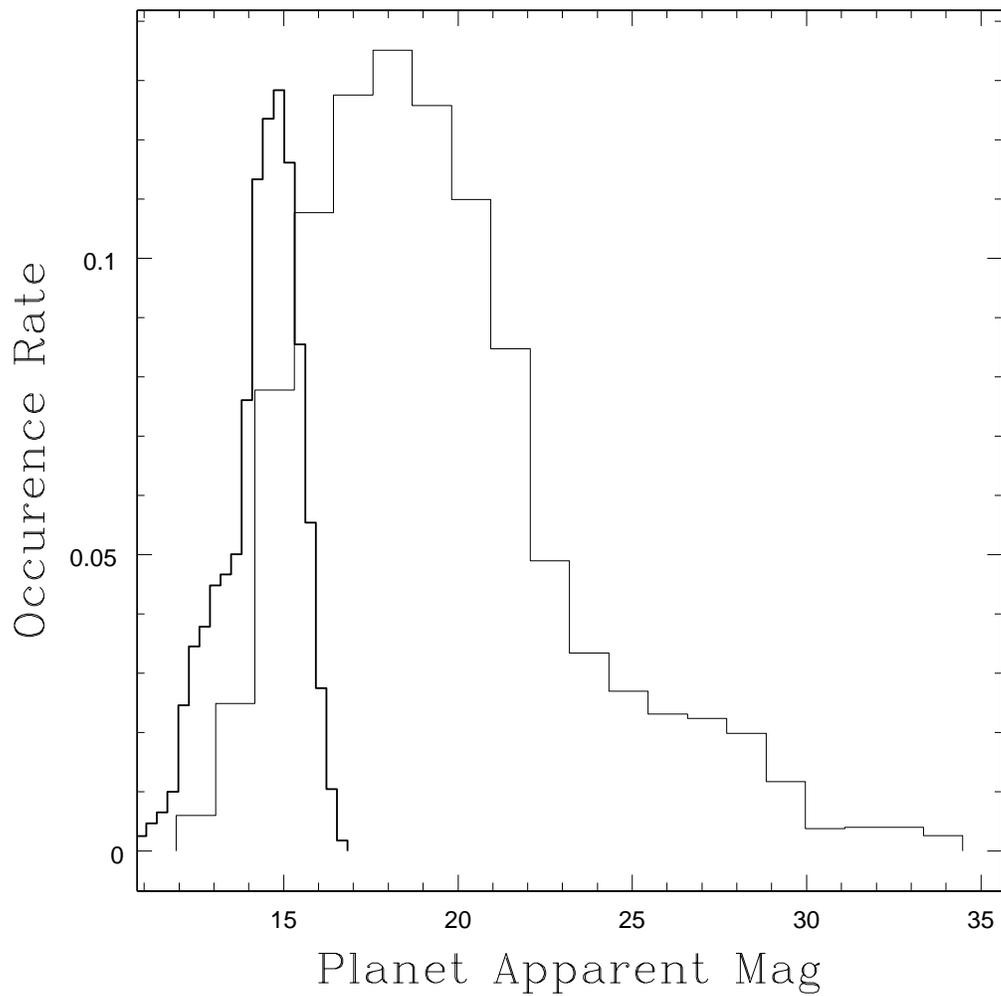


Figure 4.20 Apparent magnitude histogram for the planets in the Monte Carlo simulation based on the Burrows et al. (2003) models. The thin line histogram is the magnitude distribution of all the planets in the simulation; the heavy line gives the distribution of those that were detected.

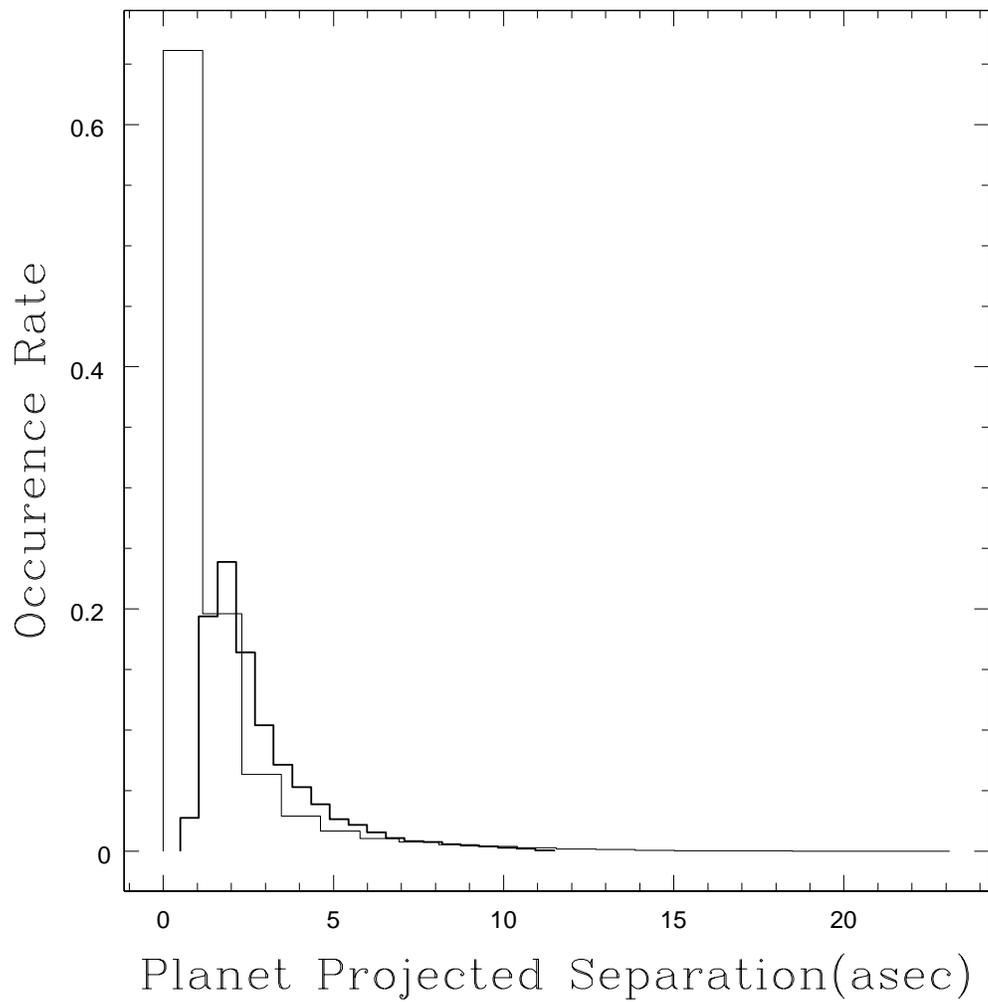


Figure 4.21 Separation histogram for the planets in the Monte Carlo simulation based on the Burrows et al. (2003) models. The thin line histogram is the separation distribution of all the planets in the simulation; the heavy line gives the distribution of those that were detected.

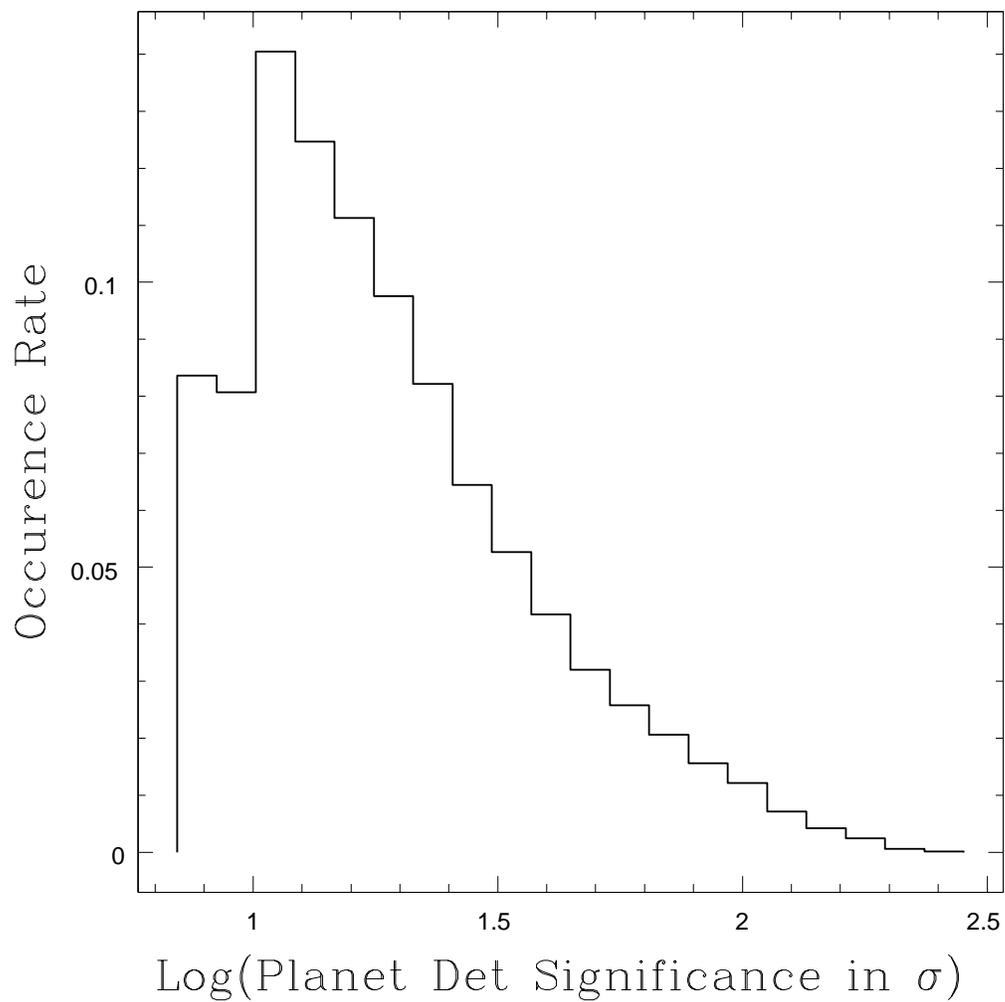


Figure 4.22 Detection significance histogram for detected planets in the Monte Carlo simulation based on the Burrows et al. (2003) models.

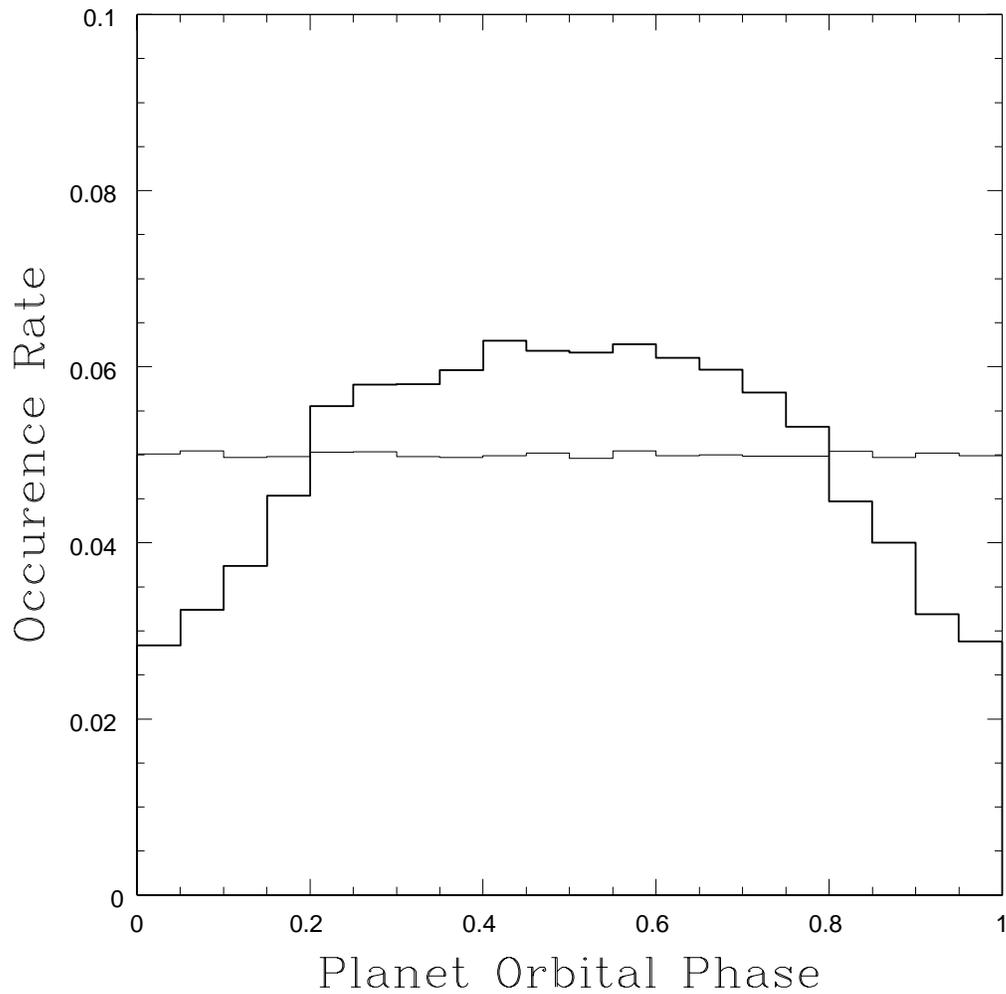


Figure 4.23 Orbital phase histogram for detected planets in the Monte Carlo simulation based on the Burrows et al. (2003) models. The thin line histogram is the phase distribution of all the planets in the simulation; the heavy line gives the distribution of those that were detected.

These tables indicate that although the best two or three stars do have very high probability of hosting detectable planets relative to the rest of the sample, they do not dominate the survey. In fact, the likelihood of hosting detectable planets is fairly well spread out. This true of simulations using both model sets, but the distribution is considerably better for the one using the Baraffe et al. (2003) models.

These tables do indicate some deadweight: seven stars had no detectable planets at all in 50,000 survey realizations. Thus, for the purposes of these Monte Carlo simulations, we surveyed 43 stars rather than 50. Why were these stars included at all? If the question relates to why they were observed, the answer is that in some cases observational issues such as bad seeing or an unexpected need to stop observing caused the data to be less good than anticipated, and sensitivity to planet-mass objects would have been obtained if all had gone well. In other cases observational constraints precluded imaging more promising targets, and so observations were made with the goal of obtaining sensitivity at least to low-mass brown dwarfs. In the case of GJ 3876, an error in the Gliese catalog and an erroneous young age led us to think it was a promising target when it was not.

If the question relates to why these stars have been included in the Monte Carlo simulations, the answer is that, first, some of the stars with no detectable planets using the distributions of planets in this Section may well have some with the more optimistic distributions to be considered below in Section 4.7. Secondly, these simulations could easily be extended to brown dwarf masses, and then the seven stars would no longer be deadweight.

A final interesting question the Monte Carlo simulations can address is how important the  $M$  band observations were to the survey results. Observations at  $M$  were appropriate for bright, nearby stars of particular interest. The stars

needed to be bright because the background limited sensitivity in  $M$  band is always worse than  $L'$ , so  $M$  band only gives an advantage for contrast-limited observations because the planet/star flux ratio for  $M$  band is even better than for  $L'$ . Of course, improving the sensitivity for contrast-limited observations is most important around bright stars where most of the image is contrast-limited. The stars should be nearby, so that the sensitivity at  $M$  extends down to cool objects with red  $L' - M$  colors. And, of course, the stars should be of particular interest for the collection of an additional data set to be worthwhile. Only 4 of the 50 stars were observed and finally analyzed at  $M$  band:  $\epsilon$  Eri, GJ 820A, and the two components of  $\xi$  Boo. These stars met all the criteria listed above:  $\epsilon$  Eri needs no introduction; GJ 820A, otherwise known as 61 Cyg A, is a famous, very nearby star, and the  $\xi$  Boo system is one of the youngest stellar systems within 10 pc of the sun. Table 4.11 gives the fraction of total planets for each star that were detected at  $L'$  only, at  $M$  band only, and in both bands, for the two simulations. For the simulation using the Burrows et al. (2003) models, the  $M$  band observations made a significant contribution to the number of detected planets around each star, especially  $\epsilon$  Eri, the most interesting star in our survey. For the Baraffe et al. (2003) models,  $M$  was substantially less important, although the likelihood of having two bands on any detected planets is still a significant benefit. This greater usefulness of the  $M$  band under the Burrows et al. (2003) models than under Baraffe et al. (2003) is not surprising if one considers Figures 4.13 and 4.14: the Baraffe et al. (2003) models predict substantially brighter planets at  $L'$ , while at  $M$  band the predictions are very close. Thus it is not surprising that  $M$  band should be less important relative to  $L'$  for the simulation using the Baraffe et al. (2003) models, and more important for the one using the Burrows et al. (2003) models. It is important to note, however, that in both cases the most

likely outcome for all stars was detection in both bands. This means including  $M$  band observations is more likely to increase the information gathered from detected planets than to greatly increase the likelihood of planets being detected at all. The usefulness of such a two-band detection should not be underestimated, however, since it would immediately supply a color and probably convincingly rule out background stars and brown dwarfs, leaving the planetary hypothesis the only viable one even before proper motion confirmation.

Table 4.9. Percentage of Detected Planets Found Around Each Star

Star Name	% Detected Planets (Burrows Models)	% Detected Planets (Baraffe Models)
$\epsilon$ Eri	15.1731	9.6042
GJ 117	10.2885	8.5977
GJ 625	8.5774	8.5292
HD 29391	6.2465	5.5624
GJ 519	5.9627	5.8051
GJ 879	4.1468	3.7949
GJ 5	3.9460	3.7549
BD+60 1417	3.6580	3.5978
GJ 349	3.3306	3.4008
GJ 355	3.0119	2.8040
GJ 410	2.9508	2.9668
GJ 450	2.9115	3.7235
GJ 159	2.5492	2.3871
GJ 354.1A	2.2742	2.2244
GJ 820A	2.1345	3.1181
HD 220140A	2.0341	1.9845
GJ 216A	1.9774	2.1159
GJ 564	1.8421	1.8503
GJ 820B	1.6806	3.2095

Table 4.9—Continued

Star Name	% Detected Planets (Burrows Models)	% Detected Planets (Baraffe Models)
GJ 3860	1.6282	1.9760
HD 78141	1.4710	1.5619
GJ 311	1.4274	1.4991
GJ 278C	1.2135	1.4734
BD+20 1790	1.1611	1.5277
GJ 211	1.1480	1.4820
GJ 282A	1.0782	1.3792
HD 96064A	0.8730	1.1450
GJ 896A	0.7857	1.0194
HD 113449	0.7683	1.0137
BD+48 3686	0.7683	1.2564
GJ 896B	0.7071	0.6796
HD 220140B	0.5937	0.6482
ξ Boo B	0.4234	0.4597
GJ380	0.3143	1.1650
ξ Boo A	0.3056	0.3112
HD 139813	0.1921	0.5540
HD 1405	0.1877	0.4312
GJ 166B	0.0917	0.2684

Table 4.9—Continued

Star Name	% Detected Planets (Burrows Models)	% Detected Planets (Baraffe Models)
HD 96064B	0.0917	0.1999
GJ 166C	0.0349	0.1542
GJ 505B	0.0175	0.1399
$\tau$ Ceti	0.0175	0.4826
GJ 505A	0.0044	0.0514
HD 77407B	0.0000	0.0000
GJ 3876	0.0000	0.0000
GJ 659A	0.0000	0.0000
GJ 659B	0.0000	0.0000
GJ 702B	0.0000	0.0000
GJ 702A	0.0000	0.0000
HD 77407A	0.0000	0.0000

To finish off the analysis of the planets detected in the Monte Carlo simulations on a somewhat recreational note, we present Figures 4.6.2 through 4.6.2. These show the orbit and position of the first detected planets around most of the stars in our survey, in the Burrows et al. (2003) Monte Carlo simulation. Note that the probability of finding a planet varies significantly from star to star, and in some cases, notably GJ 505AB, it was almost zero. However, these plots do give a striking visual impression of the types of planets that could have been detected in our survey — and that may be detected in similar surveys in the near future.

Table 4.10. Stars Binned by Fraction of Detected Planets

Detected Planet Fraction Bin	Number of Stars In this Bin (Burrows Models)	Number of Stars In this Bin (Baraffe Models)
0% - 0%	7	7
0.00% - 0.25%	8	4
0.25% - 0.50%	3	5
0.50% - 1.00%	6	3
1.00% - 2.00%	10	14
2.00% - 4.00%	10	12
4.00% - 8.00%	3	2
8.00% - 16.00%	3	3

Table 4.11. Importance of the  $M$  Band Data

Star	Simulation	Planets	Planets	Planets
		Detected in $L'$ Only	Detected in $M$ Only	Detected in $L'$ and $M$
$\epsilon$ Eri	Burrows	5.1%	35.2%	59.8%
$\epsilon$ Eri	Baraffe	9.2%	13.9%	76.9%
GJ 820A	Burrows	21.7%	16.2%	62.1%
GJ 820A	Baraffe	43.8%	2.3%	53.9%
$\xi$ Boo A	Burrows	9.3%	21.5%	69.2%
$\xi$ Boo A	Baraffe	19.6%	2.2%	78.2%
$\xi$ Boo B	Burrows	13.9%	14.6%	71.5%
$\xi$ Boo B	Baraffe	26.6%	1.2%	72.2%

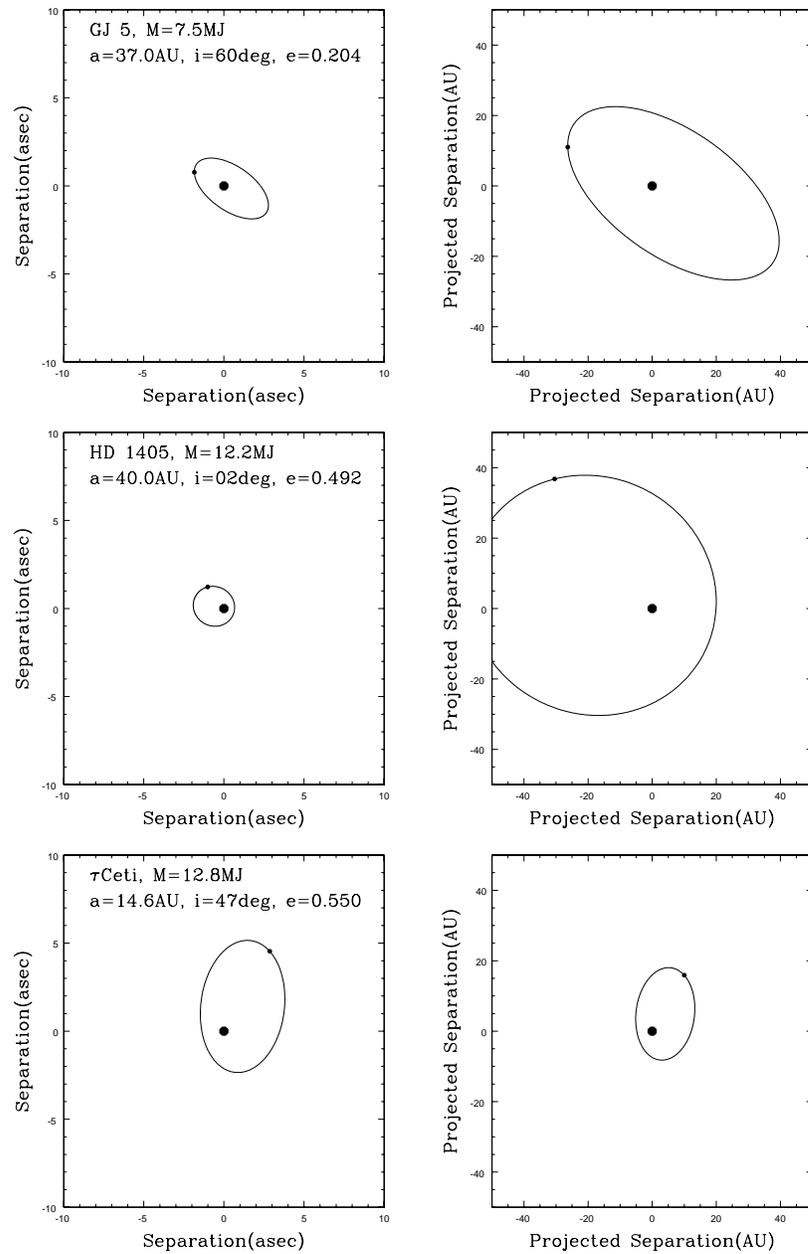


Figure 4.24 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a  $10 \times 10 \text{ asec}$  box; projected physical orbits on right in a  $50 \times 50 \text{ AU}$  box.

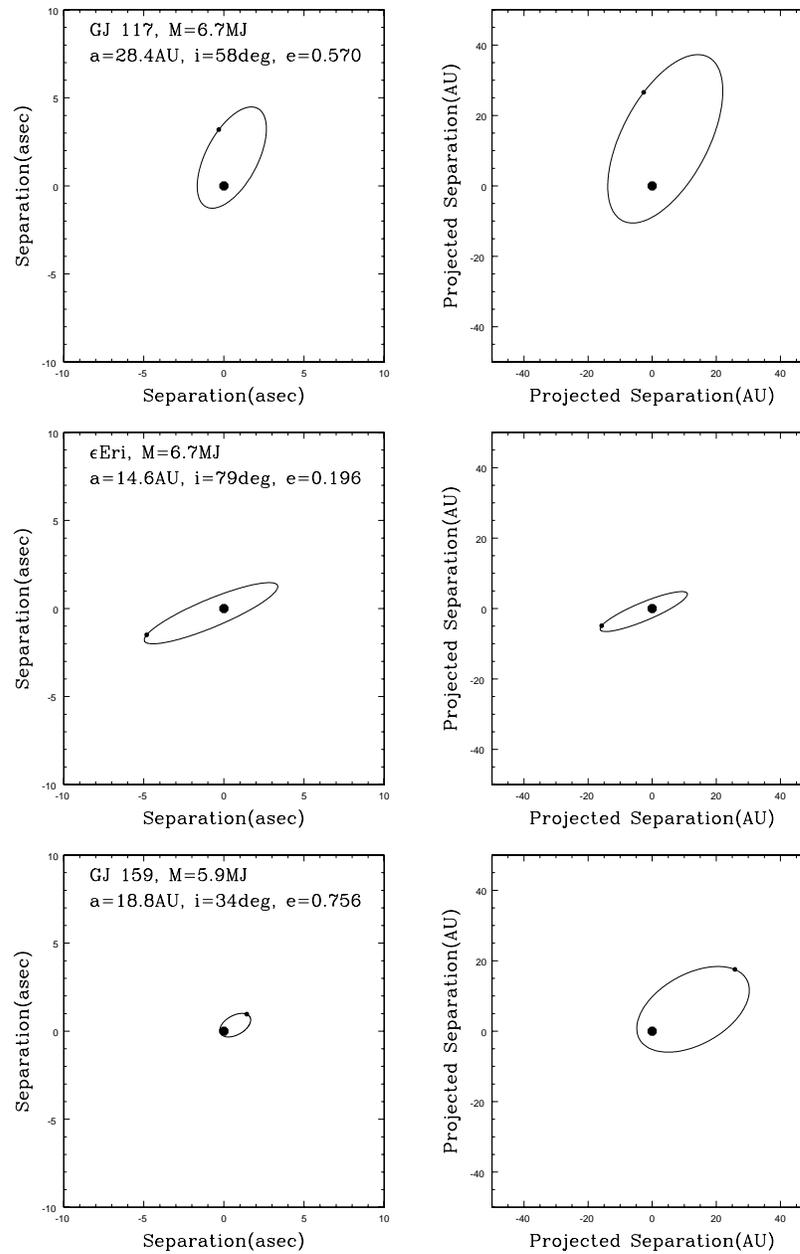


Figure 4.25 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

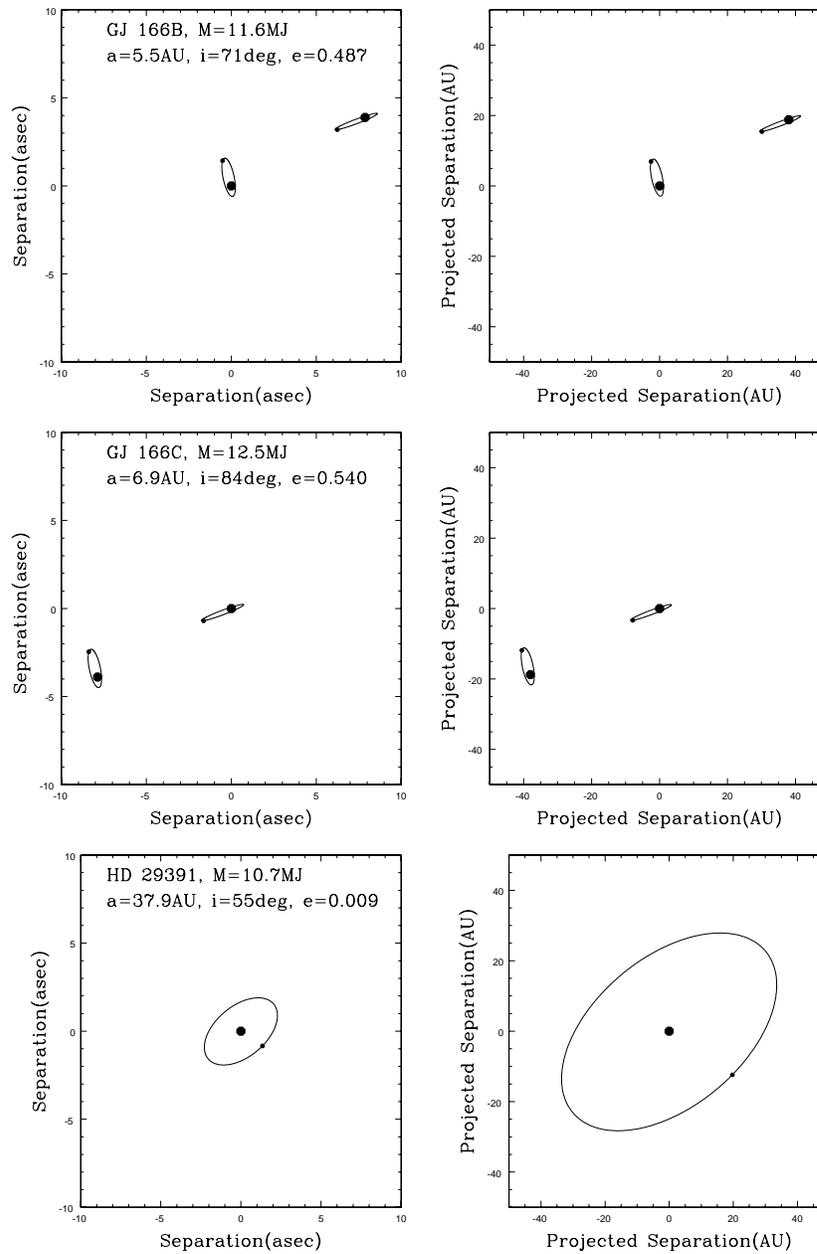


Figure 4.26 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a  $10 \times 10 \text{ asec}$  box; projected physical orbits on right in a  $50 \times 50 \text{ AU}$  box. Note the binary star GJ 166BC.

## 4.7 Monte-Carlo Simulations: The Larger Picture

### 4.7.1 Excluded Regions of Parameter Space

Figures 4.40 through 4.43 give the probabilities of finding no planets as a function of the power law exponent of the distribution of orbital semimajor axes and the outer truncation radius of the distribution. We define the power law slope as  $\gamma$ , where  $P(a) \propto a^\gamma$  is the probability distribution of the semimajor axis  $a$  of planetary orbits. Note that even for  $\gamma = -0.5$ , the integral of the distribution from  $a = 0$  to  $\infty$  does not converge, so the outer, truncation radius to the distribution is a very important parameter. Note also that because planets may be found near aphelion in highly eccentric orbits, the truncation radius for the distribution of semimajor axes does not correspond to a hard truncation in the distribution of physical planet-star distances at any given time. Slight raggedness in the contours of the figures is due to small number statistics in the Monte-Carlo simulations used to make them. Fixing the raggedness by running more survey realizations would have required a prohibitively large investment of computer time; each of the current figures required 5000 realizations of our survey for each of the 1147 pixels, and took about 8 hours to make on a fast pc.

Figure 4.40 gives the contours that apply if we set the slope of the planet mass power law distribution to its adopted value  $\alpha = -1.44$ , and use the Burrows et al. (2003) models to convert planet masses to magnitudes. The black cross marks the parameters for the Monte Carlo simulations that were discussed in detail in Section 4.6.

Figure 4.41 gives the result using the best fit mass slope with the Baraffe et al. (2003) models. According to this model set we exclude a significantly larger area of the parameter space.

Figures 4.42 and 4.43 give the contours using the Burrows et al. (2003) models

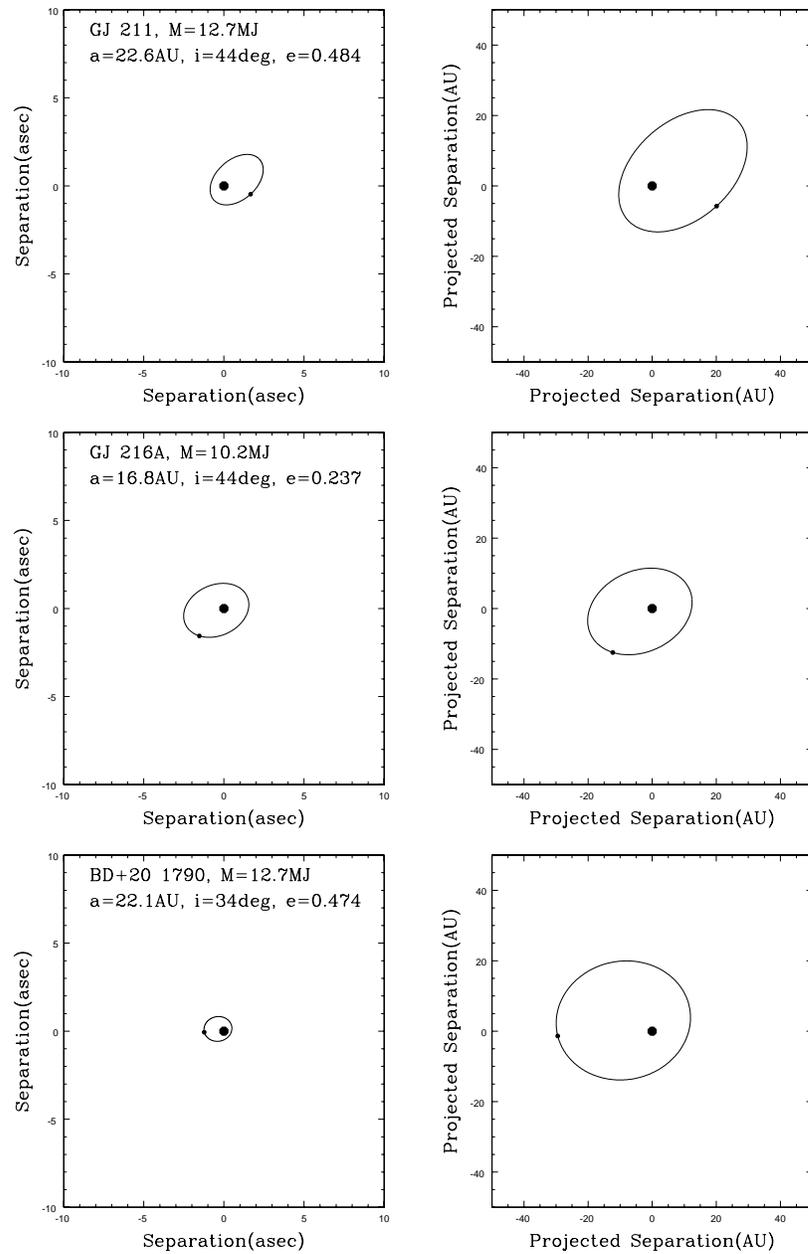


Figure 4.27 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

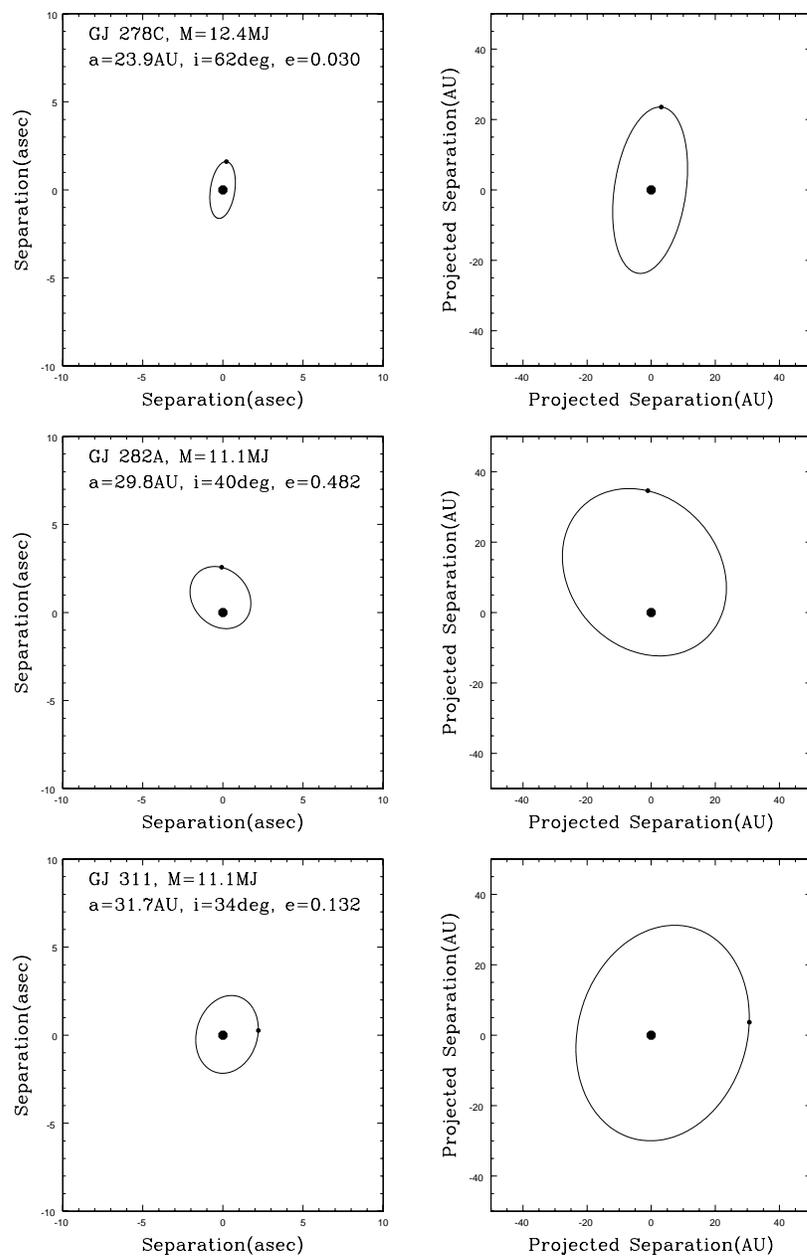


Figure 4.28 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

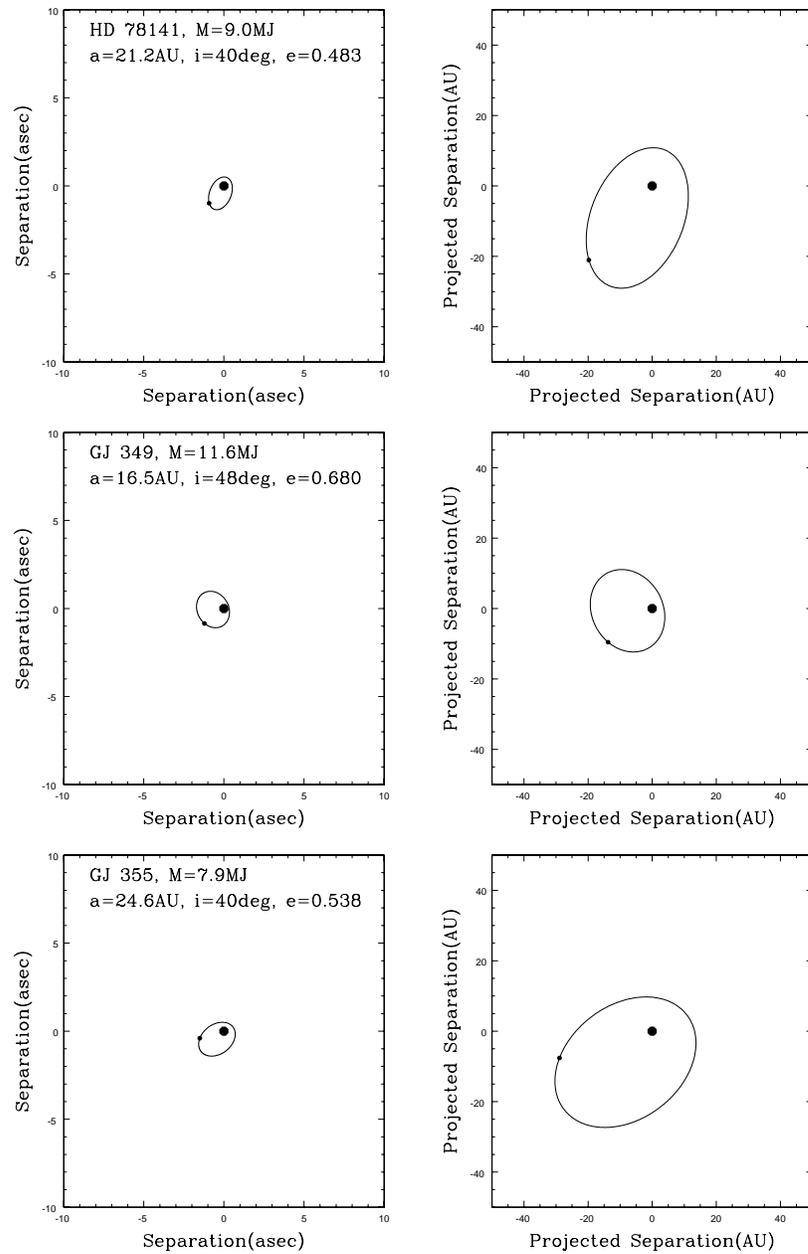


Figure 4.29 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

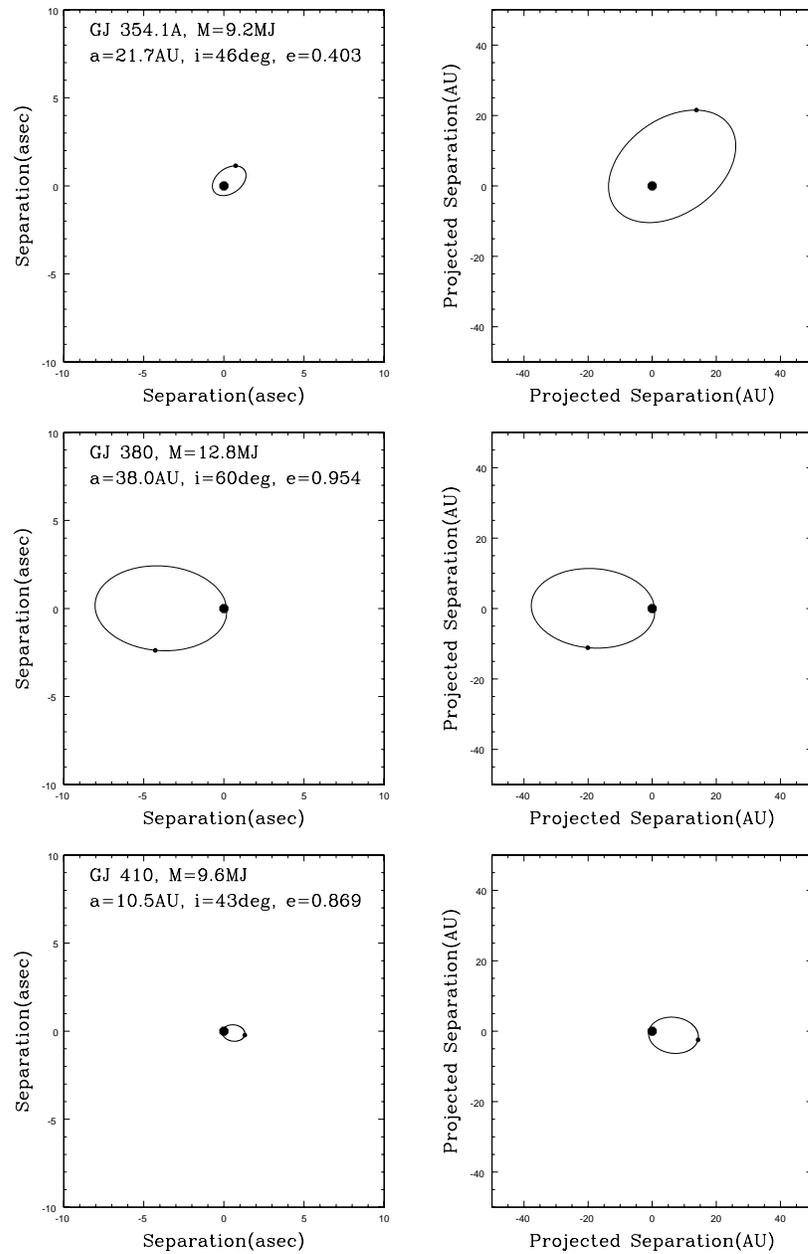


Figure 4.30 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

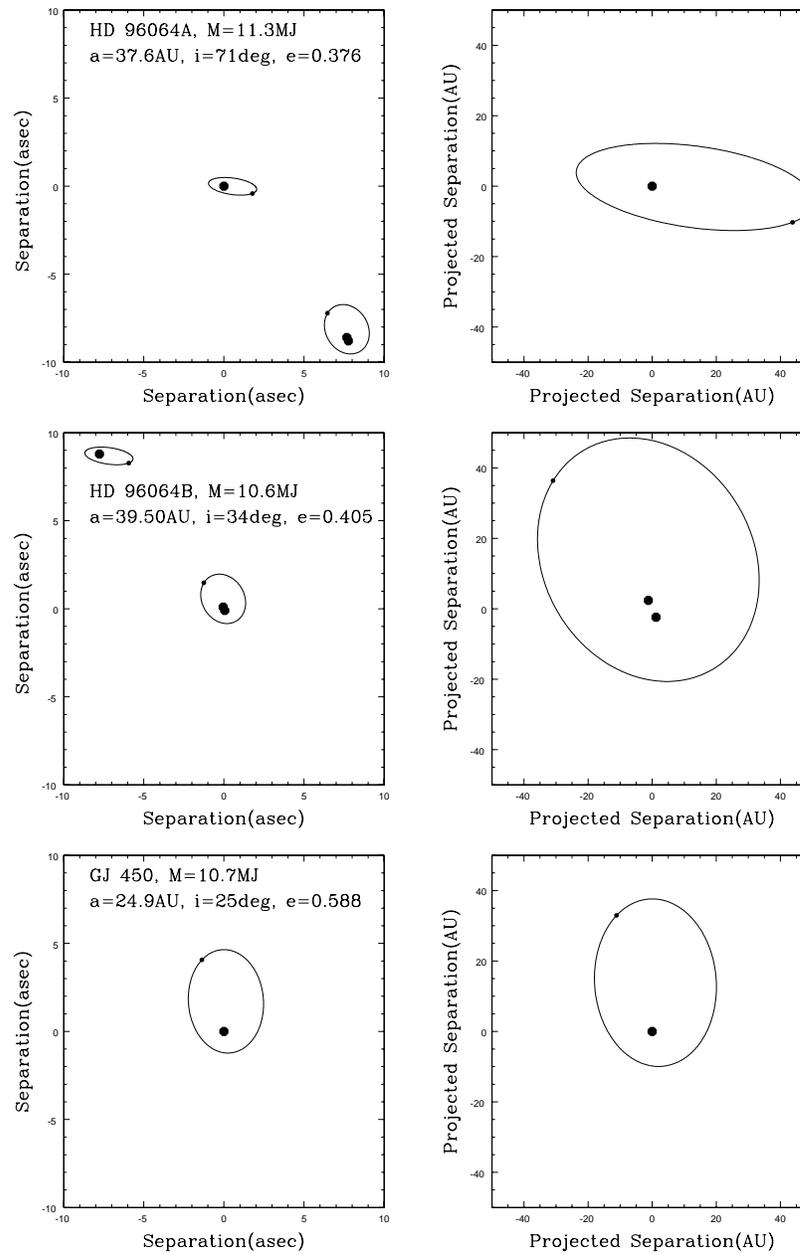


Figure 4.31 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box. Note the binary star HD 96064AB; the B component is itself a close binary.

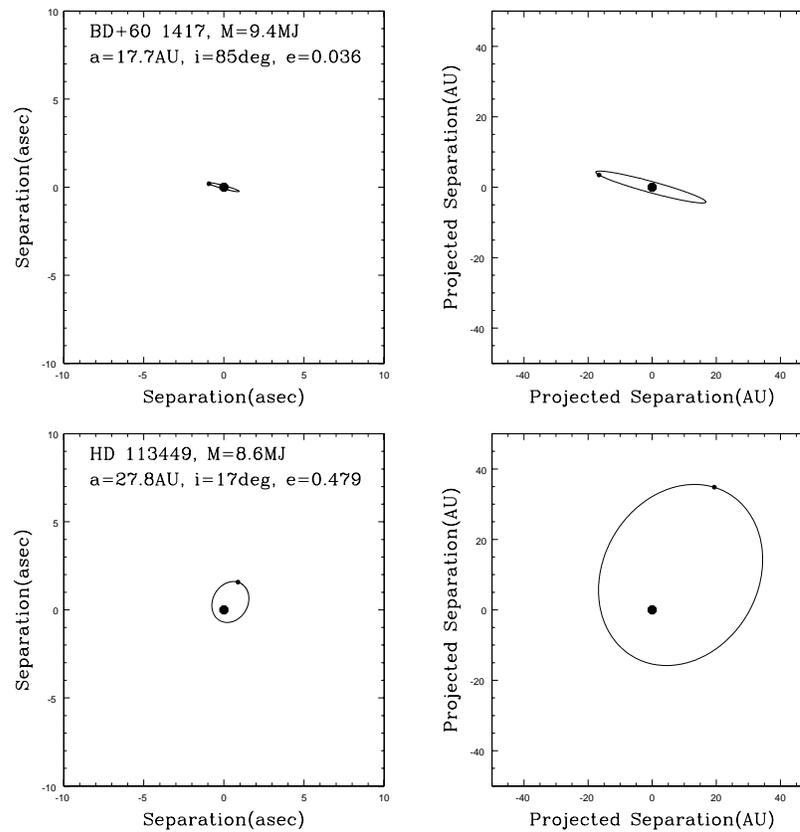


Figure 4.32 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a  $10 \times 10$  asec box; projected physical orbits on right in a  $50 \times 50$  AU box.

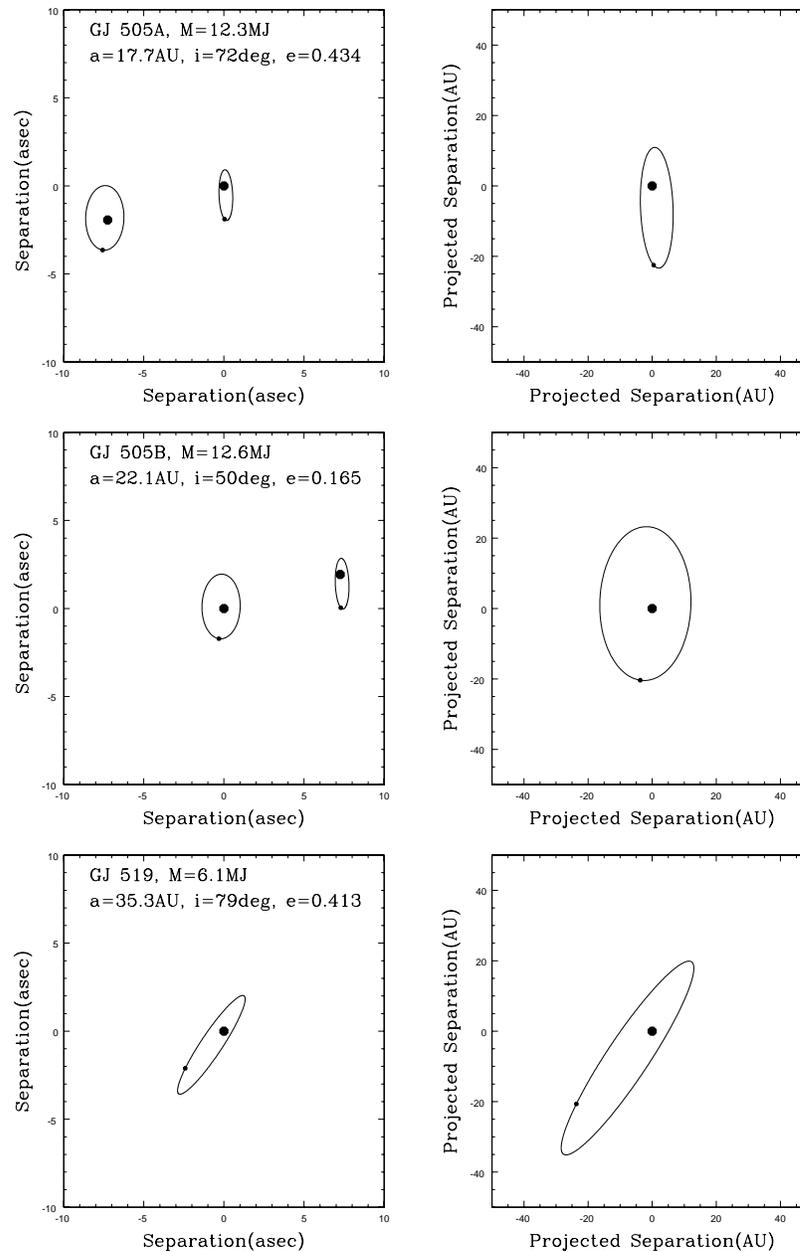


Figure 4.33 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box. Note the binary star GJ505AB.

