

PLANETARY SYSTEMS IN POLARIZED LIGHT: DEBRIS DISK OBSERVATIONS AND
INSTRUMENTATION

by

Maxwell A. Millar-Blanchaer

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Astronomy & Astrophysics
University of Toronto

© Copyright 2016 by Maxwell A. Millar-Blanchaer

Abstract

Planetary Systems in Polarized Light: Debris Disk Observations and Instrumentation

Maxwell A. Millar-Blanchaer

Doctor of Philosophy

Graduate Department of Astronomy & Astrophysics

University of Toronto

2016

Understanding planet formation is one of the major challenges of modern astronomy. Polarimetry is a powerful tool with which we can confront this challenge. In particular, polarimetric observations can be useful for imaging debris disks and characterizing exoplanet atmospheres. With that in mind, this thesis has been constructed with two main aspects: i) observational studies of two debris disk systems, β Pic and HD 157587, using the Gemini Planet Imager and ii) the characterization and testing of a new type of diffraction grating, called a polarization grating, that we plan to use for future observations of exoplanet atmospheres.

The Gemini Planet Imager is a high-contrast imager that includes a polarimetry mode designed to image circumstellar disks. Here we detail the development of new data analysis techniques that reduce systematics and noise in processed GPI data. We apply these techniques to observations of the β Pic and HD 157587 debris disks and then fit each disk image to a geometric disk model. The β Pic disk model's morphology cannot be explained by interactions with the planet β Pic *b*, and the presence of a second planet could be invoked to explain the discrepancy. In the case of HD 157587, the disk model's geometric centre is offset from the location of the star, which could be explained by a perturbing planet. Characterization of the planets' interactions with their debris disks is a critical method to gain more information about these two systems.

The second component of this thesis focuses on polarization gratings, thin film optical devices that can simultaneously act as polarizing beam splitters and as spectral dispersive elements. Moreover, they can be designed for high diffraction efficiency across a broad wavelength range. These features make polarization gratings useful for many types of astronomical observations. We have carried out laboratory and on-sky test observations using a polarization grating optimized for visible wavelengths. The laboratory tests confirm the expected diffraction efficiency and beam splitting capabilities of the grating. Our on-sky observations demonstrate the grating's ability to measure linear polarization fraction and position angle, and recover spectra in an astronomical setting. In the future we plan to use a near-infrared polarization grating to search for spectropolarimetric features in exoplanet atmospheres.

To Fabienne, Mom, Dad, Mar, Rainer and Wallis.

Acknowledgements

This thesis could not have been possible without the support of a strong community that includes many mentors, family, friends and loved ones.

Mom and Dad, you have always encouraged me to follow my dreams and have provided me with the opportunities to do so. Mar, you have been like another parent to me and have always been there to prop me up when I'm down. All three of you have always believed in me and I am truly privileged to have grown up in such a positive and nurturing environment. Rainer and Wallis, I love how close we've grown over the past number of years. The two of you always keep me grounded and give perspective to my life.

Dae-Sik, we have now known each other for over 8 years during which you have given your time and effort to make sure that I succeed to the best of abilities. You have taught me to expect more from myself and have drawn out the best in me. Throughout the years I have never doubted that you were always looking out for me and had my best interests at heart. It was working with you that I experienced my first real introduction to instrumentation and I wouldn't be where I am now without you.

James, you have provided invaluable guidance over the years and though much of our collaboration was from a distance, you have had a profound influence on my studies. You have set an excellent example as a scientist and a mentor. I sincerely hope to continue to work together into the future.

Paul, it has been a great pleasure getting to know you and working with you over the past few years. I have treasured our many conversations, both academic and not, and I expect to have many more in the coming years.

To my many colleagues at U of T, it has been fantastic working with you over the past five years and you have all had an influence on both my work and my personal well-being. Steph, I loved sharing the office with you and I am certain I'll never have another officemate as good as you. I miss you and the office hasn't been the same since you left. In addition, I'd like to specifically thank Nick, Heidi, Stephen, Lauren, Liam, Ari and Mubdi for all the good times over the years.

Elexa, Sam, Thomas, Carly, Max, Alex, Gergely, Simon, Patrick, Jeff and Kyle, you are and will forever be an important part of me. I love you all and I will miss you terribly as I move away. I expect visits from all of you soon.

Most importantly, Fabienne, I could not have done this without you and I want to thank you with all my heart for the support and warmth that you've provided to me over the past 4 years. I hope that in the next year I can provide that same support to you as you finish your thesis. I love you.

Contents

1	Introduction	1
1.1	Debris Disk Observations with GPI	1
1.1.1	Debris Disk Formation and Evolution	3
1.1.2	Observations of Debris Disks	4
1.1.3	The Gemini Planet Imager	6
1.2	Polarization Gratings	7
1.3	This Work	9
2	Gemini Planet Imager Observational Calibrations: Polarimetric Contrasts and New Data Reduction Techniques	10
2.1	Chapter Summary	10
2.2	Introduction	10
2.3	Data Reduction and Contrast Measurement in Polarimetry Mode	11
2.4	Assembling Datacubes from Raw Data	14
2.4.1	Persistence	15
2.5	Flat Fielding	16
2.6	Instrumental Polarization	17
2.6.1	Unpolarized Standard Stars	18
2.6.2	Subtracting Apparent Stellar Polarization from Polarization Datacubes	21
2.7	GPI Exoplanet Survey Contrasts	23
2.8	Conclusions	26
3	β Pictoris' Inner Disk in Polarized Light and New Orbital Parameters for β Pictoris <i>b</i>	28
3.1	Chapter Summary	28
3.2	Introduction	29
3.3	Observations and Data Reduction	31
3.3.1	Polarimetry Mode Observations	32
3.3.2	Spectral Mode Observations	34
3.4	Disk Results	36
3.4.1	Disk Modeling	37
3.5	Planet Results	41
3.5.1	Astrometry in Spectroscopy Mode	41
3.5.2	Astrometry in Polarimetry Mode	41

3.5.3	Orbit fitting	42
3.5.4	Planet Polarization	43
3.6	Discussion	45
3.6.1	The Debris Disk	45
3.6.2	β Pic b	48
3.6.3	The Disk-Planet Interaction	49
3.7	Conclusion	52
4	Imaging an 80 AU Radius Dust Ring Around the F5V Star HD 157587	54
4.1	Chapter Summary	54
4.2	Introduction	54
4.3	Observations and Data Reduction	56
4.4	Results	57
4.4.1	Polarized Intensity Image	57
4.4.2	Total Intensity Image	58
4.5	Disk Modeling	60
4.6	Discussion	64
4.7	Conclusion	66
5	Polarization Gratings for Visible and Near-infrared Astronomy	68
5.1	Chapter Summary	68
5.2	Introduction	68
5.3	Polarization Gratings	69
5.4	Laboratory Characterization	71
5.4.1	Lab Set-up	71
5.4.2	Lab Measurements	72
5.5	Scientific Application of Polarization Gratings	75
5.5.1	Polarization Grating-Based Spectrograph Designs	75
5.6	Future Plans and Conclusions	76
6	On-sky Observational Tests of a Polarization Grating for Visible Astronomy	77
6.1	Chapter Summary	77
6.2	Introduction	77
6.3	Polarization Grating Instrument Design	78
6.4	Data Reduction	80
6.4.1	Linear Polarization Extraction	84
6.5	Observational Test Results	86
6.5.1	Twilight Sky Observations	86
6.5.2	Polarized Standard - HD 183143	87
6.5.3	Spectrophotometric Calibration	89
6.6	Conclusion	89