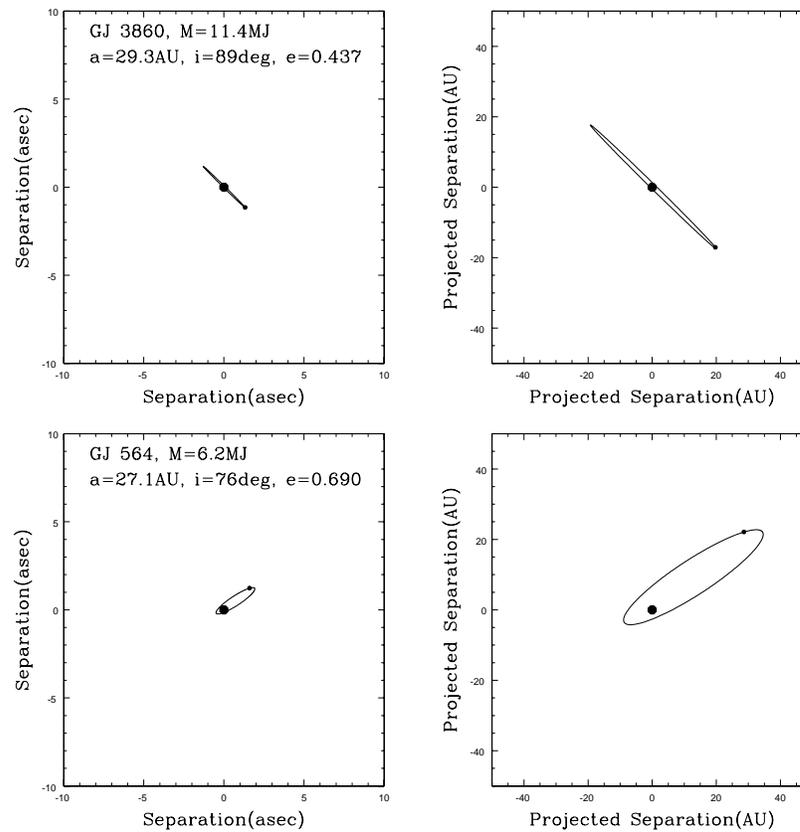


Figure 4.34 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

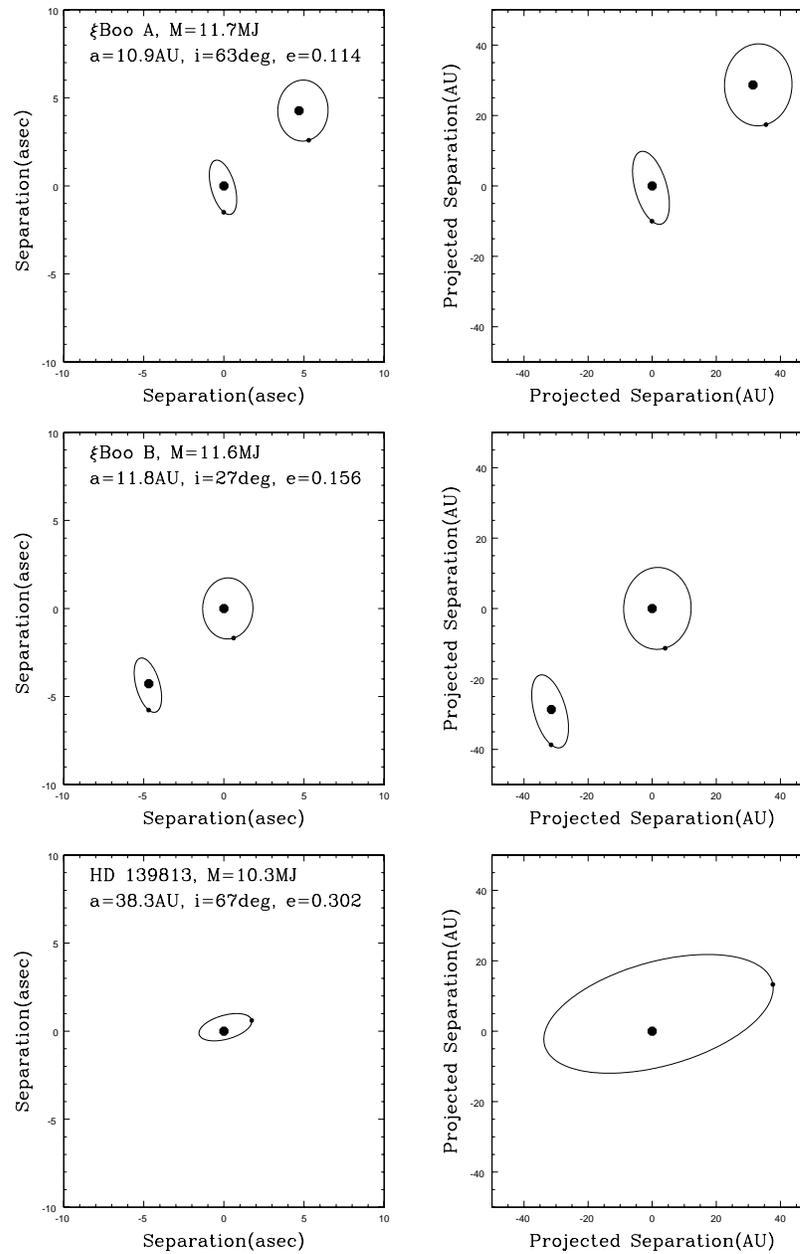


Figure 4.35 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box. Note the binary star  $\xi$  Boo AB.

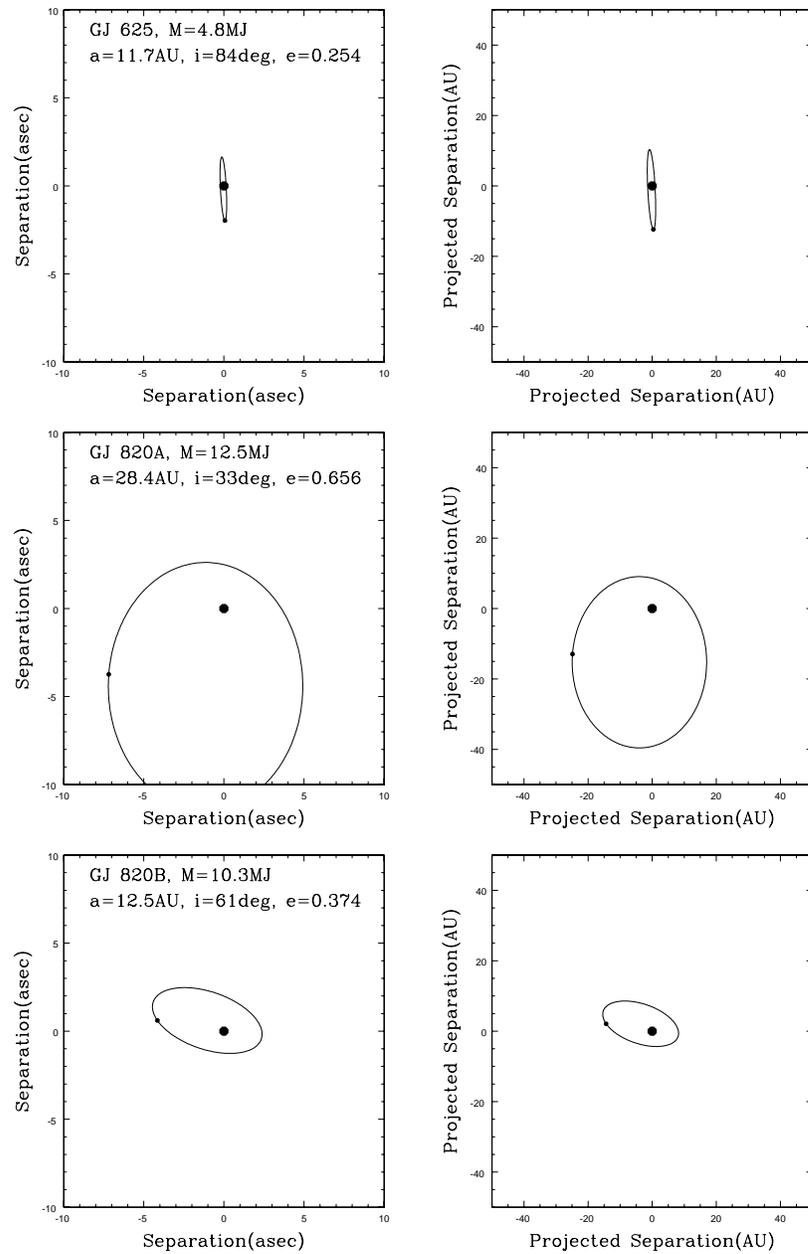


Figure 4.36 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

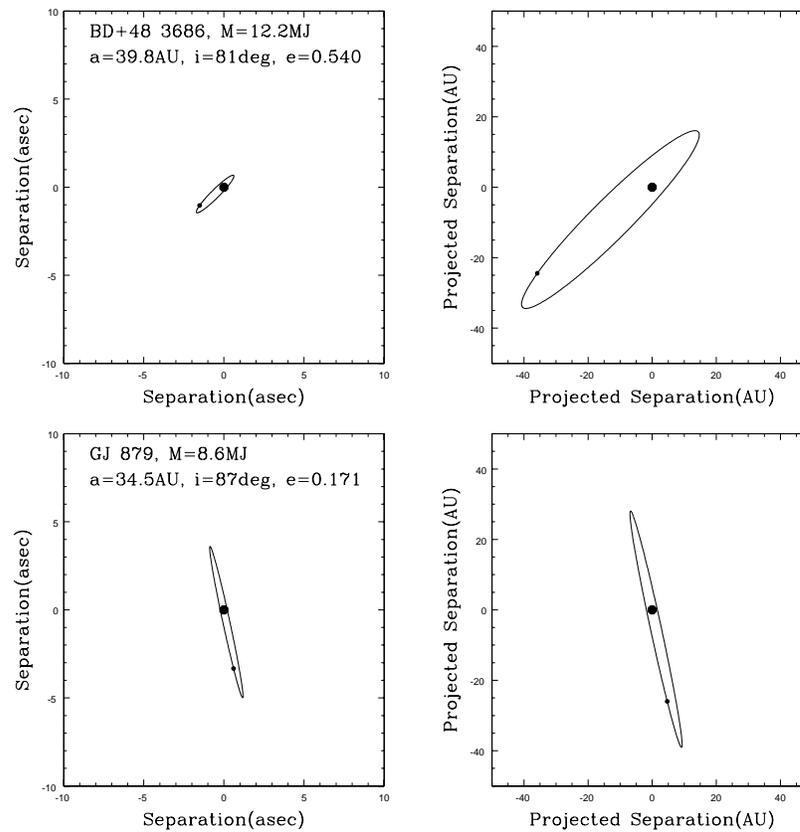


Figure 4.37 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box.

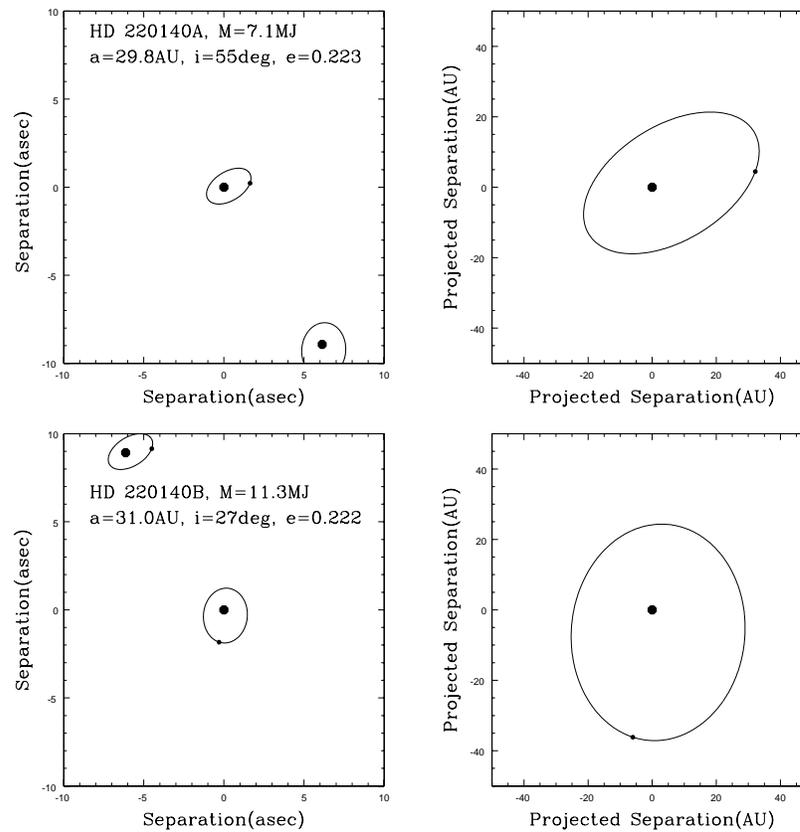


Figure 4.38 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box. Note that HD 220140 is a binary star.

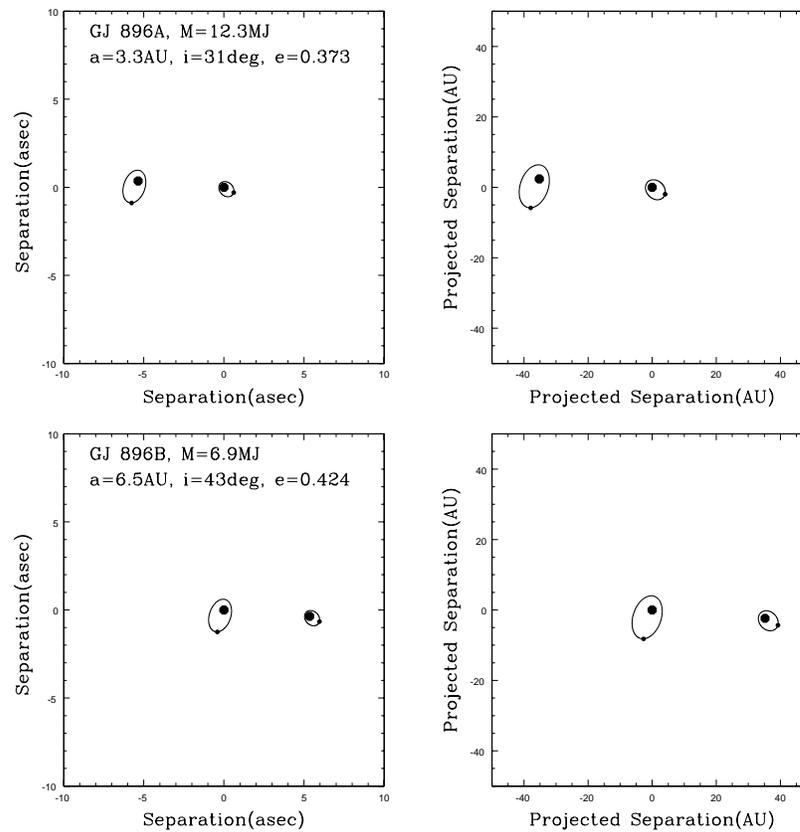


Figure 4.39 Orbits of the first detected planets around each star from the simulation using Burrows models. Sky orbits on the left in a 10x10 asec box; projected physical orbits on right in a 50x50 AU box. Note that GJ 896 is a binary star. These were some of the closest-in orbits detected in the simulation.

with mass power law slopes of  $\alpha = -1.881$  and  $\alpha = -1.089$ , respectively. These are the extreme pessimistic and optimistic edges of the 95% confidence interval on  $\alpha$  from our KS test fit to the RV planet statistics. Of course, we exclude a considerably larger area of parameter space on Figure 4.43, using the  $\alpha$  value that predicts a larger proportion of giant planets, than on Figure 4.42.

It is interesting to note that even on Figure 4.43, with the Burrows et al. (2003) models and the most optimistic value of  $\alpha$ , we do not exclude as much of the parameter space as in Figure 4.41, with the adopted value of  $\alpha$  and the Baraffe et al. (2003) models. This is a measure of how strongly the choice of model sets affects our results — and by extension, of how unconstrained and uncertain giant planet models currently are.

#### 4.7.2 Comparing Our Results to Other Surveys

. Different surveys have described the regions of parameter space they excluded in different ways. The VLT/NACO  $L'$  imaging survey of Kasper et al. (2007) provided a contour plots analogous to ours, using the Baraffe et al. (2003) models only. They exclude more parameter space than we do. We note, however, that they have used  $5\sigma$  limits in their simulation, as opposed to our  $7\sigma$  and  $10\sigma$  limits, and they used a sensitivity estimator that did not correctly handle correlated noise.

In their abstract they give an example of a distribution that they have ruled out with 90% confidence: one with a semimajor axis power law slope  $\gamma = +0.2$  and a truncation radius of 30 AU. With our best fit value of  $\alpha$  and our standard scheme with objects detected at 7 and  $10\sigma$ , we rule this distribution out only at the 84% confidence level. However, we tried an experiment in which we set our detection limits to  $5\sigma$ . With this (unrealistic!) sensitivity threshold, we ruled out the distribution with a 93% confidence level. We note from Chapter 3 that process-

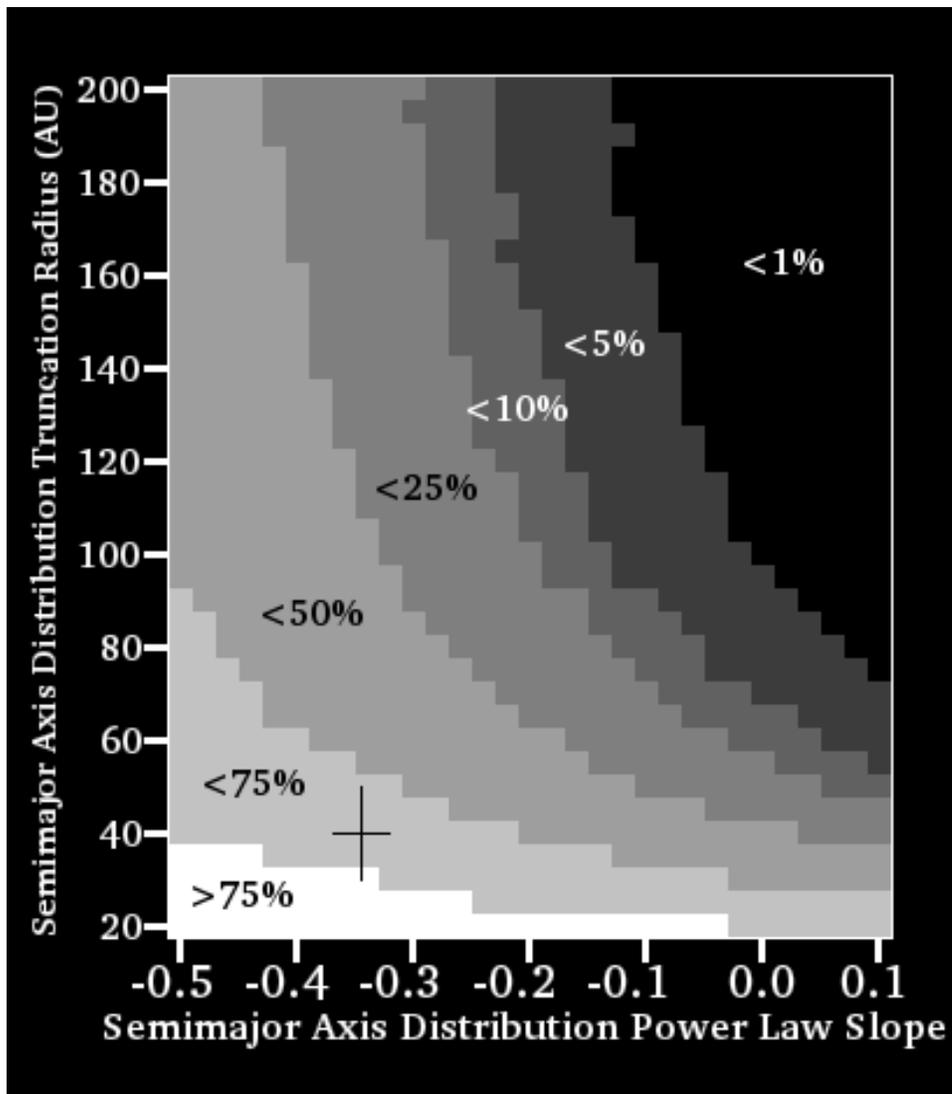


Figure 4.40 The probability of detecting no planets in our survey as a function of the power law slope  $\gamma$  for the distribution of the planets' orbital semimajor axes and the outer truncation radius for this distribution. Here we have used the Burrows et al. (2003) planet models and the adopted slope of -1.44 for the mass power law slope  $\alpha$ . From black to white, the contours are for a 0–1%, 1–5%, 5–10%, 10–25%, 25–50%, 50–75%, and more than 75% probability of detecting no planets. The black cross marks our simulation from Section 4.6.

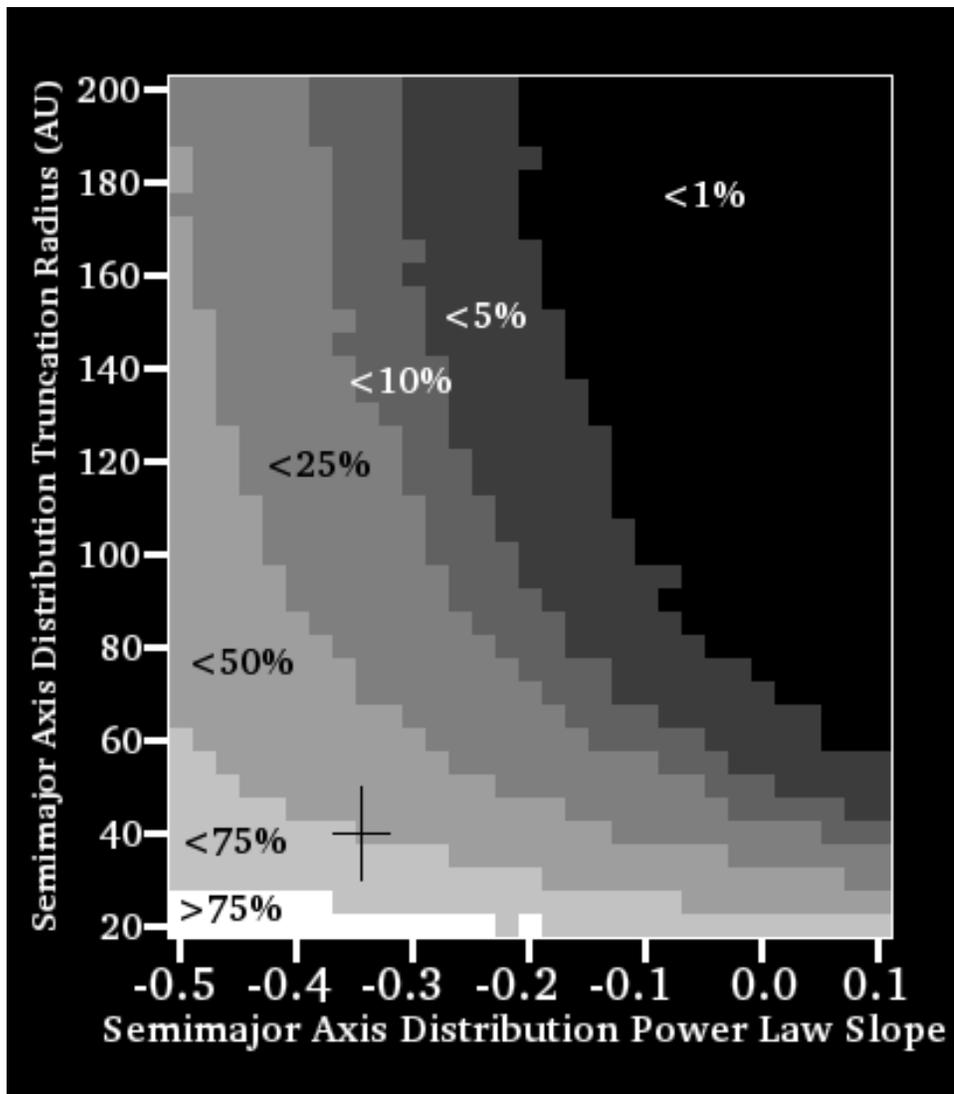


Figure 4.41 The probability of detecting no planets in our survey as a function of the power law slope  $\gamma$  for the distribution of the planets' orbital semimajor axes and the outer truncation radius for this distribution. Here we have used the Baraffe et al. (2003) planet models and the adopted slope of -1.44 for the mass power law slope  $\alpha$ . From black to white, the contours are for a 0–1%, 1–5%, 5–10%, 10–25%, 25–50%, 50–75%, and more than 75% probability of detecting no planets. The black cross marks our simulation from Section 4.6.

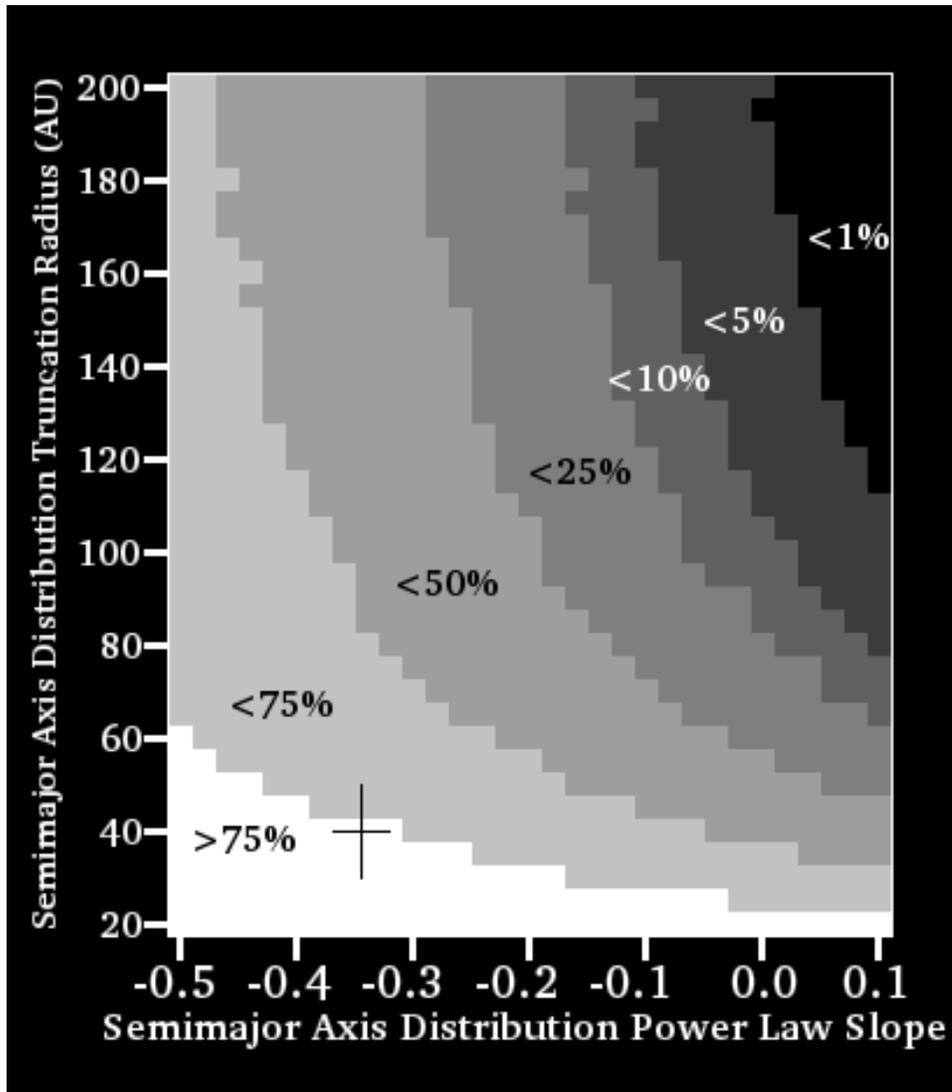


Figure 4.42 The probability of detecting no planets in our survey as a function of the power law slope  $\gamma$  for the distribution of the planets' orbital semimajor axes and the outer truncation radius for this distribution. Here we have used the Burrows et al. (2003) planet models and have set  $\alpha = -1.881$ , the most pessimistic slope for the mass power law that is permitted by the 95% confidence interval. From black to white, the contours are for a 0–1%, 1–5%, 5–10%, 10–25%, 25–50%, 50–75%, and more than 75% probability of detecting no planets. The black cross marks our simulation from Section 4.6.

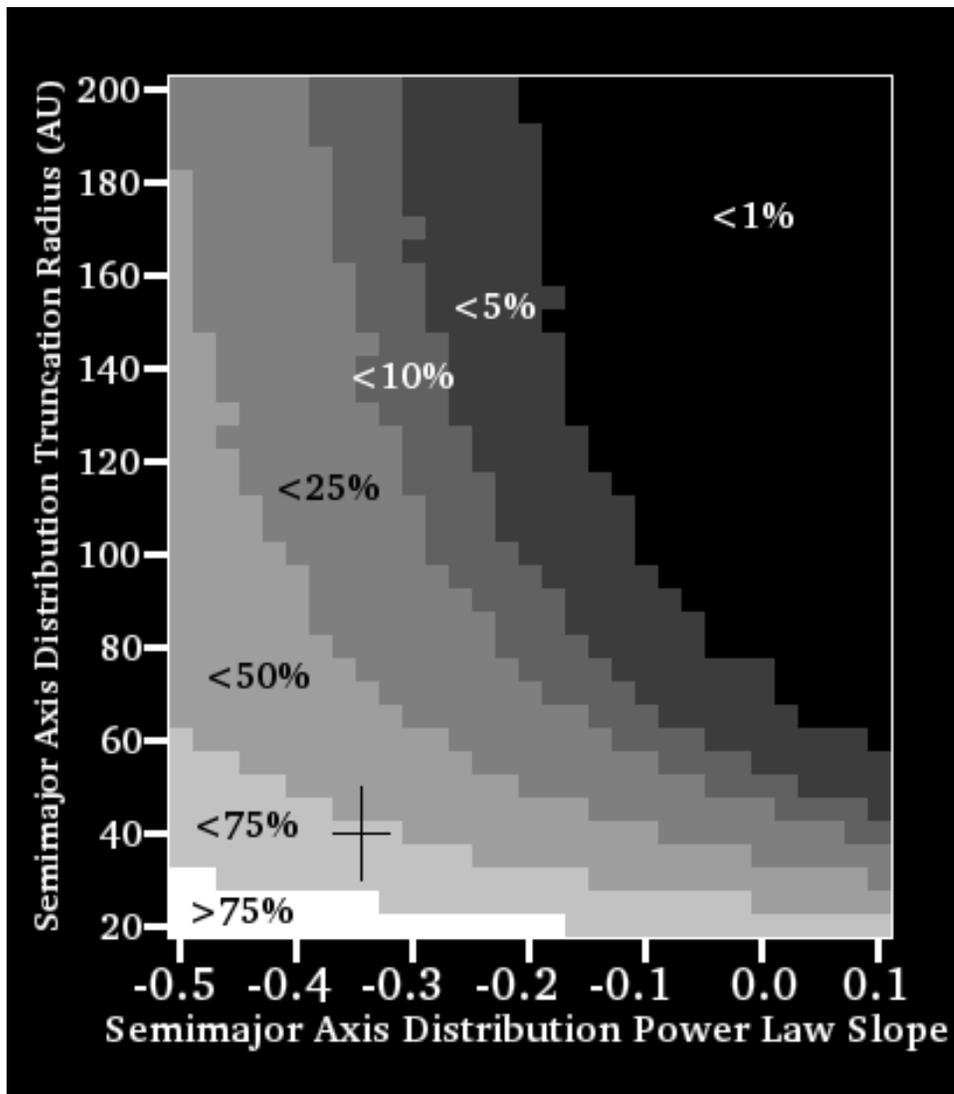


Figure 4.43 The probability of detecting no planets in our survey as a function of the power law slope  $\gamma$  for the distribution of the planets' orbital semimajor axes and the outer truncation radius for this distribution. Here we have used the Burrows et al. (2003) planet models and have set  $\alpha = -1.089$ , the most optimistic slope for the mass power law that is permitted by the 95% confidence interval. From black to white, the contours are for a 0–1%, 1–5%, 5–10%, 10–25%, 25–50%, 50–75%, and more than 75% probability of detecting no planets. The black cross marks our simulation from Section 4.6.

ing losses mean that our nominal  $5\sigma$  limits are formally more like  $4\sigma$  using our sensitivity estimator. However, Kasper et al. (2007) probably have uncorrected processing losses as well, and they may have used a sensitivity estimator that overestimates the sensitivity in the presence of correlated noise. All things considered, the Kasper et al. (2007)  $5\sigma$  limits are probably no more realistic than ours, and the comparison is probably roughly fair. We conclude that the two surveys in fact attained similar power to exclude regions in the  $\alpha, R_{truncation}$  parameter space.

Nielson et al. (2007), using the very sensitive small-separation planet imaging data from the Biller et al. (2007) survey, set, in general, stronger constraints than we can. They state that they can reject at the 95% confidence interval the hypothesis that more than 20% of stars have planets of more than 4 MJ orbiting between 20 and 100 AU, using the Burrows et al. (2003) models for mass-magnitude conversions. Assuming that these planets are distributed according to our adopted values of the mass slope  $\alpha$  and the semimajor axis slope  $\gamma$ , our 95% limit on the frequency of planets of at least 4 MJ in orbits between 20 and 100 AU is 32%.

Nielson et al. (2007) further state that they can reject a distribution with  $\alpha = -1.16$ ,  $\gamma = 0.0$ , and truncation radius 17 AU, or  $\alpha = -1.16$ ,  $\gamma = -0.5$ , and truncation radius 46 AU with 95% confidence, using the Burrows et al. (2003) models. We can reject these two models only at the 24% and 38% confidence levels. Since no distribution would ever be rejected because the values matching the properties of a data set were not drawn from it 24% or 38% of the time, these are not rejection levels at all. This is an illustration of how much better sensitivity the powerful SDI method employed by the Biller et al. (2007) survey obtained at small separations. Our observations lag less far behind theirs in excluding distributions extending out to planets in more distant orbits, but still there is no doubt

this survey set stronger constraints on the planet distributions than ours. This is not surprising when one considers the relative sensitivities plotted in Chapter 3, for all stars except  $\epsilon$  Eri.

Nielson et al. (2007) compare the regions of parameter space excluded by their results with those excluded by the Lafrenière et al. (2007b) survey. In every case Lafrenière et al. (2007b) did better. Thus, of course, our constraints lag even further behind those of Lafrenière et al. (2007b) than behind Nielson et al. (2007). The probable corrections to the sensitivities of Lafrenière et al. (2007b) and Biller et al. (2007) that should be applied to make their limits entirely realistic would close this gap a little, but not much. The  $H$  regime surveys have constrained the distribution of extrasolar planets much more strongly than we have.

Our survey has probed the older stars, however. Since giant planets are expected to experience considerable orbital evolution for at least 0.1 Gyr after their formation, we have probably set limits on a somewhat different distribution of planet orbits than the other surveys. Nonetheless the greatest importance of our survey in constraining the planet distribution is that we used a very different wavelength regime than most other surveys, thereby diversifying the investment of planet-search efforts for the whole astronomical community. Our ability as a community to place some constraints on the distribution of extrasolar planets is thereby protected against the possibility that theoretical models severely overpredict the  $H$  band brightness of giant planets. As discussed above, the theoretical models are uncertain for the longer wavelengths as well; the important thing is that we can now constrain the planet distributions based on null results across a wide range of wavelengths.

## 4.8 Conclusion

Our survey did not detect any planets. It had a good chance of doing so based on the Monte Carlo simulations of Section 4.6, which used what we have considered a reasonable guess at the true distribution of extrasolar planets. However, the chance of detecting no planets was sufficiently high with this distribution that the survey null result does not rule out the distribution. More sensitive surveys may in fact have ruled out this distribution at high confidence, but this is not clear since they have not quoted rejection confidence levels for distributions closely resembling ours. If the distribution we have used in Section 4.6 does in fact describe the population of mature extrasolar planets, a more sensitive survey of our sample set would detect a planet with very high probability. Such a survey could be carried out with the LBT, when it becomes fully operational with adaptive secondaries in place, and is fitted with an  $L'$  and  $M$  band imaging camera.

Since we have focused on an older set of stars than all other planet imaging surveys to date, we have likely probed a population of planets in evolved, stable orbits, while the other surveys probe a population that is still undergoing dynamical evolution. We may thus have set unique constraints on the distribution of the orbits of old planets. The constraints placed by Lafrenière et al. (2007b) are so much better than ours, however, that it is quite possible that a subset of the oldest stars from the Lafrenière et al. (2007b) survey would set just as tight limits as we can.

Our survey was worthwhile in another sense, though: at a time when theoretical models of extrasolar planets differ substantially from one another and are still largely unconstrained observationally, we have set interesting constraints on the distribution of extrasolar planets in a wavelength regime far from the  $H$

band. Thus, if it turns out that current models substantially overpredict the  $H$  band brightness of extrasolar planets, our survey and that of Kasper et al. (2007) will insure that some constraint on the distribution of old and young extrasolar planets, respectively, is still available to the astronomical community. It is easy to imagine that unexpected atmospheric chemistry, cloud structure, or evolution histories could render theoretical planet models substantially inaccurate. The  $H$  band might be very sensitive to model errors because its usefulness depends on strongly super-blackbody flux due to a spectral window at this wavelength and strong line-blanketing over the  $K$  band region. The  $L'$  and  $M$  bands, closer to the peaks of planet-temperature blackbodies, should intuitively be more robust. There is a possible counterexample in the fact that brown dwarfs with  $T_{\text{eff}}$  from 700 K to 1300 K are 0.2-0.7 mag fainter at  $M$  band than models predict (Leggett et al., 2007; Reid & Cruz, 2002). However, this does not affect the  $L'$  band at which most of our observations were carried out, and the likely cause (CO upwellings from deep in the planet (Leggett et al., 2007)) may not apply to colder planets to which most of our  $M$  band observations were sensitive. The main concern is not that the model fluxes at any wavelength are certain, but that consistent results from a wide range of wavelengths allow us to place greater confidence in our conclusions.

The next generation of giant telescopes will be sensitive to fainter, cooler planets around every star system. Such planets are predicted to be substantially redder in  $H - L'$  and  $H - M$  color than their hotter, more massive counterparts that can be detected with current telescopes. Thus, the planet/star flux ratio advantage delivered by longer wavelengths relative to the  $H$  band regime will become increasingly large with the increasing sensitivity delivered at every wavelength by upcoming large telescopes. The  $L'$  and  $M$  bands will likely be more and more

used; indeed, they may well be substantially superior to the *H*-band regime for, for example, a planet-imaging survey of stars within 15 pc carried out by GMT. The investments being made now to develop these wavelengths will likely have a large payoff with future giant telescopes. Space based observations at *L'*, *M* band, and longer wavelengths using JWST will attain background-limited sensitivity well beyond what is possible with GMT, and will thus be good for finding planets in large orbits far from the parent star. Twenty to thirty meter ground-based telescopes will deliver better sensitivity close-in.

## CHAPTER 5

WHERE FROM HERE? LESSONS LEARNED FROM THE CLIO SURVEY, AND THE  
FUTURE POTENTIAL OF THE  $L'$  AND  $M$  BANDS

## 5.1 Introduction

Here we consider what we have learned that can usefully guide future work. First, in Section 5.2 we discuss some observational issues: specifically, what quality of data we have obtained using Clio under various conditions. This should help in developing optimized observing strategies, and choosing the most reasonable targets under sub-optimal observing conditions such as poor seeing. We also determine whether the big picture from our data supports the conclusion that sensitivity goes up as  $\sqrt{t}$ , where  $t$  is the exposure time. This is an important question because of its relevance to the feasibility of obtaining increased sensitivity using extremely long integrations.

In Section 5.3 we consider extensions of our survey using the same or similar target lists and the same wavelength bands, but larger telescopes and/or longer exposures. We compare the power of such hypothetical surveys to the Lafrenière et al. (2007b) survey. Our conclusion from this comparison is that  $L'$  and  $M$  band imaging surveys using sample sets like ours (stars ranging out to 25 pc distances) place weak constraints on planet populations compared to the survey of Lafrenière et al. (2007b) and possible future surveys using their method.  $L'$  and  $M$  band imaging surveys using a sample set like ours are useful mainly to diversify planet-imaging effort against the possibility that current models over-predict the  $H$  band brightness of extrasolar planets.

However, as exemplified by our  $\epsilon$  Eri observations discussed in Chapters 2 and 3,  $L'$  and  $M$  band observations do much better than  $H$ -regime techniques

for the nearest, brightest stars. At present they are substantially more powerful than  $H$  regime methods only for star systems within about 3-4 pc. As larger telescopes are built, this limit will move outward. For a 50 meter telescope the  $L'$  and  $M$  bands will deliver far superior planet-imaging sensitivity for all stars within about 16 pc. This limit is conservative.

In Section 5.4 we provide recommendations for future observing projects in the  $L'$  and  $M$  bands.

## 5.2 Clio Data Quality Under Various Conditions

We will begin with an approximate estimate of Clio's background limited sensitivity. The read noise is about 700 electrons; the gain is roughly 85 electrons per ADU. A typical sky background at  $L'$  is 10,000 ADU/sec. Dark current usually saturates the detector's approximate 40,000 ADU full well in about 7 seconds. A 2 second exposure therefore picks up 20,000 ADU from the sky and 11,500 ADU from dark current. This corresponds to 2.7 million electrons, with an expected  $\sqrt{N}$  noise of 1600 electrons. Adding this to the 700 electron read noise yields a total noise of 1730 electrons, or 22 ADU.

In a 1-hour exposure there are 1800 of these 2 second frames. The noise in the stack should go down by the square root of the number of images ( $\sqrt{N}$  again). Nod subtraction presumably increases the noise by  $\sqrt{2}$ . This indicates the pixel-to-pixel RMS in a final stacked frame should be  $20\text{ADU} \times \sqrt{2} \div \sqrt{1800} = 0.67\text{ADU}$ .

We will assume 40% of the flux from a point source falls in an aperture of radius 2.5 pixels. The noise in the aperture is  $\sqrt{\pi \times 2.5^2} \times 0.67\text{ADU} = 3.0\text{ADU}$ . Correcting this for the fact that the chosen aperture contains only 40% of the flux, we get that the  $10\sigma$  limit for a point source should correspond to an integrated flux of  $3.0\text{ADU} \times 10 \div 0.4 = 74\text{ADU}$ , or 37 ADU/sec for a 2 second exposure.

Given a typical photometric calibration of 15,000 ADU/sec for a 10th mag star, the sensitivity should be  $L' = 16.5$ . While we usually don't get a background limited sensitivity quite this good in an hour, we do get fairly close. Residual noise from column variations may contribute to reducing our sensitivity slightly below this level. It seems clear, however, that there are no large, unexpected noise sources.

A similar calculation may be done for the  $M$  band. We will use the specific example of the April 2006 Vega observations. 110 images were taken, each consisting of 90 coadded frames, for a total of 9900 frames. The true exposure was about 0.2596 sec. The average sky level on each frame was about 36,700 ADU, or about  $3.12 \times 10^6$  electrons, resulting in a  $\sqrt{N}$  noise of 1770 electrons. Adding in the read noise gives 1900 electrons, or 22.3 ADU per pixel, and the equal noise from the nod subtraction frame raises it to 31.5 ADU. 56% of the flux was contained within an aperture of 2.8 pixel radius. We multiply 31.5 ADU/pixel times the square root of the number of pixels contained in a 2.8 pixel radius aperture, and divide by 0.56 to get a true single-frame  $1\sigma$  point source sensitivity of 280 ADU. Dividing this by  $\sqrt{9900}$  gives a  $1\sigma$  point source sensitivity of 2.8 ADU on the final image. This corresponds to a  $10\sigma$  limit of 28 ADU, or 108 ADU/sec. Using the photometric calibration of 4317 ADU/sec for a 10th magnitude star, we obtain a  $10\sigma$  background limited point source sensitivity of  $M = 14.0$ . Our actual  $M$  band sensitivity on this data set peaked at around 13.7. As with the  $L'$  observations, we can say that some additional sources of noise may exist, but they are not large, and our sensitivity is about what we would expect.

Optimizations to Clio, our image processing techniques, and the MMT AO system may result in considerable sensitivity improvements. First, a new detector for Clio (or a new  $L'$  and  $M$  band camera) with lower read noise and dark

current would deliver approximately 0.3 mag better sensitivity. Planned fine-tuning of the MMT AO system could increase the Strehl ratio enough to improve the background-limited sensitivity by 0.5 magnitudes and the contrast-limited sensitivity by a larger amount. Improved data processing and source recovery methods would likely add an additional 0.3 magnitudes. More sophisticated PSF subtraction based on the LOCI algorithm of Lafrenière et al. (2007a) might add additional sensitivity close to the star. In any case it would appear that an upgrade to Clio operating on a fine-tuned MMT AO system with better image processing could deliver 1.1 magnitudes better sensitivity, setting background limits near  $L' = 17.6$  or  $M = 14.5$  in a 1.5-2 hour exposure. The gains would likely be greater close to the star.

It is interesting to consider on what variables our sensitivity may depend. We will consider exposure time, PSF FWHM, sky brightness, and parallactic rotation. We have too little  $M$  band data for meaningful statistics so we analyze only the  $L'$  data. There is no reason to think the  $M$  band would behave qualitatively differently.

In Figure 5.2 we show how sensitivity varies with increasing exposure time. The question of whether sensitivity goes up as the square root of the exposure time as expected is extremely important if very long integrations (ie, to detect the planet of  $\epsilon$  Eri) are contemplated. On the left panel of Figure 5.2, the peak apparent  $L'$  sensitivity and the median  $\Delta$ -mag sensitivity at 0.5 and 1.0 arcsec separations are plotted. The right panel shows the same values scaled to a constant 1 hour exposure under the assumption that sensitivity goes as  $\sqrt{t}$ .

The background-limited sensitivity shows a clear positive slope in the left hand plot, and actually has a slight positive slope even in the right hand panel. This indicates the background sensitivity goes up faster than the expected square

root law. This is actually not surprising, since longer exposures were more likely to have greater parallactic rotation, smearing ghosts and nod-subtraction artifacts out of existence. The contrast limited sensitivity values are so scattered that even the 1.0 arcsecond values barely have a statistically significant slope on either plot. However, in both cases the slope is positive on the left, and still positive but statistically indistinguishable from zero on the right. These data thus support the expectation that our sensitivity goes up as the square root of the exposure time at both large and small radii from the star. It appears to go up even faster in the background limited case for 0.5-1.5 hour exposures. We note that tests comparing halves or other small chunks of the image stack on a given target to the full image stack have also shown a roughly  $\sqrt{t}$  dependence of the sensitivity.

In Figure 5.2 we plot the 1-hour scaled background-limited sensitivity, and the similarly scaled  $\Delta$ -mag sensitivities at 0.5 and 1.0 arcsec, against parallactic rotation. The background limited sensitivity clearly goes up until about 50 degrees of rotation, when it seems to level off. This is not surprising, as 50 degrees should be enough to remove most ghosts and the incremental improvement from further rotation will be small. The close-in sensitivities do not show a clear trend.

In Figure 5.3 we plot the 1-hour scaled peak background limited sensitivity against the sky brightness in ADU/sec. The scatter remains large, but a fairly clear dependence can be seen. Although the noise will be higher at higher sky backgrounds, acceptable data can be taken right up to the highest plotted sky levels, which correspond to about 14,000 ADU/sec. We note that a striking trend is seen in the sky brightness values from different runs: December, January, and early April yield brightnesses in the 9000 ADU/sec range; June and July are more like 12,000 ADU/sec or even more. It seems obvious that this is to be attributed to the higher temperature in the summer contributing to the thermal glow from at-

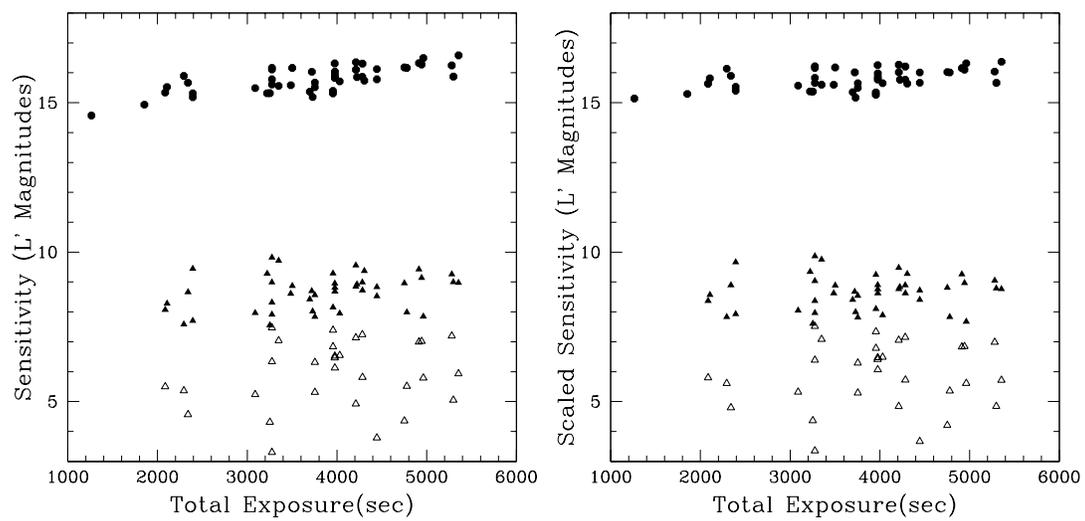


Figure 5.1 Sensitivity as a function of exposure time. On the left, raw sensitivities are plotted. On the right, the sensitivities have been scaled to a 1-hour exposure assuming a  $\sqrt{t}$  dependence. Filled circles are the peak background-limited sensitivity. Filled triangles are the median  $\Delta$ -mag sensitivity at 1.0 arcsec, and open triangles are the same at 0.5 arcsec.

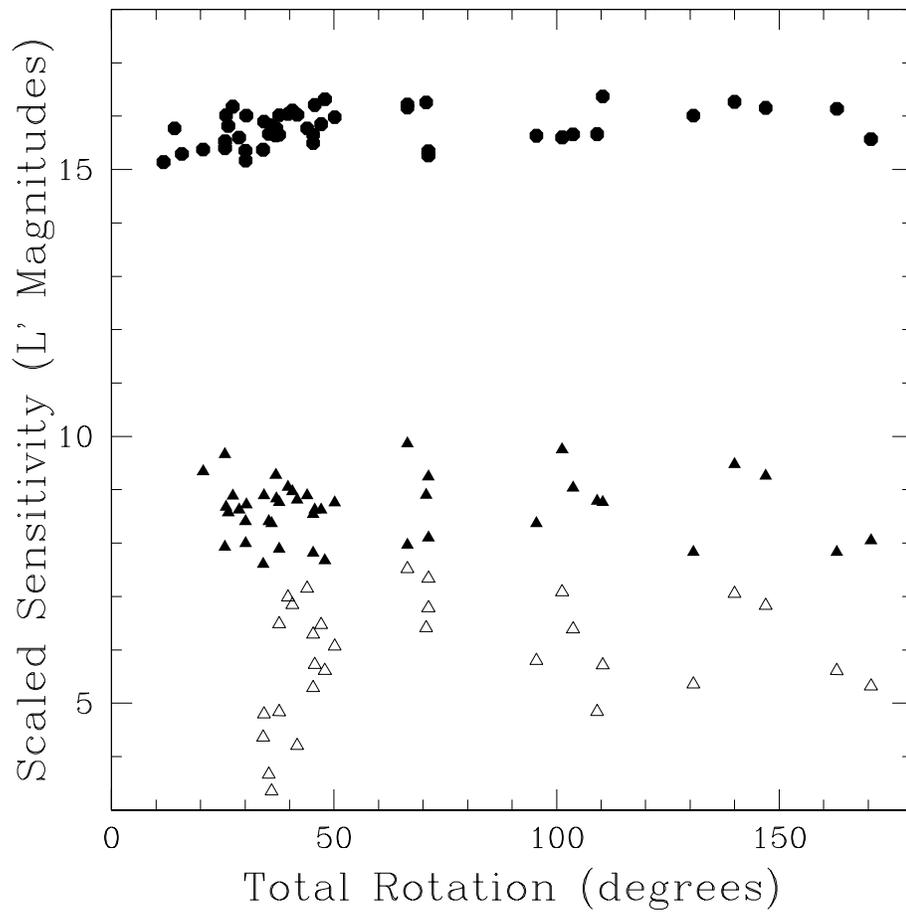


Figure 5.2 Sensitivity as a function of parallactic rotation. Filled circles are the peak background-limited sensitivity. Filled triangles are the median  $\Delta$ -mag sensitivity at 1.0 arcsec, and open triangles are the same at 0.5 arcsec.

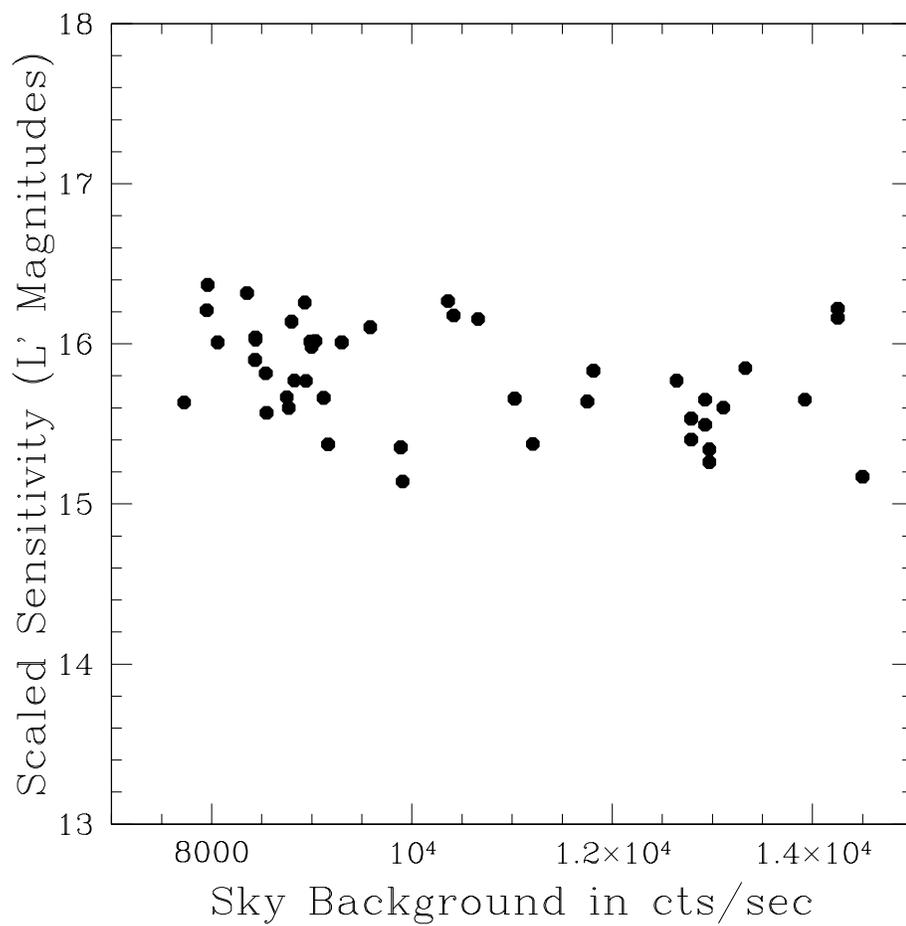


Figure 5.3 Peak background-limited sensitivity as a function of sky brightness in ADU/sec.

mosphere and optics. While our cold-weather observations do tend to be slightly more sensitive, sky brightnesses up to 14,000 ADU/sec are no cause for concern: excellent data can still be taken.