7	Conclusions	92
7.1	Thesis Summary	92
7.2	Future work	93
7.2.1	GPI Data Analysis and Debris Disk Observations	93
7.2.2	Polarization Gratings	94
	Bibliography	95

List of Tables

2.1	Unpolarized Standard Stars Observed in Direct Mode as Polarimetric Calibrators	20
3.1	Summary of GPI Observations of β Pic	30
3.2	Model Parameters and Prior Distributions for the β Pic Disk Model	37
3.3	Orbit Fit Parameters and Prior Distributions for β Pic <i>b</i>	43
4.1	Properties of the Candidate Point Sources Around HD 157587	59
4.2	Best-Fit Disk Model Parameters for the HD 157587 Disk	62
5.1	Design Parameters for the Polarization Gratings Tested at the University of Toronto. . .	70
6.1	Telescope and PSP Design Parameters	79
6.2	PSP Observations on 2016 April 27 UT	87
6.3	Polarization Fraction and Position Angle Measurements of HD 183143	87

List of Figures

2.1	Polarimetric Contrasts for Different Datacube Assembly Methods	13
2.2	Persistence in a Raw Detector Image	14
2.3	Persistence Over Time in a Polarization Datacube	15
2.4	Polarized Flat-Field Signal	16
2.5	Polarized Flat-Field Correction in a Polarization Datacube	16
2.6	Polarized Flat-Field Correction in an Extended Observation Sequence	17
2.7	GPI Total Intensity and Linear Polarized Intensity Images Demonstrating Instrumental Polarization.	18
2.8	Instrumental Polarization in the Q_r and U_r Bases	19
2.9	Instrumental Polarization of GPI Estimated with Unpolarized Stars	21
2.10	Polarized Intensity Contrast After Instrumental Polarization Subtraction	22
2.11	Polarization Datacube Contrasts from GPIES	23
2.12	Contrasts from GPIES Observations of HD 10472	24
2.13	Contrast as a Function of Exposure Time for HD 10472	25
2.14	Polarized Intensity Contrasts at $0.25''$, $0.4''$ and $0.8''$ for all GPIES Disk Observations	25
3.1	The β Pictoris Debris Disk in Polarized Light with GPI	32
3.2	Best-Fit Disk Model Residuals and Polarized Systematics	33
3.3	Total Intensity Image of β Pic b	35
3.4	Best-Fit Disk Model Posterior PDFs	39
3.5	Best-Fit Disk Model Image for β Pic	40
3.6	Best-Fit β Pic b Orbital Posterior PDFs	44
3.7	Measured and Best-Fit R.A. and Dec. Offsets for β Pic b	45
3.8	Measured and Best-Fit Radial Velocity of β Pic b	46
3.9	Determining the Inner Radius with the Disk Model	47
3.10	The β Pic system in Three Orthogonal Projections	50
4.1	The HD 157587 Debris Disk in Polarized Light with GPI	55
4.2	Signal to Noise Map of the HD 157587 Debris Disk	57
4.3	Regions Used to Calculate the Brightness Asymmetry in the HD 157587 Disk	58
4.4	Total Intensity Image of the Point Sources Around HD 157587	60
4.5	Best-Fit Disk Model for HD 157587	63
5.1	A Photo of a Polarization Grating	70
5.2	A Raw Detector Image from the Test Setup and a Photo of the Test Setup	72

5.3	Diffraction Efficiency of the Polarization Grating	73
5.4	Modulation Efficiency of the Polarization Grating	74
5.5	Three Possible Polarization Grating Implementations	75
6.1	Photo of the PSP Instrument Taken In-lab	79
6.2	Cartoon Schematic of the Optical Layout of the PSP	80
6.3	Photo of the PSP Mounted on the 16-inch UofT Telescope	81
6.4	Example of a Single Raw Data Frame from the PSP	82
6.5	Reduced Spectra from a Single PSP Observation	83
6.6	Retardance of the Thorlabs QWP	85
6.7	Measured Position Angle of the Twilight Polarization	86
6.8	Polarized Spectrum of HD 183143	88
6.9	Flux-calibrated Spectrum of HD 106252	90

Chapter 1

Introduction

The understanding of planet formation processes is one of the main goals of modern astronomy. Since the discovery of the first exoplanet orbiting a main-sequence star, 51 Pegasi *b* (Mayor & Queloz 1995), our knowledge of the types of planetary systems around other stars has grown exponentially. Using techniques such as the transit method, radial velocity technique, microlensing and direct imaging, over 3300 confirmed exoplanets have now been detected¹. This newfound knowledge has led to many exciting discoveries. For example, we now know that on average every star in our galaxy hosts at least one exoplanet (Cassan et al. 2012) and that the occurrence rate of habitable zone Earth-size planets around Sun-like stars is approximately 1 in 5 (depending on the exact definition of the habitable zone; Petigura et al. 2013). In addition, we are now able to directly image exoplanets themselves (e.g. Chauvin et al. 2004).

This thesis seeks to advance the field of exoplanet science through the innovative use of polarimetry as an observational technique. It is divided into two main parts: a) the development of data analysis techniques for the Gemini Planet Imager’s (GPI’s) polarimetry mode, directly coupled to observations of the β Pic system and the debris disk around HD 157587, and b) in-lab and on-sky characterization and performance tests of a novel grating technology, called polarization gratings, that can be effectively used in future observations of exoplanetary atmospheres. Below I provide relevant background on debris disks, both theoretical and observational (§1.1), and a technical description of polarization gratings (§1.2), followed by the structure of this thesis (§1.3).

1.1 Debris Disk Observations with GPI

The plethora of recent exoplanet discoveries have allowed us to begin the calibration of planet formation models, but much still remains unknown. Notably, the formation pathways of Jovian-mass planets remain ambiguous and measurements of their mass and semi-major axis distributions are required for further progress. In addition, the existing atmospheric models appear to be incapable of fully recreating the observed spectra of directly imaged planets (e.g. Chilcote et al. 2015). For planets at these semi-major axes (~ 5 –100 AU), a small relative inclination with respect to the observer will prevent an observable transit, and the timelines required to confirm a detection (both for transit and doppler methods) are prohibitively long. Therefore, direct imaging may be the only viable discovery method for most of these

¹<http://www.openexoplanetcatalogue.com/>

further out massive planets. However, the imaging and characterization of planets at these semi-major axes is challenging due to both their close angular proximity to their host star and the extreme contrast ratios (10^{-6} and greater) between a given star and its planets. Instrumental point-spread functions (PSFs) and atmospheric turbulence in ground-based observations extend the stellar emission out to angular separations where planets are found, masking their signal and making them difficult to detect.