In Figure 5.4 we show the 1-hour scaled  $\Delta$ -mag sensitivities at 0.5 and 1.0 arcsec plotted against the FWHM of the unsaturated PSF. The dependence one would expect is there: the sharper PSFs yield higher  $\Delta$ -mag sensitivities. The dependence is not extremely strong, however. The plot suggests good data can still be taken in substandard seeing with PSFs of FWHM = 3.7 pixels (0.18 arcsec). The author feels, however, that if the unsaturated FWHM gets beyond 4 pixels (0.2 arcsec), planet imaging observations are probably a waste of time. Under such conditions, lucky-imaging observations of bright targets such as close binaries are a better use of observing time. In lucky-imaging hundreds or thousands of short exposures are taken, and all but the sharpest are rejected. Lucky imaging with Clio+MMTAO on bad seeing nights can yield excellent results, especially in the  $M$  band where the exposures are shortest and the seeing most forgiving.

We note that windshake can be a more serious problem than bad seeing. As an example of this, excellent parallactic rotation was obtained on the HD 78141 data set, but windshake made the images of the primary mildly elliptical at a variable position angle, and as a result the PSF subtraction was unusually poor. Under conditions of severe windshake, abandoning a planet-imaging program in favor of one that can use lucky imaging is probably the best plan.

In conclusion, we find that the sensitivity of Clio images in both the contrast-limited and background limited regimes appears to follow the  $\sqrt{t}$  law. We note also that useful planet-search data may be obtained with Clio under a wide range of conditions, including high sky brightness and substandard seeing. Windshake or very bad seeing (producing an AO-corrected PSF with a FWHM of more than

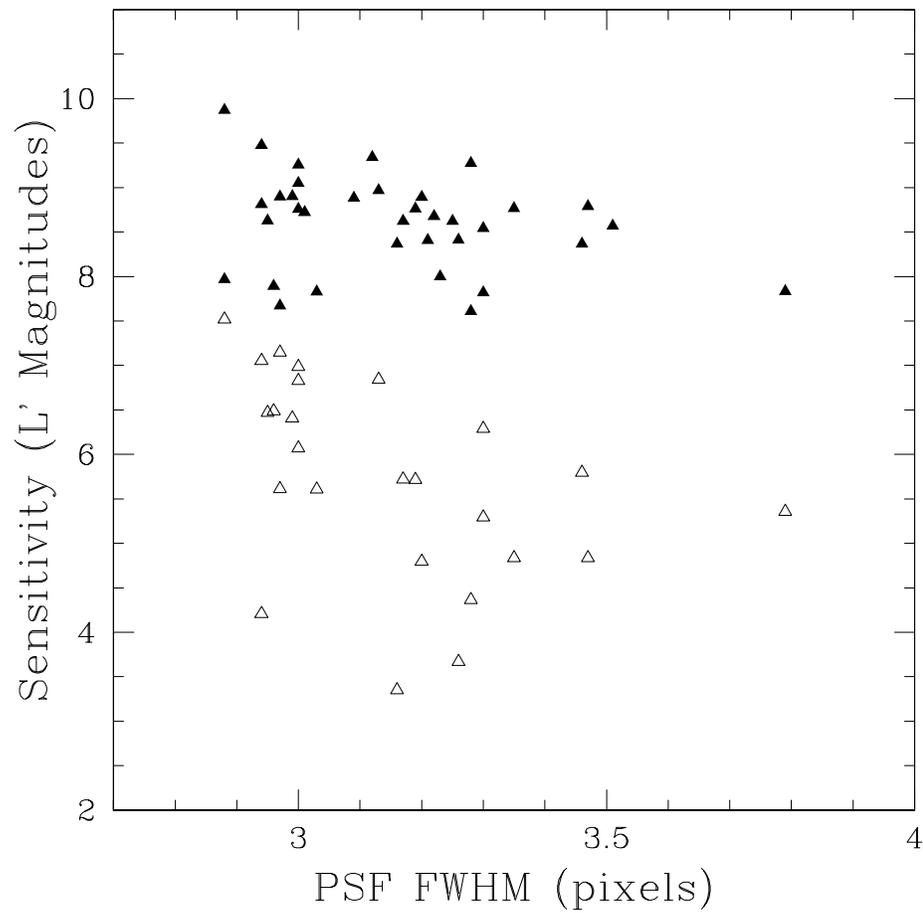


Figure 5.4 1-hour scaled  $\Delta$ -mag sensitivity at 0.5 and 1.0 arcsec as a function of the PSF FWHM in pixels. Open triangles are for 0.5 arcsec; filled ones for 1.0 arcsec.

4 pixels, or 0.2 arcsec), are two examples of conditions that probably do warrant ceasing planet-search observations and switching to lucky-imaging targets to avoid spending time gathering planet-search data that will not be useful. Clio observations attain slightly higher sensitivity in cold weather, so if possible science targets requiring sensitive observations should be observed in the winter or early spring. The effect is not as important as having the target well placed for observing, of course: summer objects will still be better observed in the summer.

### 5.3 Comparing the Future Potential of Long and Short Wavelength Regimes

Table 5.1 compares the Lafrenière et al. (2007b) survey with our survey and with hypothetical  $L'$  and  $M$  band surveys of our target sample with larger telescopes and/or longer exposures. The telescopes considered are the 6.5 meter MMT we actually used; one aperture of the 8.2 meter LBT; both apertures of the LBT combined incoherently (ie, in post-processing after data acquisition with two separate cameras); the two apertures of the LBT combined as a Fizeau imaging interferometer; a hypothetical 15 meter telescope; and finally the 24 meter GMT. For all telescopes we assumed that the flux of the faintest detectable object was proportional to the inverse-square of the telescope aperture. We assumed that the Fizeau interferometer concentrates the light from the LBT by a factor of 2 relative to incoherent combination of the two apertures; this is conservative. For each telescope, results are provided for a survey using the same exposures we did and one using exposures 3 times longer.

The data in the table were produced using Monte Carlo simulations exactly like those discussed in Chapter 4 on altered versions of the sensitivity files used for the Chapter 4 simulations. For the simulation of an MMT survey with 3x longer exposures, we simply increased the sensitivity at every point in our real

sensitivity files by  $\sqrt{3}$ , or 0.5964 mag. For the other telescopes, we assumed that a larger format, cleaner detector with fewer ghosts was in use in a hypothetical next-generation version of Clio on the telescope, and that greater flexibility in the image nod sequences eliminated artifacts from negative nod-subtraction star images. It made sense to assume that with this hypothetical camera our 90th percentile sensitivity vs radius curves would apply throughout the field. We adopted the peak sensitivity as the background limit. We scaled all separations inside the separation at which the peak sensitivity was attained by the ratio of the MMT's 6.5 m aperture to the aperture of the larger telescope, to pull in a given sensitivity to a smaller radius, simulating the smaller diffraction pattern given by the larger telescope. We separated a 'star' and 'background' contribution from each noise measurement under the usual assumption that noise adds in quadrature. We improved the background-limited sensitivity using the assumption that sensitivity is proportional to the telescope aperture squared, and added the resulting scaled-down noise to the star noise to get a new total noise from the improved background limit. We preserved the Clio maximum field radius, but set the sensitivity at every separation beyond that where the peak value was obtained to the peak value. For the 3x longer exposures, we simply increased the sensitivities by 0.5964 mag as before.

Finally, we ran Monte Carlo simulations using the same code as for Section 4.6 above, and the same planet distributions with  $\alpha = -1.44$ ,  $\gamma = -0.35$ , and a truncation radius of 40 AU. Table 5.1 records the confidence with which this model distribution could be rejected by a null result from the different hypothetical surveys. The median mass and median semimajor axis for detected planets are also recorded. For the simulated surveys by the full LBT and larger telescopes, the results are biased in the direction of weaker constraints by the fact that our sim-

ulations did not include planets less massive than 1.0 MJ; models suggest that such planets would be detectable in large numbers using the GMT. We note that according to this table the GMT, or at least a 15 meter telescope, would be required to rule out the model distribution at the 95% confidence level using an  $L'$  and  $M$  band survey with our target sample

That it could far more easily be done using an  $H$  band regime survey is indicated by Table 5.2. In this table, we consider a distribution of planets used as a test case in the Lafrenière et al. (2007b) survey. This distribution had  $\alpha = -1.2$  and  $\gamma = -1.0$ . No truncation radius was adopted; instead Lafrenière et al. (2007b) considered the upper limits they could place on the fraction of stars with planets between 0.5 and 13 MJ and orbital semimajor axes in 3 different intervals, assuming that within the intervals the planets followed the power laws above. With 95% confidence, they concluded that no more than 29% of stars have a planet with mass between 0.5 and 13 MJ and orbital semimajor axes between 10 and 25 AU. For orbits between 25 and 50 AU the corresponding limit was 13%, and for 50-250 AU it was 9%.

Our simulations do not consider planet masses below 1.0 MJ, so we are unable to follow Lafrenière et al. (2007b) in considering a full range from 0.5-13 MJ. Accordingly, we have adjusted their values by the ratio of the number of planets between 1.0 and 13 MJ to the number of planets between 0.5 and 13 MJ, given a power law slope  $\alpha = -1.2$ . This factor is:

$$\left[ \int_{1.0}^{13.0} M^{-1.2} dM \right] \div \left[ \int_{0.5}^{13.0} M^{-1.2} dM \right] = 0.7296.$$

Thus, we transform the Lafrenière et al. (2007b) 0.5-13 MJ limits of 29%, 13%, and 9%, respectively, for their three regions of semimajor axis space into 1-13 MJ limits of 21%, 9%, and 7%. These can then be directly compared in Table 5.2 to the limits set by the null result of our real survey, and null results for various

hypothetical surveys using the same bands and sample set with larger telescopes and/or longer exposures. We note that our Monte Carlo code does not work with a slope of exactly -1.0, so we have used  $\gamma = -0.99$ ; the difference should be negligible.

Table 5.1. Performance of Bigger Telescopes

Telescope Used	Exp. Time	Model Set Used	Rejection Confidence	Median Mass	Median Semimajor Axis
MMT	1x	Burrows	36.6%	8.92 MJ	26.60 AU
MMT	1x	Baraffe	50.3%	8.51 MJ	26.07 AU
MMT	3x	Burrows	49.9%	8.24 MJ	26.13 AU
MMT	3x	Baraffe	62.5%	7.62 MJ	25.80 AU
Half LBT	1x	Burrows	57.4%	8.13 MJ	25.73 AU
Half LBT	1x	Baraffe	69.9%	7.43 MJ	25.48 AU
Half LBT	3x	Burrows	69.4%	7.41 MJ	24.99 AU
Half LBT	3x	Baraffe	79.1%	6.60 MJ	25.18 AU
Full LBT	1x	Burrows	61.5%	7.82 MJ	25.95 AU
Full LBT	1x	Baraffe	73.1%	7.38 MJ	26.00 AU
Full LBT	3x	Burrows	71.7%	7.15 MJ	25.17 AU
Full LBT	3x	Baraffe	81.4%	6.30 MJ	25.48 AU
Fizeau LBT	1x	Burrows	63.1%	7.54 MJ	26.07 AU
Fizeau LBT	1x	Baraffe	74.7%	6.86 MJ	26.08 AU
Fizeau LBT	3x	Burrows	73.7%	6.93 MJ	25.33 AU
Fizeau LBT	3x	Baraffe	82.9%	6.07 MJ	25.70 AU
15 M	1x	Burrows	84.3%	6.69 MJ	23.31 AU
15 M	1x	Baraffe	91.5%	5.67 MJ	23.01 AU
15 M	3x	Burrows	90.0%	6.05 MJ	22.94 AU

Table 5.1—Continued

Telescope Used	Exp. Time	Model Set Used	Rejection Confidence	Median Mass	Median Semimajor Axis
15 M	3x	Baraffe	95.6%	5.07 MJ	22.89 AU
GMT	1x	Burrows	94.6%	5.54 MJ	22.12 AU
GMT	1x	Baraffe	98.0%	4.53 MJ	21.91 AU
GMT	3x	Burrows	96.7%	5.08 MJ	21.62 AU
GMT	3x	Baraffe	98.9%	4.09 MJ	21.51 AU

Note. — The rejection confidence is the confidence level with which a copy of our survey using the telescope and model set in question would reject the reasonable planet distribution we used in Section 4.6 above if it obtained a null result. Note that this is equal to the probability of finding at least 1 planet. ‘Half LBT’ refers to a single unit-aperture on the LBT, ‘Full LBT’ to simultaneous observations using both apertures with the data digitally combined after acquisition, and ‘Fizeau LBT’ refers to using both apertures interferometrically combined in the Fizeau imaging mode. At ‘Full LBT’ and beyond the results are pessimistically biased by our simulations’ failure to include planets less massive than 1 MJ.

Table 5.2 shows that it would take the GMT to better the Lafrenière et al. (2007b) constraints using an  $L'$  and  $M$  band survey of our target set. Even then, a better constraint is obtained only for the innermost semimajor axis interval. The Lafrenière et al. (2007b) survey constrained the planet distribution far more strongly than we did. It almost, but not quite, ruled out our best-guess distribution from Section 4.6 (with  $\alpha = -1.437$ ,  $\gamma = -0.345$ , and truncation radius 40 AU. A survey of our target list using the Lafrenière et al. (2007b) methods would have obtained much better sensitivity than we did on all but a handful of bright, very nearby stars. An extended survey using the Lafrenière et al. (2007b) techniques on Gemini or a more powerful telescope such as the LBT would detect a planet or rule out our Section 4.6 distribution at the 95% confidence level. Any future  $L'$  and  $M$  band surveys with similar target samples to ours using currently existing telescopes, should be carried out with the idea of diversifying the investment of planet-imaging effort against the possibility that current models severely overpredict the  $H$  band brightness of extrasolar planet, not of obtaining uniquely good sensitivity to planets around stars beyond about 5 pc.

With very nearby stars, or with very large telescopes, the situation is quite different. For a given size of telescope with given sensitivities in the  $H$ ,  $L'$ , and  $M$  bands, there will be a breakpoint distance for each band pair,  $H - L'$ ,  $H - M$ , and  $L' - M$ . We define this breakpoint as the distance such that observations in the **longer** wavelength band will deliver better **background limited** sensitivity to planets orbiting stars **nearer than** the breakpoint distance. For systems more distant than the breakpoint, the shorter wavelength band will be sensitive to lower mass planets. **As telescope aperture increases, sensitivity at any distance extends to fainter, redder planets, and the breakpoint distance moves outward.** We note that our experience with  $\epsilon$  Eri suggests that whenever the  $L'$  and  $M$

Table 5.2. Performance of L' Surveys vs Lafrenière et al. (2007b)

Survey	10-25 AU Limit	25-50 AU Limit	50-250 AU Limit
GDPS	21%	9%	7%
MMT 1x	89%	46%	68%
MMT 3x	62%	34%	48%
Fizeau LBT 1x	47%	24%	28%
Fizeau LBT 3x	36%	20%	22%
GMT 1x	16%	13%	18%
GMT 3x	14%	12%	16%
<i>Model</i>	9%	7%	0%

Note. — ‘GDPS’ means the Lafrenière et al. (2007b) survey. The limits quoted for various separation ranges are the 95% confidence upper limits on the fraction of stars with planets in the specified semimajor axis interval, using the Baraffe et al. (2003) models and the values  $\alpha = -1.2$  for the mass power law and  $\gamma = -1.0$  for the semimajor axis power law, following Lafrenière et al. (2007b). These results are for planets with masses between 1 and 13 MJ. The last, italicized line of the table gives the occurrence rate for planets in each interval from our ‘best guess’ model distribution of Section 4.6.

bands deliver background-limited sensitivity comparable to that of  $H$  band methods, they deliver a considerable advantage in the contrast-limited regime. We are thus using the breakpoint distances as a sort of proxy for the distance within which the longer wavelengths become unambiguously superior at all separations from the star. We note that JWST will likely be launched before the GMT becomes operational, and for truly background limited observations far from bright stars will offer unparalleled sensitivity at the  $L'$  and  $M$  bands. The niche for  $L'$  and  $M$  band observations with giant ground-based telescopes such as the GMT will be the detection of planets at smaller separations where JWST is strongly contrast limited. Many or most giant extrasolar planets may be found at just such rather small separations.

Table 5.3 gives the  $H - L'$ ,  $H - M$ , and  $L' - M$  breakpoint distances for a variety of telescopes under our two favored model sets. Note that in the construction of this table a limiting magnitude of  $H = 23.0$  for an 8 meter telescope has been adopted. This is better than broadband  $H$  observations on an 8 meter telescope can typically do, but is set roughly equal to the broadband  $H$  magnitude equivalent to the peak sensitivities obtained to the narrower, optimized-band observations of Lafrenière et al. (2007b).

Table 5.3. Distances at which various wavelength bands become superior

Apert.	Mod.		$H - L'$	$H - M$	$L' - M$	$H - L'$	$H - M$	$L' - M$
	Set	Age	Break	Break	Break	Mass	Mass	Mass
8 m	Bur	0.5 Gyr	0.96 pc	3.60 pc	5.84 pc	1.83 MJ	2.47 MJ	4.34 MJ
8 m	Bar	0.5 Gyr	2.94 pc	1.34 pc	<0.5 pc	1.78 MJ	1.39 MJ	...
8 m	Bur	1.0 Gyr	1.11 pc	3.38 pc	5.64 pc	2.99 MJ	3.81 MJ	6.22 MJ
8 m	Bar	1.0 Gyr	2.94 pc	1.49 pc	<0.5 pc	2.76 MJ	2.24 MJ	...
8 m	Bur	5.0 Gyr	0.96 pc	2.63 pc	4.19 pc	6.97 MJ	8.28 MJ	12.60 MJ
8 m	Bar	5.0 Gyr	3.24 pc	1.87 pc	<0.5 pc	6.55 MJ	7.45 MJ	...
LBT	Bur	0.5 Gyr	1.41 pc	5.10 pc	8.24	2.48 MJ	1.85 MJ	4.34 MJ
LBT	Bar	0.5 Gyr	3.97 pc	1.89 pc	<1 pc	1.39 MJ	1.75 MJ	...
LBT	Bur	1.0 Gyr	1.41 pc	4.69 pc	7.97 pc	3.76 MJ	2.93 MJ	6.23 MJ
LBT	Bar	1.0 Gyr	4.27 pc	2.11 pc	<1 pc	2.24 MJ	2.81 MJ	...
LBT	Bur	5.0 Gyr	1.26 pc	3.73 pc	5.94 pc	8.29 MJ	6.86 MJ	12.61 MJ
LBT	Bar	5.0 Gyr	4.86 pc	2.63 pc	<1 pc	8.29 MJ	7.66 MJ	...
15 m	Bur	0.5 Gyr	2.02 pc	6.72 pc	10.77 pc	1.88 MJ	2.47 MJ	4.30 MJ
15 m	Bar	0.5 Gyr	5.05 pc	2.47 pc	<1 pc	1.77 MJ	1.38 MJ	...
15 m	Bur	1.0 Gyr	2.02 pc	6.27 pc	10.59 pc	2.98 MJ	3.79 MJ	6.25 MJ
15 m	Bar	1.0 Gyr	5.45 pc	2.78 pc	< 1 pc	2.76 MJ	2.24 MJ	...
15 m	Bur	5.0 Gyr	1.72 pc	5.02 pc	7.89 pc	6.93 MJ	8.38 MJ	12.65 MJ
15 m	Bar	5.0 Gyr	6.34 pc	3.38 pc	<0.5 pc	7.59 MJ	6.45 MJ	...
GMT	Bur	0.5 Gyr	3.08 pc	10.75 pc	17.50 pc	1.86 MJ	2.46 MJ	4.33 MJ

Several things are noticeable from Table 5.3: **With an 8 meter telescope,  $L'$  and  $M$  band observations are preferred for sources within about 3-3.5 pc.** For larger telescopes, the breakpoints move outward rapidly. **The longer wavelengths are better within 4-5 pc for the LBT, 8-10 pc for the GMT, and 16-23 pc for a 50 meter telescope.** The breakpoints are only weakly sensitive to age, but it is a consistent trend that the Burrows et al. (2003) models show them moving inward with age and the Baraffe et al. (2003) models show them moving outward. **A far more important difference between the two model sets is that in every case the Burrows et al. (2003) models strongly prefer  $M$  band observations while those of Baraffe et al. (2003) strongly prefer  $L'$ .** As in similar discussions above, we note that the filter sets used to obtain magnitudes from theoretical spectra are slightly different for the Burrows et al. (2003) and the Baraffe et al. (2003) models. While this may explain some of the differences, inherent differences in the models do exist and are probably the dominant effect.

Table 5.3 analyzes the background limited case. It makes it clear than an observation of, e.g., a star at 5 pc distance with a 50 meter telescope would obtain better sensitivity far from the star in the  $M$  band than the  $H$  band. It does not answer the question of which wavelength would provide best sensitivity in the contrast-limited regime close to the star. We can surmise based on our  $\epsilon$  Eri data that whenever the background-limited sensitivity of longer and shorter wavelengths is comparable, the longer wavelengths give better sensitivity in the contrast-limited regime by a large factor. This rule is very likely to apply for the new generations of giant telescopes as well.

We note that **the conclusion is inescapable that the  $L'$  and  $M$  bands will be needed for efficient imaging of the coolest planets with the next generation of giant telescopes.** At any wavelength the coolest planets will only be detectable

Table 5.3—Continued

Apert.	Mod.		$H - L'$	$H - M$	$L' - M$	$H - L'$	$H - M$	$L' - M$
	Set	Age	Break	Break	Break	Mass	Mass	Mass
GMT	Bar	0.5 Gyr	8.44 pc	4.10 pc	<1.5 pc	1.75 MJ	1.41 MJ	...
GMT	Bur	1.0 Gyr	3.38 pc	10.00 pc	17.00 pc	3.02 MJ	3.76 MJ	6.25 MJ
GMT	Bar	1.0 Gyr	9.02 pc	4.39 pc	<1.5 pc	2.79 MJ	2.21 MJ	...
GMT	Bur	5.0 Gyr	2.78 pc	8.00 pc	13.00 pc	6.94 MJ	8.35 MJ	12.89 MJ
GMT	Bar	5.0 Gyr	10.17 pc	5.55 pc	<1.5 pc	7.59 MJ	6.52 MJ	...
50 m	Bur	0.5 Gyr	6.00 pc	22.94 pc	36.75 pc	1.83 MJ	2.50 MJ	4.33 MJ
50 m	Bar	0.5 Gyr	18.06 pc	8.38 pc	<4.0 pc	1.76 MJ	1.39 MJ	...
50 m	Bur	1.0 Gyr	7.00 pc	21.31 pc	35.12 pc	3.00 MJ	3.80 MJ	6.16 MJ
50 m	Bar	1.0 Gyr	18.06 pc	9.49 pc	<4.0 pc	2.74 MJ	2.26 MJ	...
50 m	Bur	5.0 Gyr	6.00 pc	16.44 pc	27.00 pc	6.97 MJ	8.26 MJ	12.78 MJ
50 m	Bar	5.0 Gyr	21.14 pc	11.16 pc	<4.0 pc	7.55 MJ	6.43 MJ	...

Note. — This table assumes the following  $H$ ,  $L'$ , and  $M$  band background limited sensitivities in magnitudes for the different telescopes: 8 meter: 23.0, 16.5, and 13.5; 2x8.2 meter LBT: 23.75, 17.25, and 14.25; 15 meter: 24.36, 17.86, and 14.86; 24 meter GMT: 25.39, 18.89, and 15.89; 50 meter: 27.0, 20.5, and 17.5. The ‘Break’ distance with each aperture and band pair is the distance within which the longer wavelength band delivers better sensitivity. The masses given are the minimum detectable masses at the break distances.

around the nearest stars, and as Table 5.3 shows, the  $L'$  and  $M$  bands provide the best sensitivity around the nearest stars. The hotter planets that will probably be discovered at larger distances at  $H$ -regime wavelengths will be very interesting and may be lower mass than those found around nearer, older stars with  $L'$  and  $M$ . However, the very coolest planets are interesting because they will represent planets which are either old, low mass, or both. As such they will be the most comparable to the planets in our own solar system. Also, the nearest stars will be those around which the Terrestrial Planet Finder (TPF) will be most likely to find an Earth-like planet. It is the  $L'$  and  $M$  bands on new giant telescopes that offer the best hope of finding the giant neighbors of such new earths from the ground.

#### 5.4 Future Observing Plans

$L'$  surveys with sample sets like ours or that of Kasper et al. (2007) should be carried out with current-generation telescopes, only with objective is to diversifying observing effort against possible model errors. If the models of Burrows et al. (2003) and Baraffe et al. (2003) are approximately correct,  $H$ -regime surveys such as those of Biller et al. (2007) and Lafrenière et al. (2007b) can set far stronger constraints on the distributions of extrasolar planets than  $L'$  surveys of stars ranging out to 25 pc distance.

In case it is thought desirable to carry on more work similar to our survey for reasons of diversification, we list in Table 5.4 a few promising stars from our target sample that we never observed.

For the present, and assuming the theoretical models are correct,  $L'$  and  $M$  band observations are sensitive to lower-mass planets than  $H$ -regime observations only around the nearest bright stars. The promising targets accessible from the northern hemisphere are  $\epsilon$  Eri,  $\tau$  Ceti, and 61 Cyg A and B. Adding the south-

Table 5.4. Distances, Ages, and Spectral Types of Survey Targets

Star	Age(Gyr)	Dist(pc)	Type	Binary?
GJ 171.2	0.15	16.37	K5	No
GJ 3305	0.020	29.8	M1	No
HD 35850	0.020	26.84	F7V	No
GJ 216B	0.5	8.01	K2V	No
GJ 3453	0.2	25.44	M1	No
GJ 282B	0.5	13.46	K5	No
GJ 3615	0.15	20.83	K1V	No

ern hemisphere brings in  $\alpha$  Cen A and B, and  $\epsilon$  Indi. These and Sirius are the only non-M-stars within 4 pc listed in Cox (2000). Sirius is interesting, but many potential planetary orbits would be destabilized by Sirius B, and the great brightness of the primary would strain even the  $M$  band's ability to bring out faint substellar objects near a bright star. The Clio phase plate coronagraph being pioneered by Matt Kenworthy (Kenworthy et al., 2007) might give the close-in sensitivity required to reach potentially stable orbits for planetary-mass objects in the Sirius system.

Taking deep, several hour exposures at both the  $L'$  and  $M$  bands on the four stars mentioned above that are visible from the northern hemisphere would be a good observing project with Clio. Five-hour exposures would attain background-limited  $10\sigma$  sensitivities in the 6-9 MJ range for  $\tau$  Ceti, the 4-5 MJ range for 61 Cyg A and B, and the 1.8-2.0 MJ range for  $\epsilon$  Eri, where we have assumed ages of 5 Gyr,

2 Gyr, and 0.56 Gyr for the 3 systems, respectively.

If a Clio-like camera could be operated with an adaptive secondary AO system in the southern hemisphere, the background-limited sensitivity near  $\alpha$  Cen would be 4-5 MJ assuming the system is 5 Gyr old and 2-3 MJ assuming it is 2 Gyr old. These sensitivities are quite interesting, and there is no doubt that observations with Clio would obtain unparalleled sensitivity around each of the star systems mentioned. However, the small sample size and the fact that sensitivity to 1 MJ objects is not reached means that it is statistically unlikely that a detectable planet would be present around one of these very nearby stars.

A Clio survey of all stars within 4 pc, including the M-stars not included in the list above, would be very interesting. However, we note that massive planets are not expected to be found orbiting  $M$  stars, and that because of the old age of most of these systems, sensitivity to 1 MJ objects would not be reached.

Finally, **Clio+MMTAO is the only currently operating system that seems to have a chance of directly imaging the known planet of  $\epsilon$  Eri.** The attempt would need to be made in the  $M$  band, near the 2010 apastron. The required exposure length corresponds to 5-10 good nights. It is possible that applying the LOCI algorithm described in Marois et al. (2006) to Clio data could lessen the time to detect  $\epsilon$  Eri b considerably, as this algorithm appears to be extremely powerful for subtracting away the bright PSFs of stars. Another possibility is the Kenworthy et al. (2007) phase plate coronagraph. Still, at least 3 nights would need to be scheduled for the project to obtain a reasonable likelihood of making the detection despite the vagaries of weather and equipment.

**An  $M$  band detection of  $\epsilon$  Eri b would be extremely interesting.** It would provide a much needed observational constraint on theoretical planet models, as well as a spectacular verification of the viability of planet imaging. It would

also shrink the error bars on the orbit of the planet and hence on the mass of both planet and star. It would point the way toward successful planet imaging strategies for the future.

The future of the  $L'$  and  $M$  bands is brighter than the present. Without  $L'$  and  $M$  band capability the next generation of giant telescopes will be blind to the coolest, lowest mass objects in nearby star systems. These faint, cool planets are very interesting because they represent the Jupiters (and possibly Saturns too) of solar systems in which Terrestrial Planet Finder-type missions may find Earths. To take advantage of the excellent sensitivity  $L'$  and  $M$  band observations will offer for very cool planets orbiting nearby stars, **all of the next generation of giant telescopes should be built with adaptive secondary mirrors**. Although JWST will provide unequalled background-limited sensitivity at the  $L'$  and  $M$  bands, its contrast limited performance is uncertain. Giant ground-based telescopes such as the GMT will certainly provide better sensitivity to planets at small separations, and current null results (Kasper et al. (2007), Nielson et al. (2007), Lafrenière et al. (2007b), and this work) may indicate it will be at these separations that the majority of extrasolar giant planets will be found.

Care should be taken to equip the new generation of giant telescopes for sensitive  $L'$  and  $M$  band observations. Failing to do this would cause an interesting and useful wavelength regime for planet imaging studies to fade into the thermal noise.

## APPENDIX A

THE ALGORITHMS USED BY THE LEGOLAS04 PROCESSING PACKAGE FOR CLIO  
IMAGES

**Abstract:** We describe here our data reduction pipeline for images from the Clío camera. This pipeline was largely developed by the author, but Suresh Sivanandam made important contributions. The processing makes extensive use of routines from Press et al. (1992), and from the `cfitsio` library. This is not the only pipeline for reducing Clío data; others have been developed by Daniel Apai, Phil Hinz, and Matt Kenworthy. The one we discuss here is probably the most versatile and sophisticated, and certainly the most thoroughly used. Its development was begun immediately following Clío's first light on the MMT in June 2005, early than that of any other Clío pipeline except the one Phil Hinz used for the Hinz et al. (2006) images. It retains one characteristic that was well matched to the early Clío data but probably would not have been chosen for a pipeline designed based on the slightly different data produced by Clío's more sophisticated operation during the April 2006 and subsequent observing runs. This is the `colfudge` routine for removing column pattern noise, which turns out to be noticeably sub-optimal for later Clío data.

We begin with a brief discussion of Clío's scientific reason for being, and its capabilities and limitations. This leads naturally into a discussion of the optimal observing strategy, and the type of data that are acquired. With this groundwork laid, we proceed to a detailed discussion of how we process the data.

### A.1 The Clio Camera, its Capabilities and Limitations

Clio is an  $L'$  and  $M$  band optimized AO imager. These wavelengths are interesting for extrasolar planet imaging surveys because they deliver a more favorable planet/star flux ratio than shorter wavelengths. They are challenging because the thermal background is much higher than at shorter wavelengths. Also, detector technology is less advanced and the available detectors have smaller formats.

The MMT adaptive optics (AO) system delivers lower thermal background than other AO systems because its unique deformable secondary mirror feeds AO-corrected light into the science camera after only two reflections off non-cryogenic optics. AO systems where the deformable mirror is not the secondary require several more warm reflections, each of which adds to the thermal background. Thus, Clio and the MMTAO system represent a capable combination for imaging in the  $L'$  and  $M$  bands. The obtainable sensitivity in the  $M$  band is probably the best in the world at present.

Even with the low instrumental thermal background delivered by MMTAO, the thermal background from the sky at the  $L'$  and  $M$  bands is very bright. A detector with a deep full well and fast read time is necessary to avoid saturating on this bright background. The Clio detector has a full well of about  $3.4 \times 10^6$  electrons, and a read time of about 59.6 msec. The gain of the detector is about 85, resulting in a dynamic range of 40,000 analog-to-digital units (ADU). The bias level is typically around 10,000 ADU, so the detector saturates at 50,000 ADU.

The deep full well and fast read time come at the price of high dark current and read noise with current technology. The dark current is small relative to the thermal background: for example, the dark current takes 7-10 seconds to saturate the detector, but the background will do it in 1.7-3 seconds at  $L'$  and in 0.15-0.3 seconds at  $M$ . The read noise is about 700 electrons. This makes a very small

contribution to the noise, of order 10% or so, but only if the exposure is long enough that the thermal background fills a large proportion of the full well capacity. For example, a 2 second exposure at  $L'$  under good conditions gets about 20,000 ADU ( $1.7 \times 10^6$  electrons) from the thermal background and 8,000 ADU ( $0.68 \times 10^6$  electrons) from the dark current. The  $\sqrt{N}$  noise from photon/electron statistics is then 1300 electrons from the background alone and 1540 from the background plus dark current. Adding the 700 electron read noise in quadrature gives a total noise of 1700 electrons; the contribution of the read noise is only 10%. However, for an 0.1 second exposure the  $\sqrt{N}$  statistical noise from background and dark current together would be only 340 electrons, and the 700 electron read noise would make the dominant contribution. The high read noise thus forces us to use long exposures, typically set so that the background plus dark current take 70% of the full well capacity. This insures that the observations are basically background limited, and gives us the best possible sensitivity against the bright background. Even for these long exposures, the saturation radii of all but the brightest stars are only a few pixels, a small fraction of an arcsecond.

## A.2 Clio Data

As we have already discussed, Clio planet search data is typically taken with a long exposure, set so that the thermal background plus dark current fill about 70% of the detector full well. The nominal exposure that meets this condition ranges from 1200-2000 msec in the  $L'$  band and from 70-200 msec in the  $M$  band. Because of the Clio read time, in the integrate-while-reading (IWR) mode always used for efficient science observation the true exposure is always about 59.6 msec longer than the nominal exposure. The IWR mode, and many other features critical to Clio's performance as an efficient science camera, are due to Andy Bre-

uningering, who wrote the software that runs the detector.

To increase efficiency and reduce data volume, Clio is normally set to take coadded frames. This means that a number of individual exposures are co-added by the software that runs the detector, and only the coadded stack is saved as a 2-D FITS image. The individual frames that made up the coadd are not recoverable. The coadd is a simple sum, not an average. An 10-coadd image has a bias level of 100,000 ADU, a dynamic range of 400,000 ADU, and a saturation level of 500,000 ADU, as opposed to the respective values of 10,000 ADU, 40,000 ADU, and 50,000 ADU for a one-coadd frame. Andy Breuninger also provided a mode that saves 3-D FITS cubes instead of 2-D coadded images. Individual frames can then be obtained by slicing the FITS cubes. FITS cubes offer the same good observing efficiency as coadded images, but the data volume is larger by a factor equal to the number of coadds. This results in longer processing times and can even result in such large data volumes that storage and transfer become difficult. For these reason we have usually taken coadded images, reserving FITS cubes for special situations such as 'lucky' imaging, where they are very useful.

The Clio detector is linear to within a few percent over most of its dynamic range. For very low count levels, perhaps within 1,000 ADU of the bias level, it is significantly nonlinear. The nonlinearity is in the sense that a given number of photoelectrons deposited into a very empty well are translated into a larger change in ADU than the same number deposited into a fuller well. This can render photometry based on very short exposures highly inaccurate. However, for all of our planet search and photometric standard star data the dark current and thermal background produced enough ADU to lift the whole image out of the regime where nonlinearity was a concern.

Raw Clio images show an intense, irregular pattern of column variations which

appear to be due primarily to variations in the bias level from one column to the next. There are also numerous cosmetic defects involving single pixels or small groups of pixels with deviant bias levels or sensitivity levels. The relative brightness over cosmetic defects, columns, and broad regions of the image is not in general perfectly linear with the input sky background flux. We solved this problem using the traditional method of astronomical IR imaging: taking the data in nodded sets.

In nodding observations, a number of images are taken with the telescope in one position (or 'beam'). The telescope is then moved a small amount (typically 5.5 arcsec with Clio) and the same number of images are taken again. The images taken in one position can then be subtracted from the images taken in the other position. Many types of detector artifacts vanish completely from the nodded images, while real celestial objects leave a bright and a dark signature, separated by 5.5 arcsec (or whatever nod amplitude was actually chosen).

A large nod amplitude could be selected such that the images from one beam showed the target star while those from the other beam showed only clean sky. We have chosen a smaller amplitude, so that the target is in the field for images from both beams. This has the disadvantage that a dark, negative image of the star appears in each nod-subtracted frame 5.5 arcsec away from the bright stellar image. It has the more important advantage, however, that images from both beams deliver good sensitivity to faint objects near the star. Thus the data are acquired twice as efficiently as they would be using a strategy in which one nod position showed only blank sky.

Our nodding strategy is illustrated in Figures A.1 through A.3. Figure A.1 shows a raw image of the star  $\epsilon$  Eri from nod beam 1. The star is centered vertically but displaced to the left. Figure A.2 shows a raw image of  $\epsilon$  Eri from nod

beam 2, with the star to the right of center. Figure A.3 shows the image produced by subtracting Figure A.2 from Figure A.1. The intense column striping and many other detector artifacts have vanished, so that away from the star the sky looks cleaner than in Figures A.1 and A.2, even though the display stretch has been increased *by a factor of twenty*.

In Figure A.3, the bright star image is on the left and the negative image is on the right. In a standard processing sequence, we would also make an image by subtracting Figure A.1 from Figure A.2. This image would be as clean as Figure A.3, but would have the bright star image on the right and the negative image on the left. The bright data it contained on any real astronomical objects would be independent of that in Figure A.3. We would then shift Figure A.3 and the other nod-subtracted image to register the bright star images, and add them together. The result would be a central positive image of the star, combining data from both nod positions, flanked by two dark negative images, one on the right and one on the left.

For Clio planet search imaging, we set the nominal exposure to 1200-2000 msec in the  $L'$  band, or 70-200 msec in the  $M$  band, depending on the intensity of the thermal background, which varies according to the ambient temperature and thus the time of year. We typically set the coadds to 10-20 in the  $L'$  band or 50-200 in the  $M$  band, so that a single FITS image output by Clio represents 10-30 sec of actual integration. We take the images in nod sets each consisting of 5-10 frames in each nod position. A typical data set consists of 15-25 such nod cycles, and when processed and stacked yields a cumulative true exposure time of 1-1.5 hours. When possible, a few nod cycles of shorter exposures, which do not saturate the primary star, are also acquired. These serve to give a measurement of the average PSF during the science observations, which in turn helps calibrate



Figure A.1 Raw Clío image of  $\epsilon$  Eri from nod position 1. Note the strong column striping from the detector, and other artifacts.

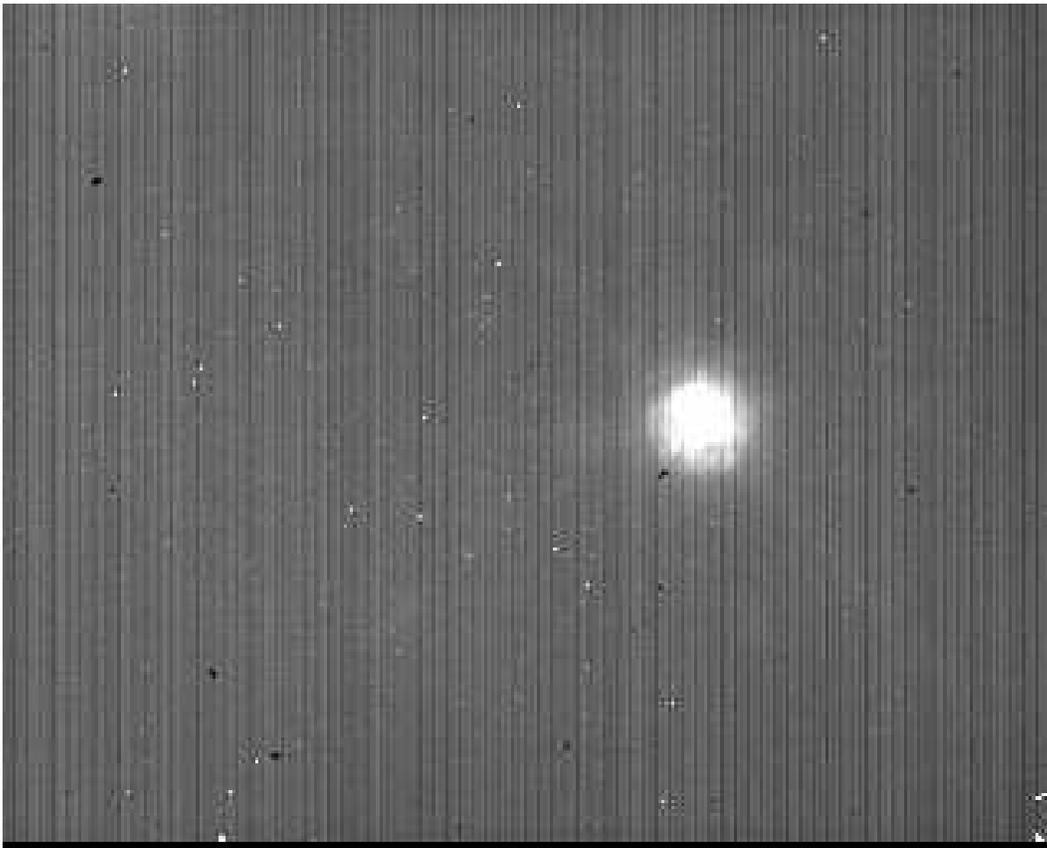


Figure A.2 Raw Clío image of  $\epsilon$  Eri from nod position 2. Note the strong column striping from the detector, and other artifacts.

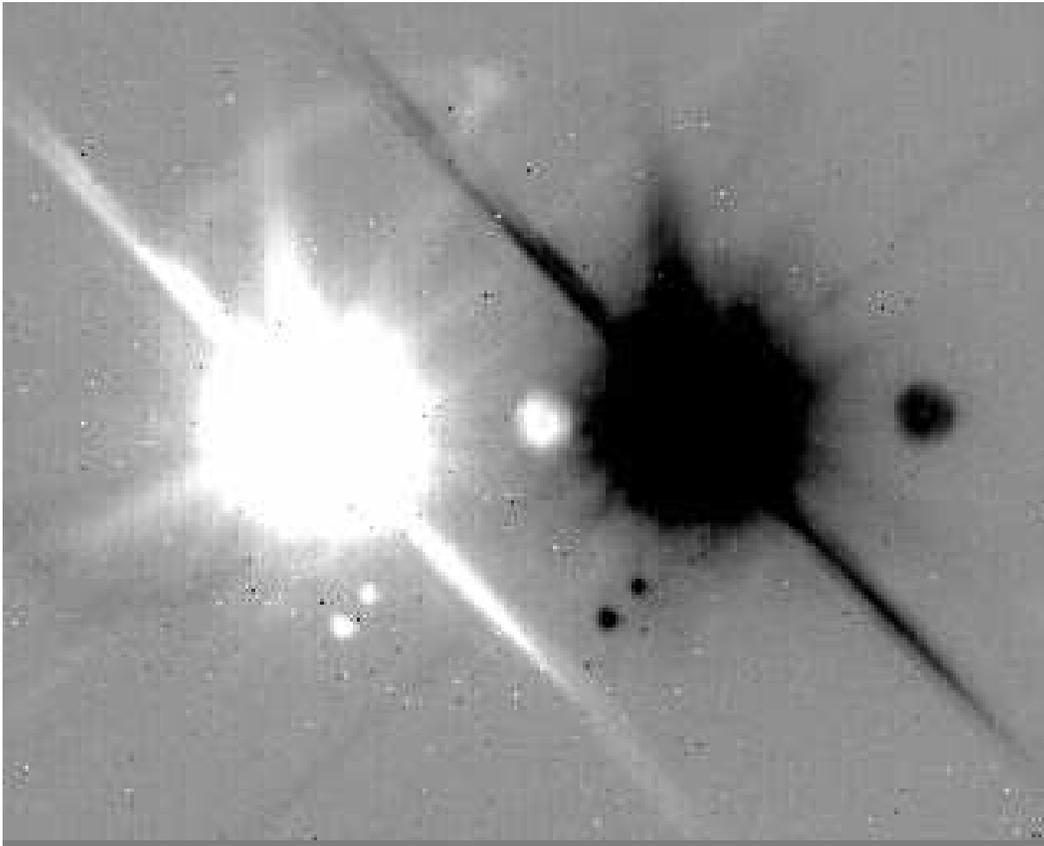


Figure A.3 Nod-subtracted image of  $\epsilon$  Eri made from Figures 3.4 and 3.5. Many detector and sky-glow related artifacts have vanished, revealing faint structure in the outer halo of this very bright star.. This image is stretched *20 times* more than the previous two figures, yet it looks cleaner apart from the bright artifacts from the star itself.

point-source sensitivity estimates and photometry of faint objects.

One further detail concerning the data is necessary for a good understanding of our image processing pipeline: we do not use the MMT instrument rotator, so the Clio instrument remains fixed with respect to the telescope, not with respect to the sky. Thus diffraction spikes, ghosts, and all other artifacts, whether they originate from the telescope or from the instrument, appear at a consistent location on all Clio images from a given data set. Real celestial objects, however, rotate from one Clio image to the next, according to the parallactic rotation of the sky. Thus, in processing, each Clio image must be rotated through an angle that depends on the instantaneous parallactic angle. Only then will real astronomical sources far from the star be properly registered in the final stack.

### A.3 Overview of the Processing Steps for Clio Images

A processing run on Clio images begins with the construction of the nod subtraction pairings. Each individual frame is then processed, and finally the processed frames are stacked to make a master image. Saving the descriptions of the pairing and the final stacks until later, we present here an overview of the processing of an individual frame.

We present images representing seven different processing stages from the moderately faint star GJ 450 ( $L' = 5.40$ ) in Figures A.4 through A.10 and from the very bright star  $\epsilon$  Eri ( $L' = 1.72$ ) in Figures A.11 through A.17.

As we note above, the raw images are literally coadds of a number of individual frames, and the ADU counts reflect this. This can be seen in the image display stretch values quoted for Figures A.4 and A.11. The first step in our processing of a given raw image is to divide it by the number of coadds, thereby obtaining an image which is normalized to a single coadd.

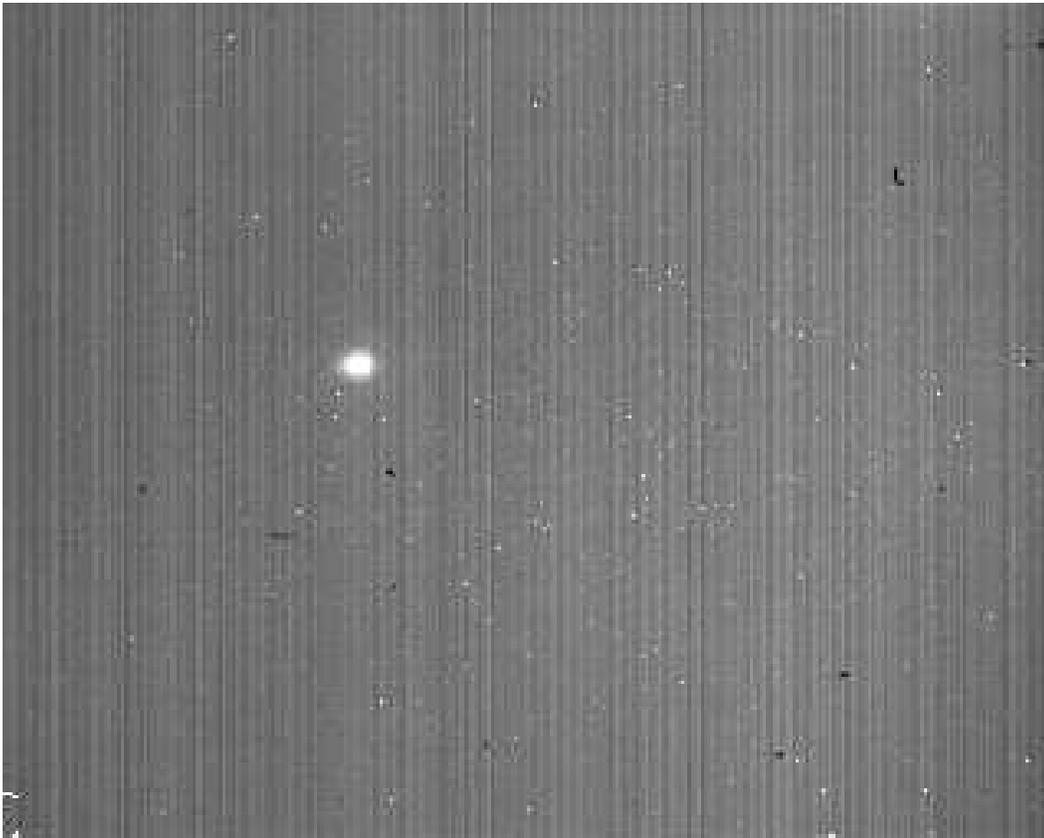


Figure A.4 Raw image of GJ 450. The display stretch for this 10-coadd image is from  $2.5 \times 10^5$  to  $4.5 \times 10^5$ , a range of 200,000 ADU. This stretch is equivalent to 20,000 ADU for a single-coadd frame.

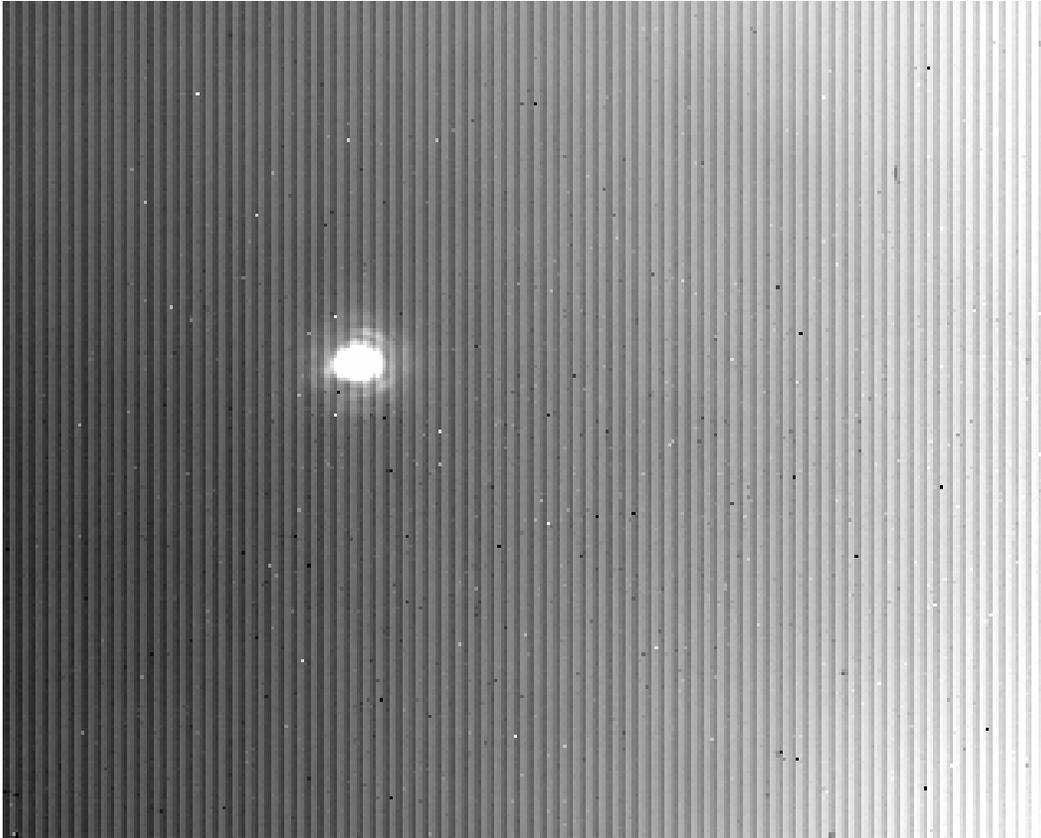


Figure A.5 Dark subtracted, flat fielded image of GJ 450. The display stretch for this image, which has been normalized to one coadd, is from 15,000 to 18,000, a range of 3,000 ADU. The dark subtraction has eliminated much of the column striping.

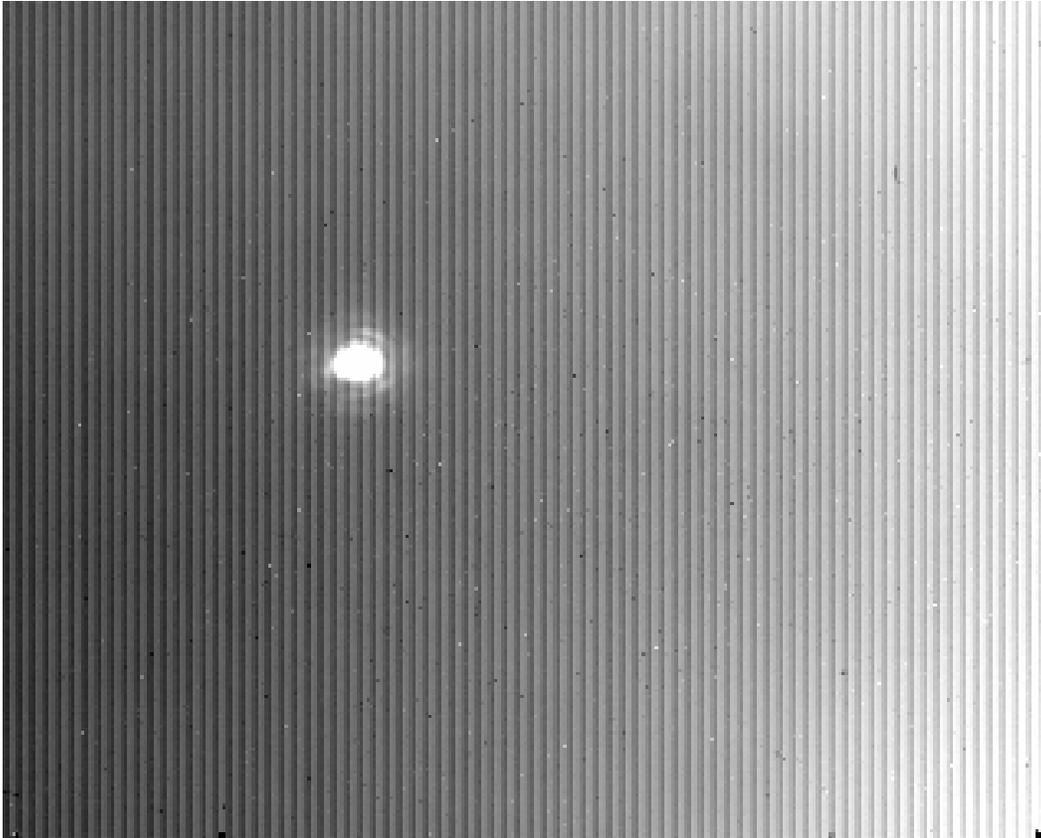


Figure A.6 Image of GJ 450 which has been dark subtracted, flat fielded, and subjected to a first round of bad-pixel fixing. The display stretch for this image is from 15,000 to 18,000, a range of 3,000 ADU. Comparison with Figure A.5 makes it clear many bad pixels have been rejected.