Despite these challenges, direct imaging searches for further-out planets have yielded a number of significant discoveries. Perhaps the most well known is β Pic *b* (Lagrange et al. 2009), a $12.7 \pm 0.3 M_{Jup}$ planet (Morzinski et al. 2015) embedded in a bright debris disk (Smith & Terile 1984). Detailed observational work combined with dynamical modelling have attributed a significant warp in the debris disk profile to the misalignment of the planet’s orbital plane with the main disk plane, suggestive of a complicated dynamical history (Dawson et al. 2011; Lagrange et al. 2012b; Nesvold & Kuchner 2015). Another compelling discovery is that of the multi-planet system HR 8799, the only multi-planet directly imaged system so far (Marois et al. 2008).

Recently, a new generation of direct imaging instruments have come online, employing high-order adaptive optics systems coupled to advanced coronagraph designs in order to mitigate the challenges associated with high-contrast imaging (e.g. GPI, Macintosh et al. 2014; Spectro-Polarimetric High-contrast Exoplanet REsearch, a.k.a SPHERE, Beuzit et al. 2008). One of the first major discoveries coming out of these new instruments is that of 51 Eri *b*, a 2-10 M_{Jup} planet at a separation of 14_{-3}^{+7} AU, discovered by GPI in Dec 2014 (Macintosh et al. 2015; De Rosa et al. 2015a). 51 Eri *b* distinguishes itself as the first directly imaged planet to have a near-infrared (NIR) spectrum showing deep methane absorption, much like the NIR spectrum of Jupiter. In the coming years, ongoing surveys being carried out with these new facilities, such as the GPI Exoplanet Survey (GPIES) and its SPHERE counterpart, hope to expand these discoveries to even lower masses and smaller semi-major axes.

A complimentary method to the direct imaging of planets for studying planetary systems at large semi-major axes is the imaging and characterization of debris disks, extrasolar analogs to our solar system’s asteroid and Kuiper belts. Debris disks are remnants of the planet formation process and their morphologies and compositions are intimately linked with the planet formation and evolutionary history of the systems in which they reside (see Wyatt 2008; Matthews et al. 2014, for recent reviews). Indeed, the mass and location of the planet β Pic *b* were predicted using images of the β Pic debris disk (e.g. Mouillet et al. 1997). Further, many directly imaged planets are known to reside within debris disks (e.g. β Pic, 51 Eri, HR 8799, Fomalhaut and HD 95086).

In scattered light, debris disks are of similar brightness to directly imaged planets and suffer from the same observational challenges. As a result, many debris disk studies have relied on space-based observations due to the lack of atmospheric turbulence and the stability of the PSF. Although recent improvements in the field of adaptive optics have led to an increase in ground-based observations of debris disks, advanced observing techniques (e.g. angular differential imaging, or ADI, Marois et al. 2006) and data processing methods (e.g. Locally Optimized Combination of Images, or LOCI, Lafrenière et al. 2007; Karhunen-Loève Image Processing, or KLIP, Soummer et al. 2012) are still required to recover debris disk emission obscured by the residual stellar PSF. However, these techniques can be subject to biases such as self-subtraction and over-subtraction of the disk signal, which can make it difficult to recover the true disk morphology of a given disk (Milli et al. 2012).

The newest generation of ground-based high-contrast imagers (such as GPI and SPHERE) have included polarimetry modes designed specifically to take advantage of the polarized nature of scattered

light in order to achieve further suppression of the uncorrected (unpolarized) PSF halo in observations of debris disks (Thalmann et al. 2008; Langlois et al. 2014; Perrin et al. 2015). This technique, sometimes known as polarimetric differential imaging (PDI), has the advantage that it is not subject to the same biases as the ADI-based methods mentioned above. In addition, measurements of the degree of polarization of the scattered-light can be used in conjunction with total intensity measurements (when available) to distinguish between different grain properties (e.g. Graham et al. 2007).

The following subsections provide a brief summary of the formation and evolution of debris disks, and a description of different types of debris disk observations, alongside a technical introduction to GPI.

1.1.1 Debris Disk Formation and Evolution

Debris disks are formed through the collisional cascade of larger-than-km-sized planetesimals that grinds the disk bodies down to μm and sub- μm sizes (e.g. Strubbe & Chiang 2006). The formation of these parent planetesimals is very similar to that of planets, requiring growth by many orders of magnitudes from the sub- μm sizes of ISM dust. In the early stages of this process, the growth of μm sized dust to decimeter- or meter-sized particles happens via collisional agglomeration as dust settles to the mid-plane of the disk (see Dominik et al. 2007, and references therein). Growth beyond a meter requires a different mechanism as relative velocities between meter-sized bodies reach levels such that nearly all collisions become destructive, and gas drag causes meter-sized bodies to spiral into their host protostars within a few 100 years of formation (Benz 2000; Weidenschilling 1977). There is mounting evidence that this barrier can be overcome through the gravitational collapse of over-dense regions of mm-cm sized pebbles (Johansen et al. 2007; 2014), such as those inferred to be forming in the dust gaps of the HL Tau protoplanetary disk (Zhang et al. 2015).

Grain growth continues throughout the protoplanetary disk phase, which can last up to 3–10 Myr, until a point at which most of the primordial gas is depleted (e.g. Haisch et al. 2001; Panić et al. 2013). The debris disk phase then begins when the constituent gas and dust are predominantly produced through collisions (e.g. Matthews & Kavelaars 2016), which begin once orbital eccentricities are stirred to greater than 10^{-2} – 10^{-3} (e.g. Kenyon & Bromley 2008). To achieve this level of stirring, several mechanisms have been suggested. Turbulence, Brownian motion, differential settling and radial drift can act on dust during the protoplanetary stage, leaving mm-sized grains with eccentricities of $\sim 10^{-3}$, a mechanism known as ‘pre-stirring’ (Brauer et al. 2008; Wyatt 2008). Stirring can also be induced from within the disk (‘self-stirring’), once planetesimals have grown to sizes of ~ 1000 km (Kenyon & Bromley 2004; 2010; Kennedy & Wyatt 2010). Eccentric planets can also induce stirring via secular perturbations from up to tens of AU away from the location of the planetesimals (Mustill & Wyatt 2009a). Finally, dynamical interactions with a passing star can also excite the required eccentricities, though such events are statistically rare (Kenyon & Bromley 2002).

Once sufficiently stirred, planetesimals participate in a collisional cascade. The cascade produces a size distribution of particles usually described as a power law $n(D) \propto D^{-\alpha}$, where in most models α ranges between 3 and 4 depending on the composition and the critical energy for fragmentation. In this range most of the mass is trapped in larger grains, but the surface area (and hence the emission) is dominated by smaller grains. If the fragmentation energy is scale-independent, then $\alpha = 3.5$ (Dohnanyi 1969). The lower end of this size distribution is usually governed by the blow out size, the size at which radiation pressure forces dust grains into hyperbolic orbits, on the order of $\sim 1 \mu\text{m}$ for Sun-like stars.

As debris disks age, their brightness decreases due to the depletion of the biggest planetesimals in the cascade (Wyatt et al. 2007). For a comprehensive explanation of all observational results, however, stochastic processes must be evoked to account for the unusual brightness of some disks (e.g. Siegler et al. 2007).

A number of different studies have tried to identify a correlation between the presence of planets and the presence of a debris disk. A preliminary study of the nearest 60 G-type stars by Wyatt et al. (2012) indicated a tentative correlation between the presence of a debris disk and the presence of low-mass planets, and was later confirmed by Marshall et al. (2014). On the other hand, after rejecting stars with unknown ages, those with ages less than 10 Gyr, and those with binary companions at less than 100 AU, Moro-Martín et al. (2015) find no such correlation. However, their sample size is only large enough to detect such a correlation if debris disks were four times more abundant in systems with low-mass planets than those from their control sample. Moro-Martín et al. (2015) note that such efforts are limited by the detectability of both the debris disks and the low-mass planets, both of which could have significant populations below detection limits. Thus the question of whether or not such a correlation exists remains unanswered.

In the presence of a nearby planet, dynamical interactions with the parent planetesimals population can lead to a range of morphological features in debris disks. If a planet’s orbital plane is misaligned from that of a nearby disk, it can cause dynamic perturbations whereby the inclinations of the planetesimals in the disk oscillate about the planet’s orbital plane in a wave that propagates outwards through the disk over time (Dawson et al. 2011; Nesvold & Kuchner 2015). Seen at a given moment in time this can appear as a warp in the the disk when viewed edge-on, as is the case for β Pic (Burrows et al. 1995; Mouillet et al. 1997). The razor sharp inner edge and the eccentricity of the Fomalhaut debris disk is thought to be caused by an eccentric Neptune- to Saturn-mass planet, with a semi-major axis of 119 AU (Quillen 2006). Likewise, the sharpness of the outer edge of the HR 4796A ring has been attributed to the presence of a planet (e.g. Lagrange et al. 2012a), though Thébault & Wu (2008) demonstrate that the observed surface brightness decrease could also be due to a dynamically cold ring. An eccentric planet can also cause a stellocentric offset of the center of symmetry of a disk through secular forcing of the constituent dust particles, resulting in an asymmetric brightness in the disk know as a pericenter glow (Wyatt et al. 1999). Indeed, planetary interactions have now be implicated in the swept-out morphology of HD 61005 (Esposito et al. 2016), the origin of which has also been attributed to either interaction with the ISM via ram pressure affecting the small grains (Hines et al. 2007; Debes et al. 2009) or via gravitational interaction with low density ISM gas (Maness et al. 2009). An effort by Lee & Chiang (2016) to unify the many morphologies observed in scattered light has resulted in five broad categories of debris disk morphologies: rings, needles, ships-and-waves, bars and moths (or fans). All of these morphologies can be explained with their simple model of a single planet secularly perturbing a narrow ring of parent bodies.

1.1.2 Observations of Debris Disks

Debris disks were first detected via infrared observations obtained with the Infrared Astronomical Satellite (IRAS) that revealed flux exceeding the level expected from stellar photospheres alone (e.g. Aumann et al. 1984). This emission was posited to originate from the thermal emission of circumstellar dust, a supposition that was confirmed shortly thereafter with the scattered-light imaging of the disk around β Pic (Smith & Terrile 1984). Since then the number of stars with infrared excesses has grown to

many hundreds thanks to surveys using space-based telescopes such as the Infrared Space Observatory, AKARI, the Spitzer Space Telescope and the Herschel Space Observatory. Results from this work, in particular observations from Herschel and Spitzer, have allowed us to put constraints on the incidence of debris disks: 25 % – 30 % around A-stars (Su et al. 2006; Thureau et al. 2014) and ~ 20 % around FGK stars (Eiroa et al. 2013). Modelling of infrared spectral energy distributions (SEDs) can lead to preliminary estimates of debris disk separations and radial extents (e.g. Wyatt 2008). However, dust location and grain emission properties are degenerate, requiring spatially resolved images to understand the true spatial distributions of the grains.

Debris disks have been resolved in both thermal emission in the IR and mm regimes and also in scattered starlight in the visible and near-infrared. These two regimes provide complimentary information; thermal emission traces the locations of mm-sized dust grains and scattered light traces the gains of μm and sub- μm size. Longer-wavelength thermal emission observations are useful because the distribution of mm-sized dust grains follows that of the parent planetesimal population and is independent of the scattering properties of the grains (Wyatt 2006). As a result, observations of thermal emission are particularly helpful in measuring, for example, stellocentric offsets (e.g. Telesco et al. 2000).

On the other hand, better spatial resolutions can be achieved with shorter-wavelength scattered-light observations, which can lead to more detailed morphological characterization. However, the smaller dust grains probed at these wavelengths can be affected by stellar winds, radiation pressure and Poynting-Robertson drag, which affect the location of the observed dust. The Atacama Large Millimetre Array (ALMA) is now able to achieve comparable spatial resolutions to those of visible and NIR observations, but these high-resolution observations by ALMA at longer wavelengths are only feasible for some of the brightest disks.

Upwards of 80 debris disks have now been resolved (e.g. Choquet et al. 2016), largely at longer wavelengths, with only ~ 25 disks seen in scattered light. A significant number of the disks seen in scattered light have been imaged with the Hubble Space Telescope (HST; e.g. Schneider et al. 2014), but ground-based observations are becoming increasingly important as advances in AO have allowed for studies at smaller inner working angles with higher contrasts. Both space- and ground-based scattered-light observations of debris disks have led to a number of significant discoveries, for example: (1) the already-discussed warped disk around β Pic, used to predict the presence of β Pic *b*; (2) the large disk with a notably sharp inner edge around Fomalhaut (Kalas et al. 2005), also a planet-host, although Fomalhaut *b*'s orbit appears to cross the disk and the exact relationship between the planet and the disk is unclear (Graham et al. 2013; Beust et al. 2014); (3) the remarkable swept-back morphology of HD 61005, a.k.a “The Moth” (Hines et al. 2007); and too many others to summarize here. A recent noteworthy result is that of Boccaletti et al. (2015) who combined archival HST images with newly obtained SPHERE observations to reveal numerous clumps of material in the AU Mic debris disk that appear to be moving at speeds that indicate highly eccentric or possibly unbound trajectories, a phenomenon unexplained by standard disk theories.