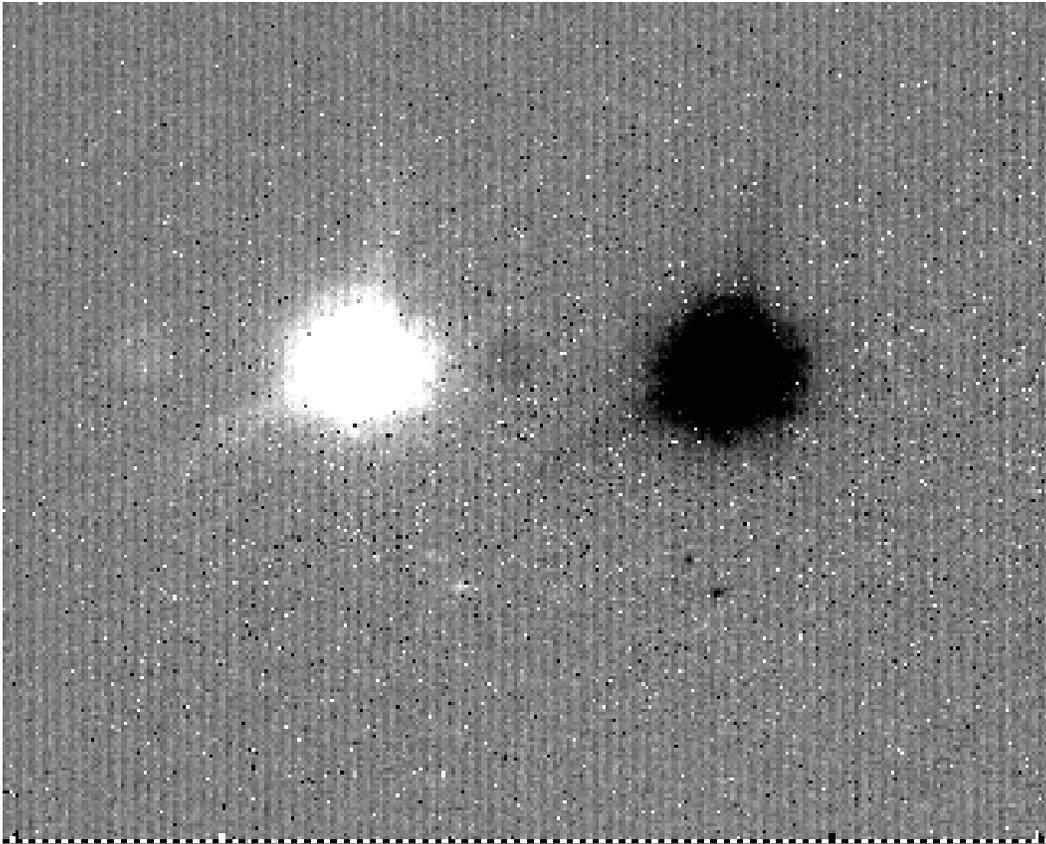


Figure A.7 Image of GJ 450 which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, and nod-subtracted. The display stretch for this image is from -100 to +100, a range of 200 ADU. Nod subtraction removes many artifacts which are left by dark subtraction, flatfielding, and pixel fixing.

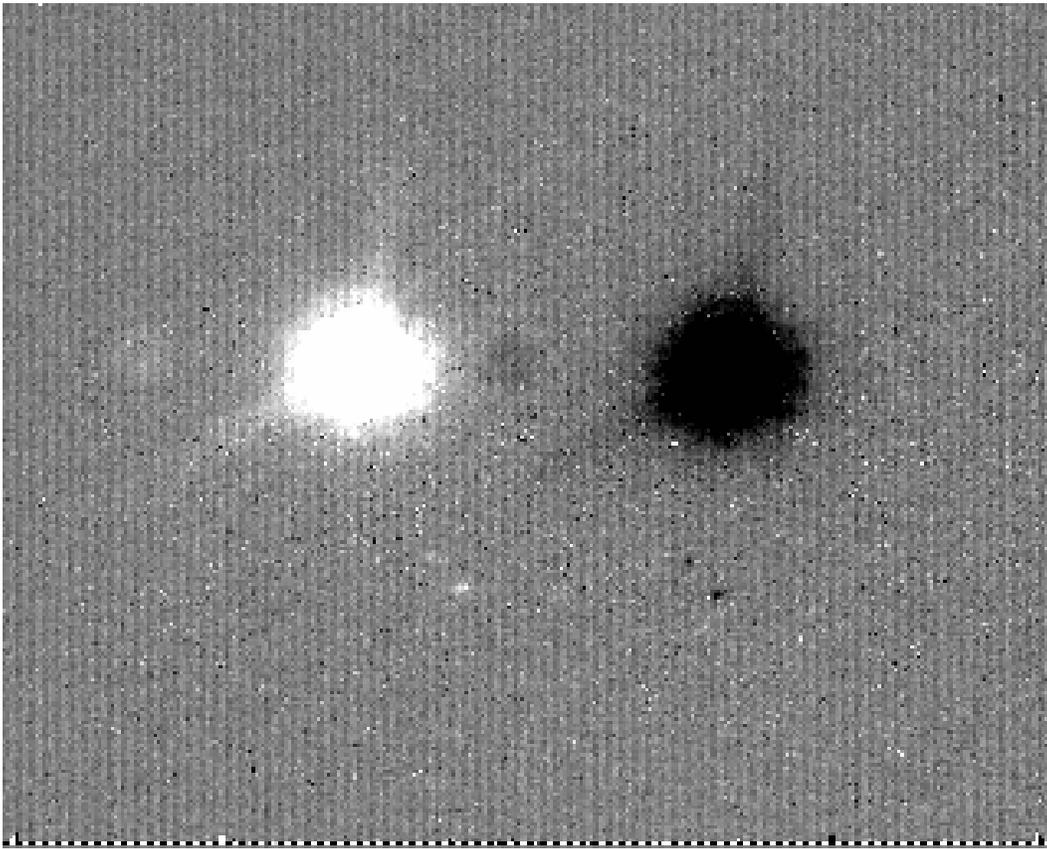


Figure A.8 Image of GJ 450 which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, nod-subtracted, and then subjected to a second round of bad-pixel fixing. The display stretch for this image is from -100 to +100, a range of 200 ADU. Comparison with Figure A.7 makes it clear many bad pixels have been rejected.

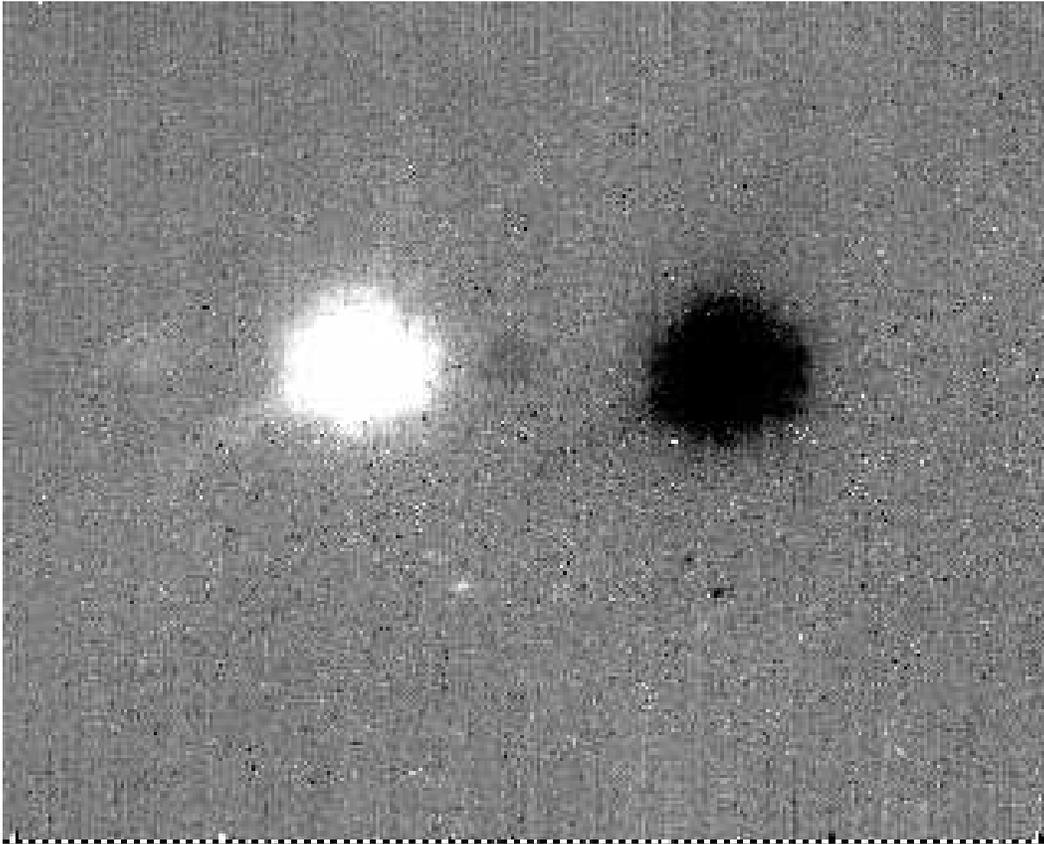


Figure A.9 Image of GJ 450 which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, nod-subtracted, subjected to a second round of bad-pixel fixing, and had column variations corrected using the colfudge routine. The display stretch for this image is from -100 to +100, a range of 200 ADU. For this relatively faint star image with significant intrinsic column variations, colfudge was very helpful.

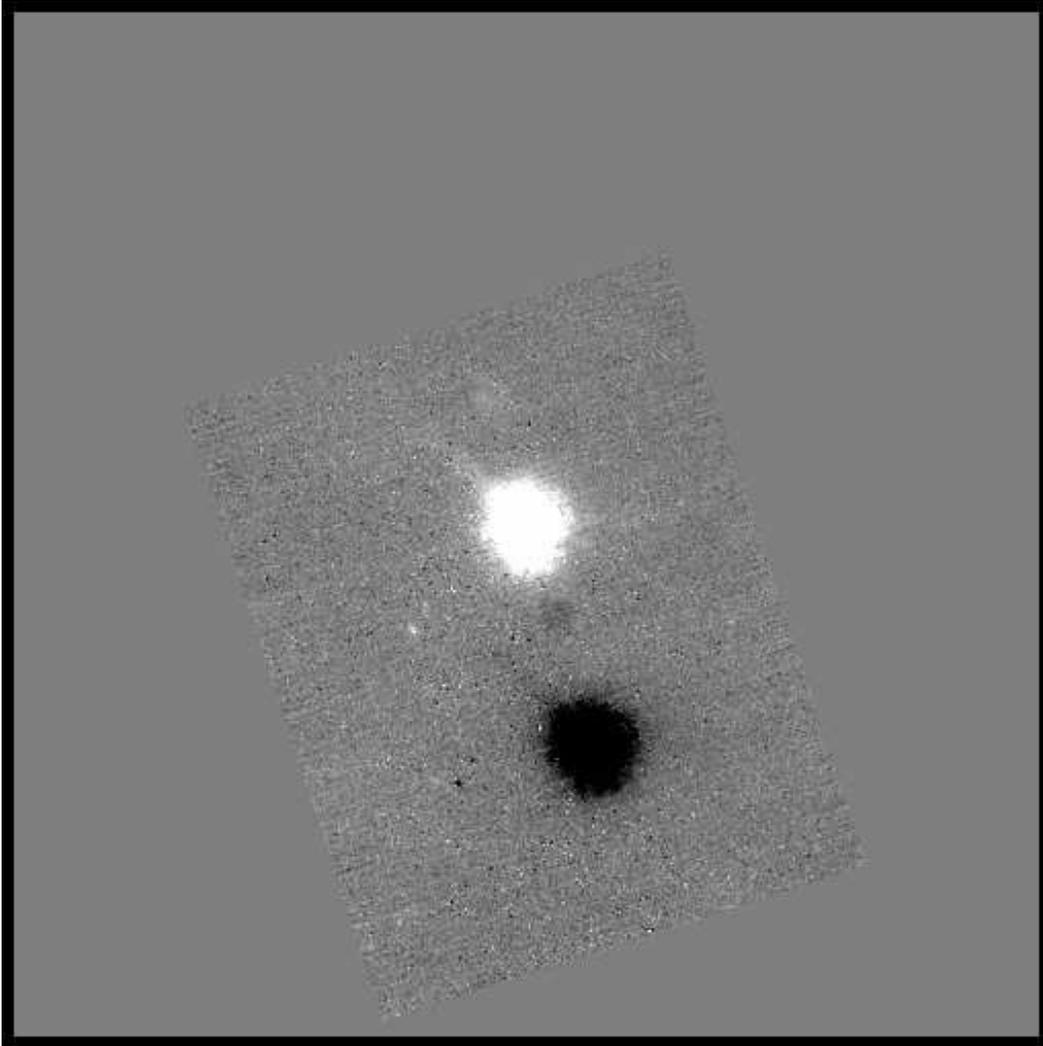


Figure A.10 Image of GJ 450 which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, nod-subtracted, subjected to a second round of bad-pixel fixing, had column variations corrected using the colfudge routine, and been zeropadded, shifted, and rotated in a single bicubic spline operation. The display stretch for this image is from -100 to +100, a range of 200 ADU.

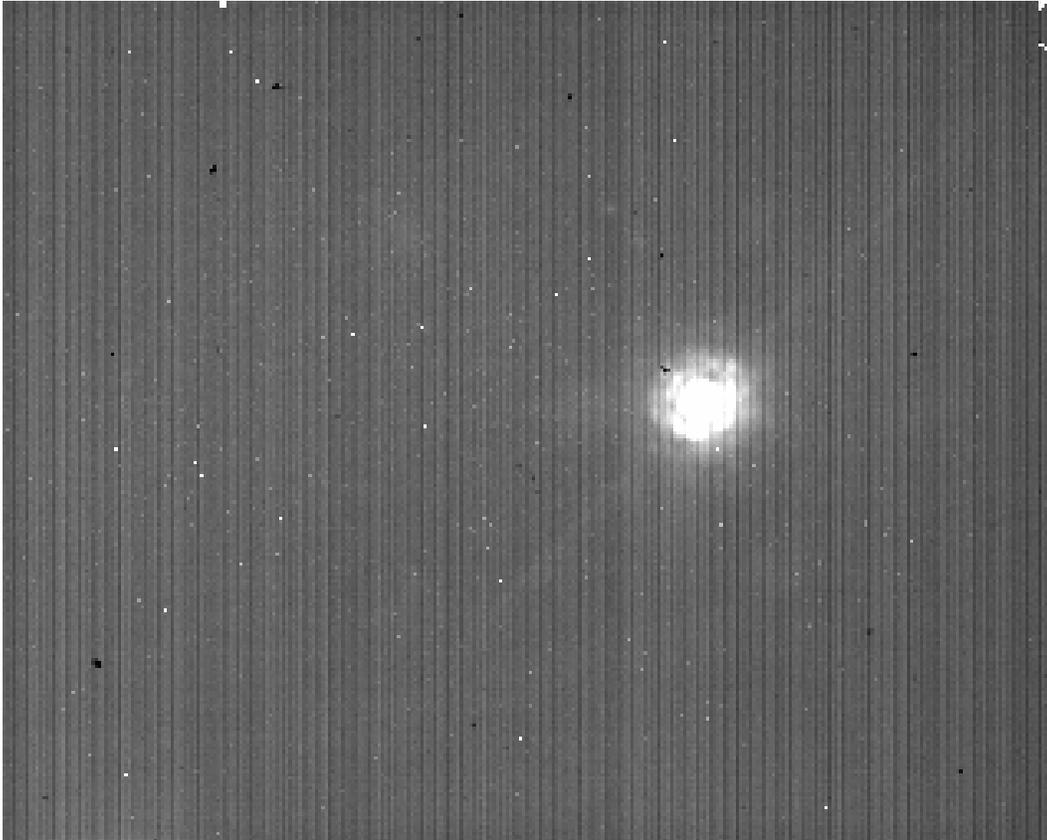


Figure A.11 Raw image of  $\epsilon$  Eri. The display stretch for this 15-coadd image is from  $4 \times 10^5$  to  $7 \times 10^5$ , a range of 300,000 ADU. This stretch is equivalent to 20,000 ADU for a single-coadd frame.

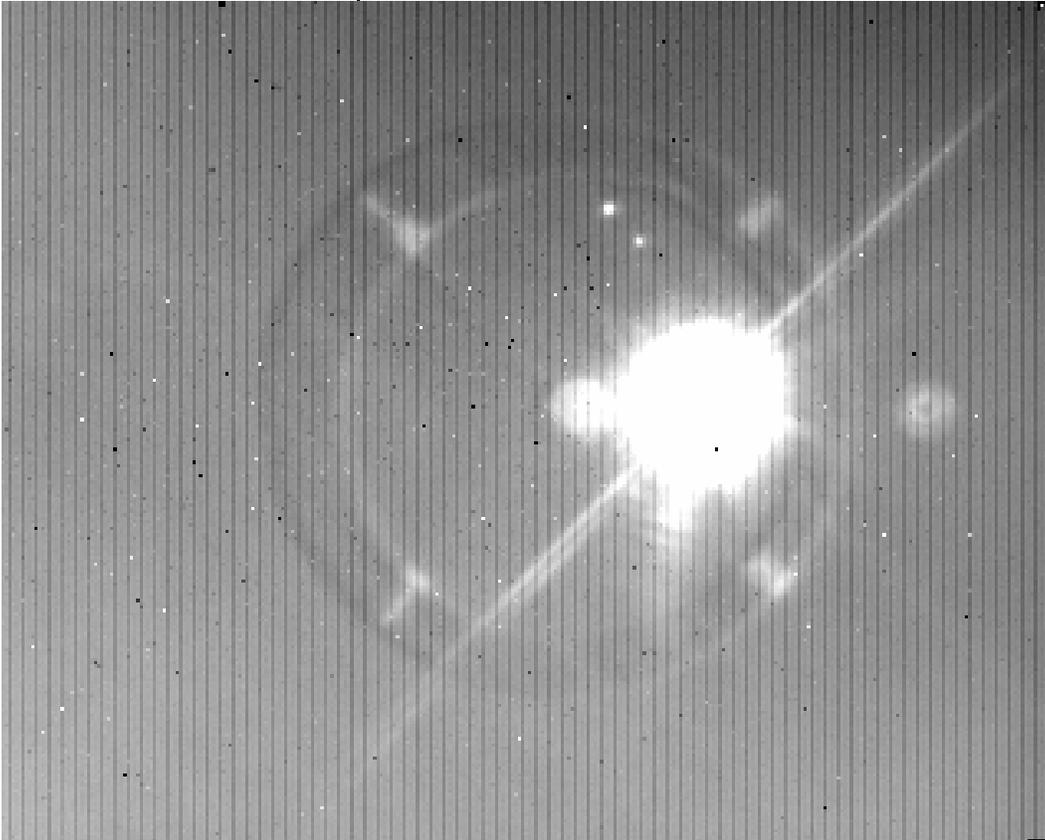


Figure A.12 Dark subtracted, flat fielded image of  $\epsilon$  Eri. The display stretch for this image, which has been normalized to one coadd, is from 17,500 to 19,500, a range of 2,000 ADU. The dark subtraction has eliminated much of the column striping.

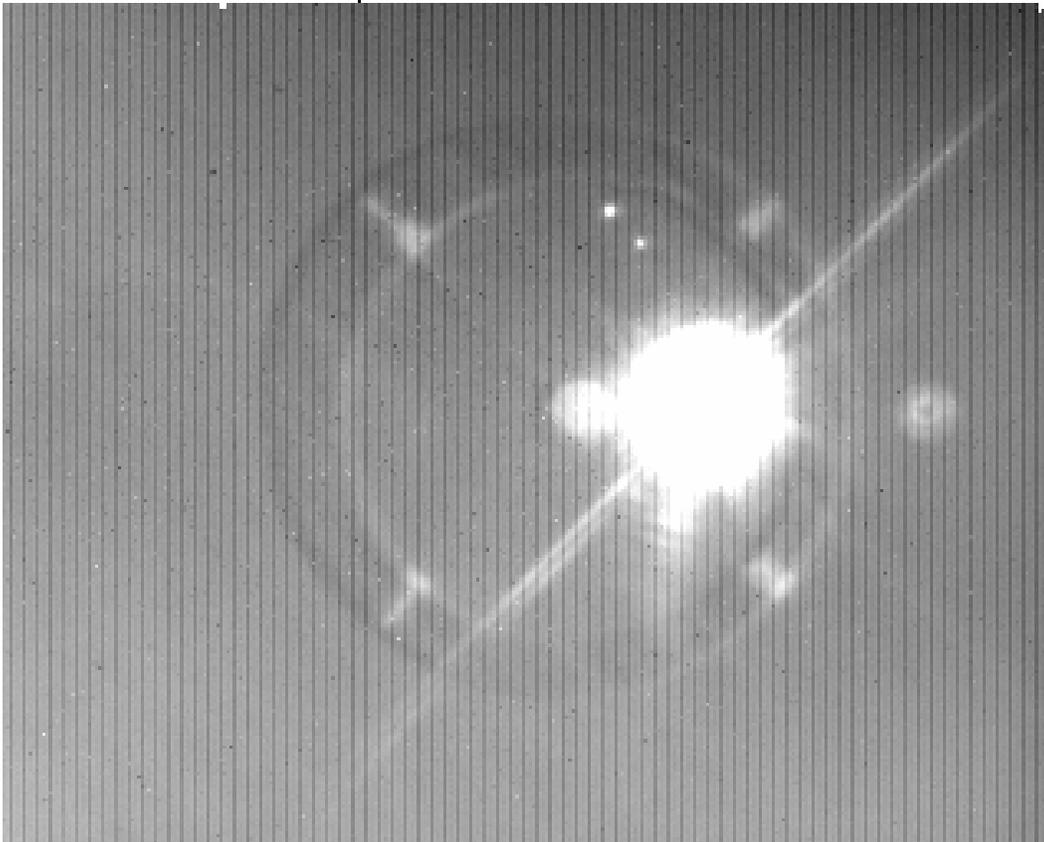


Figure A.13 Image of  $\epsilon$  Eri which has been dark subtracted, flat fielded, and subjected to a first round of bad-pixel fixing. The display stretch for this image is from 17,500 to 19,500, a range of 2,000 ADU. Comparison with Figure A.12 makes it clear many bad pixels have been rejected.

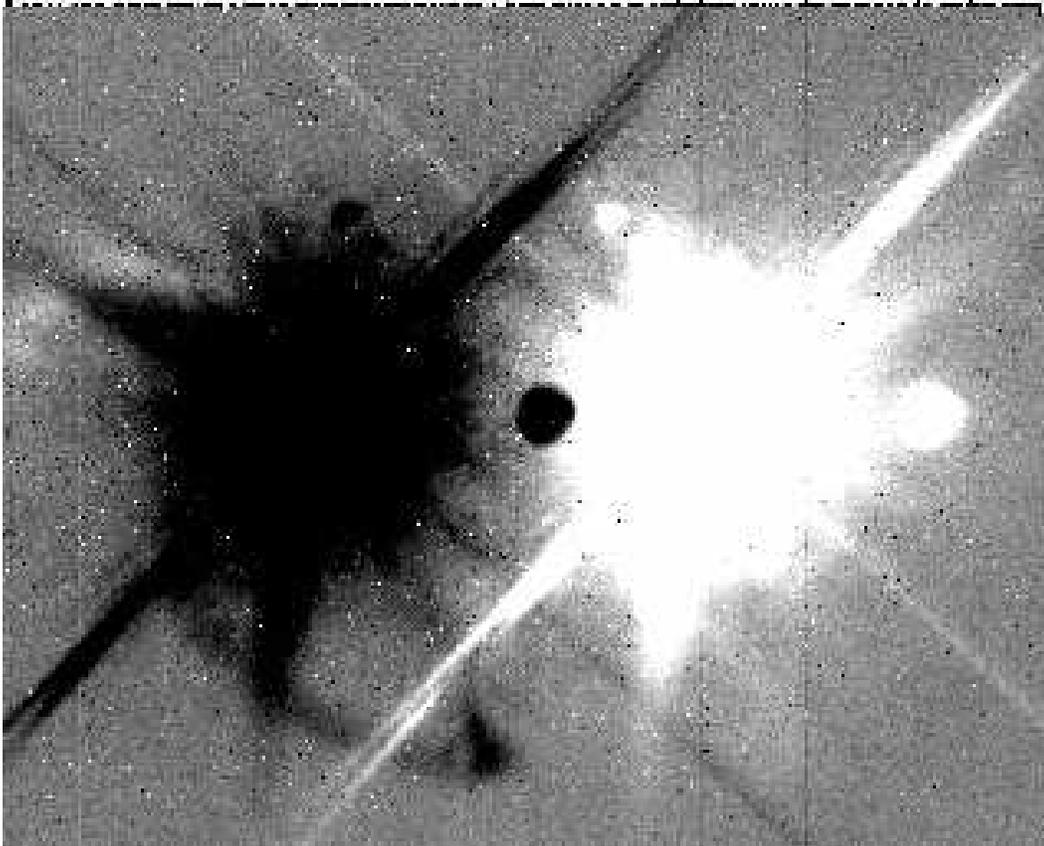


Figure A.14 Image of  $\epsilon$  Eri which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, and nod-subtracted. The display stretch for this image is from -100 to +100, a range of 200 ADU. Nod subtraction removes many artifacts which are left by dark subtraction, flatfielding, and pixel fixing.

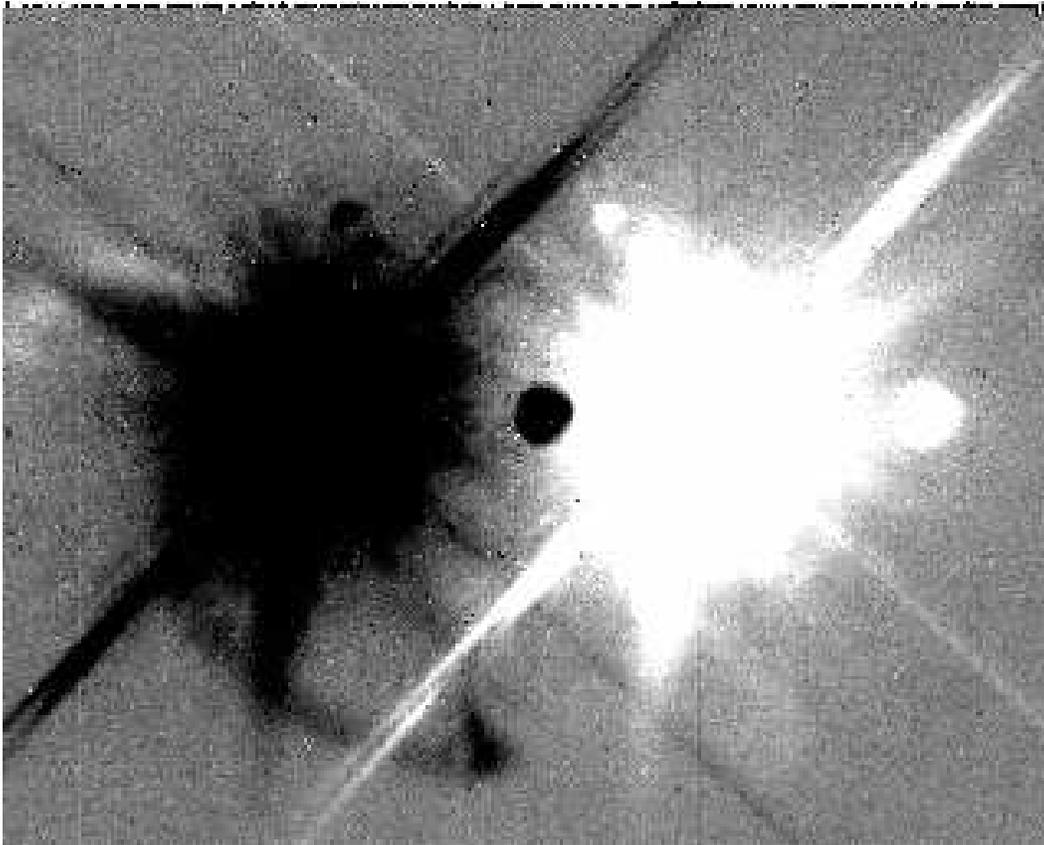


Figure A.15 Image of  $\epsilon$  Eri which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, nod-subtracted, and then subjected to a second round of bad-pixel fixing. The display stretch for this image is from -100 to +100, a range of 200 ADU. Comparison with Figure A.14 makes it clear many bad pixels have been rejected.



Figure A.16 Image of  $\epsilon$  Eri which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, nod-subtracted, subjected to a second round of bad-pixel fixing, and had column variations corrected using the colfudge routine. The display stretch for this image is from -100 to +100, a range of 200 ADU. For this bright star image with low intrinsic column variations, colfudge may have added nearly as many artifacts as it removed.

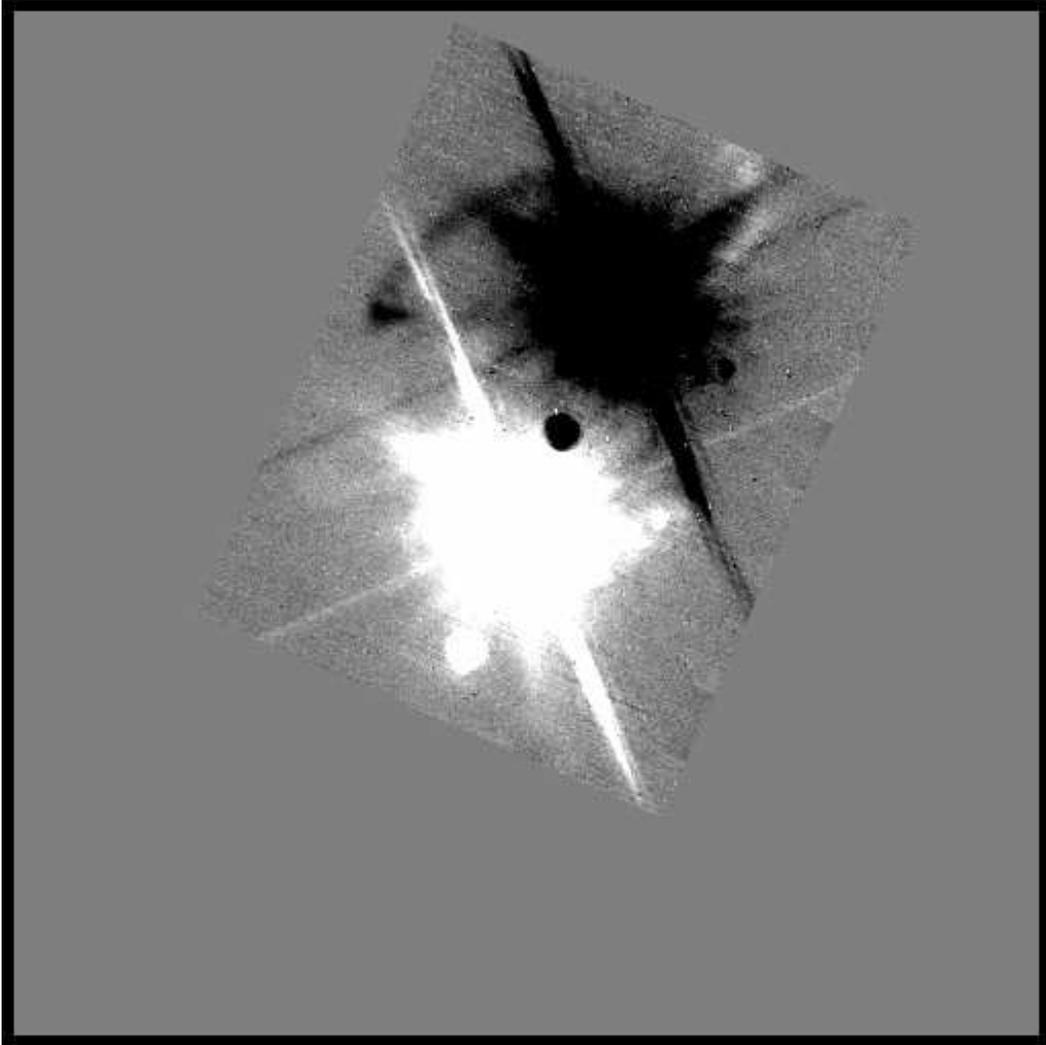


Figure A.17 Image of  $\epsilon$  Eri which has been dark subtracted, flat fielded, subjected to a first round of bad-pixel fixing, nod-subtracted, subjected to a second round of bad-pixel fixing, had column variations corrected using the colfudge routine, and been zeropadded, shifted, and rotated in a single bicubic spline operation. The display stretch for this image is from -100 to +100, a range of 200 ADU.

Next, we subtract a dark frame made from images with the same exposure time as the science image. We divide by a normalized flat frame made from images of a lens cap, telescope cover, or blank sky. The images in Figures A.5 and A.12 represent this stage of the processing. The dark frame removes much of the column striping that is so prominent in the raw images. At this point we apply three different types of bad pixel fixing and one type of bad-column fixing. These will be explained in detail later. Figures A.6 and A.13 show images at this point in the processing.

The next step is nod subtraction. The subtraction image is scaled slightly, based on the mean pixel values over a region of blank sky on both it and the science image. This scaling is the reason the subtraction of a dark image is not made wholly irrelevant by the nod subtraction. Figures A.7 and A.14 show images after nod subtraction. Note how powerfully the nod subtraction removes artifacts, allowing images stretched much more than in the previous Figures to appear at least equally clean well away from the star.

After the nod subtraction, the three types of bad pixel fixing and one type of bad column fixing are applied again. Many more bad pixels are removed that were not revealed until the nod subtraction. Figures A.8 and A.15 show the images at this point in the processing.

Some column variations typically remain after nod subtraction. Their intensity varies greatly from image to image in a given data set, and even more from one data set to another. In the examples given here, the GJ 450 image in Figure A.8 shows significant column variations, while the  $\epsilon$  Eri image in Figure A.15 shows much smaller ones, with only a few columns noticeably affected. The rather complicated 'colfudge' routine is our solution to all types of residual column variation. It powerfully removes column variations, but also introduces some artifacts

and dims faint point sources typically by about 15%. The post-colfudge images shown in Figures A.9 and A.16 show that for the GJ 450 image colfudge has been very helpful, while for the  $\epsilon$  Eri image it may have caused nearly as many artifacts as it removed.

Finally, the images are zero padded, shifted, and rotated in a single bicubic spline operation. The resulting image has celestial north up and the primary star centered. This final outcome is shown in Figures A.10 and A.17.

#### A.4 Constructing the Nod-subtraction Pairings

The assignment of nod subtraction pairings is an apparently simple task. Some complexities, however, must be addressed if an algorithm is desired that will robustly produce a reasonable result in a variety of different circumstances.

The point of nod subtraction is to remove artifacts by subtracting two images that are identical except for the positions of any real celestial objects. The problem of constructing nod pairings is, for a given image in a given beam of a nod sequence, to pick the best image from the opposite beam to subtract from it. To simplify the discussion, we will refer to the given image as the ‘science’ image and to the image that is subtracted from it as the ‘subtraction’ image. The reader should note that every image in a Clio nod sequence should serve as the science image once, and typically will also serve once as a subtraction image. The designations ‘science’ and ‘subtraction’ therefore refer to the roles the images play in a given processing operation and not to their inherent properties.

Many artifacts may manifest themselves in a way that varies nonlinearly with the level of the background flux. For this reason it makes sense to choose a subtraction image that is as near to the science image in the observing sequence as possible, so that if the thermal background is slowly varying (as is almost

Table A.1. Nod Subtraction Pairing Option 1

Science Beam	Subtraction Beam	Sci. Img.		Sub. Img.
Beam 1	Beam 2	1	—	6
...	...	2	—	6
...	...	3	—	6
...	...	4	—	6
...	...	5	—	6
Beam 2	Beam 1	6	—	5
...	...	7	—	5
...	...	8	—	5
...	...	9	—	5
...	...	10	—	5

certainly the case), the backgrounds on the science and subtraction images will match as well as possible. Thus, in the simple case of a single nod cycle with 5 images in beam 1 followed by 5 images in beam 2, we might choose the pairing illustrated in Table A.1.

The Table A.1 pairing does indeed pair each science image with the closest possible subtraction image. However, it has a serious problem with noise independence. A positive version of the specific background noise pattern on each image will contribute exactly once to the final stack. However, the noise patterns on images 5 and 6 will contribute negatively *5 times each*. The final images will be

Table A.2. Nod Subtraction Pairing Option 2

Science Beam	Subtraction Beam	Sci. Img.		Sub. Img.
Beam 1	Beam 2	1	—	6
...	...	2	—	7
...	...	3	—	8
...	...	4	—	9
...	...	5	—	10
Beam 2	Beam 1	6	—	1
...	...	7	—	2
...	...	8	—	3
...	...	9	—	4
...	...	10	—	5

noiser than necessary because the negative noise, coming from only two images, will not average down as well as the positive noise, coming from ten different images. A wider range of subtraction images will produce more independent noise in the individual frames, and lower total noise in the final stack. This consideration suggests the pairing illustrated in Table A.2.

The Table A.2 pairing seems like a good compromise between having science and subtraction images as close together as possible and keeping the noise independent. However, in the general case some images from one beam or the other may be unusable due to bad seeing or other problems. The strategy in Table A.2,

unlike that in Table A.1, does not suggest a natural way of dealing with sequence gaps.

The pairing algorithm we have actually adopted is a compromise between the strategies illustrated by Tables A.1 and A.2. It is not the product of any sophisticated optimization, but it does robustly produce reasonable pairings with good noise-independence for any input data set, even if there are many gaps and/or irregularities in the nod cycles. It has no problem with cases in which there are more images in one beam than in another, though it does require that the nod beam with more images have no more than twice as many as the nod beam with fewer.

Our algorithm pairs a given science image with the nearest image in the opposite beam that has not already been paired with another science image, and that lies fewer than a user-specified maximum number of images away in the observing sequence. We refer to this maximum number of images as the 'doubling distance,' because if there is no unused subtraction image within the doubling distance, our algorithm chooses the nearest image that has been used only once and therefore 'double-pairs' that subtraction image. If there is no image within the doubling distance that has not already been used twice, the program reports failure to construct a valid pairing set and instructs the user to select a larger doubling distance. The code will not pair the same subtraction image with more than two science images: it will not 'triple-pair' a subtraction image.

If a pairing set is successfully constructed with no double-pairings, our program will immediately begin processing the images. If there are some double pairings, the program reports this to the user, who can then decide whether to go ahead with processing or to construct another pairing set using a larger doubling distance, in the hopes of improving noise-independence by reducing the number

Table A.3. Actual Nod Subtraction Algorithm, Example # 1

Science Beam	Subtraction Beam	Sci. Img.		Sub. Img.
Beam 1	Beam 2	1	—	6
...	...	2	—	7
...	...	3	—	8
...	...	4	—	9
...	...	5	—	10
Beam 2	Beam 1	6	—	5
...	...	7	—	4
...	...	8	—	3
...	...	9	—	5
...	...	10	—	4

of double-pairings.

In the example we have been discussing, the pairing will fail unless the doubling distance is at least 6. Table A.3 gives the pairing for a doubling distance of 6.

This doubling distance gives 2 double pairings. Images 1 and 2 are never used as subtraction images, because by the time the algorithm gets to pairing images 9 and 10 they are too far away. Setting the doubling distance to 10 can fix this. The resulting pairings are given in Table A.4.

In this example there are no double-pairings and each image is used as a sub-

Table A.4. Actual Nod Subtraction Algorithm, Example # 2

Science Beam	Subtraction Beam	Sci. Img.		Sub. Img.
Beam 1	Beam 2	1	—	6
...	...	2	—	7
...	...	3	—	8
...	...	4	—	9
...	...	5	—	10
Beam 2	Beam 1	6	—	5
...	...	7	—	4
...	...	8	—	3
...	...	9	—	2
...	...	10	—	1

traction image exactly once.

We note, however, that for typical science data sets involving many nod cycles rather than just one, the problem illustrated by Table A.3 exists and cannot be solved by any reasonable value of the doubling distance. The first few images in the first beam are typically never used for nod-subtraction, and this results in a few double-pairings at the end of the sequence. A more sophisticated algorithm could surely fix this. However, having a few percent of the images in a long science data set double-paired has a negligible effect on the final noise. Our algorithm is simple, and robustly produces reasonable, if not strictly optimal, pairing results.

#### A.5 Dark Subtraction and Flatfielding

We now begin a detailed discussion of the processing applied to each individual image. The first step following the initial division by the number of coadds is dark subtraction and flatfielding. These are fairly self-explanatory, but we briefly comment on the nature of the darks and flat frames that we use, and why we consider these processing steps desirable. **We note that for Clio data taken before September 2006, a right-to-left mirror flip of each image was carried out between coadd-normalization and dark subtraction.** Andy Breuninger altered the Clio software so that the raw images were correct and no mirror-flip in processing was required for images taken during or after September 2006.

Our dark frames are taken with the Clio detector covered by an cold opaque disk in one of the internal filter wheels. We use the same exposure time as for the science observations, and usually the same number of coadds. The dark frames are taken as close in time to the science observations as possible, because the bias level of the Clio detector drifts slowly with time. Often we take the darks for a

given science target while the telescope is slewing to the next target.

Typically only 3-10 dark frames are taken to make a single master dark. The master is produced by dividing each individual frame by the number of coadds, and then stacking them using a creeping mean combine (see below). This master, coadd-normalized dark is then subtracted from the science image (and from the subtraction image with which the science image has been paired).

Our flat frames are made from exposures of the lens cap, telescope cover, or blank sky. We typically take flat frames once per observing run, in contrast to the dark frames which we take several times each night. The main purpose of our flat frames is to correct high spatial frequency variations in the gain and sensitivity of Clio images. Low frequency variations (caused, ie, by vignetting), do not appear to be large, and in any case are mitigated by the shifting and coadding of images with different nod positions and rotation angles. For this reason we consider lens cap or telescope cover flats acceptable when sky flats are not available. This situation sometimes arises because science observations are sometimes prioritized above sky flats whenever the dome is open during a given run. With  $L'$  and  $M$  band observing, in contrast to optical work, science observations are possible whenever sky flats are.

Our flats are coadd-normalized and dark subtracted; then known regions of bad pixels are interpolated away. Finally, the flats are normalized to have an overall image mean of 1.0, and combined using a creeping mean stack. Both the science and the subtraction images are divided by the normalized master flat before the nod-subtraction is performed.

Using flat frames makes sense with nod-subtracted data, since the flat will have an effect wherever there was any flux from an astronomical source. Subtracting dark frames from images that are about to be nod-subtracted is com-

pletely useless for simple nod subtraction, since the result is:

$$(I_1 - D) - (I_2 - D) = I_1 - I_2 - D + D = I_1 - I_2.$$

Where  $I_1$  is the science image,  $I_2$  the subtraction image, and  $D$  the master dark frame. The dark frame cancels; one might as well have simply subtracted the unaltered images.

This simplistic picture does not exactly apply. The flat frame would not keep it from applying, since the flat simply divides every term in the equation above. We perform bad pixel fixing before the nod subtraction, and this is rendered more effective by the fact that the dark has reduced the amplitude of the column striping, causing a greater number of bad pixels to stand out and be removed. Most importantly, however, our nod subtraction is not simple subtraction, but involves a scaling of the subtraction image so that its mean flux exactly matches the science image over a user-selected region of clean sky. The scaling factor is always very near unity, but its small deviations from unity map small changes in the sky brightness from image to image, and insure good subtracted images with zero-mean sky backgrounds despite the changing background. **If we did not subtract a dark, the scaling would be performed on the sum of bias, dark counts, and thermal background. Since we do use dark subtraction, the background level alone is scaled, and variations in the background are properly tracked and removed.**

## A.6 Bad Pixel Fixing

We apply bad pixel fixing both before and after the nod subtraction. Applying it before probably has little effect on the final images, since a much larger number of bad pixels appear against the very clean sky background after nod subtraction. However, fixing bad pixels pre-subtraction does at least mean that the nod-

subtraction scaling is performed over a region with fewer bad pixels. The exact same types and sequence of bad pixel fixing are performed before and after nod subtraction, so our discussion here applies to both.

The first type of pixel fixing is designed to remove isolated bad pixels in a given image. These are expected to be detector pixels that are normally well behaved, but for some reason are deviant on a particular image. Up to one thousand such pixels may appear in a given nod-subtracted Clio image. The author does not know their physical origin. Our algorithm for removing them considers every pixel of the image in turn to see if it is bad. The mean in a 3x3 box centered on, but excluding, the pixel itself is computed. The RMS variation in a 5x5 box centered on, but excluding, the pixel is also found. If the pixel differs from the 3x3 mean by more than a set threshold times the 5x5 RMS, the pixel is considered deviant and is replaced with the 3x3 mean. We typically set the threshold to 5.0: that is, pixels that are deviant by  $5\sigma$  or more are considered bad and replaced with the 3x3 mean.

For undersampled data sets it is clear that this algorithm might identify the central pixels of real point source images as bad pixels. The result would be a large reduction in the fluxes of point sources. Since our data are well sampled, with typical point-source FWHM values ranging from 2.9 to 3.7 pixels, we would not expect any point source centers to be rejected. The PSF image extends to surrounding pixels sufficiently to insure that the 5x5 RMS calculated by our algorithm is far too high to call for the rejection of the central pixel. A careful test in which both bright and faint simulated objects were inserted in a 1000-image stack using real stellar PSFs from Clio images has shown that, indeed, our bad pixel fixing has no appreciable effect on the recovered flux of either bright or faint point sources.

While this first algorithm can be relied upon not to mistake real astronomical sources for bad pixels, it can be prevented from recognizing bad pixels if they are involved in the haloes of bright stars. The radial gradient from the bright stellar halo keeps the 5x5 pixel RMS too high for bad pixels in this region to be rejected. To alleviate this problem, we apply another round of bad pixel fixing using a modification of the same algorithm. In this version, we make a copy of the input image and unsharp mask it. Unsharp masking is a method of removing low spatial frequencies from an image. It works by subtracting from the image a blurred version of itself. The unsharp masking used in our bad pixel fixing algorithm creates the blurred image by convolving the original with a unit box, typically of size 15x15 pixels.

The unsharp masking substantially reduces low spatial frequency features such as stellar haloes and diffuse ghosts, allowing bad pixels to stand out more clearly. The bad pixels are then identified on the unsharp-masked image **but then fixed on the original image using a 3x3 mean of original-image pixels**. This second round of unsharp masking removes some additional bad pixels that escaped the first round through involvement in stellar haloes or other diffuse image features. We apply *both* methods of pixel fixing to every image, rather than just the second version, because in tests we have found that just as the second version finds some pixels that are not removed by the first version, so also the first version removes some bad pixels that would be neglected by the second.

We perform a third type of bad pixel fixing to remove known regions of pixels that are bad on the detector, and thus consistently give deviant results. These pixels are cataloged in a file, and removed by interpolating across the bad regions. We always interpolate along columns, because of the column-to-column variations common in Clio images.

Our fourth and final type of bad pixel fixing is intended to deal with situations in which an entire column is deviant on a particular image. Such deviant columns appear on a significant minority of nod-subtracted images, though they are rarely if ever sufficiently deviant to be identified and fixed in the images prior to nod subtraction. To identify bad columns, the mean of every column is first calculated over range of rows specified by the user, supposed to be free from bright features of real astronomical sources on all the images. Then each column is considered in turn to see if it is bad. The method is a 1-D analog of our 2-D algorithm for identifying isolated bad pixels. The RMS of the column means is calculated over a region extending 5 columns on either side of the column under consideration, with the column itself not included in the calculation. The mean of the two adjacent columns is calculated. If the column under consideration deviates from the mean of its two neighbors by more than a set threshold times the RMS from the larger 10-column set, it is replaced by interpolation between its immediate neighbors.

We remind the reader that the four types of bad pixel fixing described here are applied to the science and subtraction images before the nod subtraction, and then again to the science image after the nod subtraction is complete. In general far more bad pixels are corrected in the post-subtraction image than in the two pre-subtraction images, because the great reduction in the column pattern and other artifacts after nod subtraction causes more bad pixels to stand out so that they can be identified and removed.

## A.7 Nod Subtraction

Nod subtraction is the single most important step in eliminating artifacts on Clio images so that faint sources can be detected. We have described in Section A.4

how each image is processed as the science image exactly once in the processing sequence, but most images also serve once as the subtraction image. We have described in Sections A.5 and A.6 how the science and subtraction images are processed in parallel. Lest this confuse the reader, we note here that in all cases our image processing leaves the original raw image unaltered, and merely makes additional processed versions of it. Thus all the processing we describe being performed on the subtraction image is performed on a copy made only for subtraction use, and does not affect the ability of the original of the subtraction image to be used as an unaltered science image itself when the time comes to process it.

To perform our nod subtraction, we compute the mean on the science and subtraction images within a region specified by the user, which is supposed to contain only clean sky on both images, with no real astronomical sources, rays, ghosts, or other artifacts. The subtraction image is then scaled by the ratio of means, so that the mean value on the scaled version is identical within the specified region to the mean value on the science image. We note that the means are calculated using a creeping mean with rejection fraction 50%, so that if some stellar artifacts happen to be included in the region they will bias the scaling as little as possible.

The scaled version of the subtraction image is subtracted from the science image, leaving a clean image with mean zero in regions of sky away from the bright positive and dark negative images of the primary star.

## A.8 The Colfudge Algorithm

After nod subtraction and the bad pixel fixing that follows it, some column-to-column variations remain in the image. We fix these using a rather complicated routine called colfudge.

### A.8.1 History

As mentioned in the abstract above, colfudge is the one aspect of our processing developed based on the June 2005 Clio data that might not have been chosen based on the slightly different characteristics of later Clio data. In June 2005, Clio's first light on the MMT, we read out the detector using a single amplifier. Nod subtracted images showed column-to-column variations without any clear pattern. In fact, columns could be brighter than average near the top of the image and darker than average near the bottom, or vice versa. Colfudge was developed to remove or reduce the significant image noise resulting from these unpredictable column variations. It caused some artifacts and a typically 15-18% reduction in the brightness of faint point sources, but these were a small price to pay for the removal of the column noise.

Although Clio runs were planned for September and December 2005, they were lost due to problems at the MMT and to weather, so no more Clio data were acquired until April 2006. In the mean time, Andy Breuninger developed the integrate-while-reading (IWR) mode for Clio, reading out the detector through four different amplifiers to obtain very high observing efficiency. The column variations still existed, and were still effectively removed by colfudge, but their character had changed. The author simply checked that colfudge still removed them, and did not consider the possibility that a simpler algorithm that would not dim point sources or introduce artifacts might now be possible.

In 2007 Matt Kenworthy began work on a method for processing Clio images taken with a phase plate coronagraph he developed. The author's pipeline was not suitable for images of this type. The removal of column variations was important for phase plate data too, and Dr. Kenworthy developed an alternative method. It would not have worked with Clio first light data, but it was simple, el-

egant, and very effective for data from all subsequent Clio runs. Unlike colfudge, it introduced no significant artifacts and did not dim faint point sources. Unfortunately this discovery was not made until all of the data discussed in Chapters 2-4 of this dissertation had already been processed using a pipeline that included colfudge, and thoroughly analyzed. As this was an investment of several months of full-time effort, reprocessing and reanalysis were not feasible.

However, the author did implement Dr. Kenworthy's method and incorporated it into his own pipeline. Side-by-side tests on selected real data sets showed no significant difference in the final image quality between the two algorithms. It appears, therefore, that the artifacts introduced by colfudge are of very little significance, and its only real disadvantage is the dimming of point sources. This has already been considered in the final sensitivity comparisons in Chapters 2-4 above.

### A.8.2 Colfudge

We will now describe the colfudge algorithm because it was used to process so much data, and because similar algorithms may be of use for data from other IR cameras, if they resemble Clio first light data rather than later Clio data. After this we will describe Dr. Kenworthy's algorithm, which has been incorporated into the author's pipeline and will likely be used for all future processing runs.

The parameters of colfudge are a block size and a limiting pixel value. For Clio planet search images taken in the  $L'$  band, we typically set the block size to 36 pixels and the limit value to 35.0 ADU. The limit value is the maximum deviation from zero that a pixel may have and still be considered by colfudge; it is intended to insure that colfudge determines its column noise corrections mostly based on clean, zero-mean sky regions of the image, and does not create strong artifacts around bright sources.

Colfudge begins by making a copy of the image and unsharp masking it. The blurred image subtracted in the unsharp masking is created by convolving the original image with a unit box of size 21x21 pixels. The objective of the unsharp masking is to reduce low spatial frequency variations on the image, so that the high frequency column-to-column variations can more easily be seen.

Colfudge divides the unsharp masked image up into square blocks, usually 36x36 pixels as mentioned above. It carefully allows for the possible existence of partial blocks at the right hand edge and top edge of the image, in case the block size does not divide evenly into the image size. It calculates the mean of all pixels within each block and partial block. It also calculates the mean of all the pixels in each column in each block. At this point for each block there is an overall block mean, and a mean for the portion of each column that is contained within that block. The means are calculated only over pixels that deviate from zero by less than the given limit, typically 35 ADU. They are calculated using a creeping mean with 50% rejection, to reject point sources and sharp rays and ghosts, thus minimizing the introduction of artifacts. If 75% or more of the pixels in a given column in a given box lie outside the limit, the colfudge algorithm considers the remaining pixels (ie, 9 or fewer out of 36) too few for a valid calculation. It simply assigns the overall block average as the column average for that column. This is equivalent to applying no correction to that column.

Having calculated the block averages, and the column averages within each block, colfudge proceeds to make an image designed to be added to the original image to remove the column variations. This image is constructed by linear interpolation between blocks. Let individual pixels in the image be indexed by  $i, j$ , where  $i$  corresponds to the  $x$  coordinate and  $j$  to  $y$ . Let blocks be indexed similarly by  $k, l$ . Then we have the block means  $B_{k,l}$  and the column means within

each block  $C_{i,l}$ . Note that we do not give  $k$  (the block  $x$  coordinate) for the column means  $C_{i,l}$ . This is because the pixel  $x$  coordinate  $i$  uniquely determines the block  $x$  coordinate  $k$ .