Polarimetry has emerged as a useful tool for studying debris disks in scattered-light, both because of its diagnostic power in constraining grain characteristics, and because it can be used to achieve increased speckle suppression in high-contrast applications. In the visible and NIR, the emissivity (j_λ) at wavelength λ of a collection of dust grains due to scattered light can be written as

$$j_\lambda = n \frac{dC_{sca}}{d\Omega} F_{*,\lambda}, \quad (1.1)$$

where n is the number density of dust particle, $dC_{sca}/d\Omega$ is the differential cross-section and $F_{*,\lambda}$ is the incident stellar flux at wavelength λ . The differential cross-section is a function of the 2×2 scattering matrix (S) that describes the transformation of the incident electric field (E_{in}) to the scattered electric field (E_{scat}) and depends on the shape, size and composition of the dust grains:

$$E_{scat} = S \cdot E_{in}, \quad (1.2)$$

where E_{scat} and E_{in} are two-dimensional quantities describing the two components of the E-field (Bohren & Huffman 2004). When considering randomly oriented dust grains and incident unpolarized light (as is often assumed planetary systems) only the diagonal elements of the matrix remain, and the scattering phase function and the degree of polarized light, P_{lin} , can be expressed as:

$$\frac{dC_{sca}}{d\Omega} = \frac{1}{2}(|S_2|^2 + |S_1|^2), \quad P_{lin} = \frac{|S_2|^2 - |S_1|^2}{|S_2|^2 + |S_1|^2}. \quad (1.3)$$

Thus, combined measurements of the total intensity and polarized intensity can tell us both about the spatial distribution and the scattering properties of the grains in a debris disk.

The archetype of such an analysis is that of the AU Mic debris disk by Graham et al. (2007), who used HST/ACS polarimetric observations to constrain the grain properties in the disk. Their analysis combines total intensity measurements with polarization fraction measurements in order to fit a geometric disk model using a selection of different grain scattering models. The results indicate that their observations are best fit by the highly porous aggregate model of Kimura et al. (2006), providing evidence for grain growth in the system. A similar analysis has recently been performed on GPI polarimetric observations of the HR 4796A disk by Perrin et al. (2015), who found that the grains must be relatively large to produce the observed polarization fraction, though their exact composition remains unclear. One interpretation of these results is that the disk is optically thick and that the dust in the ring is shepherded by nearby planets, analogous to the shepherding of some of Saturn's rings by nearby moonlets. While this scenario still requires further study, such a situation appears to be similar to the dynamically cold disk suggested by Thébault & Wu (2008), as Saturn's rings are also dynamically cold.

The inherently polarized nature of scattered light can be taken advantage of in high-contrast imaging observations to probe beneath the residual unpolarized speckles, left uncorrected by the AO (e.g. Kuhn et al. 2001; Perrin et al. 2008). By employing a polarimetry mode to analyze the polarized intensity, the host star's unpolarized light can be rejected and, subsequently, the disk revealed beneath. This method can be particularly effective for face-on disks, such as the TW Hya or V4046 Sgr (protoplanetary) disks (Rapson et al. 2015a;b), where the morphology is azimuthally smooth and therefore difficult to recover in total intensity using ADI. The disk observations presented in this work (Chapters 3 and 4) indeed only recover the disks in polarized intensity. While the lack of total intensity information precludes any modelling of the dust grain composition, these polarized intensity observations can nonetheless be used to analyze the disks' morphologies.

1.1.3 The Gemini Planet Imager

GPI is a NIR high-contrast imager, installed on the Gemini South telescope in the fall of 2013. Its optical system consists of a high-order AO system, an apodized-pupil Lyot coronagraph, and an integral field

spectrograph (IFS), which employs both spectral and polarimetric dispersing prisms. The combination of the AO system and the Lyot coronagraph provides diffraction-limited resolutions between 0.9 and $2.4 \mu\text{m}$ down to inner working angles of $\sim 0.15''$. The IFS samples the focal plane at 14 milliarcsecond intervals and provides a $2''.8 \times 2''.8$ square field of view. Raw GPI data is converted to calibrated science-ready datacubes using a dedicated software package called the GPI Data Reduction Pipeline (Perrin et al. 2014a; Maire et al. 2012a).

GPI is equipped with an imaging polarimetry subsystem designed specifically for the imaging and characterization of debris disks (Perrin et al. 2015). The polarimetry system is composed of a rotatable half-wave plate modulator and a Wollaston prism analyzer, making GPI sensitive to linear polarization. On-sky calibrations have revealed that GPI’s instrumental polarization is $\sim 0.6\%$ and that the polarimetry mode can suppress unpolarized speckles by up to a factor of 200 (Wiktorowicz et al. 2014a).

Since its commissioning, GPI has produced many interesting scientific results, including (but not limited to) the discovery of the planet 51 Eri *b* (Macintosh et al. 2015), the first images of the debris disk around the planet-hosting HD 106906 (Kalas et al. 2015) and images of the TW Hya protoplanetary disk that show a gap in the disk consistent with the presence of a $\sim 0.2 M_{Jup}$ planet (Rapson et al. 2015b). GPI is now available to the Gemini community as a standard observing instrument. In addition, GPI is being used for both the GPI Exoplanet Survey (GPES; PI - Bruce Macintosh) and for a Gemini Large and Long Program (LLP) with the primary goal of detailed debris disk characterization (PI - Christine Chen).

1.2 Polarization Gratings

While polarimetry has now been established as a powerful tool for the study of circumstellar material, a number of ongoing efforts are also attempting to detect polarization from exoplanets themselves. Light reflected off close-in Jupiters is predicted to create a polarized signal on the order of tens of parts per millions in integrated stellar light due to Rayleigh scattering in the upper atmosphere (e.g. Lucas et al. 2009). Such a phenomenon has been reported for HD 189733 *b* (Berdyugina et al. 2008; 2011), a transiting Jupiter-sized planet (Bouchy et al. 2005). However, this measurement has been called into question, as it has not been reproduced by other groups (Wiktorowicz 2009; Bott et al. 2016). The detection of reflected polarization is difficult because of the large number of photons required to achieve the necessary precision as well as the stringent requirements on the stability of the instrumental polarization.

On the other hand, thermal emission of directly imaged planets is expected to have polarized fractions of up to 1% (e.g. Marley & Sengupta 2011). Polarized thermal emission can be induced in directly imaged planets when the disk-integrated light experiences some level of asymmetry, either due to rotationally induced oblateness or to patchy cloud cover (Marley & Sengupta 2011; de Kok et al. 2011). Initial steps in characterizing GPI’s ability to detect broadband polarization of point sources have been carried out by Jensen-Clem et al. (2016), while polarized observations of directly imaged planets are ongoing. Spectropolarimetry of directly imaged planets has the potential to act as a probe of surface gravity and can be used to distinguish between hazy and cloudy atmospheric models in situations where these properties can appear invariant in a planet’s spectrum (Marley & Sengupta 2011; de Kok et al. 2011). Although no detection of a polarized exoplanet has yet been confirmed, the potential diagnostic power of polarimetric observations of exoplanets is undeniable.

An exciting new technology for carrying out future spectropolarimetric observations of exoplanets is a type of diffraction grating called a polarization grating (Crawford et al. 2005; Escuti et al. 2006; Provenzano et al. 2006). Polarization gratings (PGs) combine polarized beam splitting capabilities with wavelength dispersion to enable spectropolarimetry using only a single thin-film optical element. Furthermore, these gratings can be designed for extremely high throughput across a broad wavelength range. These unique characteristics make polarization gratings an appealing technology not only for exoplanet studies, but for many other interesting fields of astronomy (e.g. Packham et al. 2010).

PGs are designed to diffract incident light into only two orders, $m = \pm 1$. The amount of light diffracted into each of the orders depends on the circular polarization state of the incoming beam. The regular patterning of the liquid crystal polymers (LCPs), a material commonly used for the fabrication of PGs, causes them to disperse light spectrally as a diffraction grating. In contrast to conventional surface relief gratings, which affect the phases of orthogonal polarizations equally, PGs are able to control the two phases separately. As a result, PGs can readily overcome the fundamental limitations on the diffraction efficiency of conventional ruled gratings (Wyrowski 1991). Conventional gratings may be blazed to achieve relatively high diffraction efficiencies, but only over a small wavelength range, and diffraction efficiency must often be traded for a spectral coverage. PGs, on the other hand, are able to achieve near perfect diffraction efficiency over a near unity bandwidth, $\Delta\lambda/\lambda_0 \approx 1$ (Oh & Escuti 2007). This highly desirable property makes PGs relevant as a dispersing element not only for polarimetric applications, but for normal spectroscopy as well.

The development of these devices has been spearheaded by Dr. Michael Escuti at North Carolina State University (NCSU), where the applications to date have mostly been in the telecommunications technology and liquid crystal display fields. The PGs developed at NCSU are based on LCP thin film coatings. Within the thin film layer the LCP orientation rotates both in one direction along the grating surface and along the optical axis. The periodic rotation pattern in the grating plane cause PGs to follow standard diffraction theory, where a 180° rotation defines the grating period. In particular, PGs diffract light as a function of wavelength following the grating equation:

$$m\lambda = \sigma (\sin \alpha + \sin \beta), \quad (1.4)$$

where $m = \pm 1$ for polarization gratings, σ is the grating period, and α and β are the angle of incidence and angle of diffraction. In general, polarization gratings are used in a face-on configuration ($\alpha = 0^\circ$).

The polarimetric properties of the gratings stem from the birefringent nature of the liquid crystals. The rotation in the plane of the grating induces the beam splitting behaviour and the rotation along the optical axis works to extend the wavelength range. The result is a polarization sensitivity (Γ_{\pm}) in each order that follows the Jones matrix:

$$\Gamma_{\pm} = \frac{\sin(\pi\Delta n d/\lambda)}{2} \begin{bmatrix} \mp 1 & 1 \\ 1 & \pm 1 \end{bmatrix}, \quad (1.5)$$

where Δn is the birefringence of the material, d is the thickness of the LCP film. The \pm indicates the the behaviour of the PG's two orders.

An important design consideration of PGs is their use of detector space. To fully realize both the polarimetric capabilities and the high diffraction efficiency, both the $m=+1$ and $m=-1$ orders must be recorded simultaneously. The downside to this is that there is considerable space between the two orders

where no light falls. If a single detector is used to capture both orders, then a significant amount of the detector space between the two orders may go unused. Another important design consideration is that the beam-splitting angle is determined by the grating period and, therefore, is directly connected to the spectral resolution and linear dispersion of the grating at the detector plane. It is possible that these effects can be resolved with a sophisticated optical design, but it remains as a challenge in optical designs adopting PGs.

1.3 This Work

The work presented herein is composed of two major aspects, both centred on the theme of using polarimetry to further exoplanet science. The first aspect involves the development and characterization of advanced data analysis techniques for GPI's polarimetry mode, which are successfully applied to the data from GPI observations of two debris disks. The second aspect involves the laboratory characterization and on-sky performance testing of a polarization grating with the primary goal of demonstrating the viability of this new technology for use in future observations of exoplanet atmospheres.

The GPI data analysis techniques presented in this thesis (Chapter 2) are now adopted as standard practice in the reduction of GPI data, in particular in the GPIES campaign and the Debris Disk LLP, and they have been included in the most recent release of the GPI Data Reduction Pipeline. In the work presented here these tools are applied to observational studies of the debris disks around the stars β Pic (Chapter 3) and HD 157587 (Chapter 4), where the disks are detected only in polarized intensity. The polarized intensity images are fit to a disk model using Markov Chain Monte Carlo techniques to obtain best-fit model parameters. For each disk the best-fit model indicates potential signs of interactions with unseen planets in the systems. In the case of β Pic, the disk morphology is inconsistent with having been sculpted by the known planet β Pic *b*, contrary to expectations. This discrepancy could possibly indicate the presence of an unseen perturber. The best-fit model of the disk around HD 157587, on the other hand, shows an offset between the geometric centre of the disk and the location of the host star, consistent with interactions with a planet. In both cases the mass of the hypothetical planetary perturbers may be too small for the planets to be detected directly, rendering the study of their interactions with the debris disks one of the few ways to characterize them.

The second component of this work consists of in-lab (Chapter 5) and on-sky (Chapter 6) characterization and performance tests of a PG optimized for operation between 500 and 900 nm. The in-lab tests confirm the expected high diffraction efficiency properties and the polarimetric properties of the grating. Encouraged by the confirmation, we developed a simple PG-employed spectropolarimeter that we used to conduct test observations of the polarized twilight sky, polarimetric standard stars and spectroscopic standard stars. The test observations were able to recover the expected polarization and spectroscopic signatures of our targets, confirming the performance of the PG in astronomical spectropolarimetry. This work sets the stage for the broader use of PGs in the astronomical community.

Chapter 2

Gemini Planet Imager Observational Calibrations: Polarimetric Contrasts and New Data Reduction Techniques

A version of this work titled “Gemini Planet Imager Observational Calibrations XIV: Polarimetric Contrasts and New Data Reduction Techniques” by Millar-Blanchaer et al. has been submitted to be published in the proceedings of the 2016 SPIE Astronomical Telescopes + Instrumentation conference.

2.1 Chapter Summary

The Gemini Planet Imager (GPI) has been designed for the direct detection and characterization of exoplanets and circumstellar disks. GPI is equipped with a dual channel polarimetry mode designed to take advantage of the inherently polarized light scattered off circumstellar material to further suppress the residual seeing halo left uncorrected by the adaptive optics. We explore how recent advances in data reduction techniques reduce systematics and improve the achievable contrast in polarimetry mode. In particular, we consider different flux extraction techniques when constructing datacubes from raw data, division by a polarized flat-field and a method for subtracting instrumental polarization. Using observations of unpolarized standard stars we find that GPI’s instrumental polarization is consistent with being wavelength independent within our errors. In addition, we provide polarimetry contrast curves that demonstrate typical performance throughout the GPI Exoplanet Survey.