Let the pixel coordinates of the center of block  $k, l$  be  $X_k, Y_l$ . Then we can assign an interpolated block mean  $MB_{i,j}$  and an interpolated column mean  $MC_{i,j}$  to every pixel in the image.

Let pixel  $i, j$  be above and to the right of the center of block  $k, l$ . Then simple bilinear interpolation gives the interpolated block mean  $MB_{i,j}$  as:

$$MB_{i,j} = B_{k,l} + \left( (B_{k+1,l+1} - B_{k,l+1}) \times \left( \frac{i - X_k}{X_{k+1} - X_k} \right) - (B_{k+1,l} - B_{k,l}) \times \left( \frac{i - X_k}{X_{k+1} - X_k} \right) \right) \times \left( \frac{j - Y_l}{Y_{l+1} - Y_l} \right).$$

The interpolated column mean  $MC_{i,j}$  for the same pixel is:

$$MC_{i,j} = C_{i,l} + (C_{i,l+1} - C_{i,l}) \times \left( \frac{j - Y_l}{Y_{l+1} - Y_l} \right).$$

Then we can estimate the correction that must be added to pixel  $i, j$  to remove the local column variation. The correction is:

$$MB_{i,j} - MC_{i,j}.$$

Colfudge places this value into the correction image, and, when this image is complete, adds it to the original nod-subtracted science image to remove column variations.

We note that since nod subtraction is supposed to yield a zero-mean sky background, and we further unsharp mask the image from which colfudge derives its

corrections, all of the block means  $B_{k,l}$  should be very close to zero. It may not even be necessary to take them into account as carefully as we have done. Still, we would expect this to reduce artifacts in cases where there are residual sky variations, ie from the outer halo of a bright star, even after the unsharp masking.

Colfudge interpolates carefully up to the center of the partial blocks on the upper and right-hand edges of the image. For all pixels closer to the image edge than the center of the outermost block, interpolation is done only in the direction parallel to the image edge. For corner pixels nearer the corner than the center of the corner block, no interpolation is performed and the correction is simply set to the block and column average.

### A.8.3 Matt Kenworthy's Algorithm

When Dr. Kenworthy examined the nod-subtracted images from Clio runs after June 2005, he saw images substantially different from those the author had seen during the June 2005 run. The earlier images had shown unpredictably varying columns. The sign of a column's difference from the mean could change along a column. The later images showed a very clear pattern, repeating every four columns. The sign of a given column's deviation from the mean was consistent along that column: a column that was too bright at any point was too bright by a consistent amount from top to bottom.

The origin of the repeating 4-column pattern was obvious: it was the four amplifiers used to read out the Clio detector, with every 4th column assigned to a given amplifier. These amplifiers all had slightly different gain values. For some reason, probably very small nonlinearities in the electronics, flat fielding and nod subtraction did not entirely eliminate the amplifier-to-amplifier differences. Matt Kenworthy simply calculated the mean offset on a given image of the columns assigned to each amplifier, and applied the resulting corrections to the whole im-

age. These offsets changed unpredictably from image to image. Sometimes they were near-zero, sometimes they were very significant. In any case they were the dominant, and often the only, form of column variations. The complicated interpolations of colfudge could be replaced with four numbers. No artifacts were introduced. The column correction was not biased by point sources and therefore did not dim them. The column noise was eliminated at least as thoroughly as by colfudge.

#### A.8.4 Examples Images from Both Algorithms

Figures A.18 through A.23 give examples of the effects of both algorithms. In Figure A.18 we show an image from our data set on the star GJ 450 (the same image that was used as an example in Section A.3), just before the colfudge step in our processing. Fairly strong column pattern noise is evident, and repeats with the 4-column cycle from the readout amplifiers. No column artifacts other than the 4-column cycle are evident.

In Figure A.19, we show the image after processing by the author's colfudge algorithm. The column noise is effectively removed, but some artifacts appear. Faint point sources away from the star would be dimmed by 15-18% by the colfudge algorithm.

Figure A.20 shows the same image processed using Dr. Kenworthy's algorithm. The column noise is removed as effectively as by colfudge, but no artifacts are introduced, and point sources would not be dimmed.

Figures A.21 through A.23 give an example of the effects of the two algorithms from our  $\epsilon$  Eri data set. Note that the image used is not the same one as was used in Section A.3; that image had too little column variation to be a good example, so another frame taken a few minutes later has been used. In Figure A.21 we see an image with mild column noise. The 4-column pattern is easy to see, and there

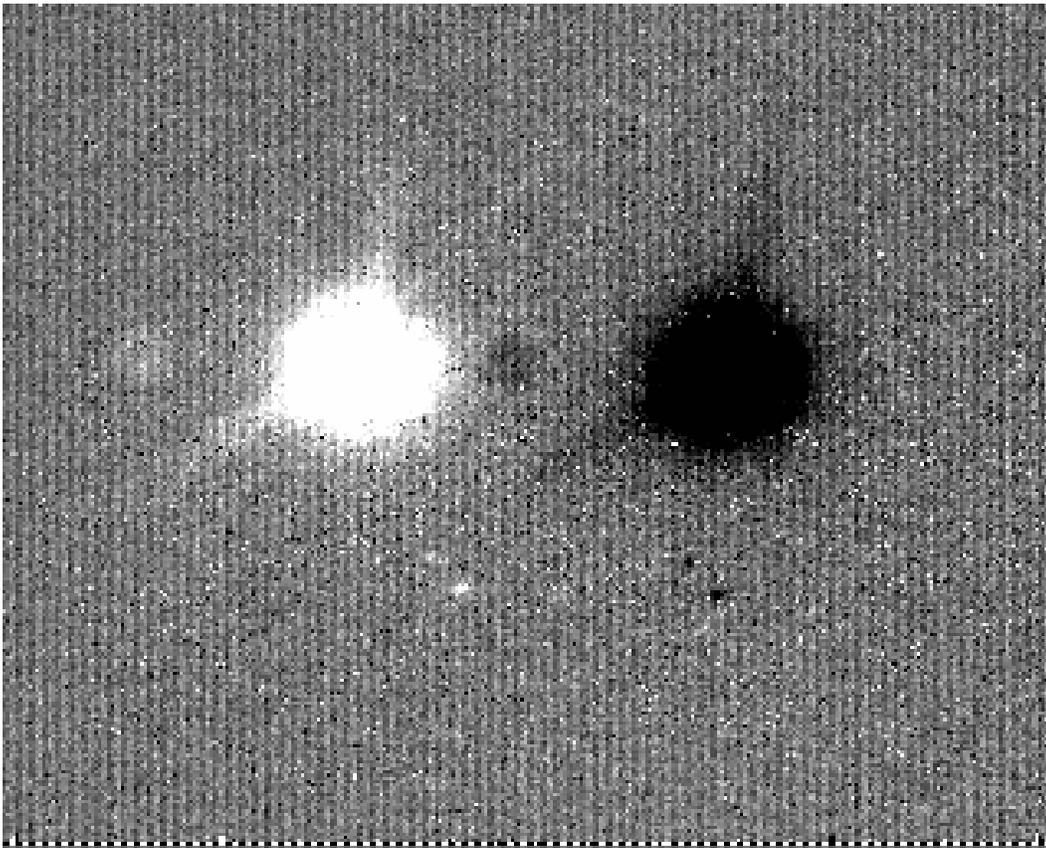


Figure A.18 Image of GJ 450 just before processing to remove column variations. Four-column pattern noise is evident; no other column variations appear.



Figure A.19 Image of GJ 450 with column noise removed by the author's colfudge algorithm. The 4-column pattern is effectively removed, but some artifacts are introduced.

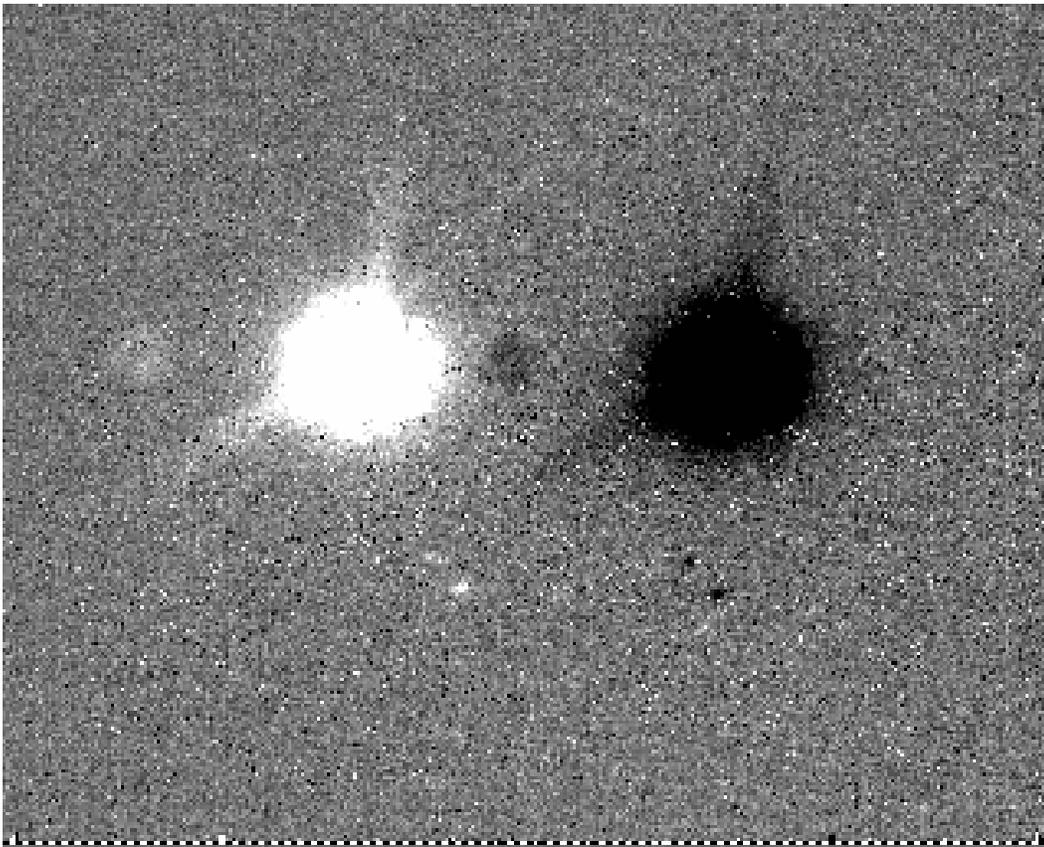


Figure A.20 Image of GJ 450 with column noise removed by Matt Kenworthy's 4-column pattern algorithm. The 4-column pattern vanishes, and no artifacts are introduced.

is also an anomalously bright column about 80% of the way from the left to the right edge of the image. This column has escaped the bad column removal code discussed in Section A.6, probably because it did not deviate strongly enough to be detected against the ubiquitous rays and ghosts produced by this very bright star.

In Figure A.22, the image has been processed by colfudge. The 4-column pattern and the single bright column are effectively removed. However, fairly significant artifacts from the halo, rays, and ghosts from this bright star have been introduced, and, as before, point sources in clear sky would be dimmed by 15-18%.

Figure A.23 shows the result of Dr. Kenworthy's 4-column algorithm. The pattern noise has vanished, and no artifacts have been introduced. Point sources would not be dimmed. The algorithm can only correct variations that follow the 4-column cycle from the readout amplifiers, so the single bright column is not removed.

As a final comparison between colfudge and the 4-column pattern algorithm, we show in Figures A.24 and A.25 the final, stacked images of  $\epsilon$  Eri made from data processed using the two images. Figure A.24, processed with colfudge, shows no more evidence of noise increased by residual or introduced column artifacts than does Figure A.25, processed with Dr. Kenworthy's 4-column algorithm. The two images are essentially identical. As  $\epsilon$  Eri is a bright star, and thus especially prone to colfudge artifacts, this comparison suggests that colfudge artifacts rarely if ever introduced significant noise into our final images. The main disadvantage of colfudge relative to the 4-column algorithm thus appears to be the fact that it dims faint point sources.

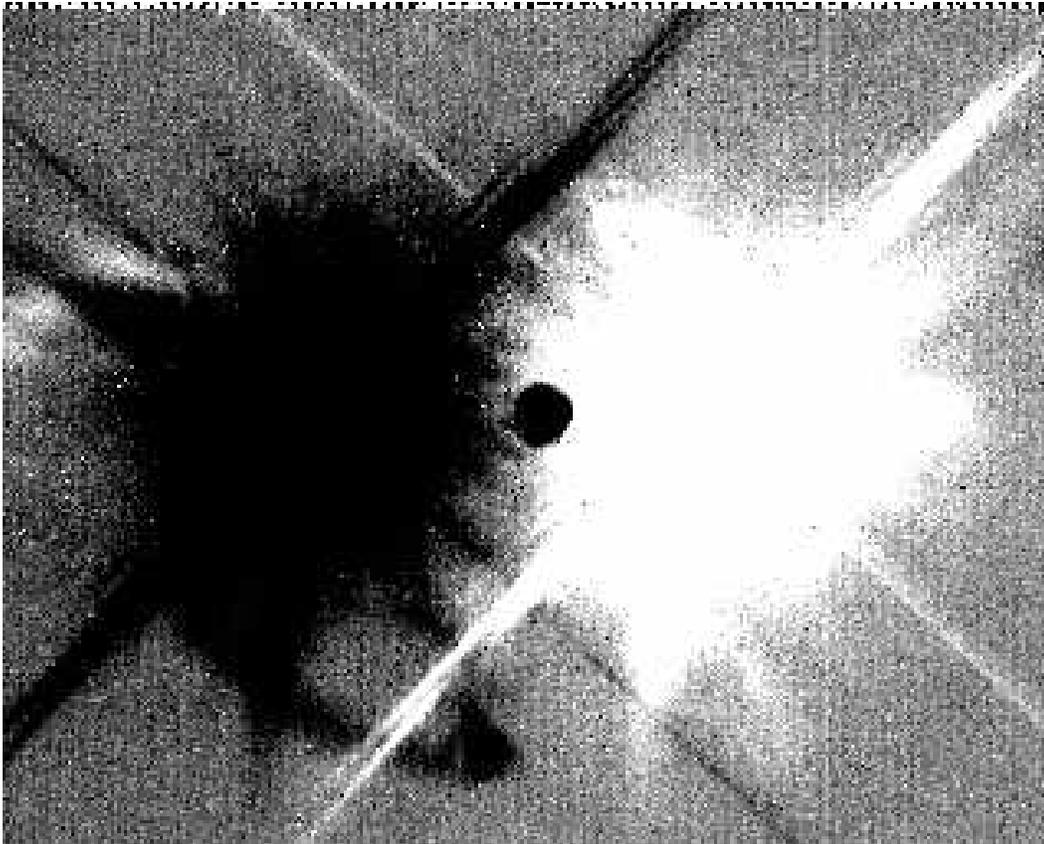


Figure A.21 Image of  $\epsilon$  Eri just before processing to remove column variations. Weak 4-column pattern noise appears; there is also a single anomalously bright column on the right hand side of the image, about 4/5 of the way across from left to right.

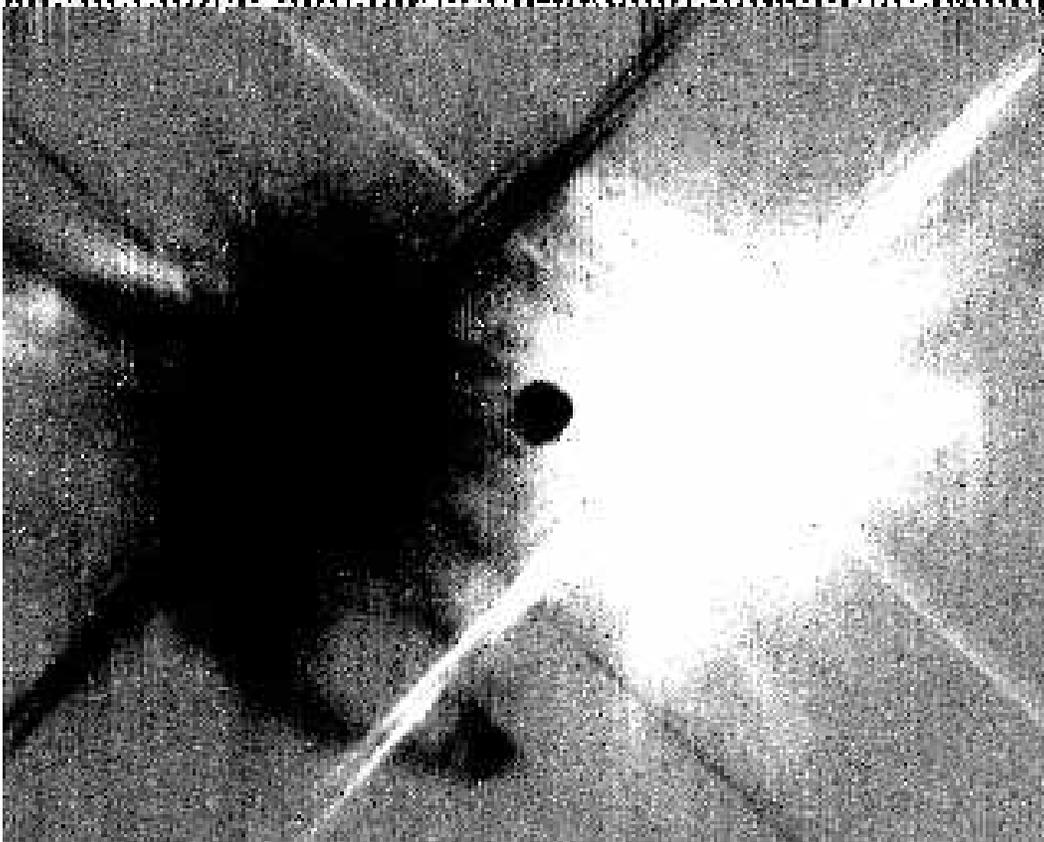


Figure A.22 Image of  $\epsilon$  Eri with column noise removed by the author's colfudge algorithm. The 4-column pattern and the single bright column are removed, but fairly significant artifacts are introduced due to the halo, rays, and ghosts from this bright star.

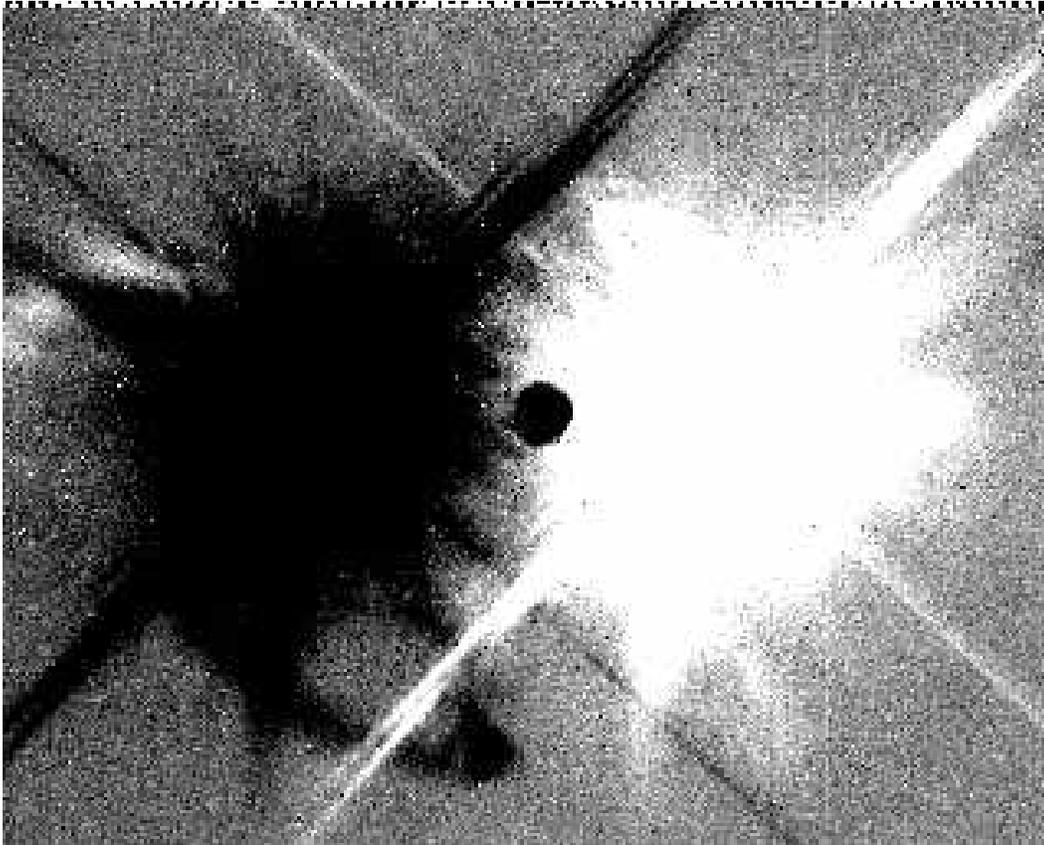


Figure A.23 Image of  $\epsilon$  Eri with column noise removed by Matt Kenworthy's 4-column pattern algorithm. The 4-column pattern is removed without the introduction of any artifacts. The single bright column is not removed.

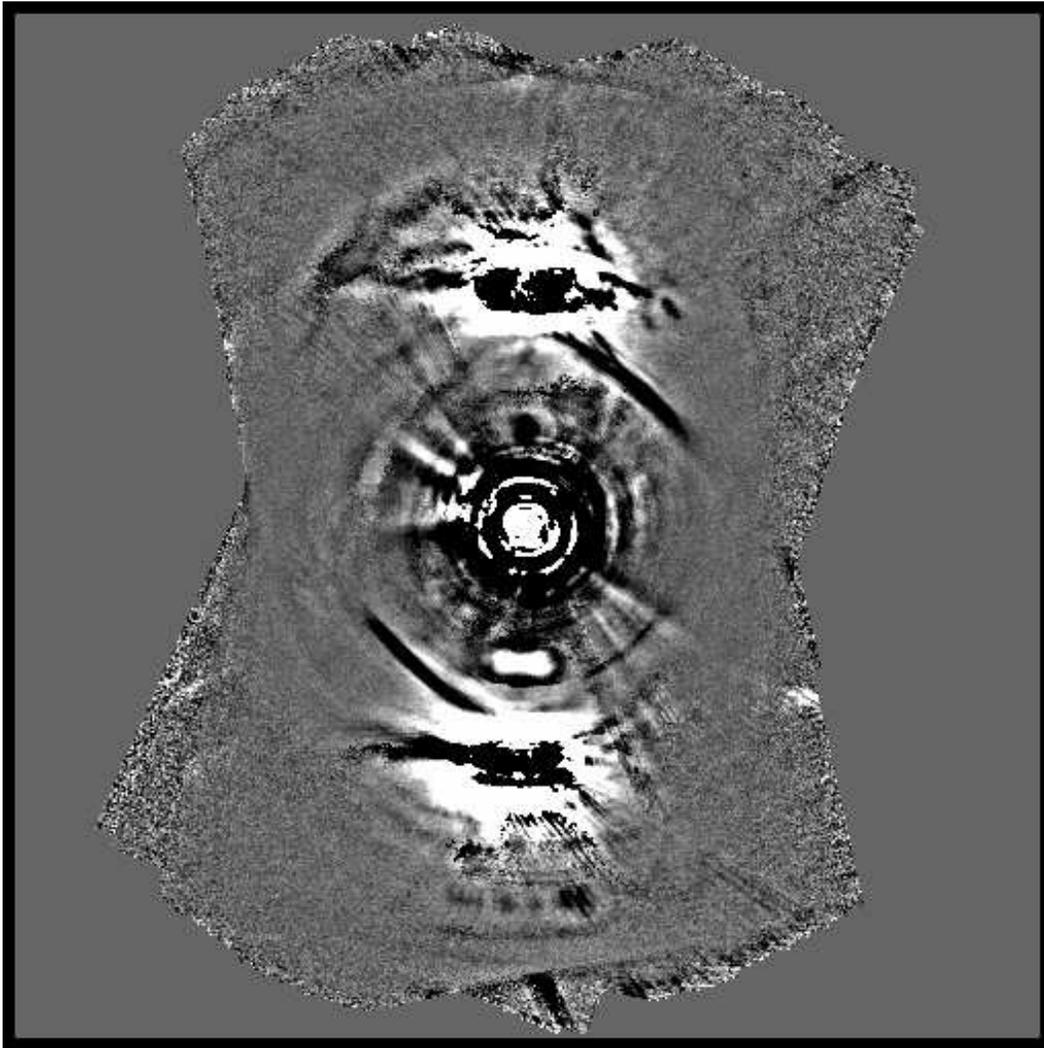


Figure A.24 Final stack of  $\epsilon$  Eri images processed with the colfudge algorithm. It is essentially identical to Figure A.25, which indicates that colfudge does not in general introduce significant artifacts into a final stacked image.

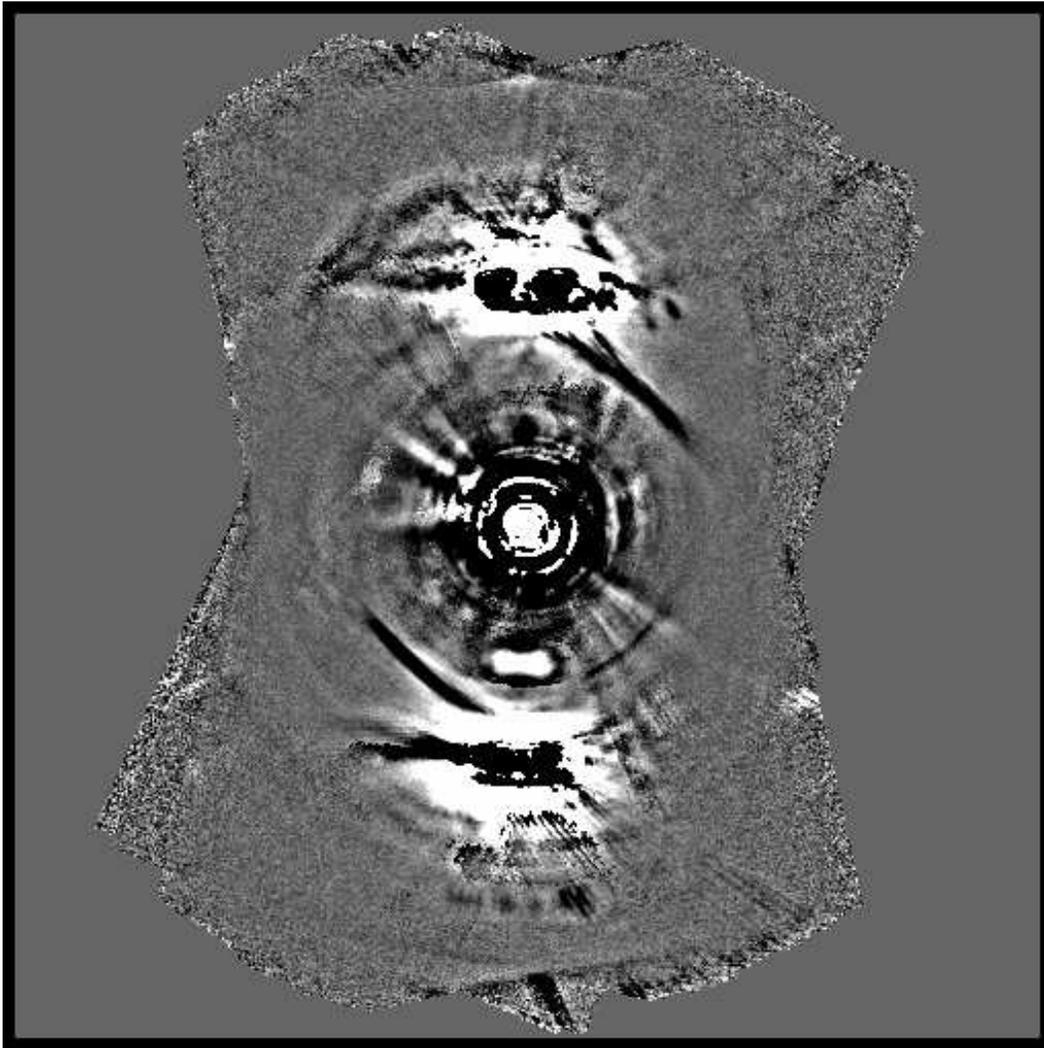


Figure A.25 Final stack of  $\epsilon$  Eri images processed with Dr. Kenworthy's 4-column pattern algorithm. It is essentially identical to Figure A.24, which indicates that the two algorithms produce very similar final results.

### A.8.5 Dimming of Faint Sources By Colfudge

We typically run the colfudge algorithm with a block size of 36 pixels. The FWHM of a typical PSF from our data is about 3.2 pixels. A faint point source does not appear against the noise in a single frame, so the attempt of the colfudge algorithm to reject it (via creeping mean pixel averaging, as explained above) does not succeed. The unseen source biases the column correction, and colfuge produces dark artifacts around it for this reason.

The 50% creeping mean rejection used by colfudge should throw out the 18 most deviant pixels in every average over a 36 pixel column segment. We would hope that even very faint point sources are partially rejected since they increase the odds that positive noise peaks at their location will be rejected. However, if we assume that colfuge averages over 18 pixels, including 3 pixels biased by a point source, we can estimate that colfudge would apply a spurious correction equal to  $1/6$  the flux from the point source. The point source would then appear nested in a dark trough on the image, and its flux would be reduced by  $1/6$ , or 16.7%. Note that this view of the operation of colfudge is simplistic because it does not take into account the interpolations between blocks. Also, because of parallactic rotation our final images do not actually show dark troughs extending from point sources.

Does colfudge really dim point sources by about 17%?. To find out, we performed the experiments discussed in Chapter 3 above. We will reiterate the important aspects relating to colfudge here for easy reading.

For our first experiment, we constructed 1000 simulated images. Each had a pure gaussian noise background with simulated sources inserted based on the PSF from some of our real star images. We used images of GJ 5, which were unusually sharp because we observed this star under conditions of excellent seeing.

We inserted one ‘fake star’, scaling the PSF to a very high flux, so that it was obviously seen on each individual image, and five ‘fake planets’, for which we made the PSF too faint to see against the noise in an individual image, but easy to see in the final stack of all 1000 images. Figure A.26 shows the result of stacking the images without prior processing, though in the figure the images serve a second purpose of illustrating the creeping mean combine. We ran colfudge on each of the 1000 images, and then stacked them. No rotation was used, so the fake planets appeared superimposed on dark troughs due to colfudge.

Colfudge dimmed the ‘fake planets’ by 14.6%. The fake star was not dimmed at all because the creeping mean combine effectively rejected it on every frame.

Thinking that colfudge might affect a less sharp PSF more severely, and that rotation might affect the dimming in an unpredictable way, we performed an additional experiment in which we inserted fake planets into a real data set. We chose the real images and PSF of the star GJ 354.1A, because the PSF was unusually blurry, with a FWHM of about 3.7 pixels. The measured dimming was 18.5%, of which all but 1-2% is most likely properly attributed to colfudge. Thus we concluded that colfudge dims sources by about 15-18%.

One can imagine ‘unlucky’ cases in which significantly more flux than this is removed. This arises from the fact that a limiting pixel brightness is one parameter of colfudge, and it does not attempt to derive a column correction from pixels that deviate from zero by more than this limit. Consider a block close to the central star, either above or below it. The strong flux from the star might put all the pixels in the block beyond the colfudge threshold, except for a small strip of pixels at the bottom. If the strip were only 3 pixels wide, a real source in it would bias the colfudge correction not by 1/6 but by almost 100% of its own flux level. Colfudge would therefore remove it entirely!

Fortunately this is not possible for several reasons. First, colfudge has a built-in limit, discussed above: it will not apply any correction to a given column in a given block if fewer than 9 pixels lie within the threshold. The hypothetical 3-pixel strip considered above would actually result in no column correction at all, and thus no dimming of point sources. Second, because colfudge uses interpolation based on the center coordinates of each block, a source near the edge of a block will be subject to less severe dimming than one at the center. Third, because of the parallactic rotation of all our data sets, it is unlikely (in many cases impossible) that a source could remain in the ‘unlucky’ situation we have considered throughout an entire data sequence. There is no doubt that sources exactly the wrong distance above or below a star in some of the images in a given set would be dimmed more than the 15-18% we have quoted for sources in clean sky, but the sources would not be caused to vanish. We can probably set an upper limit to their dimming at 50%, with 30% a more reasonable estimate for sources not carefully positioned at a point of ‘maximum unluckiness’. We remind the reader that even this 30% dimming could occur only in a small fraction of any given data set. Also, sources that occur *close enough* to a star that the star’s halo counts lift them above the threshold do not get dimmed by colfudge at all.

We note that sources can also be severely dimmed due to narrow ‘partial blocks’ at the top edge of an image. Our typical block size of 36 pixels results, unfortunately, in a strip of blocks only 6 pixels high along the top of the image. The 9-pixel limit applies only in the case of 36 pixel-high blocks, and so will not protect point sources in this region. Interpolation will not protect them well either, since they will be near the center of their narrow block. However, colfudge is applied before zero trimming. The zero trimming for images in September 2006 and later (covering half or more of all Clío data acquired to date) trimmed 6 pix-

els off the top of each image. Before Sept 06, due to the mirror-flip, only 3 pixels were trimmed from the image tops. So either a 3 or a zero pixel strip along the top of the final useful image area would be affected. In this strip, faint sources would be severely dimmed or even erased. However, because this is far from the star, parallactic rotation would very quickly move sources through this regime. Thus we would expect them to be missing only from a very few images. The integrated flux in the final stack would not be severely affected.

Based on Figures A.24 and A.25, we can conclude that any artifacts colfudge introduces do not in general survive to affect the final images significantly. The loss of some sensitivity due to the dimming of point sources is thus the only drawback of our accidental choice of a sub-optimal algorithm. We can set this loss at 15-18% in general in clear sky, 0% for bright sources or sources in the bright inner halo of the star, and no more than 30-50% for the most unfortunately placed sources above and below a bright star.

Our blind sensitivity tests, discussed in Chapters 2, 3, and 4 above, were sensitive to the dimming of point sources by colfudge and by all other stages of our processing. We note that the randomly placed fake planets in these tests were concentrated around the central star to probe close-in sensitivity. Some of them were therefore certainly subjected to the ‘unlucky’ dimming we have discussed above. It is on sensitivity estimates from these very realistic tests that we based the science results of this thesis.

## A.9 Rotation and Shifting

### A.9.1 Calculation the Required Rotation

We rotate images to get celestial north up. The required rotation depends on the instantaneous parallactic angle at the time the image was taken and on an offset

that is determined by how the instrument is bolted on to the telescope: more precisely, on how the detector chip is rotated relative to the axes of the telescope mount.

Normally, Clio image headers contain the parallactic angle read from the telescope control system. This is probably very accurate since it is used to determine the telescope pointing. However, our algorithm calculates the parallactic angle independently using the UT from the image headers. This has two advantages.

First, it sometimes happens that due to an observer error images are taken without querying the telescope system. These headers of these images do not have the parallactic angle. They do not have the accurate UT from the telescope control system either, but they do have the Clio computer UT. If the error is caught early enough, it is easy to find the offset between the telescope and system UT and patch the images so that our pipeline can process them easily.

Secondly, for very precise astrometry the amount of rotation that happens *during the taking* of a given image can be relevant. This can be true even for targets not very close to the zenith, which are not undergoing extremely rapid parallactic rotation. The author discovered this problem when comparing astrometry from long and short exposures of two different double stars. There was a statistically significant systematic offset in the measured position angles of each binary between the long and short exposures. Investigation showed it was due to the different amount of rotation that occurred during the long and short exposures. The UT time stamps in the image headers were from the time when the original coadded image was written to disk, while the astrometry sampled the parallactic rotation from the midpoint of the exposure. We corrected the UT by subtracting half the observing time per image, before using it to calculate the parallactic angle. This correction removed the offset between short and long exposures.

Calculating the parallactic angle from the time stamp, rather than merely reading the header value, makes it easier to correct a particular type of observer error, and enables exceptionally accurate astrometry when it is desired.

The constant rotation that must be applied to the images in addition to the parallactic angle is found by observing known double stars. We have a list of standard double stars, and observe one or more of them during each Clio run. We thus get a unique rotational offset for each run. The offset changes by about  $0.1^\circ$  from one run to another due to effects such as the slop of the mounting bolts before they are tightened down, or (if Clio has been disassembled since the previous run) slightly different placement of the detector and other internal components. The internal precision of bright-star position angle astrometry with Clio is perhaps  $0.01^\circ - 0.02^\circ$ . Unfortunately the position angles of our standard binaries are not known to better than about  $0.1^\circ - 0.2^\circ$ , so at present we are unable to obtain the offset with sufficient accuracy to take full advantage of Clio's potential astrometric capability.

#### A.9.2 Calculation the Required Shift

The shift is calculated based on the position of the primary star. This is found by five iterations of centroiding. The user sets the centroiding radius for each iteration based on the brightness of the star. For bright stars with large saturation radii the radius on the final iteration should of course still be well outside the saturation radius. For faint stars with unsaturated cores the radius should get quite small, perhaps down to 2-3 pixels, to get a good centroid against the background noise.

We use simple center-of mass centroiding on a nod-subtracted version of the image. Since the image is nod-subtracted, sky subtraction is not needed or used. We have confirmed that Clio star images yield very accurate centroids even if

they are severely saturated. The most stringent test cases of this are our sharp images of the white dwarf Procyon B, and of faint background stars near 61 Cygni A and B. In both cases images were shifted and stacked solely based on centroiding a highly saturated primary star, and yet the PSFs of faint companions were sharp. In the case of 61 Cygni A, especially, the accuracy and consistency of our rotational calculation was stringently tested as well because the background objects were far from the primary and the data set had a great deal of parallactic rotation.

The centroiding is actually performed on a nod-subtracted copy of the science image that is made exclusively for centroiding purposes, *before the main nod subtraction step*. This is because accurate knowledge of the centroid of the primary can be useful before the nod subtraction, as will be explained below. The centroid of the star on the *subtraction image* is also obtained.

### A.9.3 Zero-Trim, Zero-Pad, Shift, and Rotate: A Single Operation

Zero-trimming, in which strips along the edges of an image are set to zero, is necessary because the Clio detector has a 2-pixel wide overscan region along one edge. In any case it would be desirable because the extreme edges of the detector are in general artifact-prone and do not have normal sensitivity. We typically trim 2 pixels off the left and right edges of the image, 6 pixels off the top edge (which has the overscan region), and 3 pixels off the bottom edge. This is for the September 2006 and following runs; prior to this the overscan was on the bottom and the top and bottom trim depths were flipped accordingly.

Zero padding is required to avoid large, otherwise useful areas of images from being lost off the edge of the working area in the shifts and rotations required to register the data.

A general shift-and-rotate operation will not result in new pixels that directly

overlay the integer coordinates of the old pixels, so interpolation of some type is required. Every interpolation operation blurs the image slightly, so it is very desirable to accomplish shifting and rotation in a single operation. The math is slightly more complicated; other than that there is no disadvantage.

The simplest interpolation method is bilinear interpolation. An image can be zero trimmed and then zero padded in two separate operations; then a bilinear shift and rotation can be accomplished in one operation. The result is generally acceptable. Edge effects exist at the data/zero-pad boundary, but they are only one or two pixels in thickness.

Unfortunately, even a single bilinear interpolation operation blurs an image significantly. Bicubic spline interpolation is much better. However, performing a bicubic spline interpolation on a zero-padded image causes extensive edge effects at the data/zero-pad boundary. We developed a method to perform the zero-trim, zero-pad, shift, and rotation in a single operation.

For the bicubic spline, we use the `splie2` and `splin2` routines from Press et al. (1992). The `splie2` routine takes in an image and the  $xy$  pixel grid and creates a bicubic spline interpolation matrix. The routine `splin2` is then called using this matrix and a floating point  $x, y$  position as input. The interpolated pixel value at supplied  $x, y$  location in the original image will be output.

This makes our task conceptually simple. We call `splie2` with the original image and pixel grid. Next, we allocate an empty matrix to hold the larger, zero-trimmed, zero-padded, shifted, and rotated image. We loop over every pixel in this new, empty image. Using the shift and rotation we want to apply, we calculate the position this new-image pixel corresponds to in the old image. If it is beyond the image edges, the new-image pixel remains zero. If it is within the old image but falls in an area that is to be removed in the zero trim, the new-image

pixel is likewise left at zero. If, however, the new-image pixel corresponds to a point on the old image that lies in the good data region, the value of the new pixel is easily assigned by calling `spline2` with the appropriate old-image coordinates and the interpolation matrix output by `spline2`. In the end, a shifted, rotated, zero-trimmed version of the old image appears in the new matrix. Since the new image is larger, the zero-padding has effectively happened by construction. The edge effects are not just minor, they are non-existent. Every nonzero pixel in the new, shifted and rotated image represents valid data.

#### A.10 Image Stacking and the Creeping Mean

Our standard zero-trim, zero-pad, shift, and rotate operation takes a 320x258 pixel image, trims it down to 316x249, and leaves it shifted, rotated, and zero-padded to 500x500 pixels. Celestial north is up and the primary star is precisely centered on pixel 250,250. The task of our image stacking is simply to add a few dozen to a few hundred such images together in the most reasonable way.

The standard method for combining astronomical images is the median. It is simple and effective at rejecting inconstant artifacts. However, we have chosen the more complicated method of the creeping mean combine.

A creeping mean uses a single parameter: the rejection fraction. A set of values is considered – here, the brightness of a given image pixel through the image stack. The values are averaged, then the one that deviates most from the mean is rejected. The remaining data are averaged again, and the new most deviant point is rejected. Repeated iterations of this process are carried out until the required fraction of the points has been rejected. While we discuss here the use of the creeping mean in stacking images, it is also useful for averaging the pixel values over regions of a single image. It is used, ie, in nod subtraction to average the sky

regions on which the subtraction image scaling is based (Section A.7) and for the column and block averages used in the colfudge algorithm (Section A.8).

The creeping mean combine was first suggested to the author by his father, Dan Heinze. A PhD in geophysics, Dr. Heinze has a great deal of experience in the field of seismic processing for oil exploration. He has found the creeping mean very useful in analyzing seismic data, because of its ability to cause large chunks of deviant data to vanish from the final mean. Consider, for example, a bimodal distribution in which 60% of the data are drawn from a gaussian distribution which is the real 'signal' that is being measured. The other 40% of the data are affected by an error source and represent another gaussian distribution strongly offset from the first. A median will be biased by the large amount of spurious data; the value selected will be offset from the peak of the 'signal' gaussian in the direction of the 'error' gaussian. A creeping mean with 50% rejection, however, will reject points from the deviant distribution until they are all gone. The final mean will be carried out only over the signal, and the bias will be very small. The creeping mean is very effective at rejecting strongly deviant chunks of data, provided the fraction of data affected is less than the rejection fraction of the creeping mean. If more than half the data are deviant, of course, neither the creeping mean nor the median can help.

Since we observe with the MMT instrument rotator off, ghosts, rays, and other artifacts in our images rotate slowly with respect to real celestial sources. The intense artifacts from the negative nod subtraction images also rotate. This creates exactly the situation where the creeping mean performs best: when a large, high-amplitude artifact rotates through a region, it makes the pixels there strongly deviant. However, provided the parallactic rotation was sufficient, a sufficient fraction of images will exist with good data in that region that the creeping mean

will reject the artifact. Figure A.26 compares a 20% rejection creeping mean with a median stack with a single iteration  $5\sigma$  sigma clip. Along the bottom of the images are simulated real planets present in 100% of the data. At top left is a simulated bright star image, also present in 100% of the frames. The other objects are simulated ghosts. In order from left to right and top to bottom they are present in 50%, 20%, 19%, 18%, 15%, 10%, and 5% of the data. As expected, neither the creeping mean nor the median stack helps against the 50% ghost; if half the data is bad, we're out of luck. For the other ghosts the creeping mean clearly does a better job. The median sigma clip removes the cores but is more strongly biased by the remaining bright haloes. Because data has literally been thrown away and was not then available to average over, the creeping mean has a slightly higher background noise level than the median. This is a small price to pay for effective removal of ghosts.

We comment, however, that side-by-side tests of the creeping mean against the median on real Clio data do not necessarily show a strong, clear advantage to the creeping mean. The median does better in some areas, the creeping mean in others. What is clear is that the creeping mean is a reasonable alternative to the median, and has worked well for us. Observers wanting the best possible sensitivity for heavily haunted or artifact-ridden images would do well to test the creeping mean against more conventional image combination methods to see if it might be the best choice for their particular data set. A creeping mean combine with an extreme rejection fraction of 50%, in particular, sometimes seems to work magic on very bad data sets, at the cost of increased background noise since half the data has been rejected.

Our implementation of the creeping mean algorithm ignores pixels that were equal to zero. This is very important, because the nod offsets and the parallactic

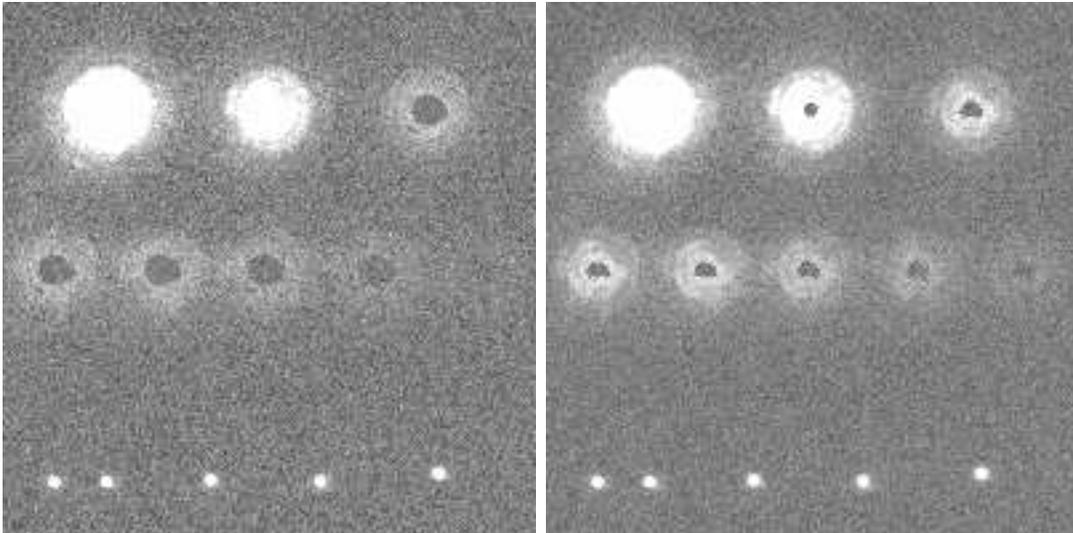


Figure A.26 Comparison of different image stacking methods. On the left, a 20% rejection creeping mean. On the right, a median combine with a single-iteration  $5\sigma$  sigma-clip. The faint images at bottom are simulated planets; the bright upper left image is a simulated star. In order from left to right and top to bottom the other artifacts simulate ghosts present in 50%, 20%, 19%, 18%, 15%, 10%, and 5% of the data. The creeping mean clearly removes most of the ghosts more effectively, while the median removes only the inner cores and is more strongly biased by the halo regions. The sky noise is somewhat lower in the median combined image.

rotation in our data sets make it inevitable that valid data regions on some images would overlay zero padded regions on others. It is essential to keep the zero pads from being averaged into the real data. Since our sky backgrounds have mean zero, zero is a possible data value. However, the creeping mean stack is performed on fully processed images with FITS floating point pixel values. Valid data pixel values of exactly zero are vanishingly improbable.

### A.11 Unsharp Masking

After the creeping mean stack, we unsharp mask the final image to remove the bright stellar halo. This also helps with dark haloes left over from the negative nod subtraction images, and with other diffuse, low spatial frequency features. Unsharp masking is a very significant step to reveal faint point sources for easier detection by both manual and automatic methods.

Our procedure for unsharp masking is to create a version of the image convolved with a gaussian kernel of  $\sigma = 5\text{pix}$ , and then subtract this from the original image. Our gaussian has a FWHM of 11.8 pixels, as opposed to about 3 pixels for a typical PSF, so the unsharp masking has little effect on point sources.

A few innovations in our masking strategy have substantially improved our results. First, we do not use Fourier methods in the convolution, as these would have caused artifacts at the zero pad boundaries. Rather, we simply implemented the mathematical definition of a convolution on a pixel by pixel basis. That is, we set each pixel in the convolved image equal to the integral over all the pixels in a region surrounding that pixel on the original image times the normalized gaussian kernel centered on the pixel. This is much slower than a Fourier convolution, but it gives us very useful flexibility. First, we define the region over which the integration is performed as a disk with radius set so that the gaussian value at its

edge is 1% of the central value. Second, we reject zero-valued pixels in the main image from consideration in the integral. This prevents artifacts at the edge of the zero pad. Second, we perform 3 iterations of high-rejection sigma-clipping *on the set of pixels contained in the disk on the original image*. The clipping threshold is set to a very low  $2.5\sigma$  to reject even faint point sources as completely as possible. Only pixels that do not get rejected are included in the integral. This means that sufficiently bright point sources are rejected from the convolved image. Therefore they are *totally unaffected* by the unsharp masking operation, rather than being only slightly affected. This innovation also greatly reduces the tendency for dark halos to form around bright point sources as a result of the unsharp masking.

The unsharp masking is the final step in our image processing pipeline. Figure A.27 shows its effect on the final image.

## A.12 Specialty Processing

Up to this point we have described our baseline processing method. There are 8 different variations on this method, which we will now describe.

### A.12.1 PSF Subtraction: The ‘b’ Reduction Method

We have referred to our baseline processing, described above, as the ‘a’ processing method. The ‘b’ method involves an attempt to subtract the bright stellar PSF to show faint sources near the star.

Our method for this makes use of the fact that most features of the stellar PSF remain fixed with respect to the telescope and the instrument rather than the sky. Importantly, this includes not only diffraction rays and ghosts but also the persistent spurious point sources known as ‘super speckles’. Super speckles probably arise from slight aberrations in the primary mirror or other telescope optics that have too high a spatial frequency for the AO system to correct them. They typ-

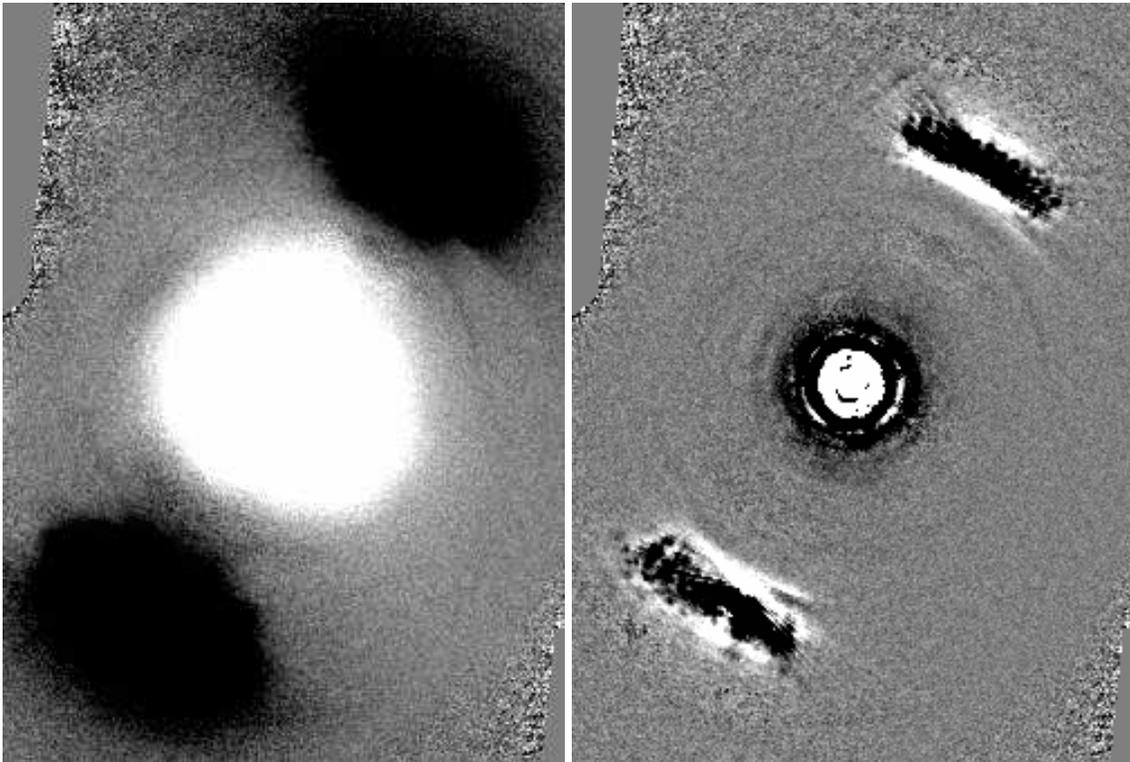


Figure A.27 Example of our final unsharp masking. On the left is a final stacked image of GJ 211, processed with our standard processing but not unsharp masked. On the right is the same image, unsharp masked using our standard method. The stretch on these image is -10 to +10 ADU.

ically remain approximately constant during observations of a given star, and then change when the telescope slews to another target. They may be due in part to gravitational flexure of the primary mirror.

In principle, we can make a master PSF image by processing a data set exactly the same as we would for regular science data, except for the parallactic rotation. Then in the final, shifted images, diffraction rays, ghosts, and superspeckles will all occupy consistent locations, but real sources will rotate through the stack. A creeping mean combine of such images *rejects any real sources, rather than the ghosts and artifacts*. A properly rotated version of the resulting master PSF image could then be subtracted from every science image before the science stack. Stellar features such as the stellar halo, ghosts, diffraction rings, and rays would vanish. Real sources, since they did not appear in the master PSF image, would not be subtracted away.

In practice we encountered a difficulty. Sometimes we were unable to schedule the observations of a given star optimally. This could be due to other high-priority targets, bad weather, or equipment problems. We might start or end our observations near the time of the star's transiting the meridian, instead of getting numbers of images on either side of transit. In this case, if the star was at a favorable declination so that it passed close to the zenith, one half of the data might contain the vast majority of the parallactic rotation. The other half of the data might have very little rotation, as the star set or rose more-or-less straight down or up. In that case, a real source could appear fairly brightly in the master PSF image, at its consistent location in the half of the data with little rotation. The source would then be mostly subtracted away from that half of the data.

To avoid this, we made two master PSFs, one from the first half of the data, and the other from the second half. We subtracted the second-half PSF from sci-

ence images in the first half of the data, and vice-versa. Thus, we minimized the self-subtraction of real sources in data sets with sub-optimal rotation patterns.

Changes in the seeing can change the amount of light going into the central PSF of a star, causing more or less to be diffracted into the broader stellar halo. Also, small changes in transparency can cause the total flux from the star to change. For these reasons, we thought it necessary to scale the PSF to match each image. We also normalize the individual images in the master PSF construction step to an average flux before stacking them, so that they match as well as possible. Since in general the cores of the star images are saturated, we scale or normalize the PSFs in an annulus surrounding the core. The user selects the inner and outer radii of this annulus based on the saturation radius of the star.