2.2 Introduction

GPI is a high-contrast instrument on the Gemini South 8-m telescope designed for the direct detection and characterization of Jupiter-like planets and dusty debris disks around young nearby stars (Macintosh et al. 2014). Its optical design combines an extreme adaptive optics system (Poyneer et al. 2014) and an apodized-pupil Lyot coronagraph (Soummer et al. 2011) with a lenslet-based integral field spectrograph

(IFS; Larkin et al. 2014). GPI also includes a polarimetry sub-system composed of a rotatable half-wave plate (HWP) modulator and a Wollaston prism analyzer that when deployed replaces the spectrograph’s prism (Perrin et al. 2015). The Wollaston prism disperses incident light such that each lenslet produces two spots of orthogonal linear polarization on the detector. For a detailed description of GPI’s polarimetry mode we direct readers Perrin et al. (2015). In addition to being available to the Gemini community as a standard observing mode, GPI’s polarimetry mode is currently being used to carry out scattered-light (*H*-band) debris disk observations as part of the GPI Exoplanet Survey (GPIES). Target stars with known infrared excess are observed with a short ‘snapshot’ observing sequence and stars with previously resolved disks, or disks discovered in a snapshot are observed for an hour long sequence. GPI is also being used to carry out a detailed study of debris disk composition as part of a Gemini Large and Long Program (PI: Christine Chen). However, the results of the large program study are not discussed in this work.

Here we report on recent updates to the data reduction process and summarize the typical performance of the polarimetry mode during the GPIES survey, building upon the early characterization work presented by Wiktorowicz et al. (2014a) and Perrin et al. (2015). Throughout this work we will attempt to quantify performance improvements based on the 5-sigma contrast measured in the total linear polarized intensity, $P = \sqrt{Q^2 + U^2}$.

In §2.3 we summarize the reduction of polarimetry data and describe the method used to calculate contrast. In §2.4 and §2.5 we discuss the methods used to convert the raw data to a polarization datacube and the use of polarized flat fields, respectively. We present observations of unpolarized standard stars that we use to assess GPI’s instrumental polarization in §2.6, where we also present the current method used to subtract the instrumental polarization. In §2.7 we present the typical polarized intensity contrasts achieved in the GPIES survey, followed by some concluding remarks in §2.8.

2.3 Data Reduction and Contrast Measurement in Polarimetry Mode

A standard polarimetry mode observation sequence involves taking images with the HWP at positions of 0, 22.5, 45 and 67.5 degrees, making GPI sensitive to linear polarization (Stokes Q and Stokes U). Data are reduced using the GPI Data Reduction Pipeline (DRP) which converts raw data into analysis-ready datacubes, where each step in the reduction is known as a ‘primitive’ (Maire et al. 2012a; Perrin et al. 2014b). We briefly summarize the relevant data reduction steps here.

In a typical reduction, a raw data frame is first dark subtracted, corrected for bad pixels and cleaned for correlated detector noise. The raw data are then converted into a three-dimensional ‘polarization datacube’ where the first two dimensions are spatial dimensions and the third dimension holds two orthogonal polarization states. Therefore each spatial pixel (‘spaxel’) has two flux values. The polarization datacubes are then cleaned for bad pixels. At this point the position of the occulted star (behind the focal plane mask) is estimated using a radon-transform based method that relies on GPI’s fiducial ‘satellite spots’, four replicas of the stellar PSF that are imprinted on the focal plane by a grid located in the apodizer (Wang et al. 2014). Each polarization cube can then be summed to obtain a total intensity image or differenced to obtain the linear polarized intensity with an orientation defined by the position of the HWP during the observation (stored in the FITS header of each file). A double differencing routine is then applied to all the polarization datacubes that serves to remove any bias introduced by

non-common path errors for each lenslet. A sequence of polarization datacubes can be combined into a Stokes datacube, where the third dimension holds a Stokes vector $([I, Q, U, V])$ at each spatial location, by solving a set of equations that describe the expected response of the instrument to incident polarized light given the known HWP and parallactic angles for each frame. Finally, the Stokes cube can be transformed to the radial Stokes convention $([I, Q, U, V] \Rightarrow [I, Q_r, U_r, V])$ where a pixel in the Q_r frame holds linear polarized intensity that is oriented either parallel (negative values) or perpendicular (positive values) to a line connecting the pixel to the central star’s location. The U_r frame holds polarization oriented $\pm 45^\circ$ from that line. Under the assumption that the measured polarization is due to single scattering off relatively small dust grains, then all the flux should be in the Q_r frame and the U_r frame should just contain noise. In debris disks, the mass of the dust is sufficiently small that contributions from multiple scattering only introduce a weak signal in the U_r map, up to a few percent of the Q_r signal (Canovas et al. 2015). In GPI images this signal is expected to be much below the noise threshold in the U_r image, justifying its use as a noise map. However, when observing protoplanetary disks where the dust density is significantly higher, Canovas et al. (2015) demonstrate that multiple scattering may introduce an observable U_r signal and the assumption that it is dominated by noise may break down. Therefore caution is urged when interpreting data presented using the radial Stokes convention.

The metric we use to quantify performance in polarimetry mode is the 5-sigma polarized intensity contrast, which we have developed to be as similar as possible to the contrast measured in spectroscopy mode. In GPI’s spectroscopy mode the final data product is a spectral datacube, where the third dimension holds spectral information. For a single wavelength slice, the 5-sigma point source contrast is measured at each angular separation as five times the standard deviation of the pixel values in an annulus around the star with that same angular separation. This value is then divided by the average peak satellite spot brightness of the four satellite spots in that wavelength slice and then multiplied by the known peak-satellite-spot-brightness-to-stellar-flux ratio to obtain the contrast relative to the stellar flux.

In polarimetry mode the entire bandpass is seen in each polarization datacube and the satellite spots appear as elongated smears rather than as replica PSFs. Instead of measuring the peak satellite spot brightness, as is done in spectroscopy mode, we measure the total flux of the satellite spots using a DRP primitive called “Measure Satellite Spot Flux in Polarimetry”, that saves the flux information in the FITS header of the polarimetric datacube (Hung et al. 2016). The equivalent satellite spot peak brightness is estimated by multiplying the average of the total satellite spot fluxes of the four spots by a conversion factor. We assume a Gaussian PSF and the conversion factor is then the ratio between the the peak and total flux in a 2D Gaussian function. Thus, we are calculating the peak flux of a theoretical Gaussian PSF whose total flux is equal to the flux measured in the elongated satellite spots. The contrast is then calculated in the same manner as in spectroscopy mode, by dividing the standard deviation in concentric annuli by the equivalent satellite spot peak brightness and then multiplying by the peak-satellite-spot-brightness-to-stellar-flux ratio to obtain the contrast in units of stellar flux. The 5-sigma contrast is then calculated by multiplying by a factor of 5.