Each master PSF is made from a 50% rejection creeping mean stack of the appropriate half of the data set. The appropriate master PSF image is subtracted from each science image after the zero-padding, shifting, and rotation. This image is scaled to match the science image within the assigned annulus. Generally, full PSF subtraction is performed out to a radius of 75 pixels from the central star. To avoid hard edge-artifacts at the boundary, the PSF subtraction is faded from 100% to zero between 75 and 85 pixels out. These large radii are chosen so that even distant ghosts and diffraction rays are subtracted. PSF subtraction introduces substantial noise into the regions of the science images that are well away from the bright central PSF. This is because the background noise in the PSF image adds to the background noise in the science image. This causes no problems in the final analysis of a data set, because the images from the baseline 'a' reduction, rather than the 'b' reduction, are used to investigate the areas well away from the bright central star. The PSF subtracted 'b' images provide a useful check on the 'a' images even far from the star, at the locations of bright ghosts and

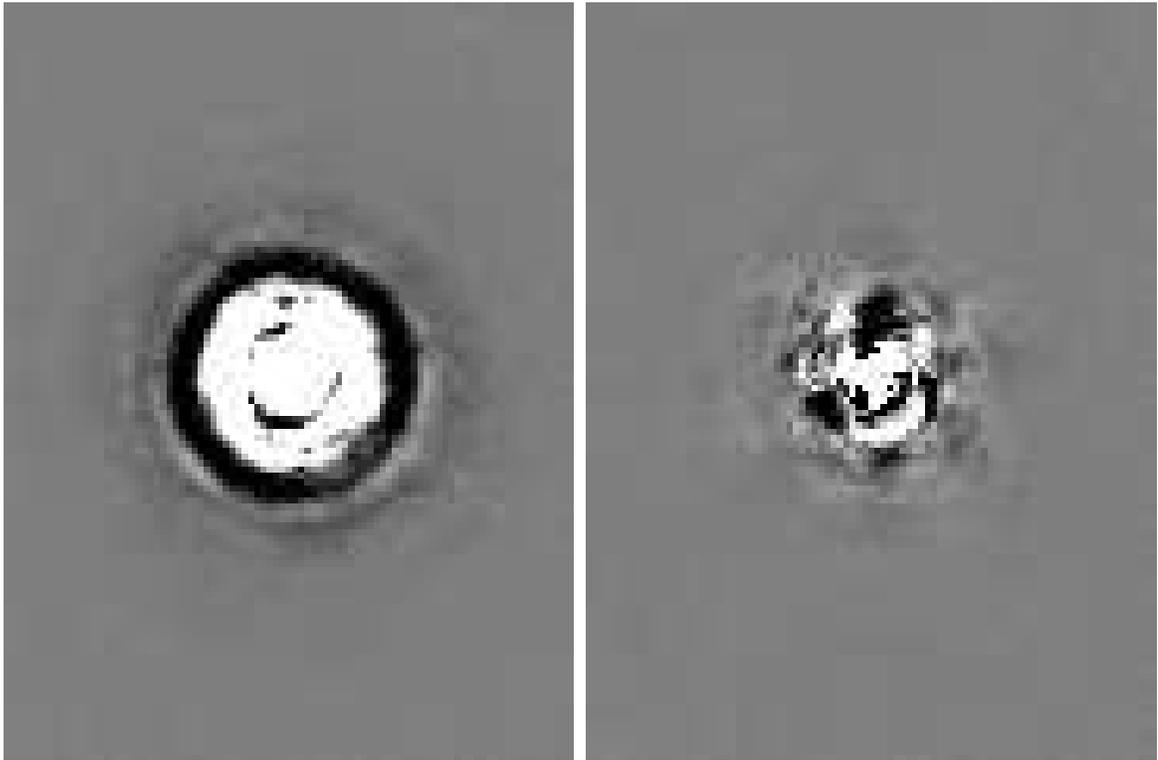


Figure A.28 Example of the effect of PSF subtraction using the 'b' processing method. On the left is our 'a' method master image of the star GJ 211. On the right is the the image made by processing the same data with 'b' method PSF subtraction. The power of the method is clearly seen. The stretch on these images is -100 to +100 ADU.

along rays. Sometimes such artifacts might be mistaken for real sources on the 'a' image, but the fact that they cleanly subtract away on the 'b' image clearly shows them to be spurious.

Figure A.28 illustrates the effect of the 'b' method PSF subtraction.

Whenever the 'b' method PSF subtraction algorithm is applied to a data set with less than about  $50^\circ$  of parallactic rotation, there is a radius from the star within which real sources would be significantly dimmed by self-subtraction.

Since the two master PSF images are made from the first and second halves of the data, and therefore the mean parallactic angles for the two images probably differ by about half the total parallactic rotation in the data set, we typically estimate this minimum radius as:

$$R_{min} = \frac{3.0\text{pix}}{\sin(\theta/2)}$$

Where 3.0 pixels is the approximate FWHM of a typical PSF, and  $\theta$  is the total amount of parallactic rotation in the data set.  $R_{min}$  is therefore the minimum distance where parallactic rotation will cause a real source to rotate by one FWHM within the stack used to make each master PSF image, and the position of the source residuals in the two master PSFs will be at least one FWHM apart. We do not consider the 'b' method PSF subtracted images valid for sensitivity calculations within  $R_{min}$ . However, it is quite possible to detect real sources well within  $R_{min}$  if they are sufficiently bright. Such sources will have a distinctive self-subtraction signature: a bright, azimuthally compressed source with dark shadows on both sides. We intentionally look for such features when manually examining images for faint companions.

The PSF subtraction method we have described here uses the same property as the ones described in Marois et al. (2006) and Lafrenière et al. (2007a). It is, however, less sophisticated and apparently less powerful. The LOCI algorithm described in Lafrenière et al. (2007a) is apparently very sophisticated and powerful when used on shorter wavelength images from the Gemini North telescope. It is not clear how well LOCI or a variant thereof would work on Clio data. One problem that the LOCI algorithm might encounter with Clio data is the larger angular size of the diffraction disk, which means that self-subtraction is more of a problem if the parallactic rotation is not very large. Another is the noise from the

bright thermal background, which means there is a high cost to subtraction because the noise from the subtraction image adds to what was already there. Still, according to Lafrenière et al. (2007a) the algorithm is both versatile and powerful. We will probably try to apply some adaptation of it to Clio data in the future.

#### A.12.2 PSF Subtraction: The ‘c’ Reduction Method

If there is less than about  $10^\circ$  of parallactic rotation in a given data set,  $R_{min}$  as calculated above becomes more than 30 pixels. Except for exceptionally bright stars, there is little PSF to be subtracted at radii beyond 30 pixels. The ‘b’ reduction method is therefore not very useful for such data sets.

The ‘c’ reduction method makes its own approximate PSF from each image. Every pixel within a specified distance of the central star is considered in turn. A circle centered on the primary star and passing through the pixel in question, is considered. An arc of length 25 pixels, centered on the pixel in question, is constructed. The creeping mean average of this arc, with 50% rejection is computed, and the resulting value is the value of the approximate PSF at that pixel.

Subtraction of this approximate PSF leaves point sources only slightly dimmed, but diffraction rings, halo residuals, and diffuse ghosts strongly suppressed. In the case of data sets for which both methods work, the ‘c’ method is inferior to the ‘b’ method because it cannot remove super speckles (they look like point sources). The ‘c’ method has the advantage at larger distances from the star that it does not introduce sky noise as the ‘b’ method does. In general with the ‘c’ method the user selects an inner radius, typically approximately equal to the saturation radius. Inside this radius no PSF subtraction is attempted and the pixels are simply set to zero. As with ‘b’ method PSF subtraction, full PSF subtraction is usually performed out to 75 pixels radius, after which the subtraction is faded linearly from 100% at 75 pixels to zero at 85 pixels.

The 'c' reduction method is the reason for an odd fact that we noted above: we calculate the centroid of the primary star on the science *and* the nod subtraction image using temporary subtracted copies made solely for the purpose, before the 'real' nod subtraction step in the processing. This allows us to apply 'c' method PSF subtraction to the science *and the nod subtraction image* before subtraction. Having the PSF partly removed on the nod subtraction image reduces the intensity of the negative nod subtraction image. This chance to remove much of the negative nod subtraction image is a sort of 'free bonus' to the 'c' PSF subtraction method; we have made use of the opportunity. This would not work as well with the 'b' method because of the sky noise it would introduce around the negative star images.

The 'c' method PSF approximation, like the centroid, is calculated based on temporary, nod-subtracted versions of both the science and the subtraction image made solely for the purpose. Once calculated, the PSF approximation is subtracted from the original, unaltered images, and then nod subtraction proceeds in the usual way.

Figure A.29 illustrates the effect of the 'c' method PSF subtraction.

In general we have not used the 'c' method on data sets for which the 'b' method produced good results. It is not as powerful as the 'b' method because it cannot remove super speckles.

### A.12.3 Pre-Stack Unsharp Masking: The 'd' Reduction Method

The 'd' reduction method was developed to combat a problem we sometimes experience with creeping mean combined images of bright stars. If the star has an extended halo, there can be parts of the image that are very bright from the stellar halo on images from one nod position, but dark because of the negative stellar halo on images from the other position. This creates a strong bimodality within

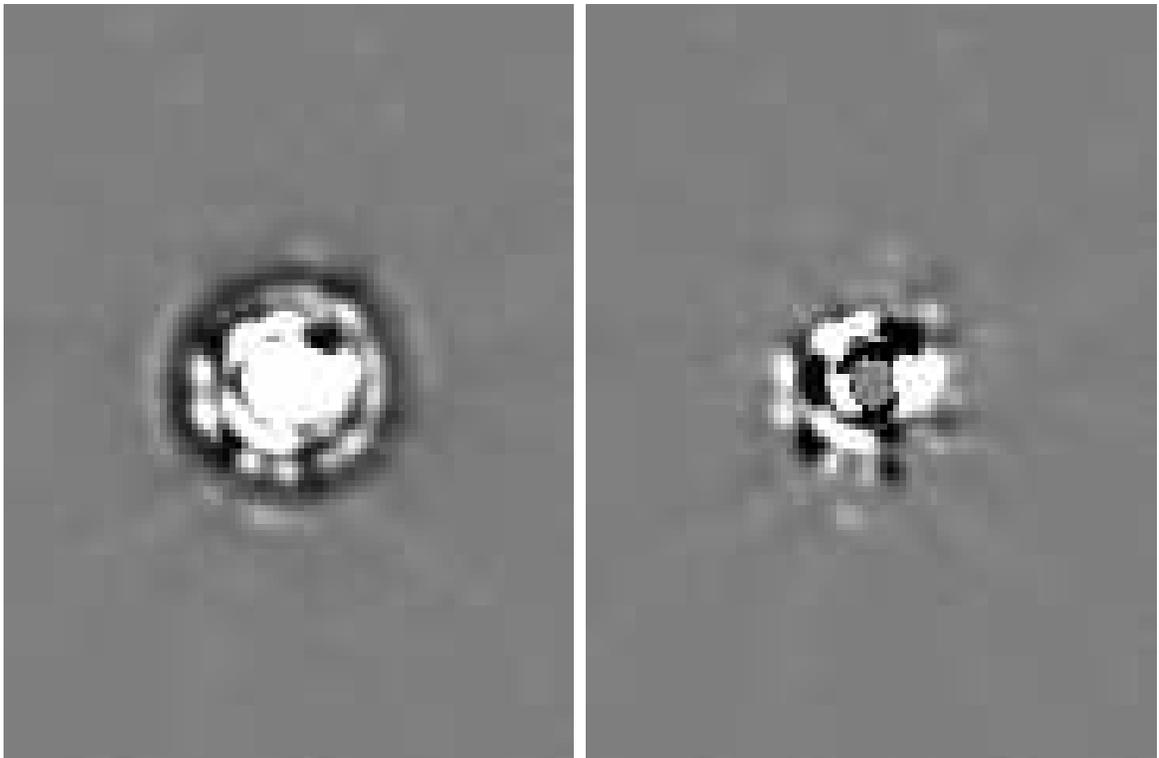


Figure A.29 Example of the effect of PSF subtraction using the 'c' processing method. On the left is our 'a' method master image of the star GJ 278C. On the right is the the image made by processing the same data with 'c' method PSF subtraction. The PSF is helpfully subtracted, but super speckles remain, in contrast to the 'b' method result shown in Figure A.28. The stretch on these images is -100 to +100 ADU.

the final image stack for pixels in certain regions. The creeping mean tends randomly to select one mode of the bimodal distribution, and to reject pixels from the other mode. The chosen mode can be different for neighboring pixels, so that one pixel can be very bright in the final stacked image, while a nearby one can be dark. The data quality in the affected regions is totally destroyed by this intense ‘bimodality noise’. We note that it is likely a median/sigma clip stack would have experienced a similar problem; this is not an exclusive weakness of the creeping mean. We have solved the problem by unsharp masking each individual image *before the final stack*. This removes the stellar haloes and insures the stack is performed on images with good, zero-mean sky values in most regions. The improvements in the data quality is enormous.

Images processed using the ‘d’ method are unsharp masked (again) after the final stack, as are images reduced using all of our methods. The same flux-loss experiments described in Section A.8.5 also tested the loss due to pre-stack unsharp masking: it was found that it removes 5-10% of the flux from faint point sources. Little flux is lost from sharp point sources; more is lost from fuzzy ones.

Even for less bright stars, the ‘d’ script images usually prove remarkably clean, and provide a useful comparison against the ‘a’ images. Spurious sources due to artifacts (especially of diffuse ghosts) in the ‘a’ images often vanish or are made much fainter on the ‘d’ images.

Figure A.30 shows how ‘d’ method processing cleans up an image.

#### A.12.4 Combining Methods: The ‘e’ Reduction Method

Our ‘e’ data reduction method is a combination of the ‘b’ and ‘d’ methods, to remove bimodality noise from and generally clean up PSF-subtracted images of bright stars. The pre-stack unsharp masking is applied after the PSF subtraction.

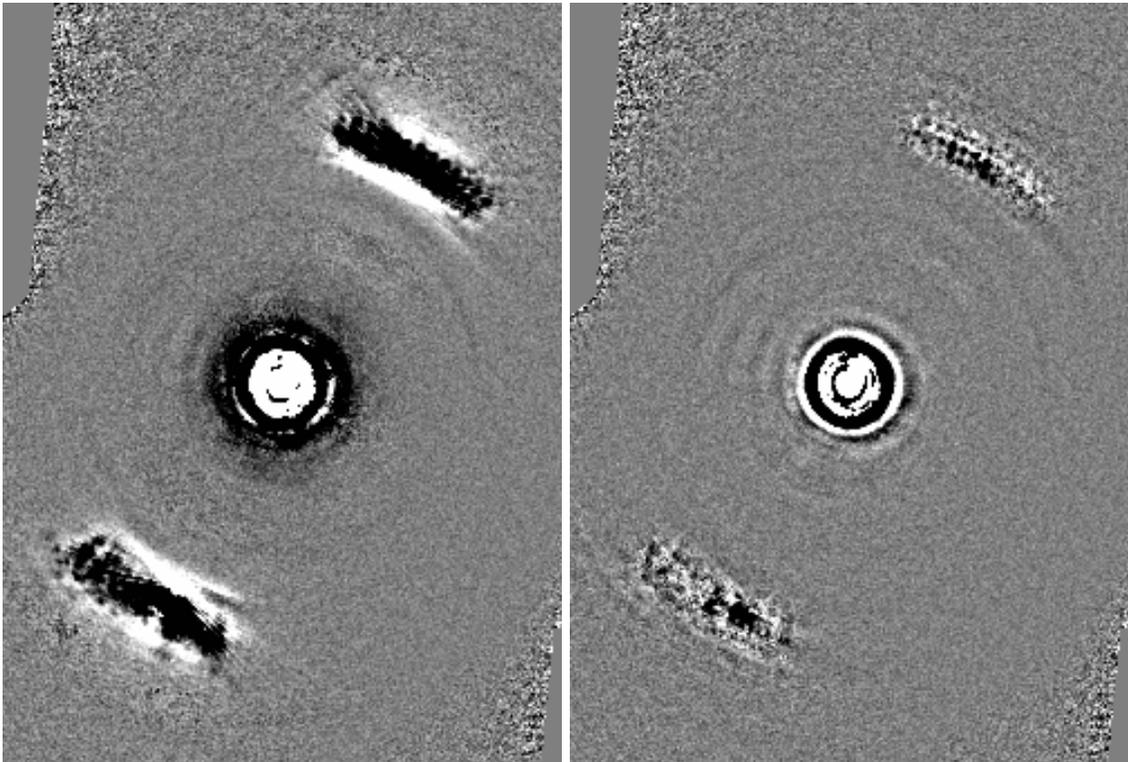


Figure A.30 Example of how the 'd' processing method cleans up an image. On the left is our 'a' method master image of the star GJ 211. On the right is the the image made by processing the same data with the 'd' method. The stretch on these images is -10 to +10 ADU.

### A.12.5 Removing Nod Artifacts: The 'x' Reduction Method

The 'x' data reduction method uses a variant on nod subtraction that avoids the dark negative images. A single master sky image is made for the whole data set, and this is subtracted from each individual science image, rather than the assigned subtraction image.

The science/subtraction image pairs are assigned as usual, but before the main processing a master sky image is made using all of the pairs. Each subtraction image is scaled to match the science image, just as it would be for a regular nod subtraction. Rather than performing the nod-subtraction, however, we then slice each image along the perpendicular bisector of the line between the position of the star in the science image and the position of the star in the subtraction image. This produces two pieces of each image: one with a star on it and one without. We discard the pieces with the stars and join the starless ones together. The result is an image that is maximally free from stellar artifacts, and shows only blank sky.

After a starless sky image has been made from each science/subtraction image pair, the sky images are scaled to have the same mean brightness match and then stacked with a 50% rejection creeping mean to produce a master sky image.

After the sky image construction is complete, the images are processed in the usual way. However, at the nod subtraction step the master sky image, rather than a single selected frame from the opposite beam, is used as the subtraction image. Thus there are no negative star images from the nod subtraction. Because the master sky image is a low-noise average of many individual frames, there are no serious issues of non-independence of subtraction noise. However, because the sky image is made from the whole data set, it does not track subtle variations in the images due to changing sky background as well as ordinary

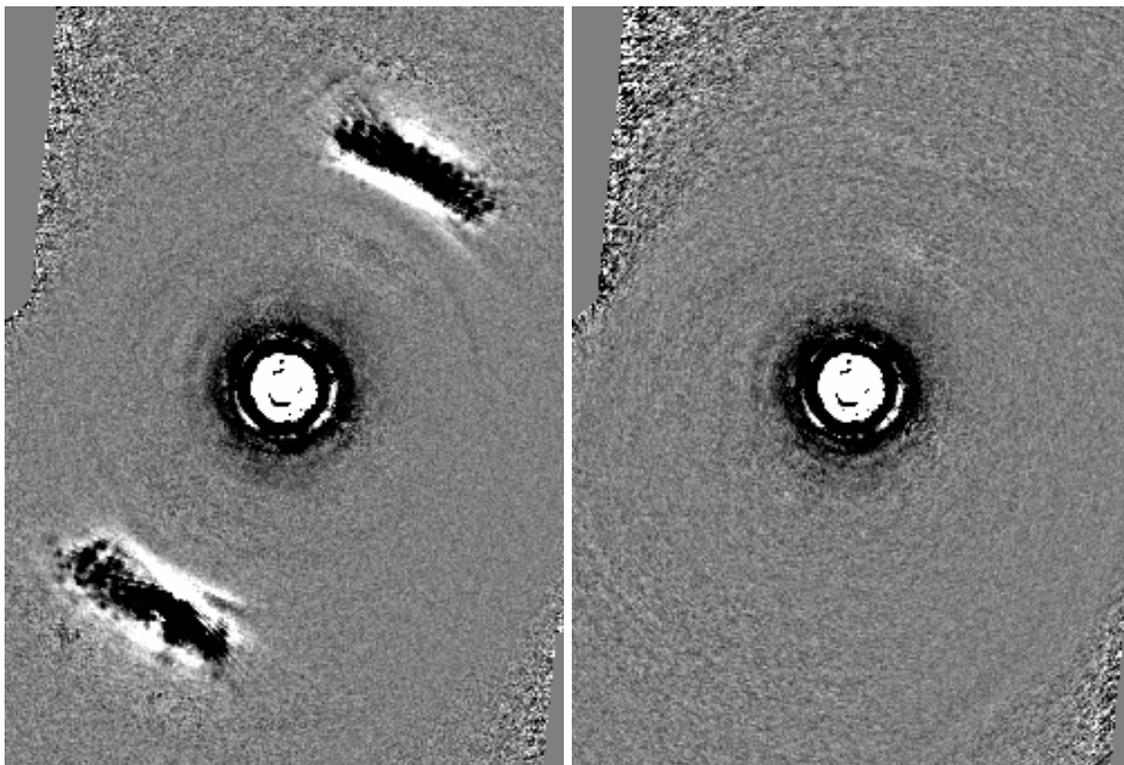


Figure A.31 Example of how the 'x' processing method removes the nod subtraction artifacts from an image. On the left is our 'a' method master image of the star GJ 211. On the right is the the image made by processing the same data with the 'x' method. The stretch on these images is -10 to +10 ADU.

nod-subtraction frames, chosen to be as close to the science image in the observing sequence as possible. This always makes the 'x' method images a bit noisier than 'a' images in places where both show clean sky.

Figure A.31 shows the effect of 'x' method processing.

The usefulness of the 'x' processing method varies enormously from one data set to another, presumably depending on how stable the sky background was. In all cases the 'x' processed image has a noisier sky background than the 'a' image. However, in some cases the differences are very slight, and the absence of the

dark nod subtraction artifacts on the 'x' image seems almost miraculous, delivering good sensitivity in regions of the image previously destroyed by nod artifacts. In other cases the 'x' image is hopelessly full of intense column noise and other artifacts. We note that colfudge is unable to handle this high-amplitude column noise because it typically violates its correction threshold. Matt Kenworthy's 4-column algorithm, on the other hand, might do better since there is no limit to the magnitude of the corrections it can measure and apply.

#### A.12.6 Combining Methods: The 'y' Reduction Method

The 'y' reduction method is a combination of the 'x' and 'd' methods. Nod subtraction is performed using a master sky image, as in the 'x' method. After the nod subtraction images are subjected to pre-stack unsharp masking, as in the 'd' and 'e' methods.

The 'y' reduced image is often much better than the 'x' image, offering a much greater improvement over it than the 'd' images offers over the 'a' image. This is because for the reason mentioned above the master sky image does not subtract away all artifacts as well as ordinary nod subtraction frames do. It often happens that fairly high-amplitude, low spatial frequency variations in the sky background remain after the subtraction of the master sky image. The pre-stack unsharp masking suppresses these very effectively, with the result that the final stack of 'y' processed images is much cleaner than the 'x' image.

Figure A.32 gives an example of how 'y' method image processing produces a cleaner image than the 'x' method.

#### A.12.7 Binary Star PSF Subtraction: the 'f' Reduction Method

Binary stars can offer a uniquely powerful method of PSF subtraction. If the stars are close enough together to appear in the same field, far enough apart that the

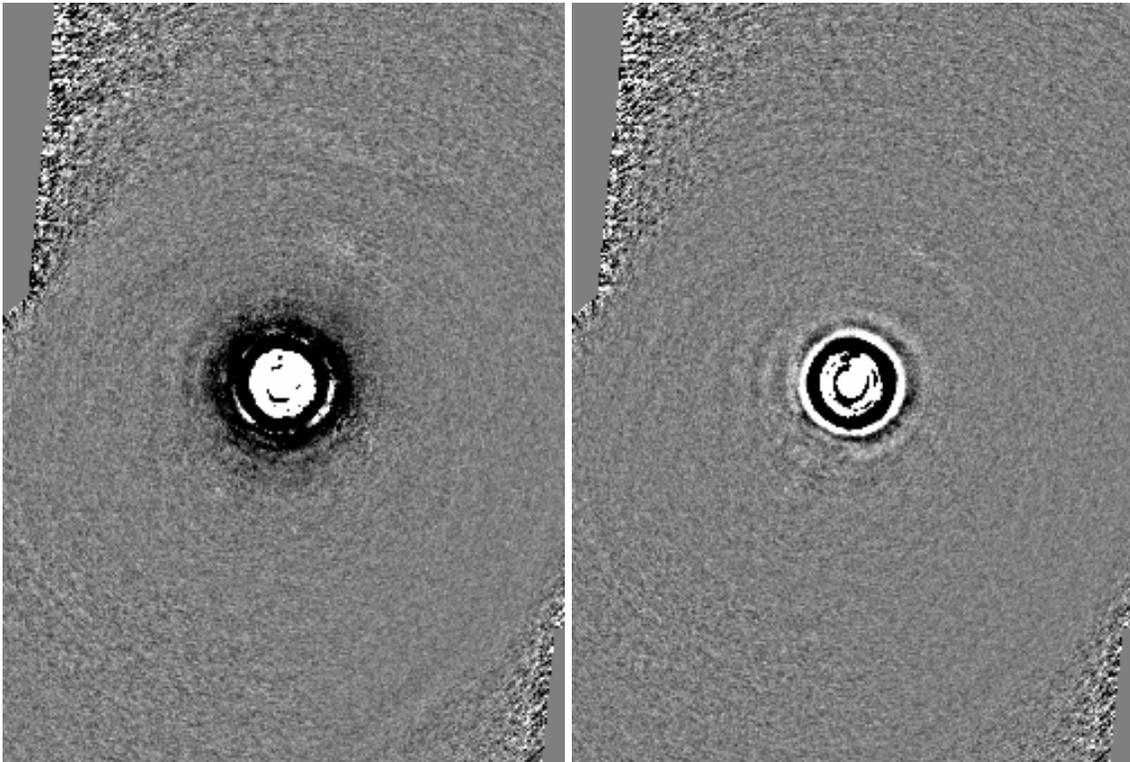


Figure A.32 Example of how the 'y' processing method produces a cleaner image than the 'x' method. On the left is our 'x' method master image of the star GJ 211. On the right is the image made by processing the same data with the 'y' method. The stretch on these images is -10 to +10 ADU.

PSFs are not confused, and similar in brightness, the PSF of each star can be used to subtract the other.

This is a very powerful method of PSF subtraction because it can be done on an image by image basis, rather than averaging over many images as is the case for the 'b' reduction method. The PSF of both stars on a given image is determined by the residual wavefront error from the AO system integrated over the exposure of that particular image. The *isoplanatic angle*, or angle over which the wavefront aberration from the atmospheric seeing is approximately the same, is larger than the Clio field at the  $L'$  and  $M$  bands. For this reason the light from the two stars has very similar residual wavefront error, and therefore the PSFs of the two stars are very similar.

Point sources are not at all dimmed in this method of PSF subtraction, since the PSF from a different star is used. The only way a real faint companion could be subtracted away is in the vanishingly improbable case that the other star had a companion with exactly the same relative position and brightness.

We have tried different methods of binary star PSF subtraction, but the one we have settled on in the end operates on the individual processed images output by an 'a' method reduction. Each 'a' reduced image is processed in turn. The locations of the two stars are determined by the usual centroiding procedure. Then the PSF of each is shifted over the other, scaled in brightness to match it, and then subtracted. We use an annular scaling region exactly as for the 'b' method reduction.

Rather than the standard 'b' method outer subtraction radius of 75 pixels, faded to 0% subtraction at 85 pixels, the outer radius for PSF subtraction for the 'f' reduction method must be set each time by the user. This is because the optimal radius depends so strongly on the brightness and separation of the two stars. For

example, if the stars are separated by less than 75 pixels, attempting to subtract the PSFs out to 75 pixels is foolish and will result in negative images of each star on the opposite side of the other. Different outer subtraction radii can be set for the primary and secondary stars.

After binary PSF subtraction has been performed on each of the individual 'a' processed images, the altered images are stacked with the usual creeping mean combine to produce the master 'f' method image.

We find that the 'f' reduction method usually produces extremely clean, well subtracted images of the secondary star. The subtraction of the primary star is generally good at very small separations, but at larger separations the sky noise from the secondary PSF, which had to be scaled up to match the brightness of the primary, becomes a dominant source of noise. The 'f' reduction method typically delivers excellent sensitivity to faint sources at a wide range of separations from the secondary, and to sources at small separations from the primary. If the stars are of equal brightness, of course, this does not apply. If the secondary is so much fainter than the primary that scaling it up for subtraction would introduce too much noise, it is possible to subtract only the secondary. In cases where the brightness of the two stars is significantly different, we often run 'b' method reduction on the primary as well as using the 'f' method. The two types of PSF subtraction are then complementary, with 'f' giving the best results at the smallest separations and 'b' working best at larger radii.

#### A.12.8 Combining Methods: the 'g' Reduction Method

The 'g' reduction method is a combination of the 'f' and 'd' methods: after the binary PSF subtraction, the images are subjected to pre-stack unsharp masking. They are then stacked in the usual way to produce the master 'g' method image. There is usually not much difference between the 'f' and 'g' master images.

### A.13 Conclusion

We have developed a very sophisticated and flexible image processing pipeline. Others faced with difficult image processing challenges, especially with astronomical AO images, may profit from adapting some of our algorithms and ideas for their own use. The following are likely to be the most use:

First, our method of removing single deviant pixels. We have shown that it has no effect on point-source images in well sampled data with a PSF FWHM of 3 pixels. It very effectively removes the majority of bad pixels. Our analogous method for removing single bad columns is also very effective.

Second, our method of performing a shift, rotation, zero trim, and zero pad in a single operation using bicubic spline interpolation. This method is remarkable for its ability to perform arbitrary shifts and rotations with no edge effects, no data loss (other than the intended trim regions), and minimal blurring.

Third, the creeping mean combine. It shows excellent ability to reject completely deviant chunks of image data that bias other stacking methods.

Fourth, our unsharp masking method, with its aggressive sigma-clip to avoid dimming point sources and introducing dark halo artifacts. Unsharp masking is a powerful and widely used technique. Our implementation removes or reduces some of its key drawbacks.

Other potentially useful innovations include our use of pre-stack unsharp masking to mitigate the effect of certain kinds of artifacts; the 'c', 'b', and 'f' methods of PSF subtraction; and the colfudge algorithm. We note that methods similar to our 'b' and 'f' PSF subtraction schemes have been used elsewhere (Marois et al., 2006; Kasper et al., 2007).

The colfudge algorithm represents a weakness in our processing method, which unfortunately was not identified until most of our data analysis was complete.

The key weakness is that it dims point sources. The dimming averages 15-18% in clear sky, but 25-30% may not be uncommon at unfavorable locations close to the primary star, and dimming as high as 50% may be possible in unlikely, 'maximally unlucky' locations in some data sets. In regions even closer to the star, the colfudge threshold is exceeded sources are not dimmed at all. Point sources sufficiently bright to appear clearly on individual images pre-stack are also not dimmed.

The undesirable characteristics of colfudge are not disastrous. Fairly extensive blind sensitivity tests have confirmed that almost all  $10\sigma$  sources and a majority of  $7\sigma$  sources are confirmed in our data. This despite the fact that colfudge and other effects typically reduce the nominal  $10\sigma$  and  $7\sigma$  sources to  $8\sigma$  and  $5.6\sigma$ , with significantly more reduction possible in some cases. Although we will perform all future processing of Clio data using Matt Kenworthy's column correction algorithm rather than colfudge, algorithms based on colfudge may be useful to other observers dealing with recalcitrant pattern noise in their images.

There is certainly room for improvement in the Clio pipeline. Besides replacing colfudge with the 4-column algorithm, we would like to implement some version of the LOCI algorithm (Lafrenière et al., 2007a) to see if it can deliver better PSF subtraction than the less sophisticated 'b' method reduction we describe above.

## APPENDIX B

## A MANUAL FOR THE LEGOLAS04 CLIO IMAGE PROCESSING PIPELINE

**Clio Data Reduction Manual***By Ari Heinze*

*Manual for using the legolas04 pipeline and related programs to analyze data from the Clio IR camera, with special emphasis on analyzing data to detect extrasolar planets. The software described here was mostly written by Ari Heinze, but Suresh Sivanandam made important contributions. Numerical Recipes in C subroutines were used extensively.*

**B.1 Disclaimer**

I have found lack of thoroughness a much more common failing in manuals than excessive detail, and so I have erred on the side of being thorough. I apologize in advance, therefore, for anything herein which seems irrelevant and/or an insult to the reader's intelligence. The reader is authorized to skip at will...

**B.2 Taking the Data****B.2.1 Planet Detection Data**

This is supposed to be a manual for data reduction, but it wouldn't be complete without a note at the beginning on what the data are expected to look like. The Clio detector has a high read noise, so in order to obtain background-limited sensitivity the background level should be more than half the detector full well capacity. Two thirds to three quarters of the full well are good levels to aim for. At typical operating conditions the bias level on raw, single frame images is about 10,000 DN, and saturation is about 50,000 DN, so the dynamic range is 40000

DN. Two thirds of full well is then about 36,700 DN, and three quarters is 40,000 DN. For best sensitivity to point sources against the background, the single-frame exposure time should be set to a value that results in background levels in this range of DN. This usually results in nominal exposures of 1200-2000 msec in the  $L'$  band, and 100-200 msec in the  $M$  band. **Note that the true exposure time for Clio is the nominal exposure time plus about 59.6 msec.**

The procedure recommended above for choosing the exposure time will result in saturated star images for most stars brighter than, say  $L' = 8.0$  or  $M = 6.0$ . We find that this is not a problem at all for centroiding purposes: even star images with saturation radii of twenty pixels can serve as accurate centroiding beacons resulting in good, sharp images of nearby faint objects (the most stringent test of this so far has been imaging the white dwarf companion of Procyon, but faint background companions to the nearby, bright K stars 61 Cygni A and B also gave a good example of highly accurate centroids based on a severely saturated primary image. If sensitivity is desired near a bright star, then large saturation radii are of course unacceptable. In some cases shorter exposures with more coadds may be a good choice, since stellar speckles rather than read noise will be the limiting factor close to the star. However, another, likely better, option is to use the  $M$  band rather than  $L'$  when good sensitivity near a bright star is desired. If, as is nearly always the case, the faint objects being sought are much cooler than the primary star, there is likely to be a contrast advantage in  $M$  over  $L'$ , and there is always a significant Strehl advantage and a much smaller saturation radius at  $M$ . If the  $M$  saturation radius is still too large, shorter exposures are the only remaining option for straight-through imaging, but the phase plate coronagraph being developed for Clio by Matt Kenworthy and others may be a better choice.

The number of coadds should be set so that the cumulative exposure time in a

single image is about 10-30 sec. This means typically 10-20 coadds in  $L'$  and 100-200 coadds in  $M$ . More coadds means fewer final data files and faster processing. However, it also means more hot pixels in each raw image, since the number of hot pixels increase with cumulative exposure time. Also, it reduces flexibility and means more time will be lost in the case of AO loop crashes, brief clouds, or other short interruptions in the data sequence. Note that the Clio software does literally co-add individual frames to produce the raw images it gives you, so if you have coadds set to 100,  $4 \times 10^6$  DN on your raw images corresponds to 40,000 DN on the individual frames, and means you are still in the acceptable range for background counts, albeit at the high end of this range.

Once the exposure time, filter, and coadds are chosen, the nod amplitude and direction for Clio must be selected. The standard procedure is to set a nod amplitude of 5.5 or 6.0 arcseconds, with the direction along the long dimension of the detector. This provides the largest obtainable clean, high-sensitivity area around the target star in the final images. In cases of double stars where good images of both components are desired, or in cases where imaging of a faint companion of specific known location is being attempted, a different nod choice may be required. Once the nod direction and amplitude have been chosen, we do not recommend changing them in response to parallactic rotation. Keeping them constant tends to lead to better elimination of ghosts and other artifacts in the final processing. If a double star is being observed, care should be taken to set the initial nod direction so that it remains as close to perpendicular to the vector between the stars as possible throughout the parallactic rotation expected during the observation. Generally, the system is set to take either 5 or 10 images in each nod position each nod cycle. Taking fewer results in more overhead, as it takes time to move the telescope. Taking more increases efficiency but could reduce the

quality of sky subtraction, and more importantly reduces flexibility and results in more time lost in case of AO loop crashes or other short interruptions in the data sequence.

Typically, we try to obtain at least one hour of integration time on each star. This can usually be accomplished with one and a half hours of actual observations. If time is available, observing a star for two hours or even more does result in significant improvements to the sensitivity. The exposure time, number of coadds, and, as mentioned above, nod amplitude and direction should be constant for the whole observing sequence if at all possible (that is, for the images that are to be used for faint source detection against the background. There is no problem with interspersing shorter exposure for different purposes, as we recommend below.) The image rotator should not be used and should be left at a setting of 0 degrees if possible.

Clio's background limited sensitivity to astrophysically realistic cool objects is usually better at  $L'$  than at  $M$  band. According to the planet models of Burrows et al. (2003), this situation reverses for stars closer than about 4-5pc. According to the Baraffe et al. (2003) models,  $L'$  gives better background-limited sensitivity even for the nearest stars. However, both model sets agree that  $M$  band often gives better sensitivity than  $L'$  for contrast-limited observations close to bright stars. It follows, therefore, that for faint stars  $L'$  data alone are sufficient, unless the stars are very nearby. For stars of intermediate brightness  $L'$  data are desirable for sensitivity to companions at large separations, while  $M$  data are good for close-in companions. For very bright stars, rays and other artifacts can prevent Clio from obtaining background limited sensitivity at  $L'$  except at very large separations, so  $M$  band observations should be prioritized over  $L'$ . As a rough guide, we suggest  $L'$  band observations alone are sufficient for stars fainter than

about  $L' = 3.0$  and more distant than 5pc; for stars brighter and/or nearer than this up to about  $L' = 1.0$ , both  $L'$  and  $M$  data are likely to be useful, and for stars brighter than  $L' = 1.0$ ,  $M$  data will likely be significantly superior to  $L'$ .

### B.2.2 Monitoring Data Quality

As mentioned above, the exposures we suggest herein will saturate the PSF cores of all but fairly faint stars. Again, as we have already said, this is not a problem for detection of planets. Good centroids can be obtained from saturated stellar images, and the images can be properly registered and stacked to reveal any faint companions. However, it is **very desirable** to have accurate knowledge of the true PSF obtained during a given science observation, and of any variations in atmospheric transparency during the observations. The former is useful as a template for source detection, and indispensable for deriving the appropriate aperture correction for photometry of any faint sources that are detected, while the latter is the only reliable way to verify photometric accuracy. For these reasons, if the science exposures saturate the primary star, interleaving a smaller number of much shorter, unsaturated exposures is **highly recommended**.

A reasonable policy is to use short exposures on every fifth nod cycle, without changing the coadd number or the number of images per nod. The overhead involved in this is small, since the short exposures are, well, short. Exposure times to obtain unsaturated data can range from several hundred msec down to 10 msec for very bright stars. For the brightest stars, even the shortest possible exposures may saturate (the Clio readout method always adds 59.6 msec to the nominal exposure time, so a nominal 5 msec exposure is really 64.6 msec, and nominal exposures shorter than 5 msec are not possible). If a 5 msec exposure saturates, we recommend using the neutral density filters in Clio's second filter wheel to obtain unsaturated exposures. We warn the reader that **the Clio detector has**

**significant nonlinearity within the first 1000 ADU above the bias level. This can bias photometry based on very short exposures.** For longer exposures the thermal background and/or the dark current raise the background above 1000 ADU. If a neutral density (ND) filter is being used to get unsaturated images of a star, **it is a good idea to choose one that produces good counts with a fairly long exposure, say 250-500 msec. This way the dark current will be sufficient to lift the image above the nonlinear regime.**

If for some reason non-saturated images of a particular star were not obtained, the unsaturated PSF from another star can be used for source detection and aperture calibration of the saturated star images. Hopefully, a suitable star for this will always be observed not too distant in time from the observations of a too-bright science target. A suitable star would mean one observed in the same filter, at similar airmass, and similar angle to the wind if the wind is not negligible.

### B.2.3 Darks

Bias and dark current can change significantly over a night, so taking new darks for each science data set is highly recommended. Without this, it will not be possible to get good sky brightness measurements, and some unnecessary noise may be introduced into the data. The darks must be taken at the same exposure times as the science images, but 5-10 darks at each exposure time should suffice. As an example, suppose you have observed a given star in the  $L'$  band for 2 hours, acquiring 25 nod-cycles with 5 images in each position using 2000 msec exposures and 10 coadds, thus totalling 250 images of 20 second coadded exposure each. The total integration time is 5000 seconds. In addition, you have taken 5 nod cycles of 5 images in each position using 100 msec exposures and 10 coadds, to obtain unsaturated data for a PSF template and atmospheric transparency monitoring. As you are slewing to the next star, you should simply take 5 dark frames

of 2000 msec exposure and 10 coadds, and 5 more at 100 msec exposure with 10 coadds. Such a sequence of darks is adequate and takes only a few minutes. In fact, there is sometimes enough time to take darks while the scope is slewing and then centering the next target, without any loss of sky time. If there is loss of sky time it is usually only 1-3 minutes.

#### B.2.4 Flats

The legolas04 pipeline can be run without flats, and the rather surprising discovery that this produces decent results has led observers not to invest much time and effort in producing flats. Using flats in all pipeline reductions is highly recommended, but it is not necessary to take new flats every night, nor to take flats at specific filters or exposure times. Once every third night is probably sufficient frequency for Clio flats, and a single flat may be made and then used for all data on adjacent nights. A reasonable procedure for taking flats is to take a series of exposures of the empty twilight or night sky at the same exposures used for science images, and then acquire appropriate darks immediately afterwards. A master dark image can then be made and subtracted from each flat frame, and the flats may then be normalized and combined – in fact, this is almost exactly the procedure one would use to make an optical flat. The programs zerocheckm01 and flatcheckm01 are part of the legolas04 suite, and are used to make darks and flats, respectively, using a slight optimization of the above procedure.

#### B.3 Introduction to Data Reduction

We will now discuss each stage of a Clio data reduction sequence. The input parameters and operations of the programs used in each stage will be described. Instructions on how to make good parameter choices will be given.

All of the programs in the Clio processing suite are written in the C program-

ming language. In most cases they are run on the command line with an input parameter file using the syntax:

```
./program.exe < param_file
```

The instructions we give here will largely relate to how the parameter file should be constructed. **I have become accustomed to calling parameter files such as the one in this example 'scripts' I have been told this is not correct terminology. However, herein 'script' shall mean any file input to (as above) or otherwise supplied to a program giving it the parameters by which it is to run. My apologies to sticklers for a precise nomenclature.**

#### B.4 Data Reduction: Darks

The program zerocheckm01 can be used to make all of the darks from one night in a single run. It is run on the command line with a script, like this:

```
./zerocheckm01.exe < zcm01script01
```

where zcm01script01 contains the names of raw images, the desired names of master darks to be constructed, and the parameters to use in their construction.

Pasted below is an example of a zerocheckm01 script. Table B.1 explains each entry:

```
320 258
2
10 310 10 248
5
GJ3860dark00001.fit
GJ3860dark00002.fit
```

GJ3860dark00003.fit  
GJ3860dark00004.fit  
GJ3860dark00005.fit  
1  
0.4 GJ3860darkamain.fits  
1  
15.00  
10  
darkGJ820A00001.fit  
darkGJ820A00002.fit  
darkGJ820A00003.fit  
darkGJ820A00004.fit  
darkGJ820A00005.fit  
darkGJ820A00006.fit  
darkGJ820A00007.fit  
darkGJ820A00008.fit  
darkGJ820A00009.fit  
darkGJ820A00010.fit  
1  
0.4 GJ820Adarkamain.fits  
1  
20.00

Table B.1. Explanation of zerocheckm01 script

Entry	Explanation
320 258	x and y dimensions of images in pixels in pixels (this never needs to be changed).
2	number of master darks that will be made.
10 310 10 248	minimum and maximum x, and minimum and maximum y pixel values for the region to be used for statistics. Here, the whole image is included except for a 10-pixel boundary.
5	number of individual images to combine to make first master dark.
GJ3860dark00001.fit	name of image 1 for first dark.
GJ3860dark00002.fit	name of image 2 for first dark.
GJ3860dark00003.fit	etc...
GJ3860dark00004.fit	etc...
GJ3860dark00005.fit	etc...
1	1 means yes, go ahead and make a master dark from these data.
0.4 GJ3860darkamain.fits	fraction of data that will be rejected, and name of master dark.
1	1 means yes, divide the master dark by a constant.
15.00	The constant to use. Should be the number of coadds, so the resultant dark will be normalized to one coadd.

Table B.1—Continued

Entry	Explanation
10	number of images to combine to make the second master dark.
darkGJ820A00001.fit	listing images.
darkGJ820A00002.fit	etc...
darkGJ820A00003.fit	etc...
darkGJ820A00004.fit	etc...
darkGJ820A00005.fit	etc...
darkGJ820A00006.fit	etc...
darkGJ820A00007.fit	etc...
darkGJ820A00008.fit	etc...
darkGJ820A00009.fit	etc...
darkGJ820A00010.fit	etc...
1	Yes, make master dark.
0.4 GJ820Adarkamain.fits	rejection fraction and name for the second master dark.
1	yes, divide by a constant.
20.00	the number of coadds.

The only judgment calls in this script are the statistics region (which does not actually affect the final outcome) and the rejection fraction. What does the rejection fraction mean?

When stacking images to make a combined master image, the code uses creeping mean rejection on a pixel-by-pixel basis. What this means is that, for the second master dark made from 10 individual images, it has 10 values for each pixel. For a given pixel, it takes these 10 values, finds their mean, and rejects the one farthest from the mean. Then, it finds the mean of the remaining 9, and rejects the one of them that is farthest from the new mean, and so forth, until it has rejected the required fraction of the values – in this case, 0.4, or 4 values for the stack of 10 images. The mean value after the last rejection is what goes into the final image. Note that this is not a smoothing of the images; it is performed purely on a pixel-by-pixel basis and its sole purpose is to reject hot pixels. These are numerous, and rejecting them is worthwhile. Therefore the rejection fraction should usually be set to reject at least one of the values for each pixel. For example, if you have 5 images, the rejection fraction must be at least 0.2 or you are just doing a straight average. In general it is better to reject at least 2 data points, but you should never reject more than half your data. Rejection fractions above 0.5 no longer eliminate bad data and simply increase the noise, since you are averaging over a smaller number of 'surviving' data points.

When run, `zerocheckm01` uses the defined statistics region to print out a bunch of data on the mean of the darks and the stability of this mean. Frequently, the user doesn't have time to check all of this information. If you do, note that the range on the means of the different images should be less than about 1% of the average mean. If it's not, something may have gone wrong, and if it's 5% or more, of course something has gone wrong and your dark sequence is no good. The

most likely cause of this would be that you have inadvertently combined darks from different exposure times and/or numbers of coadds.

Note that although the script given here made only two master darks, it is perfectly acceptable and efficient to make 10 or even more master darks at a time using `zerocheckm01`: all the master darks for a whole night can be made in a single execution.

## B.5 Flats

The program for making flats is `flatcheckm01`. It is very similar to `zerocheckm01`, and is run in the same way. Here is an example of a `flatcheckm01` script. The entries are explained in Table B.2:

```
320 258
1
10 310 10 248
1
bpfile03
1
20.0
5
GJ820Adarkamain.fits
flat1200a00001.fit
flat1200a00001e.fits
flat1200a00002.fit
flat1200a00002e.fits
flat1200a00003.fit
flat1200a00003e.fits
flat1200a00004.fit
```

flat1200a00004e.fits

flat1200a00005.fit

flat1200a00005e.fits

1

0.4 flat1200lmain.fits

Table B.2. Explanation of flatcheckm01 script

Entry	Explanation
320 258	Pixel x and y dimensions of images. Never needs to be changed.
1	The number of master flats to construct
10 310 10 248	min x and max x and min y and max y in pixel coordinates for a box in which statistics will be calculated, and over which the flats will be normalized.
1	Yes, perform bad pixel correction.
bpfile03	Name of bad pixel file. This file was made in April 2006, and seems perfectly satisfactory still.
1	Yes, divide flats by a constant.
20.0	The constant. <b>Must</b> be the number of coadds, or zero subtraction won't work properly
5	The number of individual images to combine to make the first (and in this case the only) master flat.
GJ820Adarkamain.fits	The name of the master dark image that applies to these flats.
flat1200a00001.fit	The name of the first raw flat image.
flat1200a00001e.fits	The name by which to call an edited output version of the first flat image.
flat1200a00002.fit	The name of the second raw flat image.

Table B.2—Continued

Entry	Explanation
flat1200a00002e.fits	The name by which to call an edited output version of the second flat image.
flat1200a00003.fit	etc...
flat1200a00003e.fits	etc...
flat1200a00004.fit	etc...
flat1200a00004e.fits	etc...
flat1200a00005.fit	etc...
flat1200a00005e.fits	etc...
1	Yes, go ahead and combine to make a master flat.
0.4 flat1200lamain.fits	Rejection fraction and name of master flat.

As in the case of `zerocheckm01`, the only judgement calls for a `flatcheckm01` script are the statistics box and the rejection fraction. Unless the Clio detector gets dirty or I burn it with a laser again, the statistics box given here, which includes all but a 10 pixel border around the edge of the image (really 9 pixels at left and bottom and 10 at right and top, since I have made the obvious fencepost error) should be perfectly satisfactory. The discussion of the rejection fraction in the previous section applies without alteration. The edited versions of the flats output by `flatcheckm01.c` are dark-subtracted, normalized, and bad-pixel-fixed.

Note that although for reasons of space efficiency I show a `flatcheckm01` script combining only 5 individual flats, this is not recommended for real data. Since flats do not have to be repeated frequently, there is no reason not to beat down the noise by taking 10, 20, or even 50 to combine to make a single master.

## B.6 Overview of a `legolas04` Run

The main Clio pipeline program is called `legolas04`. If you are wondering where the name comes from, I believe it can be traced to the the fact that both Andy Breuninger and I are fans of J. R. R. Tolkein. Andy called his wonderful program for actually running the Clio detector `mirkwood`, after the forbidding forest in Tolkein's [The Hobbit](#). Carrying on this theme, I called my analysis pipeline `legolas`, after the keen-eyed elf who comes from his ancestral home in Mirkwood to become one of the heroes of Tolkein's [The Lord of the Rings](#) trilogy. It has been hypothesized by the Steward grad students that Roger Angel is still looking for the One Mirror to Rule Them All.

`Legolas04` is run on the command line without a direct parameter file, simply:  
`./legolas04.exe`

The program then queries the user for a file I call the parameter script. This

file contains all of the required analysis parameters. After reading the parameter script, `legolas04` asks for the image script, a file which tells the program what images it is dealing with: the total number of images in each nod beam, and the name, number of coadds, approximate coordinates of the target star, sequence number, and MMT instrument rotator angle for each image. This allows for the possibility that the number of coadds or the instrument rotator angle could be changed during an observing sequence. This is not recommended, but can be dealt with. Trying to combine images with different **exposure** times using `legolas04`, on the other hand, is a highly dubious enterprise that should probably never be tried. If multiple data sets with different exposures must be combined, `legolas04` should be run separately on the set of images taken at each exposure, and then the processed frames or master images can be combined with appropriate scaling.

After it has read the image script, `legolas04` asks the user for input that determines how it is going to pair up the images from the two different nod positions to carry out sky subtraction. Before describing this input, I will describe basically what `legolas04` does when pairing up images. This is under the assumption that, as recommended in the section on taking the data, you have images in nod sets, with the star in good position on the detector in both beams. Thus each image will serve once as the 'science' image, and a 'subtraction' image from the other beam will be chosen and subtracted from it, and each image will also likely be used once itself as a 'subtraction' image for a 'science' image in the other beam.

Suppose you have a very short data sequence, just 1 nod cycle consisting of 5 images in each nod position, or beam, so that you have 10 images with images 1-5 in nod beam one, and 6-10 in nod beam two. When processing the beam 1 data, `legolas04` must choose a subtraction image from beam 2 for each of the

beam 1 images, after which it will process the beam 2 images, pairing them with subtraction images from beam 1. Clearly, there are multiple ways to do this. Since the point of nod subtraction is to remove sky effects, and these may be variable, it is desirable to use the closest image available in a sequence. This would mean image 6 would be used as the subtraction image for images 1-5, and image 5 would be the subtraction image for images 6-10. But this has problems because the specific noise pattern on images 6 and 5 will get strengthened in the final stack. Trying to make the noise as independent as possible is better. So one could pair image 1 with image 6, image 2 with image 7, image 3 with image 8, and so forth. It is necessary to have a general algorithm that will produce decent results with any kind of data set, even if the nod sets are of unequal length, or if there are gaps where some images are missing because of AO loop crashes or other problems.

The algorithm that legolas04 uses is fairly simple. When it is trying to find a subtraction image to pair with a given science image, it searches for an image in the appropriate beam that is no more than a user-specified distance away from the science image in the sequence, and that has not been used as a subtraction image before. If unused images exist in the specified range, legolas04 uses the nearest one available and then flags it as having been used once, so that re-using it can be avoided when pairing subsequent images. If, however, there is no unused image within the specified range, legolas04 will choose the nearest available image that has been used only once, and flag that as having been used twice. If there are no images within the range that have not been used twice, legolas04 will report that it was unable to construct a pairing set.

**Note: the next five paragraphs are mostly of theoretical interest and may be skipped if the reader is under time pressure to learn how to perform basic**

**processing with legolas04 as quickly as possible.** It is interesting to note that a given data set can be paired in many different ways. If the final, combined master images made from different pairings of the same data set had independent noise, they could be averaged together to make a final image with extremely low noise. At issue is the particular sum of noise that stacks on top of the location of a celestial object, say a planet. If this is independent from one pairing to another, averaging together many pairings can increase the SNR for detecting the object. In reality, the object always appears at the same place relative to the noise in the science images, so science-image noise does not average down in an average of different pairings. The noise in the subtraction images, however, may, and this still means in principle the noise could be reduced by a factor of  $\sqrt{2}$ . Legolas04 offers the option of taking advantage of this possibility. The user can request that any number of independent pairings be constructed and averaged together (though it will take 10 times as long to do 10 pairings!). Is this option any use?

The condition for the noise actually to be reduced is that the location of the planet **not** be consistent in pixel coordinates from one image to another. For example, suppose in pairing 1 that image 6 is used as the subtraction image for image 1, and image 7 is used for image 2. But in pairing 2, image 7 is used as the subtraction image for image 1, and image 6 is used for image 2. If the planet is at exactly the same pixel coordinates on image 1 and 2, exactly the same noise will be superimposed on it in both pairings. But suppose it is not in the same place on the two images. Say it is at location 1 on image 1 and location 2 on image 2. Then in pairing 1, the noise that falls upon it from the subtraction images will be the noise in image 6 at location 1 plus the noise in image 7 at location 2. In pairing 2, on the other hand, the noise will be the noise in image 7 at location 1 plus the noise in image 6 at location 2. Since the noise we are talking about is

photon noise, it should be entirely spatially uncorrelated, even from one pixel to the next. So if the planet moves even one pixel between images, the noise may be reduced in an average of many pairings.

Is this condition satisfied in real data? Well, not necessarily in general. For images within a given nod set, the MMT tracking holds celestial objects at a constant location to a much better accuracy than one pixel — as indeed we should hope! Because the AO loop must be opened in order to shift the telescope, and opening and closing it takes time, it is not practical to dither the telescope every image — this would probably take up more telescope time than the noise reduction could possibly be worth.

There are three other points, however, worth considering. First, if more than a few pairings are tried, `legolas04` will have to vary pairing by many images in order to avoid duplicate pairings from one set to another (and it does absolutely avoid duplicates, or else report failure to pair). Second, for objects that come near the zenith and are observed across the meridian, parallactic rotation can change the pixel coordinates of planets rapidly. Third, while it is impractical to dither between every image in a nod set, it is quite practical to change the zero positions from one set to another, and in fact Andy's program for running `Clio` explicitly supports this. If a set consists of only 5 or fewer images in each beam, when directed to construct multiple pairings `legolas04` will quickly be forced to span sets, and the noise in different pairings will be partially independent.

The conclusion, then, is that the option of asking for multiple pairings may be useful if the data set either has a great deal of parallactic rotation, or has been taken with slight pointing offsets after every nod. All of this is at present theoretical — it has not been well tested with real data. Also most `Clio` data have not been taken with any care to vary the zero position between nod sets. Certainly,

observing this way can hardly hurt and may well help, so it is recommended. The multiple pairing option, however, is certainly too slow to be useful in quick look reductions at the telescope. If it is used at all, it should be for final reductions later on.

**Here end the theoretical paragraphs.**

So, with that long introduction, here's what legolas04 asks the user after the name of the image script has been entered:

“Enter the number of distinct pairing sets you want to make.”

This is the option to take advantage of the theoretical possibility described in the paragraphs above and try multiple pairing sets. In general, you will only want one pairing set, so enter 1. For time-critical reductions at the telescope, you should always enter 1. Later on, if you think the data will support it and you have the time, trying something like 10 or 15 is always an option. Maybe you will be the first person to find out it really helps.

After this legolas04 says:

“Enter the maximum distance the nearest image must be away before doubling pairs”

This is the user-specified distance described above. For example, if you enter 4 it will be able to pair image 55 with image 59 if they are in different beams, but not image 55 with image 60 because they are 5 images apart in the sequence and you have told it the maximum distance is 4. Try something and hit enter. The full size of a nod cycle plus 3 is a good choice. If you use 5 images in each beam per nod cycle, the full size of a nod set is 10, so try 13. Legolas04 will attempt to come up with a single pairing set (or many independent pairing sets if you've asked

for them) with this constraint. If it can do so without using any of the subtraction images twice, it will proceed straight into processing. If it can do so, but it has to use some subtraction images twice, resulting in non-independence of the noise, it will warn you and let you know how many images were used twice. If it's a small fraction of the total, you can enter 1 and processing will proceed. If it's a large fraction of the total, enter 0 and you will have a chance to set the doubling distance to a larger value to reduce the number of images that are used twice. If legolas04 can't come up with the required number of independent pairings even by doubling images (it will never triple them), it will let you know and you can enter a larger doubling distance and try again.

In a data set with 5 images per nod position per cycle, with no gaps, with the doubling distance set to 13 it will always successfully pair with exactly 3 double pairings. For a typical data set of more than 100 images, this is negligible.

If all this seems complicated, don't worry. When actually running legolas04 you don't have to think about what's really going on. Just enter 1 for the pairing sets, start with something like 13 for the maximum distance, and then try different maximum distances until you get one that results in fewer than 5-10 percent of the images being double-paired.

## B.7 Making the Parameter Script

Legolas04 is a very flexible program, with many options mostly controlled by the parameter script. Flexibility can be a drawback if it implies that many judgment calls must be made and many parameters set before every analysis. To avoid this, there two are auxilliary programs that, with a minimum of input from the user, automatically write parameter scripts of 7 different typically useful forms for each star. Numerous parameters are set to reasonable default values so the

user doesn't even have to know about them. A single run of one of these programs can be used to generate scripts for all the stars in a given night.

The most stripped down, user-friendly parameter script writing program is `legoparswrite04`. A slightly more complicated but more flexible version is `legoparswrite03`. It differs from `legoparswrite04` in only two ways. First, it provides the potential for slightly more accurate position angle astrometry on binary stars or faint companions. Second, it allows the user the option of correcting the column pattern noise in the images with our old 'colfudge' algorithm. The simpler program, `legoparswrite04`, always uses Matt Kenworthy's four-column algorithm for this purpose, because it is generally more powerful and does not dim point sources or introduce artifacts.

We first discuss `legoparswrite04`, the simpler program.

As is usual with our suite of processing programs, `legoparswrite04` is run on the command line, with a script. (That's right, a script to write a script):

```
./legoparswrite04.exe < lpw04script01
```

An example of a `legoparswrite04` script is given below:

```
2
1
272.867400
2006 06 11
02
VegaM
VegaMdarkamain.fits
flat1200lamain.fits
220 310 10 248
```

```
20.0 17.0 15.0 15.0 15.0
18 36 56.3364 38 47 01.291
0.1796
3308.0
9.0 16.0
GJ519
GJ519darkamain.fits
flat1200lamain.fits
10 310 178 248
15.0 12.0 10.0 8.0 8.0
13 37 28.7684 35 43 03.947
1.5596
13135.0
5.5 11.0
```

The entries in this script are explained in Table B.3.

Table B.3. Explanation of legoparswrite04 script

Entry	Explanation
2	Number of stars for which scripts should be written
1	1 means raw images should be mirror-flipped in processing; 0 means they should not. For data from July 06 or earlier, use 1; for later data use 0.
272.867400	Clockwise rotation in degrees that must be applied to images (after mirror-correction if applicable) to get the +y axis of the image to line up with the +EL (up) direction referenced to the real sky, and, of course, the telescope mount. Values for this are documented on the Clio wiki.
2006 06 11	UT Date of the observations in yyyy mm dd format.
02	A numerical string to identify these scripts. Convention: 01 means short, PSF-monitoring exposures, 02 means long science exposures.
VegaM	The name of star 1. Will identify the scripts and the final output images.
VegaMdarkmain.fits	The dark frame for this star.
flat1200lamain.fits	The flat frame for this star.
220 310 10 248	The minimum and maximum x, and minimum and maximum y pixel values defining a

Table B.3—Continued

Entry	Explanation
20.0 17.0 15.0 15.0 15.0	<p>box for scaling images during subtraction.</p> <p>Five centroiding radii for the 5 iterations of centroiding that will be performed on this star.</p>
18 36 56.3364 38 47 01.291	<p>Epoch 2000.0 The RA and DEC of this star in hh mm ss.sss dd mm ss.ss format. The number of decimal places on the seconds is irrelevant.</p> <p><b>If the DEC is negative, all 3 DEC entries must carry the minus sign, ie -05 -39 -18.484</b></p> <p>Note also that the proper motion in Simbad units (mas per year in RA and DEC) <b>may</b> be added after the coordinates. If no proper motion is supplied, zero proper motion is assumed.</p>
0.1796	<p>The true single-frame exposure time in seconds for this star.</p>
3308.0	<p>The photometric calibration in the form of counts/sec for a star of magnitude 10.0 within a 30.0 pixel aperture. Typical values for this can be found on the Clio wiki.</p>

Table B.3—Continued

Entry	Explanation
9.0 16.0	The inner and outer radii for the scaling annulus for PSF subtraction on this star.
GJ519	The name of star 2.
GJ519darkamain.fits	The dark frame for this star.
flat1200lamain.fits	The flat frame for this star.
10 310 178 248	min and max x, and min and max y pixel coords for the scaling box for this star.
15.0 12.0 10.0 8.0 8.0	Five centroiding radii for this star.
13 37 28.7684 35 43 03.947	Epoch 2000.0 RA and DEC for this star.
1.5596	True single-frame exposure time for this star.
13135.0	Photometric calibration.
5.5 11.0	Inner and outer radii for scaling the PSF subtraction on this star.

The above is a script for long exposure, science images of two stars. Vega was observed at *M* band using a nominal exposure of 120 msec. Adding the 59.6 msec read time we obtain the true exposure from the script, 0.1796 sec. GJ 519 was observed in the *L'* band with a 1500 msec nominal exposure time, corresponding to a real exposure of 1.5596 sec. As you may have noticed, the parameters up through the numerical string identifying the type of script (02) are general parameters, applying to all observations of a given type throughout the night. Two blocks of parameters specific to two different stars follow.

The judgment calls in this script are the radii for the centroiding iterations, and the PSF scaling annulus.

The scaling region takes some care. When performing a nod subtraction, `legolas04` scales the subtraction image by the ratio of the mean within this box on the science image to the same mean on the subtraction image, to compensate for any variations in sky brightness, and to obtain images with mean zero in regions of clean sky. Although `legolas04` attempt to reject stellar artifacts from the scaling calculation, the scaling region should be as clean as free as possible of ghosts, rays, and stellar haloes, for all the images in the data set. The best way I have found to choose the scaling region is to select three to five nod pairs throughout the observing sequence, subtract the raw images in IRAF and look at the results. For example, if you have 200 images taken in sets of 5 images in each nod position, images 5 and 6 are in different beams, as are 75 and 76, 135 and 136, and 195 and 196. You could subtract each of these raw image pairs, and then display the results with a high stretch, so the bright and dark stellar haloes look huge (except in the case of very faint stars), considerable noise can be seen in the background, and any faint rays or ghosts are readily visible. Simply choose as your scaling region the largest box you can find that is well clear of both bright and dark stellar

haloes, ghosts, and rays on all images in the representative set you have chosen.

The pixel radii for the centroiding iterations aren't a sensitive parameter, in my opinion. I usually start at 15 pixels for faint stars and 25-35 pixels for bright stars, and then pull in by about a factor of two throughout the iterations. I recommend setting the radius for the last iteration no smaller than twice the saturation radius. If the star is unsaturated and bright, setting the radius of the last iteration to 6.0 pixels is probably good.. If it's faint, going as low as 3.0 pixels may be a good option, to insure accurate centroids through the sky noise.

Setting the PSF scaling annulus is easy: simply put the inner edge one or two pixels beyond the maximum saturation radius, and put the outer edge 5-6 pixels beyond that. In this example, the saturation radius was about 7.0 pixels for Vega and 4.0 pixels for GJ 519. For very bright stars with a large saturation radius you may want to use a wider scaling annulus. Of course, the inner edge of the annulus must always be beyond the saturation radius.

If you don't have a photometric calibration, entering any random number for the photometric calibration is OK. The only thing that will be affected is the photometry in a data file output at the end of the processing, and if your star images were saturated, that photometry was going to be wrong anyway.

The output of legoparswrite04 is a set of 7 legolas04 parameter scripts. The names of all of these scripts began with **legolas04**, followed by the star name supplied in the script, followed by the numerical string identifying the type of data, followed by a letter identifying the type of script. For example, the name supplied for the second star in the above script was **GJ519**, and the numerical string was **02**, so the name of the first script will be **legolas04GJ51902a**.

The letters identifying the 7 scripts that are written are **a, b, c, d, e, x, and y**. Thus the full list of legolas04 parameter scripts that legoparswrite04 would make

for the star GJ519 in this example is legolas04GJ51902a, legolas04GJ51902b, legolas04GJ51902c, legolas04GJ51902d, legolas04GJ51902e, legolas04GJ51902x, and legolas04GJ51902y.

What do these scripts do? The 'a' script is baseline processing, which includes mirror-flipping if requested, division by the number of coadds to normalize to a single frame, dark subtraction, flatfielding, bad pixel fixing, subtraction of a scaled image from the opposite nod position, further bad pixel fixing, column pattern noise correction, centroiding on the star image, zero-padding, shifting, rotation. The final processed frames have celestial North up with the primary star located at exact position 250.0,250.0 on a zero-padded image of dimensions 500x500 (note that the dimensions of a raw Clio image are 320x258).

Once all the science images have been processed this way, they are stacked using a creeping mean with 20% rejection. The resulting image is unsharp masked using a gaussian kernel with  $\sigma = 5.0$  pixels (FWHM 11.8 pixels), and with pixels of value 0.0 not considered in the masking so that there are no weird edge effects due to the zero padding. If you are worried about real data pixels randomly having a value of exactly zero, this is possible since the images are supposed to have mean 0.0 in areas of clean sky, but given the floating point precision of the c programming language and the typical rms about zero for the final science images, real data pixels with value exactly 0.0 should occur about as often as large astronomical telescopes are destroyed by meteor impact. Both the unsharp-masked and the non-unsharp-masked versions of the final image are output. The names of these images begin with the star name supplied in the parameter script, ie **GJ519**. The script letter follows, **a** in this case, followed by the numerical string, **02**. Next comes a **u** for the image that is unsharp masked, followed by a two digit integer showing which pairing this is, if multiple independent pairings have

been requested. Since this is almost never the case, this final integer is usually **01**. Thus for this example involving GJ519, a final non-unsharp-masked image called **GJ519a0201.fits** will be produced, along with an unsharp-masked version called **GJ519a02u01.fits**.

The 'b' script is exactly like the 'a' script except that it performs PSF subtraction to reveal close-in companions. In doing this, it takes advantage of the observed fact that on Clio images with MMTAO on a typical night, the point-spread function remains substantially constant over the 1-2 hour timescale of typical Clio observing sequences. During this time, while the PSF from the instrument and telescope is changing little, any real companions are rotating slowly about the primary star because of the parallactic rotation imposed by the Earth's rotation as the MMT tracks with its altazimuth mount.

Thus, an image stack may be made without taking this sky rotation into account, and the result will be a good average of the Clio+MMTAO PSF during the observing sequence, but any real companions, since they do move with the sky rotation, will smear out. If a 50 % creeping mean rejection is used, the companions should vanish, leaving a pure measurement of the system PSF. When run with the 'b' script, legolas04 splits the images into first and second halves, based on the UT values in the image headers, and makes such an unrotated PSF from both halves. It saves these PSF images, and goes on to the main science processing exactly as described for the 'a' script, except that as one of the final steps it subtracts a rotated, scaled version of one of the PSF images from each science image. In order to minimize the subtracted of even very close-in real sources, legolas04 uses the PSF image from the second half of the observing sequence for images from the first half, and the PSF image from the first half for images taken in the second half. The result typically improves the sensitivity to close-in companions

by 2 magnitudes over the 'a' script.

If the parallactic rotation during the observing sequence is insufficient, real close-in sources will be subtracted or dimmed by the PSF subtraction. Our rule of thumb is that the dimming should not be severe at radii larger than the minimum radius where half the total parallactic rotation in the data set is sufficient to move a real celestial source through an arc of at least one FWHM, or about 3 pixels. For example, if  $20^\circ$  of parallactic rotation has been obtained and you are wondering if the PSF subtraction would harm companions located 20 pixels away,  $20\text{pix} \times \sin(10^\circ) = 3.47\text{pix}$ . Since this is larger than 3 pixels, the images should be fine. Sources at radii much less than 20 pixels will be significantly dimmed by PSF subtraction, though they may still be detectable. The standard 'b' scripts output by legoparswrite04 perform PSF subtraction out to a radius of 75 pixels, and then fade the subtraction to zero over a 10 pixel annulus ending at 85 pixels, to avoid a hard edge to the PSF subtraction region. These large radii are chosen because PSF subtraction often handily removes ghosts and other artifacts at large radii, and provides a useful 'second opinion' for the 'a' script reduction even outside the close-in regions where it excels the 'a' script by a large margin. The PSF subtraction radii can be altered by manually editing the 'b' script. The final 'b' script images are stacked and masked as described for the 'a' script, with output images called, e.g. **GJ519b0201.fits** and **GJ519b02u01.fits**.

The 'c' script attempts to subtract the PSF in a way that does not depend on sky rotation. For each pixel, it simply constructs an arc centered on that pixel at fixed radius from the central star, finds the average of pixels along this arc, and subtracts this average from the pixel. This method of PSF subtraction is not as effective as the 'b' script method because it cannot measure the details of the PSF as well. It can still be quite helpful for data sets with insufficient parallactic

rotation for the 'b' script to work well. The output images are, you guessed it, **GJ519c0201.fits** and **GJ519c02u01.fits**.

The 'd' script operates exactly like the 'a' script, except that as one of the final steps before stacking the images it performs unsharp masking **on the individual images** using the same parameters as the 'a' script uses for masking after the stack. After the stack it operates exactly like the 'a' script, producing images called, ie, **GJ519d0201.fits** and **GJ519d02u01.fits**, the latter of which will in a sense be doubly unsharp masked, since the masking has been applied both to its constituent images before the stack, and to the final stacked image. On faint stars the 'd' script looks only slightly cleaner than the 'a' script, but on bright stars that show devastating artifacts from haloes, ghosts, and rays on the 'a' script the 'd' script can work like magic, pulling excellent sensitivity out of regions that looked hopeless on the 'a' script images. The double unsharp masking can reduce the measured brightnesses of point sources slightly.

The 'e' script is a combination of the 'b' and 'd' scripts, where the pre-stack unsharp masking follows the PSF subtraction. The 'e' script is designed to get the best results near bright stars. The names of the output images follow the same convention as for the other scripts.

The 'x' and 'y' scripts are a major departure from the other scripts, in that the nod subtraction is not performed in the usual way. Instead, a single master sky image is made from the whole data set, and this is then used as the subtraction image for all of the science images.

The master sky image is formed as follows: The pairings of images from opposite beams are used just as they would be for the ordinary science processing. The centroid of the star on each image is calculated. The science and subtraction images are sliced along the perpendicular bisector of the line between the star

positions on the two images. This makes two halves of each image, one containing the star and the other showing blank sky. The 'star halves' are discarded and the two starless pieces are scaled to match brightness and then pasted together to make a single starless sky image.

After such a starless sky image has been made from every pairing, these images are stacked with 50 % creeping mean rejection to make a single master sky image. Even if (as is not unlikely for bright stars) some rays, ghosts, or outer halo artifacts ended up in some of the individual sky frames, the creeping mean should reject them and output a clean image.

The master sky image is used as the subtraction image for all of the science images, which, subsequent to the nod subtraction, are handled exactly the same way as in the 'a' script (for the 'x' script) or the 'd' script (for the 'y' script).

The point of the 'x' and 'y' scripts is that their output images, unlike those of all the other scripts, do not contain areas where the data has effectively been destroyed by the negative star images from the nod subtraction. However, since the sky does vary throughout the night, the 'x' and 'y' images are always noisier than the corresponding 'a' and 'd' images. Their unique advantage lies in their ability to reveal companions 'behind' the negative star images that destroy some of the data in all the other reductions. In some cases their quality approaches that of the 'a' and 'd' scripts, while in others they are so bad as to be unusable. I am not sure yet what makes the difference, though the stability of the sky background would be my first guess. The names of the output images follow the same convention as for the other scripts.

There are two additional processing methods, which are sometimes described as the 'f' and 'g' scripts (or the 'f' and 'g' reduction methods). They are specialty methods for binary stars, in which the PSF of each star is used to subtract the

other. They typically produce results considerably better than the other PSF subtraction methods for cleanly separated binary stars with components of similar brightness.

The `legolas04` program cannot carry out the ‘f’ and ‘g’ reduction methods, and `legoparswrite04/03` do not make scripts for them. Rather, they are carried out using auxiliary programs that operate on the individual processed frames output by `legolas04` run with the ‘a’ script. In Section B.11 we describe how to carry out ‘f’ and ‘g’ method reductions using ‘a’ script output images.

The other script writing program, `legoparswrite03`, differs from `legoparswrite04` in only two ways. First, the parallactic rotations are calculated in a more sophisticated way (which requires more input in the script). The additional accuracy is not needed except for extremely precise astrometry of double stars. Second, `legoparswrite03` gives the user the option to select a different algorithm for column pattern noise correction.

An example of a `legoparswrite03` script is given below:

```
2
1
272.98
13 24 10.2849
-0.1220 -0.1286
2294
02
GJ349
GJ349darkamain.fits
flat300lamain.fits
10 310 200 250
```

35.0  
20.0 15.0 12.0 10.0 10.0  
09 29 54.8245 05 39 18.484  
2.0596  
16655.0  
6.5 12.0  
GJ564  
GJ564darkmain.fits  
flat300lamain.fits  
10 310 10 45  
35.0  
20.0 15.0 12.0 10.0 10.0  
14 50 15.8112 23 54 42.639  
2.0596  
16655.0  
6.5 12.0

The explanation of the entries is given in Table B.4.

Table B.4. Explanation of legoparswrite03 script

Entry	Explanation
1	Number of stars for which scripts should be written.
1	1 means raw images should be mirror-flipped in processing; 0 means they should not. For data from July 06 or earlier, use 1; for later data use 0.
272.98	Clockwise rotation in degrees that must be applied to images (after mirror-correction if applicable) to get the +y axis of the image to line up with the +EL (up) direction referenced to the real sky, and, of course, the telescope mount. Values for this are documented on the Clio wiki.
13 24 10.2849	Greenwich Mean Sidereal time at 00:00 UT on the date of the observations, in hh mm ss.ssss format. The number of decimal places given on the seconds is not important. These values can be obtained from The Astronomical Almanac.
-0.1220 -0.1286	Equation of the equinoxes in seconds on the UT date of the observations and on the following day. Can be obtained from The Astronomical Almanac.
2294	Integer number of days from 00:00 UT January 1, 2000 to 00:00 UT on the UT date of the observations.

Table B.4—Continued

Entry	Explanation
02	A numerical string to identify these scripts. Convention: 01 means short, PSF-monitoring exposures, 02 means long science exposures.
GJ349	The name of star 1. Will identify the scripts and the final output images.
GJ349darkamain.fits	The dark frame for this star.
flat300lamain.fits	The flat frame for this star.
10 310 200 250	The minimum and maximum x, and minimum and maximum y pixel values defining a box for scaling images during subtraction.
35.0	The power of the column noise 'colfudge' correction. If it is set to 0.0, colfudge will not be used. The 4-column pattern algorithm of Matt Kenworthy will be used instead.
20.0 15.0 12.0 10.0 10.0	Five centroiding radii for the 5 iterations of centroiding that will be performed on this star.
09 29 54.825 05 39 18.48	Epoch 2000.0 RA and DEC of this star in hh mm ss.sss dd mm ss.ss format. The number of decimal places on the seconds is irrelevant. <b>If the DEC is negative, all 3 DEC entries</b>

Table B.4—Continued

Entry	Explanation
<b>must carry the minus sign, ie -05 -39 -18.484</b>	Note also that the proper motion in Simbad units (mas per year in RA and DEC) <b>may</b> be added after the coordinates. If no proper motion is supplied, zero proper motion is assumed.
2.0596	The true single-frame exposure time in seconds for this star.
16655.0	The photometric calibration in the form of counts/sec for a star of magnitude 10.0 within a large aperture. Typical values for this can be found on the Clio wiki.
6.5 12.0	The inner and outer radii for the scaling annulus for PSF subtraction on this star.

This program requires three time-related entries that it uses to calculate the parallactic angles, where `legoparswrite04` required only the date in `yyyy mm dd` format. The three entries needed here are Greenwich Mean Sidereal time at 00:00 UT on the date of the observations, the Equation of the Equinoxes in seconds on the date of the observations, and the integer number of days between 00:00 UT Jan 1, 2000, and 00:00 UT on the date of the observations.

Of these, the GMST and the Equation of the Equinoxes are available from the *Astronomical Almanac*, Section B; and the integer day number can easily be calculated. When obtaining these numbers from the Almanac, note well that the appropriate table contains both the Greenwich Sidereal Time at 00:00 UT *and the UT at 00:00 Greenwich Mean Siderial Time*. The former quantity, the Greenwich Sidereal Time at 0h UT, is the one that `legoparswrite03` requires, **not the UT at 0h GMST**.

The only other difference between `legoparswrite03` and `legoparswrite04` is that `legoparswrite03` requires the input of a 'column fudge power' parameter for each star.

The `legolas04` processing program has two possible modes for correcting column pattern noise. I invented the first method, which is called `colfudge`. It can remove arbitrary column variations, as long as the amplitude is not too high. However, it introduces some artifacts, and more significantly it dims point sources typically by 15-18%. In the worst cases of unfortunately placed sources near a bright star, the dimming can be 30% or even more. Dimming greater than 30% can only occur in unlikely cases of 'maximally unlucky' source placement.

The other algorithm is due to Matt Kenworthy. It corrects only column variations that repeat in a 4-column cycle, but can correct variations of arbitrarily high amplitude. Because Clio is read out through 4 different amplifiers, the 4-column

pattern is the dominant source of column variations. In many or most cases no remaining column variation can be seen in a nod-subtracted image after the 4-column pattern is removed. Matt Kenworthy's algorithm does not dim point sources or introduce artifacts, and is at least as effective as colfudge at removing the dominant form of column variations in Clio images. It is therefore strongly to be preferred over colfudge in most cases. To use it, just set the column fudge power to 0.0 in the legoparswrite03 script.

If the column fudge power is *not* set to zero, my old colfudge algorithm will be used instead. The value chosen for the colfudge power is then an important and sensitive parameter. The meaning of the parameter is that as the colfudge subroutine determines slowly varying fudge values to add to all the pixels in the image to remove column noise, it determines these values using only pixels that do not deviate from the zero mean of the subtracted images by more than a set value. This restriction is designed to insure that it measures real column noise on mostly-clean patches of sky, and is restrained from trying to 'fudge away' bright star images and their artifacts. The colfudge power is this threshold – colfudge will conclude that any pixel with an absolute value larger than this threshold is affected by a star and is not giving information about column noise.

If the colfudge power is set too low, the full column noise will not be measured and residual column noise will remain in the images. Set it too high, and artifacts will streak out from stellar diffraction rings. For those concerned that colfudge will create artifacts around faint sources such as the planets we are trying to detect, since they will not be bright enough to break the threshold, this is true. However, point sources do not look very much like column noise, and so the artifacts are relatively weak and do not in general cause severe dimming of faint point sources (15-18% is the norm), while the removal of real column noise

increases the visibility of faint sources dramatically. The vertical sides of the outer diffraction rings of bright stars, on the other hand **do** look a bit like column noise, and so it is here that annoying artifacts appear if the colfudge power is set too high. In general, 30-50 for L' or Ks planet search exposures, 5-15 for L' or Ks short, PSF monitoring exposures, and 5-15 for M or M' observations, are good places to start. You can fine tune it from there if column noise or artifacts indicate the correction is too weak or too strong. At the telescope, you probably won't have time to try multiple runs with different column fudge powers. Any values in the ranges given above should produce acceptable results. For photometric calibration stars the power should be set very low or zero so that faint artifacts don't bias the photometry.

## B.8 Making the Image Script

The automated method for constructing the image scripts for legolas04 uses two different programs in succession, which puts one in the unusual position of having to write a script, for a program that will write a script, for another program that will write a script for legolas04.

The first program is called scriptwritelas02, and is run on the command line like this:

```
./scriptwritelas02.exe < swlas02script03
```

Scriptwritelas02 writes a script for the program starcenters01, which in turn writes the image script for legolas04. The information starcenters01 relays to legolas04 includes the name, number of coadds, MMT instrument rotator angle, and approximate primary star coordinates for each image. Starcenters01 also makes a postage stamp image giving a small area around the primary star on

every image in the data set. The user can examine the postage stamp image and quickly identify images that should be rejected due to bad seeing, AO loop failures, or other problems.

An example of the script for scriptwritelas02 is as follows:

```
2 5
320 258
20.0
10 3 20
10 310 178 248
GJ702
0.0
starcens01script03
dark100bmain.fits
legolas04GJ702im01
.fit
GJ702postim01a.fits
41
81
```

The explanation of this script is found in Table B.5.

Table B.5. Explanation of scriptwritelas02 script

Entry	Explanation
2 5	Number of nod sets, and images per beam in each set
320 258	x and y pixel dimensions of the images. Never needs to be changed
20.0	Number of coadds
10 3 20	Half-size of postage stamp images, distance to go away from stellar core to measure unsaturated wings of PSF, distance to go away from stellar core to measure sky background.
10 310 178 248	Defines the scaling box, exactly as previously
GJ702	Root name of star images
0.0	MMT instrument rotator angle
starcens01script03	Name of script to write for the starcenters01 program.
dark100bmain.fits	Master dark for this data set.
legolas04GJ702im01	Name of script starcenters01 will write for legolas04
.fit	postfix to raw image names.
GJ702postim01a.fits	Name of tiled postage stamp image that will be output
41	The number of the first image in the first nod set.

Table B.5—Continued

Entry	Explanation
81	The number of the first image in the second nod set.

The scriptwritelas02 script discussed here is for short, PSF monitoring exposures of the star GJ 702. What is being communicated in the script is that there are just 2 nod sets of such exposures, one going from image 41-50, with images 41-45 in beam 1 and 46-50 in beam 2. The second set goes from image 81-90, with 81-85 in beam 1 and 86-90 in beam 2. For a science data set there would be many nod sets rather than just 2, so the script would end with many lines each giving the number of the first image of a given nod cycle.

The only judgment calls for the scriptwritelas02 script that have not yet been discussed are the size of the postage stamp images (which will be output by starcenters01), the distance to go away from the stellar core to measure the wings of the PSF, and the distance to go away from the stellar core to measure the sky background. The parameters given here will work for unsaturated or mildly saturated stars. For heavily saturated stars, with a saturation radius, say, of more than 5 pixels, the parameters should be something like 10 7 20 instead of 10 3 20. If the stars are very heavily saturated, e.g. out to 10 pixels radius, 10 pixels half-size postage stamp images will be entirely saturated and therefore useless. Adopting a new parameter set such as 25 15 40 will should yield sufficiently large postage stamp images, with the star fairly well centered. The one time when fine tuning these parameters may be necessary is in the case of a double star. In this case the starcenters01 program sometimes gets confused and alternates which star it centers on. Setting the middle parameter, the distance to go away from the saturated image core to get the unsaturated PSF wings, exactly right can fix this problem. The key is to set the distance so it just barely clears the saturation radius.

I should point out that the entry labeled 'Root name of star images' in the script above has to be exactly that, whereas the star name in the legoparswrite04/03

script can be anything. That is, if you have a bunch of images called `img00001.fit`, `img00002.fit`, `img00003.fit`, etc, the root name in your `scriptwritelas02` script has to be **img**, while the star name you supply in the `legoparswrite03` script can be anything, even **Zubenelgenubi**. In that case `legolas04` will correctly process your `img00001.fit` etc images, and then produce final processed images called **Zubenelgenubia02u01.fits**, etc. (In that case the parameter scripts made by `legoparswrite03` would be called **legolas04Zubenelgenubi02a**, etc, and you might want make the entry for the name of the image script to be written for `legolas04` in your `scriptwritelas02` script match to avoid confusion. In my convention you would call it **legolas04Zubenelgenubiim02**).

The output of running `scriptwritelas02` with the script above will be the script `starcens01script03`, which looks like this:

```
320 258
20.000000
10 10
10
dark100bmain.fits
10 310 178 248
3 20
.fit
legolas04GJ702im01
GJ70200041
41
0.000000
GJ70200042
42
```

0.000000

etc...

etc...

GJ702postim01a.fits

Where the etc... entries indicate that all 20 images are listed, along with their number in the sequence and the instrument rotator angle at which they were taken. All of the images taken in beam 1 are listed first, and then all of the images taken in beam 2.

A more detailed explanation of the starcens01 script output by the previous scriptwritelas02 script is found in Table B.6

Table B.6. Explanation of starcenters01 script

Entry	Explanation
320 258	x and y pixel dimensions of images. Never needs to be changed.
20.000000	Number of coadds.
10 10	Number of images in beam 1 and number in beam 2
10	Half-size of postage stamp images.
dark100bmain.fits	Master dark frame for these images.
10 310 178 248	Scaling box.
3 20	Distance to go from the star image core to get to measure unsaturated wings, and distance to go to measure sky background.
.fit	postfix to image names
legolas04GJ702im01	Name of legolas04 image script that will be written
GJ70200041	Name of image 1 in beam 1
41	Position of this image in the observing sequence.
0.000000	MMT instrument rotator angle at which this image was taken
GJ70200042	Name of image 2 in beam 1
42	As above
0.000000	As above
etc...	Remaining 8 images in beam 1 listed

Table B.6—Continued

Entry	Explanation
etc...	in the same way All 10 images in beam 2 listed in the same way
GJ702postim01a.fits	Name of tiled postage stamp image that will be output

The program `starcenters01` is run with this script as one would expect:

```
./starcenters01.exe < starcens01script03
```

`Starcenters01` does two things. First, it makes the image script for `legolas04`, which in this case is called `legolas04GJ702im01`. (By the way, this is my standard convention for image scripts: `legolas04`, then the name of the star, then 'im' for image script, and then the numerical script index: 01 for short, PSF monitoring exposures, and 02 for long, planet search exposures.) Second, `starcenters01` makes a tiled postage stamp image. This means that, having found the approximate location of the brightest star on every image (as it had to do to write the `legolas04` image script), it trims out a small image centered on this star, and tiles these postage stamp images together, in this case to make the image `GJ702postim01a.fits`.

After running `starcenters01`, but before using the `legolas04` script it made, you should look at the postage stamp image. It should be immediately obvious if there are any bad images – ones where the AO loop failed, the seeing was especially bad, or the target was clouded. If there are bad images, you should edit the `starcenters01` script to remove them, and then run it again so it will make a `legolas04` image script that does not use any bad images. Figure B.1 shows the postage stamp image made by running `starcenters01` with the script above.

To delete images from the `starcenters01` script based on their bad appearance in the postage stamp image, you must first identify where they come in the sequence. The postage stamp tiles are ordered from left to right and from bottom to top. Thus, the lower left tile in the postage stamp image is the first image. In the example here, with 20 images total, the image in Figure B.1 is tiled in 4 rows

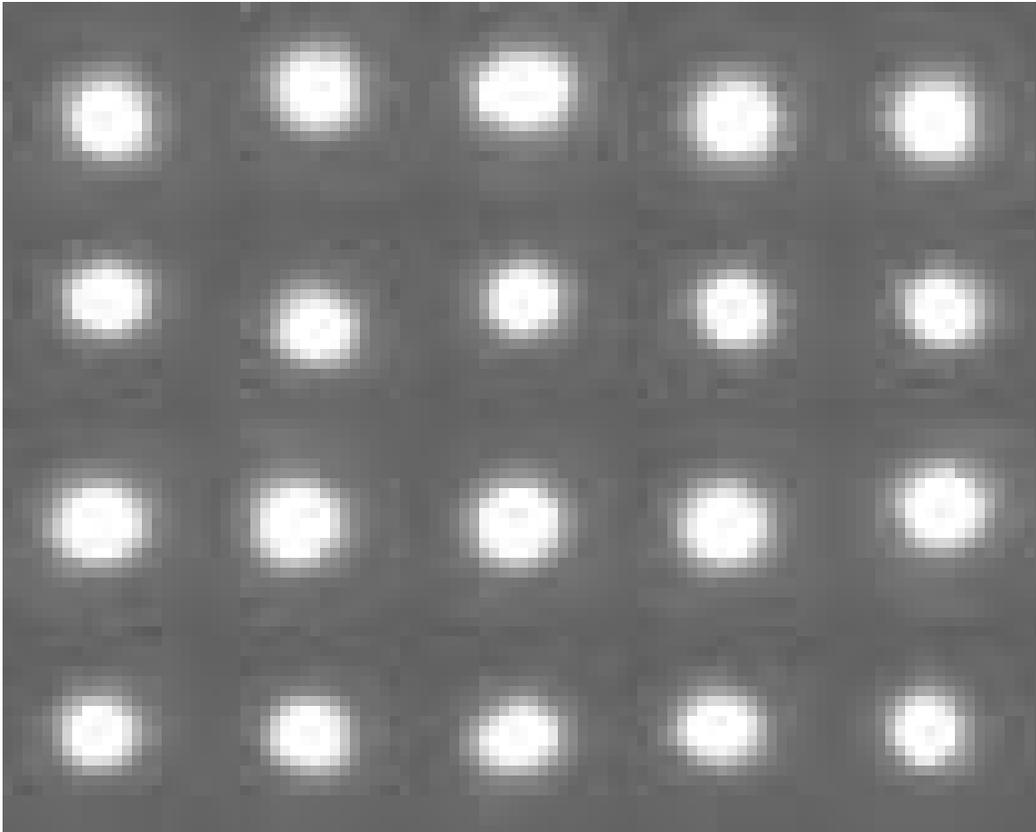


Figure B.1 An example postage stamp image, output by the `starcenters01` program when run with the script `starcens01script03` explained in Table B.6. The third image from right in both the bottom and top rows are elongated and should perhaps be rejected.

of 5 columns, or 5x4 tiles. All the images from beam 1 appear first, then all the images in beam 2 (this is, of course, the same order as in the starcenters01 script).

Figure B.1 shows images that are quite variable in appearance, suggesting the seeing during this data set was not very good. Two images stand out because of their left-right elongation, which is probably due to windshake. These are the third image from left in the bottom row, and the third image from left in the top row. They should perhaps be deleted from the starcenters01 script.

The first image is easy: since it is third from left in the bottom row, it is the third image in the starcenters01 script, or image number GJ70200043.fit. The other bad image is third in the top row. The three rows below this hold a total of 15 images, so the third image in the top row is  $15+3 = 18$ th in the starcenters01 script. Since the images in the first beam, 41-45 and 81-85, are listed first, followed by the images in the second beam, 46-50 and 86-90, the 18th image listed is GJ70200088.fit.

Each of the bad images must be found in the starcenters01 script, and it and the two lines following it that give its parameters must be deleted. Finally, we must change the numbers of images in each beam at the top of the starcenters01 script to reflect the deletions. In this case, we have deleted one image from each beam, so instead of 10 10, the line giving the number of images in each beam would be changed to 9 9.

In cases where there are many images to be deleted, editing the starcenters01 script can be tedious. I recommend identifying all the bad images, writing down their numbers in the tiling sequence, and then using the emacs ctrl-u option to page down the required number of lines and mark all of the bad images with `***`, or something of that sort. For example, if you want to delete image 101, put your cursor on the name of image 1 and then go up 3 lines. Now, since there are 3

lines for each image, if you go down  $3 \times n$  lines you will be on the name of image  $n$ . You can type ctrl-u 303 and then down arrow to go down 303 lines, which will put you on image 101, which you can mark for deletion. After you have marked all the bad images in this way, you can go back through and delete them all, without having to deal with questions like, “Let me see, I’ve deleted images 23, 28, and 34. Which image is 47 in the original order now?” Of course, since you have to edit the line giving the number of images surviving in each beam, you should keep track of how many images from each beam you are deleting. Remember, all the images from beam 1 are listed first in the starcenters01 script, followed by all the images in beam 2.

After editing the starcenters01 script, run starcenters again and look at the new postage stamp image. If you’ve done everything right, there shouldn’t be any bad images. The legolas04 image script produced by this final starcenters01 run will feed legolas04 all the good images and leave out the bad ones.

## B.9 Running legolas04 (At Last!)

Simply type:

```
./legolas04.exe
```

The result will be an immediate request:

*Enter the name of your master parameter script file.*

To perform an ‘a’ reduction of the GJ 702 short exposure data, you would enter

**legolas04GJ70201a**

legolas04 will write lots of output to the terminal as it reads the parameter script. It will then request:

*Enter the name of your image listing script file.*

In this example, you would enter:

**legolas04GJ702im01**

legolas04 will immediately say

*Got 10 images in position 1 and 10 in position 2. Enter the number of distinct pairing sets you want to make.*

The number of images it reports is, of course, a sanity check that your file is good and you are processing the right data set. In running this example we have assumed that the two elongated images that appeared in the postage stamp image have **not** been deleted.

As discussed many sections ago, 1 is usually or always the right number of pairing sets, so enter simply:

**1**

legolas04 will once again write lots of output to the terminal as it reads the image script. Finally it will say:

*Enter the maximum distance the nearest image must be away before doubling pairs.*

Try entering the full number of images in a nod set plus 3:

**13**

Legolas will output to the terminal a running report of its attempts to pair all the images. In this case, it will successfully pair all images with no doubling and proceed right into processing. If instead you had entered:

7

legolas04 would finally have reported:

*Successfully created all desired pairing sequences! However, there were 2 double pairings. If you want to try a larger doubling distance and see if you can reduce or eliminate double pairings, enter 0. Otherwise, enter 1 to go ahead with processing using the current pairings.*

And you could decide that 2 is not a disturbing number of double pairings, and enter 1 to go ahead with processing, or you could decide it should be possible to get zero double pairings, and enter 0, in which case you would get a repeat of the previous message:

*Enter the maximum distance the nearest image must be away before doubling pairs*

And you could enter 10 or 13 and find that it is indeed possible to get zero double pairings. During processing, legolas04 outputs a running log to the terminal, so you can always see which image it is working on, which processing step, etc. If you remember how many images there were in total, you can use this information to estimate how much time the run will take.

#### B.10 Output of legolas04: Processed Images

The output of final, stacked images has already been discussed in the section on legolas04 parameter scripts, but legolas04 also outputs final, processed versions of each input image. These images have the same name as the original input (for example, img00013, not Zubenelgenubi), but they have the script letter plus **e** (for 'edited') appended, and the final postfix is **.fits**, not **.fit**, as for the raw Clio images. Thus if the raw image is called **img00013.fit**, the version processed by

legolas04 using the 'a' script will be called **img00013ae.fits**. If the raw image is called **GJ820A00236.fit**, the processed version after the 'e' script reduction will be called **GJ820A00236ee.fits**.

These edited images are in the final form that legolas04 stacks to make its final outputs. They are very useful, because they can be stacked in different ways if you prefer another method to legolas04's default 20% rejection creeping mean. The images are also pretty big, since they are zero padded to 500x500 pixels, and if you process a lot of stars with all six different useful scripts ('a', 'b', 'd', 'e', 'x', and 'y') then you have made 6 500x500 pixel images from each 320x258 input raw image, so filling up disk space can be an issue. The way to deal with this, of course (if your disk is not so big that it's irrelevant) is to run legolas04, make all the alternative stacks of the edited images you want, and then delete the edited images and only keep the final stacks.

#### B.10.1 Code-Generated Files: Image Lists

Legolas04 automatically makes a lot of files that are required for its operation, and it calls these the 'cgf', or code-generated files. Some of them can be very useful after the legolas04 run is complete. They all begin with the star name, followed by cgf and the numerical string and letter identifying the script, i.e. **GJ820A**cgf02a for the 'a' script reduction on long exposure science images of the star GJ820A.

One of the most useful of the code generated files is the image combine file, or icf, called, i.e. **GJ820A**cgf02a**icf**. This is simply a list of all the names of the edited versions of the raw images output by legolas. The images are listed in the same order as legolas04 processed them, that is, all the images in beam 1 first, followed by all the images in beam 2.

### B.10.2 Post-legolas processing: cmcsets programs

Several useful programs exist for performing further analysis on Clio data after it is processed by `legolas04`. A first step might be to make different stacks of the processed images. The final stacked output of `legolas04`, as described above, is a 20% creeping mean rejection stack of the entire data set. There is a simple program to make any number of different creeping mean stacks of Clio data, and output versions with and without the `legolas04` standard unsharp masking. This program is called `cmcunsh01`. The free parameters are the subsets of images used (each subset must be listed in a file, which is input to `cmcunsh01`) and the creeping mean rejection fraction. This fraction can be set to 0.0 if desired to get a straight average of the data. The program `cmcunsh01` is run on the command line with a script:

```
./cmcunsh01.exe < cmcunsh01script01
```

A `cmcunsh01` script looks like this:

```
500 500
4
58
GJ3541Acgf02aicfq1
0.2
GJ3541Aa0201q1.fits GJ3541Aa02u01q1.fits
58
GJ3541Acgf02aicfq2
0.2
GJ3541Aa0201q2.fits GJ3541Aa02u01q2.fits
58
```

GJ3541Acgf02aicfq3

0.2

GJ3541Aa0201q3.fits GJ3541Aa02u01q3.fits

58

GJ3541Acgf02aicfq4

0.2

GJ3541Aa0201q4.fits GJ3541Aa02u01q4.fits

A brief description of what each script entry means is provided in Table B.7.

Table B.7. Explanation of cmcunsh01 script

Entry	Explanation
500 500	x and y dimensions of images
4	number of image sets to combine
58	number of images in first set
GJ3541Acgf02aicfq1	name of file listing images for first set
0.2	rejection fraction for first set
GJ3541Aa0201q1.fits GJ3541Aa02u01q1.fits	name of ouput images without and with unsharp masking for unsharp first set
58	number of images in second set
GJ3541Acgf02aicfq2	Name of image list file for second set
0.2	rejection fraction for second set
GJ3541Aa0201q2.fits GJ3541Aa02u01q2.fits	name of output images without and with unsharp masking
58	number of images in 3rd set
GJ3541Acgf02aicfq3	etc
0.2	etc
GJ3541Aa0201q3.fits GJ3541Aa02u01q3.fits	etc

Table B.7—Continued

Entry	Explanation
58	number of images in 4th set
GJ3541Acgf02aicfq4	etc
0.2	etc
GJ3541Aa0201q4.fits GJ3541Aa02u01q4.fits	etc

The `cmcunsh01` script above produces combines with 20% creeping mean rejection of the first, second, third, and fourth quarters of the images for star GJ 354.1A. There were 232 images in all, so each quarter contained exactly 58 images. The list files, such as `GJ3541Acgf02aicfq1` are just simple ASCII text files listing the full name of each image in the given set. Since the script says there are 58 images in set number one, the file `GJ3541Acgf02aicfq1` must have 58 lines each giving the full name of a Clio image after legolas processing, ie `GJ3541A00012ae.fits`.

The example script here did standard 20% creeping mean combines on different subsets of a Clio data set. The program `cmcunsh01` could, of course, also be used to make stacks of the whole data set at different rejection fractions, including zero rejection.

There are two additional programs offering simpler scripts to do standard, often-used complementary stacks of Clio data sets. The program `cmcsets01` takes the code-generated list file directly output by Clio, ie `GJ820Acgf02aicf`, and uses it to make a 50% rejection stack of the whole data set, and 20% rejection stacks of the first and second halves of the data set. Images without and with standard unsharp masking are output. Like `cmcunsh01`, `cmcsets01` is run on the command line with a script:

```
./cmcsets01.exe < cmcs01script01
```

Here is an example of a `cmcsets01` script:

```
232 500 500
GJ3541Acgf02aicf
7
GJ3541Aa0201r05.fits
```

GJ3541Aa02u01r05.fits

GJ3541Aa0201h1.fits

GJ3541Aa02u01h1.fits

GJ3541Aa0201h2.fits

GJ3541Aa02u01h2.fits

An explanation of the entries in this script is proved in Table B.8.

Table B.8. Explanation of cmcsets01 script

Entry	Explanation
232 500 500	number of images and their x and y dimensions
GJ3541Acgf02aicf	name of image list file output by legolas04
7	number of characters in the image root name
GJ3541Aa0201r05.fits	name of 50% rejection image without unsharp masking
GJ3541Aa02u01r05.fits	name of 50% rejection image with unsharp masking
GJ3541Aa0201h1.fits	name of first half 20% rejection image without unsharp masking
GJ3541Aa02u01h1.fits	name of first half 20% rejection image with unsharp masking
GJ3541Aa0201h2.fits	name of second half 20% rejection image without unsharp masking
GJ3541Aa02u01h2.fits	name of second half 20% rejection image with unsharp masking

The fact that the number of characters in the image root name is required may perhaps need some explanation. `cmcssets01` uses this to re-order the images. In the image combine file output by `legolas04`, the images are ordered by beam, all the images in beam 1 first followed by all the images in beam 2. But `cmcssets01` must divide the images into the first and second halves of the data set, meaning first and second in the ordinary time-order of observations, regardless of beam. It does this by the image numbers, using the names listed in the image combine file. In the example above, the image combine file `GJ3541Acgf02aicf` lists all 232 images from the 'a' script `legolas04` reduction of the GJ 354.1A data set. The image `GJ3541A00006ae.fits` is 6th in temporal order, but since it is in the second beam it probably appears something like 117th in the list order. `cmcssets01`, however, will read its name one character at a time, looking for a 5-digit number after the 7th character, since the script has told it the image root name is 7 characters long. Thus it will find the 5-digit number to be 00006, and correctly determine that the image comes 6th in temporal order. If the user made a mistake, and gave the root name length as 6 characters, the program would read the number as A0000, and the result is undetermined but will surely result in a scrambled order of images.

Other than this oddity, `cmcssets01` operates exactly as a special case of `cm-cunsh01`, with a simpler, effectively shorthand, script. It generates its own additional image list files, in this case `GJ3541Acgf02aicfh1` and `GJ3541Acgf02aicfh2`, giving, respectively, the first and second 116 images in the GJ 354.1A data set. These subset list files can be useful to speed the production of custom lists for `cm-cunsh01`, if they are needed.

The images output by `cmcssets01` are very useful for source detection. In the example here, the initial run of the `legolas04` 'a' script for GJ 354.1A produced the final image `GJ3541Aa02u01.fits`. This image showed an apparent faint source near

the star, but it could not be confirmed as real. Other features of equal or greater brightness were present, with shapes and locations that suggested they were artifacts, and it could not be confirmed that the more provocative source was not simply a special case of one of these. Running `cmcssets01` produced the half images `GJ3451a02u01h1.fits` and `GJ3541Aa02u01h2.fits`. On blinking between these two images, all spurious sources rotated with the changing parallactic angle from the first to the second half of the Clio data set, while the other source remained fixed and was shown to be a real celestial object. Unfortunately, further investigation showed it was a background star.

There is another program, `cmcssets02`, which again operates as a special case of `cmcsunsh01` with a shorthand script that uses the code-generated image list file from `legolas04`. It produces 20% rejection stacks of images isolated by beam. Thus it makes a complete stack of all of the beam 1 images, and also stacks of the first and second halves of the beam 1 images, and a stack of all the beam 2 images, and also stacks of the first and second halves of the beam 2 images. The script is as follows:

```
232 500 500
117 115
GJ3541Acgf02aicf
GJ3541Aa0201b1.fits
GJ3541Aa02u01b1.fits
GJ3541Aa0201b2.fits
GJ3541Aa02u01b2.fits
GJ3541Aa0201b1h1.fits
GJ3541Aa02u01b1h1.fits
GJ3541Aa0201b1h2.fits
GJ3541Aa02u01b1h2.fits
```

```
GJ3541Aa0201b2h1.fits  
GJ3541Aa02u01b2h1.fits  
GJ3541Aa0201b2h2.fits  
GJ3541Aa02u01b2h2.fits
```

This script should be fairly self-explanatory by analogy to `cmcsets01`. The second line gives the number of images in beam 1 followed by the number of images in beam 2. The user can easily determine these numbers by looking over the `legolas04`-generated file, ie `GJ3541Acgf02a.icf`. This file will contain a list of all of the beam 1 images in order of image number, followed by a list of all the beam 2 images in order, so the switch-over from beam 1 to beam 2 is the sudden drop in the image numbers. Since `cmcsets02` operates by beam, it does not need to re-order the images from the `legolas04` `icf` files. It does not, therefore, need the number of characters in the image root names.

#### B.11 Post-`legolas` Processing: ‘f’ and ‘g’ Method Reductions

The ‘f’ and ‘g’ method reductions for `Clio` images are specialty options for double stars. The PSF of each star is subtracted using a scaled version of the PSF of its companion. This is done on an image-by-image basis, rather than on the final master image. Both stars in a binary are affected almost identically in a given image by the specific set of effects from uncorrected atmospheric seeing, telescope and instrument imperfections, and even windshake that were present during the taking of that image. For this reason, better PSF subtraction is usually obtained by using the two stars of a binary to subtract each other than by any other method.

The subtraction is most powerful if the brightness of the two stars is very similar, so the required scaling factors are near unity. As with the ‘b’ and ‘e’ methods of data reduction, the PSF of each star is scaled based on an annulus.

Thus an accurate scaling can be obtained for saturated stars. If the images are not saturated, the inner radius of the annulus can be set to zero.

The program for PSF subtraction of binary stars is called `psfsub02`, and operates on single images fully processed by the 'a' script reduction method of `legolas04`. This program is run on the command line in the usual way:

```
.\psfsub02.exe < psfsub02script01
```

Here is an example of a `psfsub01` script:

```
500 500
139
250 250 101 210
10.0 8.0 8.0 6.0 6.0
40.0 45.0
4.0 5.0 11.0 50.0 60.0
4.0 5.0 11.0 50.0 60.0
GJ505A00011ae.fits
GJ505A00011fe.fits
GJ505A00012ae.fits
GJ505A00012fe.fits
GJ505A00013ae.fits
GJ505A00013fe.fits
GJ505A00014ae.fits
GJ505A00014fe.fits
GJ505A00015ae.fits
GJ505A00015fe.fits
GJ505A00021ae.fits
GJ505A00021fe.fits
```

GJ505A00180ae.fits

GJ505A00180fe.fits

The entries in this script are explained in Table B.9.

Table B.9. Explanation of psfsub02 script

Entry	Explanation
500 500	x and y pixel dimensions of input images
139	Total number of images to process
250 250 101 210	Approximate pixel coordinates of primary and secondary star. As always the primary is at 250,250. In this case the secondary is at about 101,210.
10.0 8.0 8.0 6.0 6.0	Radii for 5 centroiding iterations to be performed on each star.
40.0 45.0	Inner and outer radii for the sky subtraction annulus. If no sky subtraction is desired, enter 0.0 0.0.
4.0 5.0 11.0 50.0 60.0	Five radii defining how the PSF of the primary is subtracted using that of the secondary. Pixels inside the first radius will simply be set to zero. The second and third radii define the annulus that will be used to scale the PSF. The fourth is the outer radius for full PSF subtraction. Beyond this PSF subtraction will linearly fade from 100% to zero, reaching zero at the final radius.
4.0 5.0 11.0 50.0 60.0	Same as previous entry, but for subtracting the secondary using the

Table B.9—Continued

Entry	Explanation
	primary. Here, the same radii are used; this need not be the case.
GJ505A00011ae.fits	Name of input ‘a’ script processed image number 1.
GJ505A00011fe.fits	Name of output PSF subtracted version of this image.
GJ505A00012ae.fits	Name of input ‘a’ script processed image number 2.
GJ505A00012fe.fits	Name of output PSF subtracted version of this image.
GJ505A00013ae.fits	Name of input image number 3
GJ505A00013fe.fits	Name of output for this image
GJ505A00014ae.fits	etc
GJ505A00014fe.fits	etc
GJ505A00015ae.fits	etc
GJ505A00015fe.fits	etc
GJ505A00021ae.fits	Name of input image number 6. Note the numbering discontinuity because all images from beam 1 are listed first, and then all images from beam 2. The order doesn’t matter, but we have them ordered

Table B.9—Continued

Entry	Explanation
	like this because we simply pasted them from the legolas04 output image list ('icf') file to save time.
GJ505A00021fe.fits	Output image number 6.
etc ...	etc ...
GJ505A00180ae.fits	Input image number 139.
GJ505A00180fe.fits	output image number 139.

The only judgement calls in the `psfsub02` script are the radii for the centroiding iterations, the sky subtraction radii, and the five PSF subtraction radii.

The centroiding radii are exactly analogous to those required for `legolas04` itself. Choosing these has already been discussed above.

Sky subtraction is largely a matter of personal inclination. It shouldn't be necessary unless the data set was taken under very bad conditions, when it might not help much. On the other hand, it shouldn't hurt as long as the annular radii are large enough to clear the outer halo of the star. The ones given here should be enough for all but the very brightest stars.

For the PSF subtraction radii, the zeroing radius should probably be set approximately to the saturation radius. The inner radius of the scaling annulus should be one or two pixels beyond this. There is wide flexibility for the outer radius of the scaling annulus. The annulus should probably be at least 5 pixels wide in general, and should not be so large that the scaling is affected by sky noise or the flux from the companion star. If optimal sensitivity is desired at a specific radius from the star, a narrow annulus centered on that radius may fine-tune the subtraction to produce optimal results there.

The outer limits of subtraction are usually set by the presence of the companion. If significant companion flux lies within the outer radius for PSF subtraction, significant dark artifacts will appear around both stars. The outer limits may thus vary widely depending on the separation of the stars.

The PSF subtraction radii can be different for the two stars. An example of when this may be desirable is a case where the primary is considerably brighter than the secondary. The secondary PSF will have to be brightened a lot in scaling to subtract the primary. If the primary is subtracted out to a large radius, scaled up sky-noise from the area surrounding the secondary will dramatically reduce

the sensitivity in a wide area around the primary. It may thus be desirable to set a small outer radius for PSF subtraction around the primary, but a large one around the secondary. In extreme cases, the primary subtraction radius can be set to zero and only the secondary will be subtracted.

We note that PSF subtraction using `psfsub02` works best for binaries with separations between about 3 and 6 arcseconds. At smaller separations the halos of the stars overlap significantly and artifacts are introduced. At larger separations the PSFs of the stars begin to match less well. This is not to say binary PSF subtraction is not worthwhile for stars with separations less than 3 or more than 6 arcseconds; merely to highlight the range where the technique is most effective.

After `psfsub02` has been run, the processed images (called `GJ505A*****fe.fits` in the example above) must be stacked and unsharp masked to make master images like the usual `legolas04` final output. An easy way to do this is to make a copy of the `legolas04` 'a' script image list file and change the names of all the images in it; then use the new file as input to `cmcunsh01`. In the example above, the 'a' script image list file would be called `GJ505Acgf02aicf`. A new copy could be called `GJ505Acgf02ficf`, and all the occurrences of `ae.fits` in this new file could be replaced with `fe.fits` (using, for example, `emacs Alt-x replace atring`). Then the master images could be made using this very simple `cmcunsh01` script:

```
500 500
1
139
GJ505Acgf02ficf
0.2
GJ505Af0201.fits GJ505Af02u01.fits
```

The programs `cmcsets01` and `cmcsets02` could also be run with the image list

file GJ505Acgf02ficf as input; the results would be exactly as if the 'f' reduction method were a standard legolas04 option.

The 'g' reduction method is simply the 'f' method with pre-stack unsharp masking. That is, the 'g' method is to the 'f' exactly what the 'd' method is the 'a'. In general, 'g' method images do not differ enough from the 'f' ones to make it clear that a 'g' reduction is worthwhile. For very bright binaries it may be very helpful, however. It is carried out by applying the program unsharpm01 to the output image from psfsub02. This program is very versatile, with many possible ways to unsharp mask images. Here, however, we simply show how to set the parameters to the same standard type of unsharp masking used by legolas04, cmcunsh01, cmcsets01, and cmcsets02. The program is run on the command line in the usual way:

```
.\unsharpm01.exe < unsharpm01script01
```

Here is an example of an unsharpm01 script:

```
139
500 500
4
5.0 2.5
GJ505A00011fe.fits
GJ505A00011ge.fits
GJ505A00012fe.fits
GJ505A00012ge.fits
GJ505A00013fe.fits
GJ505A00013ge.fits
etc ...
GJ505A00180fe.fits
```

GJ505A00180ge.fits

It follows the same pattern as the psfsub02 script, with the number of images to process on the first line, their x and y pixel dimensions on the second line, processing parameters on following lines, and then finally a long set of line pairs giving the input and output names for each image. The parameter set for unsharpm01 is much simpler than for psfsub02. It consists of only two lines. Here, the entries are 4 and 5.0 2.5, respectively. The 4 chooses the unsharp masking algorithm; algorithm # 4 is the standard one used by the legolas suite. The 5.0 is the  $\sigma$  value for the gaussian kernel used in the unsharp masking, and 2.5 is a  $\sigma$  clip for rejection point sources from the mask so they do not get dimmed. The reader need not understand these parameters in any detail; suffice it to say that they are the standard ones. They are hidden in declared constants in the legoparswrite04/03, cmcunsh01, cmcsets01, and cmcset02 programs.

After unsharpm01 is run, the resulting 'ge.fits' images can be combined in the same ways we describe above for the 'fe.fits' images output by psfsub02.

## B.12 Output of legolas04: Image Data Files

After each run of legolas04, a lot of potentially useful data is stored in the data file that it produces. The name of this file begins with the star name, not the image name (ie, **Zubenelgenubi**, not **img**). Next the identifier, **datfile** is appended, and then the numerical string and the letter identifying the legolas04 parameter script. Thus the data file produced by an 'a' script reduction of the science images of the star GJ820A would be called **GJ820Adatfile02a**.

This file has 11 columns, each labeled. I have not adjusted the tabbing to get the labels to line up with the columns. You can fix this in a few seconds by adding more tabs if you open the file in emacs. If you copy the unaltered version

into Open Office Calc, the labels will also align; this is at least a good excuse for my not having edited legolas04 to insert extra tabs automatically.

The first column is more or less meaningless; it is simply the order in which legolas04 processed the images. It always processes all the images in beam 1 in order, and then all the images in beam 2.

The next column is the decimal UT in hours that legolas04 read from the image header. This can be used to reconstruct the airmass at which each image was taken, or in the unlikely event that a Clio observing sequence records a time-variable astrophysical phenomenon, to map the history of this phenomenon.

The third and fourth columns record the exact pixel coordinates of the centroid legolas04 found for the primary star on the unshifted images. If a mirror-flip in processing was commanded, these coordinates apply after the mirror-flip. They can be used to see how consistent the positioning of the primary star was from frame to frame, which may shed light on whether the multiple pairing option could work. Alternatively, it can be used to evaluate the accuracy of the MMT tracking and pointing or the legolas04 centroids. I think in general the star should be expected to move in a slow arc about some point, since the rotator is off and it is not likely the star has been fortuitously lined up with the exact center of the MMT rotator on the Clio chip.

The fifth column records the clockwise rotation that was applied to the images to get celestial North up. This rotation is of course dependent on the instantaneous parallactic angle, which is why it changes from image to image. This column is useful to see how much parallactic rotation took place during the observing sequence, and thus to evaluate how close to the star the PSF subtraction scripts may still be expected to produce good results.

The sixth column records the raw counts on each image summed within a 30

pixel radius of the primary star, with sky subtraction using an annulus extending from radius 40 to 43 pixels. These default parameters may be changed by manually editing the `legolas04` parameter script. If the star is unsaturated, this column actually represents raw aperture photometry of reasonable quality, and variations in it can be indicative of clouds or the result of changing airmass. If the star is saturated, the aperture photometry on it may still give an indication of how stable the conditions were. If the PSF and atmospheric transparency are very stable, so should be the saturated photometry, but if not, the changing PSF and/or halo brightness will make the saturated photometry unstable, and so instability in this column gives a measure of these variable atmospheric conditions.

The seventh column is the magnitude of the primary star, calculated from the photometric calibration and exposure time supplied in the `legoparswrite03` script, and the raw aperture photometry from column six. If the star is unsaturated and the photometric calibration and exposure time have been input correctly, this column supplies real photometry of the star, which is potentially quite interesting as relatively few stars have measured  $L'$  and  $M$  band magnitudes. The accuracy is not likely to be better than about 10 %, however. If the star is saturated and/or the calibration and exposure time were not input correctly, this column only presents the same stability information as column six, in, perhaps, a form slightly easier to comprehend quickly by eye.

The eighth column gives the sky brightness after dark subtraction and flat-fielding, averaged over the scaling region. The units are counts/pixel. This column can give a very useful indication of any passing clouds or other changes in the sky. Clouds always cause the sky to brighten at  $L'$  and  $M$  band wavelengths, and this brightening indicates a decrease of transparency and potentially other negative affects on the data. Also, after being normalized to the exposure time,

the sky level can be used to make sky quality comparisons from night to night and run to run, and understand how much the quality of L' and M band observations obtainable at the MMT changes with the time of year.

The ninth column is simply the pixel-to-pixel background RMS measured over the scaling region after dark subtraction and flat fielding. It is not indicative of the rms of the final processed images, since it is measured before nod subtraction and nod subtraction brings the RMS down dramatically. I'm not sure this column is good for anything, but it seems desirable to have it here for completeness.

The tenth column gives the correction applied to the parallactic angle calculation due to the precession of celestial coordinates since epoch 2000.0. The eleventh column gives the correction due to the star's proper motion, if the proper motion was not set to zero. These are relevant only for very accurate position angle astrometry, but `legolas04` applies the corrections whenever the appropriate information is supplied.

### B.13 Post-`legolas` processing: Source Detection

The standard program for source detection is named, optimistically, `planetdetect01`. It requires as input a PSF template image and a set of science images with a detection threshold in sigma for each one. A standard way of running it is to supply the full, `legolas04` output image, ie `GJ3541Aa02u01.fits`, with a threshold of 4.5 sigma, both half images made by `cmcsets01`, ie `GJ3541Aa02u01h1.fits` and `GJ3541Aa02u01h2.fits`, with thresholds of 3.0 sigma, and the 50% rejection image made by `cmcsets01`, `GJ3541Aa02u01r05.fits`, with a threshold of 3.5 sigma. The program reports as confirmed sources anything that appears at a consistent location on all images (all 4 in this example) at a significance above the detection

threshold in each case. Generally I set the match radius (ie, the definition of 'a consistent location' to 2.0 pixels). These parameters usually result in a few false positives that must be checked out manually, but they rarely or never miss a real source that appears in a region of reasonable image quality on all 4 input images. Using fewer images or widely different detection thresholds is, of course, quite possible and may produce good results.

Planetdetect01 uses two methods to create error images and detect sources. First, it takes the input PSF image and calculates the optimum aperture radius for detecting faint sources against random background noise. Random noise goes up linearly with the radius of an aperture. Thus the optimal radius for faint-source detection will be the radius where the ratio of enclosed source flux in the PSF to radius is the highest. This is usually between 2.0 and 3.0 pixels for good Clio data. Planetdetect01 takes a science image, and goes through it pixel by pixel. At each pixel, it sums the light within the optimal detection radius of that pixel, and writes the sum to that pixel location in a new image. Then, it goes through the new image pixel by pixel. At each pixel, it finds the mean and rms in an annulus centered on that pixel. I usually take 5.0 pix and 8.0 pix as the inner and outer radii of the annulus. If the pixel under consideration is above the annular mean by more than the sigma-threshold times the annular rms, planetdetect01 'detects' a source at that location on that image. After it has detected sources on all the images, it looks for sources that are consistently detected on every image, and reports them as confirmed detections using aperture photometry.

There is a special case to the above method of calculating the mean and rms in an annulus. If the pixel in question is too close to the central star, the noise characteristics may be so dependent on distance from the central star that using an annulus centered on the pixel will produce very inaccurate results. In this case,

planetdetect01 finds the mean and rms in an arc at constant radius from the central star, with a fixed length and a gap centered on the pixel under consideration. Usually, I set the length of this arc to 45 pixels, and the radius from the central star at which the rms calculation method changes from annuli to arcs at 60.0 pixels.

The second method planetdetect01 uses to detect sources is PSF fitting. Again, it goes through the image pixel by pixel, but this time, at each pixel it fits a PSF centered on that pixel to the image within a certain radius. Usually I set the radius to 6.0 pixels. The PSF fit is a least-square fit with two degrees of freedom: the amplitude of the PSF, and the background. The exact center of the PSF is not a free parameter: it is fixed centered on the pixel under consideration. Planetdetect01 produces a new image giving the amplitude of the best-fit PSF centered on each pixel, and proceeds exactly as before in detecting sources present on this image.

Finally, planetdetect outputs the lists of sources consistently detected on all images by each method.

planetdetect01 is run on the command line with a script:

```
./planetdetect01.exe < planetd01script01
```

Here is an example of a planetdetect01 script:

```
GJ3541Adetsall_a.txt
GJ3541Adetsconf_a.txt
2.0
4
500 500
GJ3541Aa0101t.fits 101 101
250.0 250.0 60.0 45
8.0
```

0.0 0.0

5.0 8.0

51 51

1.0 30.0 0.1

6.0

GJ3541Aa02u01.fits

GJ3541Aa02u01\_amp01.fits

GJ3541Aa02u01\_sky01.fits

4.5

GJ3541Aa02u01\_smooth01.fits

GJ3541Aa02u01\_aplan01.fits

GJ3541Aa02u01\_pplan01.fits

GJ3541Aa02u01h1.fits

GJ3541Aa02u01h1\_amp01.fits

GJ3541Aa02u01h1\_sky01.fits

3.0

GJ3541Aa02u01h1\_smooth01.fits

GJ3541Aa02u01h1\_aplan01.fits

GJ3541Aa02u01h1\_pplan01.fits

GJ3541Aa02u01h2.fits

GJ3541Aa02u01h2\_amp01.fits

GJ3541Aa02u01h2\_sky01.fits

3.0

GJ3541Aa02u01h2\_smooth01.fits

GJ3541Aa02u01h2\_aplan01.fits

GJ3541Aa02u01h2\_pplan01.fits

GJ3541Aa02u01r05.fits

GJ3541Aa02u01r05\_amp01.fits

```
GJ3541Aa02u01r05_sky01.fits
```

```
3.5
```

```
GJ3541Aa02u01r05_smooth01.fits
```

```
GJ3541Aa02u01r05_aplan01.fits
```

```
GJ3541Aa02u01r05_pplan01.fits
```

Don't be daunted by the apparent complexity of this long script. Almost all the parameters remain unchanged from one star to another, so once you have one good planetdetect01 script making more involves little more than copying an old one and changing the image names.

An expansion of the script above is given in Table B.10.

Table B.10. Explanation of psfsub02 script

Entry	Explanation
GJ3541Adetsall...a.txt	File to write <b>all</b> source detections
GJ3541Adetsconf...a.txt	File to write confirmed source detections
2.0	Match radius
4	Number of images to search
500 500	Dimensions of science images
GJ3541Aa0101t.fits 101 101	Name and dimensions of PSF template image
250.0 250.0 60.0 45	x y coords of primary star on all science images, followed by radius within which sensitivity should be calculated in arcs, and length of the arcs
8.0	Radius within which sensitivity should <b>not</b> be attempted (i.e. saturation radius)
0.0 0.0	Inner and outer sky subtraction radii for PSF template. Set both to 0.0 if no sky subtraction is desired.
5.0 8.0	Inner and outer radii for the sensitivity calculation annulus.
51 51	x y coords of the PSF center on the template image.
1.0 30.0 0.1	Inner and outer radii and radius step

Table B.10—Continued

Entry	Explanation
	for finding optimum aperture radius for detecting faint sources.
6.0	Fitting radius for PSF fitting.
GJ3541Aa02u01.fits	Name of science image 1.
GJ3541Aa02u01_amp01.fits	Image to contain amplitudes of PSF fit to science image 1.
GJ3541Aa02u01_sky01.fits	Image to contain sky values for PSF fit to science image 1.
4.5	Detection limit in sigma for science image 1.
GJ3541Aa02u01_smooth01.fits	Image to contain science image 1 smoothed with optimal dection radius.
GJ3541Aa02u01_aplan01.fits	Image to contain planets detected by aperture phot on science image 1
GJ3541Aa02u01_pplan01.fits	Image to contain planets detected by PSF fitting on science image 1.
GJ3541Aa02u01h1.fits	Name of science image 2.
GJ3541Aa02u01h1_amp01.fits	Other entries exactly analogous to above.
GJ3541Aa02u01h1_sky01.fits	etc
3.0	Note lower detection limit in sigma for this image, which is made from only half the data)

Table B.10—Continued

Entry	Explanation
GJ3541Aa02u01h1..smooth01.fits	etc
GJ3541Aa02u01h1..aplan01.fits	etc
GJ3541Aa02u01h1..pplan01.fits	etc
GJ3541Aa02u01h2.fits	Name of science image 3.
GJ3541Aa02u01h2..amp01.fits	etc
GJ3541Aa02u01h2..sky01.fits	etc
3.0	etc
GJ3541Aa02u01h2..smooth01.fits	etc
GJ3541Aa02u01h2..aplan01.fits	etc
GJ3541Aa02u01h2..pplan01.fits	etc
GJ3541Aa02u01r05.fits	Name of science image 4.
GJ3541Aa02u01r05..amp01.fits	etc
GJ3541Aa02u01r05..sky01.fits	etc
3.5	etc
GJ3541Aa02u01r05..smooth01.fits	etc
GJ3541Aa02u01r05..aplan01.fits	etc
GJ3541Aa02u01r05..pplan01.fits	etc

Note that the PSF template image **must** have the PSF star centered at an integer pixel value. If, as is always the case in my applications, the PSF image is an image made by legolas04 from short, unsaturated exposures of a star, this condition is easily satisfied. Legolas04 images made with standard scripts always have dimensions 500 500, with the star centered on exactly pixel 250 250. I trim these down to dimensions 101 101 to make PSF template images with the star at the exact center of the image, pixel 51 51 as in this script.

In general, when making new planetdetect01 scripts from old ones, I change nothing but the image names and the radius within which the sensitivity should not be attempted. This last should usually be set just a pixel or two beyond the saturation radius, and this, of course, is very different for stars of different brightnesses. For very faint primaries where the images are not saturated, I suppose it could be set to zero. Choosing 2.0 or 3.0 instead will keep planetdetect01 from reporting the primary as a detected source.

The output of planetdetect01 is the two files named at the beginning of the script. The file here represented by GJ3541Adetsall\_a.txt will contain all the sources reported as detections on any of the images. For the parameters I give here, there will be hundreds of false positives, so this file is not very useful. I don't usually look at it. The second file, here represented by GJ3541Adetsconf\_a.txt, will hold confirmed detections, that is, sources found at a consistent location on all images. This is the useful one. It will usually still have a few false positives. All the sources it reports should be examined carefully by hand. Then the rather difficult decision whether or not to follow up dubious sources must be made. Note that both these files explicitly report sources detected by the aperture photometry method and the PSF fitting method separately. Generally I trust PSF fitting a bit more, but there is good reason to pay attention to both. Clear, high significance

sources are always detected by both methods. Less obvious but nonetheless real ones may be detected by one and/or the other, with no consistent pattern. Very bright sources usually yield multiple detections.

#### B.14 Post-legolas processing: Sensitivity Estimation

There are two main programs for sensitivity estimation in the legolas suite: `errradplot01` and `percentileplot01`. Note that in all these cases **the magnitudes used are apparent magnitudes , not absolute magnitudes**. The apparent-to-absolute magnitude conversion, when needed to calculate planet masses from theoretical models, is done internally in each program. The user never has to work in absolute magnitudes.

##### B.14.1 `errradplot01`

The first one, `errradplot01`, takes aperture-summed and PSF-fitted images output by `planetdetect01` and uses them to create final sensitivity maps, giving the point source sensitivity at every point in the image. `Errradplot01` is designed to produce a master sensitivity image from a whole set of final stacked images from a `legolas04` run. Thus, the entire range of possible `legolas04` scripts may be run, and the output images fed into `errradplot01` to make a final sensitivity map that gives, at each location in the image, the best sensitivity obtained by any of the different reduction scripts. For example, the PSF subtraction scripts will do the best close-in; farther out the baseline ‘a’ script or the pre-stack unsharp masking ‘d’ script will do the best; in regions affected by nod subtraction artifacts the master sky image scripts ‘x’ and ‘y’ will do the best. `Errradplot01` combines all this information into a single sensitivity map giving the best sensitivity obtained at every point.

`Errradplot01` is run in the usual way, with a script:

```
./errradplot01.exe < errrdpl01script02
```

Here is an example of an errradplot01 script:

```
6
500 500
GJ3541Aa0101t.fits 101 101
250.0 250.0 60.0 45
7.0
0.0 0.0
51 51
1.0 30.0 0.1
1.0
GJ3541Aa02u01__amp01.fits
GJ3541Aa02u01__smooth01.fits
GJ3541Aa02u01.fits
0
0
0
8.0
10.0
GJ3541Aa02u01__err1.fits
GJ3541Aa02u01__err2.fits
GJ3541Ab02u01__amp01.fits
GJ3541Ab02u01__smooth01.fits
GJ3541Ab02u01.fits
0
0
0
```

8.0

10.0

GJ3541Ab02u01\_err1.fits

GJ3541Ab02u01\_err2.fits

GJ3541Ad02u01\_amp01.fits

GJ3541Ad02u01\_smooth01.fits

GJ3541Ad02u01.fits

0

0

0

8.0

10.0

GJ3541Ad02u01\_err1.fits

GJ3541Ad02u01\_err2.fits

GJ3541Ae02u01\_amp01.fits

GJ3541Ae02u01\_smooth01.fits

GJ3541Ae02u01.fits

0

0

0

8.0

10.0

GJ3541Ae02u01\_err1.fits

GJ3541Ae02u01\_err2.fits

GJ3541Ax02u01\_amp01.fits

GJ3541Ax02u01\_smooth01.fits

GJ3541Ax02u01.fits

0

0  
0  
8.0  
10.0  
GJ3541Ax02u01\_err1.fits  
GJ3541Ax02u01\_err2.fits  
GJ3541Ay02u01\_amp01.fits  
GJ3541Ay02u01\_smooth01.fits  
GJ3541Ay02u01.fits  
0  
0  
0  
8.0  
10.0  
GJ3541Ay02u01\_err1.fits  
GJ3541Ay02u01\_err2.fits

An explanation of the entries in this script is given in Table B.11.

Table B.11. Explanation of errradplot01 script

Entry	Explanation
6	Number of different legolas04 output images to process.
500 500	Dimensions of legolas04 output images.
GJ3541Aa0101t.fits 101 101	Name and dimensions of the unsaturated PSF image for this star.
250.0 250.0 60.0 45	Pixel coordinates of the primary star on these images, radius out to which sensitivity out to be calculated based on the RMS in constant-radius arcs rather than in disks, and length in pixels of the arcs to be used.
7.0	Radius from the primary star within which sensitivity calculation should <b>not</b> be attempted (ie saturation radius).
0.0 0.0	Inner and outer radii of a sky subtraction annulus for the PSF image. If no sky subtraction is desired, enter 0.0 0.0.
51 51	Pixel coordinates of the star on the PSF image. Note that the PSF <b>must</b> be centered on integer coordinates.

Table B.11—Continued

Entry	Explanation
1.0 30.0 0.1	<p>Inner and outer radii to consider in calculating the optimal aperture radius for source detection against the noise, and step size to use between these limits when searching for the optimal radius. Note that the outer radius must be the same as that used for aperture photometry of the photometric standard star images. this is typically 30 pix for <math>L'</math> and 10 pix for <math>M</math> band.</p>
1.0	<p>Strehl ratio if an idealized synthetic PSF is being used. If the PSF is from real data (recommended), enter 1.0.</p>

Table B.11—Continued

Entry	Explanation
General parameters end here, and those for the first legolas04 script ('a') images begin.	
GJ3541Aa02u01_amp01.fits	Name of PSF amplitude image output by planetdetect01 for legolas04 'a' script output image.
GJ3541Aa02u01_smooth01.fits	Name of aperture photometry image output by planetdetect01 for legolas04 'a' script output image.
GJ3541Aa02u01.fits	Name of legolas04 master image for the 'a' script..
0	Number of boxes to set to zero sensitivity due to bad data.
0	Number of circles to set to zero sensitivity due to bad data.
0	Number of annular sectors to set to zero sensitivity due to bad data.
8.0	Radius of disk in which the RMS is to be calculated for sensitivity estimation. This applies only for image regions more distant from the the star than the limit for calculating the sensitivity in arcs

Table B.11—Continued

Entry	Explanation
	that has already been set above.
10.0	Sigma level for sensitivity calculation. Here, $10\sigma$ sensitivities will be calculated
GJ3541Aa02u01_err1.fits	Name of output sensitivity map for the legolas04 'a' script, based on aperture photometry estimation method.
GJ3541Aa02u01_err2.fits	Name of output sensitivity map for the legolas04 'a' script, based on PSF fitting estimation method.
End of parameters for first ('a') legolas04 script images; beginning of those for second ('b').	
GJ3541Ab02u01_amp01.fits	Name of PSF amplitude image output by planetdetect01 for legolas04 script 'b' output image.
GJ3541Ab02u01_smooth01.fits	Name of aperture photometry image output by planetdetect01 for first legolas04 'b' script output image.
GJ3541Ab02u01.fits	Name of legolas04 master image for the 'b' script.
0	Number of boxes to set to zero sensitivity due to bad data.
0	Number of circles to set to zero

Table B.11—Continued

Entry	Explanation
	sensitivity due to bad data.
0	Number of annular sectors to set to zero sensitivity due to bad data.
8.0	Radius of disk in which the RMS is to be calculated for sensitivity estimation. This applies only for image regions more distant from the star than the limit for calculating the sensitivity in arcs that has already been set above.
10.0	Sigma level for sensitivity calculation. Here, $10\sigma$ sensitivities will be calculated
GJ3541Ab02u01_err1.fits	Name of output sensitivity map for the legolas04 'b' script, based on aperture photometry estimation method.
GJ3541Ab02u01_err2.fits	Name of output sensitivity map for the legolas04 'b' script, based on PSF fitting estimation method.
End of parameters for second ('b') legolas04 script images; beginning of those for third ('d').	
GJ3541Ad02u01_amp01.fits	etc; analagous to entries above
GJ3541Ad02u01_smooth01.fits	etc; analagous to entries above

Table B.11—Continued

Entry	Explanation
GJ3541Ad02u01.fits	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
8.0	etc; analagous to entries above
10.0	etc; analagous to entries above
GJ3541Ad02u01_err1.fits	etc; analagous to entries above
GJ3541Ad02u01_err2.fits	etc; analagous to entries above
End of parameters for third ('d') legolas04 script images; beginning of those for fourth ('e').	
GJ3541Ae02u01_amp01.fits	etc; analagous to entries above
GJ3541Ae02u01_smooth01.fits	etc; analagous to entries above
GJ3541Ae02u01.fits	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
8.0	etc; analagous to entries above
10.0	etc; analagous to entries above
GJ3541Ae02u01_err1.fits	etc; analagous to entries above
GJ3541Ae02u01_err2.fits	etc; analagous to entries above
End of parameters for fourth ('e') legolas04 script images; beginning of those for fifth ('x').	

Table B.11—Continued

Entry	Explanation
GJ3541Ax02u01__amp01.fits	etc; analagous to entries above
GJ3541Ax02u01__smooth01.fits	etc; analagous to entries above
GJ3541Ax02u01.fits	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
8.0	etc; analagous to entries above
10.0	etc; analagous to entries above
GJ3541Ax02u01__err1.fits	etc; analagous to entries above
GJ3541Ax02u01__err2.fits	etc; analagous to entries above
End of parameters for fifth ('x') legolas04 script images; beginning of those for sixth ('y').	
GJ3541Ay02u01__amp01.fits	etc; analagous to entries above
GJ3541Ay02u01__smooth01.fits	etc; analagous to entries above
GJ3541Ay02u01.fits	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
0	etc; analagous to entries above
8.0	etc; analagous to entries above
10.0	etc; analagous to entries above
GJ3541Ay02u01__err1.fits	etc; analagous to entries above
GJ3541Ay02u01__err2.fits	etc; analagous to entries above

Table B.11—Continued

Entry	Explanation
End of parameters for legolas04 script images; beginning of final output parameters.	
GJ3541A_errmain1.fits	Name of master sensitivity map based on the aperture photometry estimation method.
GJ3541A_errmain2.fits	Name of master sensitivity map based on the PSF amplitude fitting estimation method
1	number of stars on final images, around which sensitivity should be calculated.
0.048574	Image plate scale in arcsec/pixel
17.75	Distance to star in parsec.
2.0596	True single-frame exposure time in seconds.
16825	Photometric calibration: total ADU per second recieved from a 10th magnitude star within a 30.0 pixel radius aperture.
GJ3541Aradplot01.txt GJ3541Aradplot02.txt	Names of radial sensitivity plot text files for the aperture photometry and then PSF amplitude fitting estimation methods.
250.0 250.0 8.0 150.0	Coordinates of first (and only, in

Table B.11—Continued

Entry	Explanation
	this case) star around which the sensitivity should be plotted, and inner and outer radii in pixels for which the sensitivity should be plotted.

The length of this script may appear intimidating. However, like the planet-detect01 script, there are few judgement calls and few parameters that need to be changed from one star to the next. The length of the example script here is also extended by the fact that, for realism, I have included images from six different legolas04 scripts. This is the typical number for reduction of a good quality data set on a star that does not have a bright binary companion in the same Clio field.

The only judgement call in the general parameters section is the radius within which sensitivity should not be calculated. This should be set to the saturation radius, or perhaps one or two pixels beyond it. Everything else except the names and the number of different legolas04 scripts can be copied without alteration from one star to another. Note, however, that **the outer radius for calculating the optimal aperture for source detection must be equal to the radius of the aperture used to obtain the photometric calibration from the standard stars.** This is a very non-intuitive kludge to feed this information into the program without an additional script line; my apologies. The information is needed to map the flux within the small (2-3 pixel radius) optimal aperture used for the aperture photometry sensitivity calculation to the flux contained within the much larger aperture used to obtain photometric calculation from the standard stars. The small optimal aperture is designed to maximize the SNR of faint sources against the noise; the large standard star photometry aperture is designed to consistently include the vast majority of flux from a target despite variations in seeing. For  $L'$  data our standard star photometry radius is 30 pixels; for  $M$  band it is 10.0 pixels.

In the script above there are no judgement calls in the blocks dealing with images output by each of the six legolas04 reduction scripts. The 8.0 pixel radius for the RMS calculation disk and the 10.0  $\sigma$  sensitivity value should never be changed. It is very rare that there is any occasion to set the sensitivity to zero in

any boxes or annular sectors.

It does sometimes happen, however, that it is necessary to set the sensitivity to zero *on the PSF-subtracted 'b' and 'e' script images only* within a circular region centered on the star. This situation arises when the data set has insufficient parallactic rotation. The 'b' and 'e' script PSF subtraction methods will then severely dim any sources within a certain radius (see Section B.7). We usually estimate this radius as  $R_{min} = 3.0\text{pix} \div \sin(\theta/2)$ , where  $\theta$  is the total parallactic rotation over the data set. This is the radius within which parallactic rotation moves a source through an arc of length less than the twice the typical PSF FWHM. If  $R_{min}$  is greater than the saturation radius, sensitivities calculated based on the 'b' and 'e' script images inward of  $R_{min}$  will be overestimated. Then the number of circular regions to zero should be set to 1 rather than 0. In this case, the next line in the script must give the pixel coordinates of the circle center and the radius of the circle. Thus, in the above example, if parallactic rotation had been small enough that  $R_{min}$  was larger than the 7-pixel saturation radius (say 15 pixels), rather than:

```
0
0
0
```

for the number of boxes, circles, and sectors to zero, we would enter:

```
0
1
250.0 250.0 15.0
0
```

to zero no boxes, 1 circle of radius 15.0 pixels centered at pixel coords 250,250 (the location of the primary star), and no sectors.

In the case of a binary star with a saturated companion it may also be desirable to zero a circle centered on the companion with radius equal to the saturation radius.

For the final parameters of `errradplot01`, there are no judgement calls except the inner and outer radii for sensitivity plotting. These are relevant only for the radial plot files here exemplified by `GJ3541Aradplot01.txt` and `GJ3541Aradplot02.txt`. They are a matter of personal taste, except that the inner radius should be greater than or equal to the saturation radius.

These radial plot files have 9 columns. The first 3 hold the radius from the star in pixels, arcseconds, and projected AU, respectively. The next 3 hold the minimum sensitivity obtained over the best 90%, 75%, and 50% of a circle centered on the star at that radius, in units of total ADU received from a point source on a single frame. The last 3 columns hold the same sensitivities in units of magnitudes. Note that another way to describe the minimum sensitivity obtained in the best 90% of the circle is the sensitivity level that was achieved or exceeded everywhere except the worst 10% of the circle, or the 10th percentile sensitivity at the given radius. Similarly, the 75% column holds the 25th percentile sensitivity, and the 50% column the 50th percentile, or median, value. Note that 'circle' in this context means a 1-D (1 pixel wide) circle at fixed radius about the primary star, **not** a disk centered on the star. The sensitivities in the columns differ because of azimuthal variations in image quality at a given radius. These can be due to nod subtraction artifacts, rays, ghosts, or other effects.

`Errradplot01` also outputs sensitivity maps based on the images from each individual `legolas04` script, and master sensitivity maps giving the best sensitivity obtained with any `legolas04` script at each location. The sensitivity on these images is in units of total ADU received on an individual frame.

These can be converted to magnitudes as follows. Suppose the errradplot01 output image at a given point records a  $10\sigma$  sensitivity of 100 ADU for a data set in which the single frame exposure time was 1.5596 sec. This means that a point source generating  $100/1.5596 = 64$  ADU/sec at the Clio detector would be detected at the  $10\sigma$  level at that location in the final stacked image. Note this is not ADU per pixel, but total ADU integrated over the PSF. For a typical  $L'$  band photometric calibration of 15,000 ADU/sec for a 10th magnitude star, a 64 ADU/sec source has an  $L'$  magnitude of 15.9. This is typical of our peak sensitivity in regions of clean sky for a data set with a total exposure time of one hour. The sensitivities in the maps output by errradplot01 can be translated into magnitudes easily by the equation we have just implicitly used:

$$S_{mag} = 10.0 + 2.5 \times \log \left( \frac{P_{cal}}{(S_{ADU}/t_{exp})} \right)$$

Where  $S_{mag}$  is the sensitivity in magnitudes,  $P_{cal}$  is the photometric calibration in total ADU/sec for a 10th magnitude star,  $S_{ADU}$  is the sensitivity in ADU from an errradplot01 sensitivity map, and  $t_{exp}$  is the true single-frame exposure time in seconds. Note well this should not be corrected by the number of coadds: a data set with in which each image is a 15-coadd stack of 1.5596 second frames has an effective exposure time of 23.39 sec for each image, but since legolas04 divides by the number of coadds the correct  $t_{exp}$  value for the equation above is 1.5596 sec, **not** 23.39 sec.

#### B.14.2 percentileplot01

The error images and radial plot files from errradplot01 are useful; however there are additional programs to translate them into more easily interpreted forms.

First, we have percentileplot01. It takes in the master sensitivity images out-

put by `errradplot01`, and produces output plots giving any desired percentiles of the sensitivity vs radius. Plots are given with radius in arcseconds and sensitivity in magnitudes, and also radius in projected AU and sensitivity in Jupiter masses (MJ). The resulting files do an excellent job of representing the true sensitivity obtained by Clio observations on a given source. They can be plotted, or fed into Monte-Carlo simulations to predict the odds of detecting a planet given different distributions of planetary parameters.

The program `percentileplot01` is run as usual, with a script:

```
./percentileplot01.exe < percenplot01script02
```

Here is an example of a `percentileplot01` script:

```
500 500
GJ3541A_errmain1.fits GJ3541A_errmain2.fits
10
0.0
10.0
20.0
30.0
40.0
50.0
60.0
70.0
80.0
90.0
1
0.048574
17.75
```

```
2.0596
30.0
16825.0
Baraffe_Lmags01.txt
0.1
0.7
0.0
GJ3541Apercenplot01.txt GJ3541Apercenplot02.txt
250.0 250.0 8.0 240.0
```

The entries in this script are explained in Table B.12.

Table B.12. Explanation of percentileplot01 script

Entry	Explanation
500 500	Dimensions of input images
GJ3541A_errmain1.fits GJ3541A_errmain2.fits	Names of input master sensitivity maps based on aperture photometry and PSF amplitude fitting.
10	Number of sensitivity percentiles to plot at each radius
0.0	Percentile # 1: 0th percentile, or worst sensitivity obtained at each radius
10.0	Percentile # 2: 10th percentile of sensitivity obtained at each radius
20.0	Percentile # 3: 20th percentile of sensitivity obtained at each radius
30.0	Percentile # 4: 30th percentile of sensitivity obtained at each radius
40.0	Percentile # 5: 40th percentile of sensitivity obtained at each radius

Table B.12—Continued

Entry	Explanation
50.0	Percentile # 6: 50th percentile of sensitivity obtained at each radius
60.0	Percentile # 7: 60th percentile of sensitivity obtained at each radius
70.0	Percentile # 8: 70th percentile of sensitivity obtained at each radius
80.0	Percentile # 9: 80th percentile of sensitivity obtained at each radius
90.0	Percentile # 10: 90th percentile of sensitivity obtained at each radius. The 100th percentile, or best sensitivity obtained at each radius could also be plotted, but we have not done so here.
1	Number of stars around which sensitivity should be calculated. Always 1 unless the target is binary

Table B.12—Continued

Entry	Explanation
0.048574	Clio plate scale in asec per pixel.
17.75	Distance to star in parsecs.
2.0596	Single frame exposure in seconds.
30.0	Photometric aperture radius used for standard stars.
16825.0	Photometric calibration from standard stars: ADU/sec from a 10th mag star within the aperture radius above.
Baraffe_lmags01.txt	Name of file giving the mass vs magnitude conversion based on Baraffe (2003)
0.1	Age of star system in Gyr
0.7	Irrelevant dummy parameter.
0.0	Fudge value in magnitudes to be added to the sensitivities
GJ3541Apercenplot01.txt GJ3541Apercenplot02.txt	Name of ouput sensitivity plot files based on aperture photometry and PSF amplitude fitting.
250.0 250.0 8.0 240.0	Coordiantes of first (and in this case only) star

Table B.12—Continued

Entry	Explanation
	around which sensitivity should be plotted, and the inner and outer radii in pixels for the plot.

The percentileplot01 scripts are fairly simply to write. The number and values of percentiles to plot is of course simply a matter of taste. The Baraffe magnitude file should simply be obtained from me (Ari Heinze) along with the programs. The fudge value in magnitudes can be used, eg, to plot  $7\sigma$  sensitivities if errradplot01 made  $10\sigma$  images. In this case the ratio  $10/7$  in magnitudes should be added to the sensitivities. This fudge correction will be a positive 0.387 mag. The inner and outer radii for plotting sensitivities are almost irrelevant, as percentileplot handles out-of-range values gracefully.

#### B.14.3 Additional Auxiliary Programs

A few more conversions of sensitivity data may be desirable. For example, so far no provision has been made for an airmass correction to the sensitivities. Also, the magnitude+age to mass conversion is done in percentileplot01 only using the theoretical models of Baraffe et al. (2003). The models of Burrows et al. (2003) are also available and may be more accurate. To offer this flexibility there is the program convpercen01.exe. It is run in the usual way:

```
./convpercen01.exe < convp01script03
```

It takes in plot files output by percentileplot01, applies an airmass correction to the magnitudes, and re-calculates the masses using either the models of Burrows et al. (2003) or those of Baraffe et al. (2003). Here is an example of a convpercen01 script:

```
3
0
GJ659Apercenplot02.txt GJ659Apercenplot06.txt
1.011 1.113 0.086
233
```

Burrows\_Lmags01.txt

1.0

20.2

0.7

GJ3541Apercenplot02.txt GJ3541Apercenplot06.txt

1.011 1.032 0.086

233

Burrows\_Lmags01.txt

0.1 18.87

0.7

GJ450percenplot02.txt GJ450percenplot06.txt

1.011 1.031 0.086

234

Burrows\_Lmags01.txt

1.0 8.1

0.7

The entries are explained in Table B.13

Table B.13. Explanation of convpercen01 script

Entry	Explanation
3	Number of percenplot01 files to convert.
0	0 to use the Burrows models, 1 to use the Baraffe ones
GJ659Apercenplot02.txt GJ659Apercenplot06.txt	Names of first input percenplot file, and of the converted output file.
1.011 1.113 0.086	Mean airmass of photometric standard star observation, mean airmass of target observation, atmospheric extinction in mag/airmass.
233	Number of data rows in input percenplot file.
Burrow_sImags01.txt	Model file for mass conversion.
1.0	Age of star in Gyr.
20.2	Distance to star in pc.
0.7	Irrelevant dummy parameter.
GJ3541Apercenplot02.txt GJ3541Apercenplot06.txt	Names of second input percenplot file, and of the converted output file.
1.011 1.032 0.086	etc; as above

Table B.13—Continued

Entry	Explanation
233	etc; as above
Burrows_Lmags01.txt	etc; as above
0.1 18.87	etc; as above
0.7	etc; as above
GJ450percenplot02.txt GJ450percenplot06.txt	Names of third input percenplot file, and of the converted output file.
1.011 1.031 0.086	etc; as above
234	etc; as above
Burrows_Lmags01.txt	etc; as above
1.0 8.1	etc; as above
0.7	etc; as above

There are no real judgement calls here. Only the magnitudes are read from the input percenplot file. The masses are completely recalculated, and the arcsec to projected AU conversion is redone as well. Thus this program may also be used to translate a percenplot file to a different stellar distance and age than was initially used. The number of data rows, of course, refers only to the number of rows in the arcsec vs magnitude section of the input percenplot file. In the example of the GJ 354.1A percentile plot script, sensitivity would be calculated for radii from 8.0 to 240.0 pixels. It is done every integer pixel value, so accounting for fencepost error there are 233 rows. This, of course, is the value given in the convpercen01 script.

We also have a program, `pmassbar01`, which uses the same theoretical model files and interpolation methods as `percentileplot01` and `convpercen01` to convert any number of input magnitudes to planet masses using either the Baraffe et al. (2003) or the Burrows et al. (2003) models, and given an input age and distance for the star. It is simple enough to be run interactively without a script:

```
./pmassbar01.exe
```

The user can then simply enter the data the program requests, and convert a single magnitude to a planetary mass value. Afterward, to convert longer lists of magnitudes, scripts can easily be constructed.

## REFERENCES

- Aumann, H., Gillett, F., Beichman, A., de Jong, T., Houck, J., Low, F., Neugebauer, G., Walker, R., & Wesselius, P. 1984, *ApJ*, 278, L23
- Aumann, H. 1985, *PASP*, 97, 885
- Backman, D. & Paresce, F. 1993, in *Protostars and Planets III*, ed. E. H. Levy & J. I. Lunine (Tucson: Univ. Arizona Press), 1253
- Backman, D. 1996, *AAS/DPS 28.0122B*
- Baraffe, I., Chabrier, G., Barman, T. S., Allard, F., & Hauschildt, P. H. 2003, *A&A*, 402, 701
- Barrado y Navascués, D. 1998, *A&A*, 339, 831
- Benedict, G., McArthur, B., Gatewood, G., Nelan, E., Cochran, W., Hatzes, A., Endl, M., Wittenmyer, R., Baliunas, S., Walker, G., Yang, S., Kürster, M., Els, S., & Paulson, D. 2006, *AJ*, 132, 2206
- Bessell, M. & Brett, J. 1988, *PASP*, 100, 1134
- Biller, Beth A.; Close, Laird M.; Masciadri, Elena; Lenzen, Rainer; Brandner, Wolfgang; McCarthy, Donald; Henning, Thomas; Nielsen, Eric L.; Hartung, Markus; Kellner, Stephan; Geissler, Kerstin; & Kasper, Markus 2006, *Proc. SPIE*, 6272E, 74B
- Biller, B., Kasper, M., Close, L., Brandner, W., & Kellner, S. 2006, *ArXiv e-prints*, astro-ph/0601.440

- Biller, B., Close, L., Masciadri, E., Nielsen, E., Lenzen, R., Brandner, W., McCarthy, D., Hartung, M., Kellner, S., Mamajek, E., Henning, T., Miller, D., Kenworthy, M., & Kulesa, C. 2007, ArXiv e-prints, astro-ph/0705.0066
- Bryden G., Beichman, C., Trilling, D., Rieke, G., Holmes, E., Lawler, S., Stapelfeldt, K., Werner, M., Gautier, T., Blaylock, M., Gordon, K., Stansberry, J., & Su, K. 2006, *ApJ*, 636, 1098
- Burrows, A., Sudarsky, D., & Lunine, J. I. 2003, *ApJ*, 596, 578B
- Chauvin, G., Lagrange, A., Zuckerman, B., Dumas, C., Mouillet, D., Song, L., Beuzit, J., Lowrance, P., & Bessel, M. 2005, *A&A*, 438, L29
- Cohen, M., Walker, R. G., Michael, J. B., & Deacon, J. R. 2003, *AJ*, 104, 1650
- Cox, A. N. 2000, *Allen's Astrophysical Quantities* (Fourth Edition; New York, NY: Springer-Verlag New York, Inc.)
- Cutri, R., Skrutskie, M., van Dyk, S., Beichman, C., Carpenter, J., Chester, T., Cambresy, L., Evans, T., Fowler, J., Gizis, J., Howard, E., Huchra, J., Jarett, T., Kopan, E., Kirkpatrick, J., Light, R., Marsh, K., McCallon, H., Schneider, S., Stiening, R., Sykes, M., Weinberg, M., Wheaton, W., Wheelock, S., & Zacarias, N. 2003, 2MASS All Sky Catalog of point sources (The IRSA 2MASS All-Sky Point Source Catalog, NASA/IPAC Infrared Science Archive. <http://irsa.ipac.caltech.edu/applications/Gator/>)
- Deller, A. & Maddison, S. 2005, *ApJ*, 625, 398
- Favata, F., Micela, G., Sciortino, S., & D'Antona, F. 1998, *A&A*, 335, 218
- Fischer, D. 1998, PhD Thesis entitled **Lithium Abundances in Field K Dwarfs**, University of California Santa Cruz

- Fischer, D. A. & Valenti, J. 2005, *ApJ*, 622, 1102
- Freed, M., Hinz, P., Meyer, M., Milton, N., & Lloyd-Hart, M. 2004, *Proc. SPIE*, 5492, 1561
- Geißler, K.; Kellner, S.; Brandner, W.; Masciadri, E.; Hartung, M.; Henning, T.; Lenzen, R.; Close, L.; Endl, M.; & Kürster, M. 2007, *A&A*, 461, 665G
- Gillett, F., Aumann, H., & Low, F. 1984, paper presented at the Meeting on Protostars and Planets, University of Arizona, Tucson, January 1984
- Gillett, F., Aumann, H., Neugebauer, G., Low, F., & Walters, R. 1985, in *Proceedings of the First International IRAS symposium*, June 10, 1985, Noordwijk, The Netherlands.
- Greaves, J., Holland, W., Moriarty-Schieven, G., Jenness, T., Dent, W., Zuckerman, B., McCarthy, C., Webb, R., Butner, H., Gear, W., & Walker, H. 1998, *ApJ*, 506, L133
- Hatzes, A., Cochran, W., McArthur, B., Baliunas, S., Walker, G., Campbell, B., Irwin, A., Stephenson, Y., Kürster, M., Endl, M., Els, S., Butler, R., & Marcy, G. 2000, *ApJ*, 544, 145
- Hinz, P., Heinze, A., Sivanandam, S., Miller, D., Kenworthy, M., Brusa, G., Freed, M., & Angel, J. 2006, *ApJ*, 653, 1486.
- Holland, W. et al. 1998, *Nature*, 392, 788
- Hünsch, M., Schmitt, J., Sterzik, M. & Voges, W. 1998, *Astron. Astrophys. Suppl. Ser.* 135, 319

- Janson, M., Brandner, W., Henning, T., Lenzen, R., McArthur, B., Benedict, G., Reffert, S., Nielsen, E., Close, L., Biller, B., Kellner, S., Güther, E., Hatzes, A., Masciadri, E., Geissler, K., & Hartung, M. 2007, *AJ*, 133, 2442
- Johnson, J., Butler, R., Marcy, G., Fischer, D., Vogt, S., Wright, J., & Peek, K. 2007, ArXiv e-prints, astro-ph/0707.2409
- Juric, M. & Tremaine, S. 2007, ArXiv e-prints, astro-ph/0703.0160
- Karatas, Y., Bilir, S., & Schuster, W. 2005, *MNRAS*, 360, 1345
- Kasper, M., Apai, D., Janson, M., & Brandner, W. 2007, ArXiv e-prints, astro-ph/0706.0095
- Kenworthy, M., Codona, J., Johanan, L., Hinz, P., Angel, J., Heinze, A., & Sivanandam, S. 2007, *ApJ*, 660, 762
- King, J., Villarreal, A., Soderblom, D., Gulliver, A., & Adelman, S. 2003, *ApJ*, 125, 1980
- Koerner, C., Sargent, A., Ostroff, N. 2001, *ApJ*, 560, L181
- Lachaume, R., Dominik, C., Lanz, T., & Habing, H. 1999, *A&A*, 348, 897
- Lafrenière, D., Doyon, R., Marois, C., Doyon, R., Nadeau, D., & Artigau, E. 2007, *ApJ*, 660, 770
- Lafrenière, D., Doyon, R., Marois, C., Nadeau, D., Oppenheimer, B., Roche, P., Rigaut, F., Graham, J., Jayawardhana, R., Johnstone, D., Kalas, P., Macintosh, B., & Racine, R. 2007, ArXiv e-prints, astro-ph/0705.4290
- Latham, D., Stefanik, R., Tsevi, M., Mayor, M., & Gilbert, B. 1989, *Nature*, 339, 38.

- Leggett, S., Hawarden, T., Currie, M., Adamson, A., Carrol, T., Kerr, T., Kuhn, O., Seigar, M., Varricatt, W., & Wold, T. 2003, *MNRAS*, 345,144
- Leggett, S., Saumon, D., Marley, M., Geballe, T., Golimowksi, D., Stephens, D., & Fan, X. 2007, *ApJ*, 655, 1079
- Gomez, R., Levison, H., Tsiganis, K., & Morbidelli, A. 2005 *Nature*, 435, 466
- Li, A. & Lunine, J. 2003, *ApJ*, 590, 368
- Lopez-Santiago, J., Montes, D., Crespo-Chacon, I., & Fernandez-Figueroa, M. 2006, *ApJ*, 643, 1160
- Lowrance, P., Becklin, E., Schneider, G., Kirkpatrick, J., Weinberger, A., Zuckerman, B., Dumas, C., Beuzit, J., Plait, P., Malumuth, E., Heap, S., Terrile, R., & Hines, D. 2005, *AJ*, 130, 1845
- Marcy, G. & Butler, R. 1992, *PASP*, 104, 270
- Marley, M. S., Fortney, J. J., Hubickyj, O., Bodenheimer, P., & Lissauer, J. J. 2007, *ApJ*, 655, 541
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, 641, 556.
- Marsh, K., Dowell, C., Velusamy, T., Grogan, K., Beichman, C. 2006 *ApJ*, 646, 77.
- Masciadri, E.; Mundt, R.; Henning, Th.; Alvarez, C.; & Barrado y Navascués, D 2005, *ApJ*, 625, 1004M
- Mayor, M. & Queloz, D. 1995, *Nature*, 378, 355
- Montes, D., López-Santiago, J., Fernaández-Figueroa, M., & Gálvez, M. 2001, *A&A*, 379, 976

- Nielson, E., Close, L., & Biller, B. 2007, ArXiv e-prints, astro-ph/0706.4331
- Neuhäuser, R.; Brandner, W.; Eckart, A.; Guenther, E.; Alves, J.; Ott, T.; Huélamo, N.; & Fernández, M 2000, A&A, 354L, 9N
- Neuhäuser, R., Guenther, E., Wuchterl, G., Mugrauer, M., Bedalov, A., & Hauschildt, P. 2005, A&A, 435, L13
- Ozernoy, L., Gorkavyi, N., Mather, J., & Taidakova, T. 2000, ApJ, 537, L147
- Perryman, M., Lindegren, L., Kovalevsky, J., Hog, E., Bastian, U., Bernacca, P., Creze, M., Donati, F., Grenon, M., Grewing, M., Van Leeuwen, F., Van Der Marel, H., Mignard, F., Murray, C., Le Poole, R., Schrijver, H., Turon, C., Arenou, F., Froeschle, M., & Peterson, C. 1997, A&A, 323, L49
- Potter, D., Cushing, M., & Neuhauser, R. 2003, csss, 12, 689
- Press, W. H., Teukolsky, S.A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical Recipes in C (Second Edition; New York, NY: Cambridge University Press)
- Quillen, A. & Thorndike, S. 2002, ApJ, 578, L149
- Reid, I. & Cruz, K. 2002, AJ, 123, 466
- Sivanandam, S., Hinz, P., Heinze, A., & Freed, M. 2006, Proc. SPIE, 6269, 27
- Song, I., Caillault, J.-P., Barrado y Navascués, D., & Stauffer, J. 2001, ApJ, 546, 352
- Song, I., Zuckerman, B., & Bessell, M. 2003, ApJ, 599, 342
- Song, I., Schneider, G., Zuckerman, B., Farihi, J., Becklin, E., Bessell, M., Lowrance, P., & Macintosh, B. 2006, ApJ, 652, 724
- Wichmann, R., Schmitt, J., & Hubrig, S. 2003, A&A, 399, 983

Wichmann, R. & Schmitt, J. 2003, MNRAS, 342, 1021

Wilner, D. 2004, ASP Conference Series, Vol. 324.

Wilson, J., Hinz, P., Kenworthy, M., Skrutskie, M., Jones, T., Nelson, M., Woodward, C., & Garnavich, P. 2007, Proceedings of the conference In the Spirit of Bernard Lyot: The Direct Detection of Planets and Circumstellar Disks in the 21st Century. June 04 - 08, 2007. University of California, Berkeley, CA, USA.  
Edited by Paul Kalas

Wolszczan, A. & Frail, D. 1992, Nature, 355, 145

Wyatt, M. 2003, ApJ, 598, 1321

Yoichi, I., Oasa, Y., & Fukagawa, M. 2006, ApJ, 652, 1729

Zuckerman, B., Song, I., Bessell, M., & Webb, R. 2001, ApJ, 562, L87