To date no direct calibration of GPI’s total-satellite-spot-flux-to-stellar flux ratio has been carried out in polarimetry mode. This is due in part to the difficulty of finding objects that are bright enough for the AO system to close control loops, but do not saturate the detector in non-coronagraphic observations when the Wollaston prism is in place (though a planned upgrade that includes installing an neutral density filter may solve this problem in the near future). While a direct measurement of the flux ratio

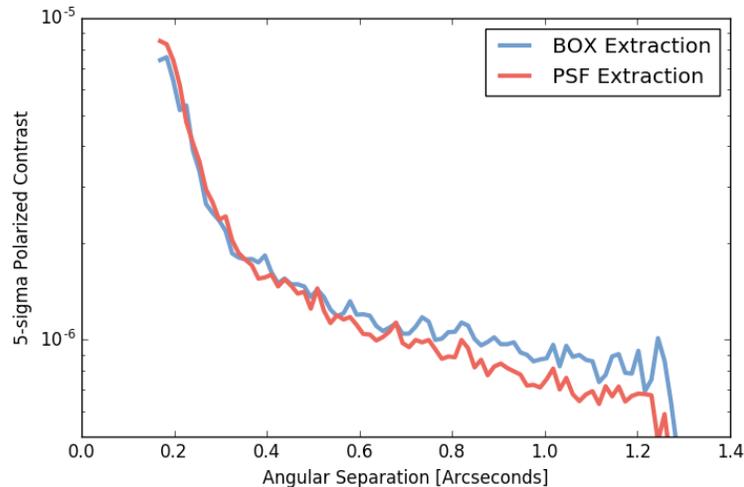


Figure 2.1: 5-sigma polarized contrast for the observations of the unpolarized standard star HD 118666 using both a BOX and PSF extraction method. The PSF extraction method shows improvement over the BOX method outside of 0.3 arcseconds where the polarized contrast is dominated by photon and read noise.

is currently unavailable, laboratory and on-sky tests comparing the total satellite spot fluxes between spectroscopy mode and polarimetry modes for the same target indicate that the satellite spots in both modes receive the same amount of flux, implying that the flux ratio is the same (Hung et al. 2016).

In polarization datacubes, contrast can be measured in total intensity as well as in polarized intensity (the difference of the two polarization states). In Stokes datacubes, the total satellite spot flux is estimated by taking the average value of the total satellite spot fluxes in all of the constituent polarization datacubes. Contrast can then be measured in any of the Stokes vector individual states (I, Q, U, or V), or in combinations of them, such as the polarized intensity ($P = \sqrt{Q^2 + U^2}$). The DRP primitive “Measure Contrast in Pol Mode” reports the total intensity and polarized intensity contrasts of both polarization datacubes and Stokes datacubes. The primitive gives the option to display the contrast immediately in a plot window and/or save the contrast information to a FITS file. Note that in spectroscopy mode, because we are typically searching for point sources, a high-pass filter can sometimes be applied to the data before measuring the contrast. Throughout this work we do not apply a high-pass filter to polarimetry mode data.

Many of GPI’s polarimetry mode targets are extended objects, such as protoplanetary or debris disks. The ability to detect extended objects above the noise floor is enhanced relative to point sources, by the fact that coherent structure can be extended across many pixels. The true detection threshold will then depend on the surface brightness and the angular extent on the sky of a given target. Thus, the point source contrast as defined in this section may somewhat underestimate the achievable detection limit for extended sources. However, as it is standard in the field, we opt to continue to use the 5-sigma point source contrast as the metric against which we test our new reduction techniques and report GPI’s sensitivity.

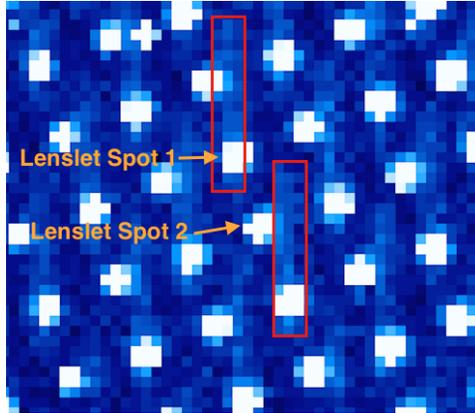


Figure 2.2: A raw detector image from a polarization observation showing persistent spectra from a previous spectroscopic observation. Two spectra are highlighted in red (though more are apparent in the image) and the labels indicate the two orthogonally polarized spots produced by the Wollaston prism for one lenslet. The persistent spectra are aligned such that the intensity measured from Lenslet Spot 1 receives a stronger bias than that measured for Lenslet Spot 2, resulting in an artificial polarization signal.

2.4 Assembling Datacubes from Raw Data

A key step in the reduction of GPI data is the conversion from raw data to datacube. After the raw data has gone through dark subtraction, bad pixel correction and the subtraction of correlated detector noise, it is converted to a polarization datacube using a DRP primitive called “Assemble Polarization cube”. The light in each lenslet (in raw detector counts, i.e. ADU/coadd) either spot is measured by summing a square aperture centered on the lenslet, known as BOX extraction or summing via a weighted PSF. The location and morphology of each lenslets’ two spots are predetermined using a polarization flat-field image obtained using Gemini’s Facility Calibration Unit (GCAL). Each lenslet in the flat-field image is fit to a 2D Gaussian function and the best fit parameters, including x and y widths, tilt and center location, are saved in the polarization calibration file, known as a ‘polcal’.

An analysis of observations of the unpolarized standard star HD 118666, presented in Perrin et al. (2015), demonstrated that for an 8 minute observation sequence, the polarized intensity contrast is limited by photon noise and read noise outside of ~ 0.3 arcseconds when using the BOX extraction. The analysis was based on observations taken as part of GPI commissioning on March 24, 2014 and consisted of eight 60s integrations, with the HWP rotating by 22.5 degrees between each observation. Since that time we have included in the “Assemble Polarization cube” primitive the option to use a weighted PSF extraction, where each pixel’s weight is determined by its relative contribution to the lenslet PSF, which is assumed to be the best fit Gaussian for that lenslet as stored in the polcal calibration file. This has the effect of reducing photon and read noise of each polarization spot measurement in the extracted polarization datacube, because pixels where the relative noise is high get down-weighted in the sum. An added advantage of the weighted PSF technique is that bad pixels in the raw data can be masked out and the surrounding pixels can still be used to provide an estimate of the counts in each lenslet spot. When using the BOX method bad pixels either have to be masked out or included in the sum, which can result in spurious values for spots that contain bad pixels.

In Figure 2.1 we present the polarized intensity contrast curve from the final Stokes cube of the HD 118666 dataset with both the BOX and PSF extraction methods. We find that the PSF extraction

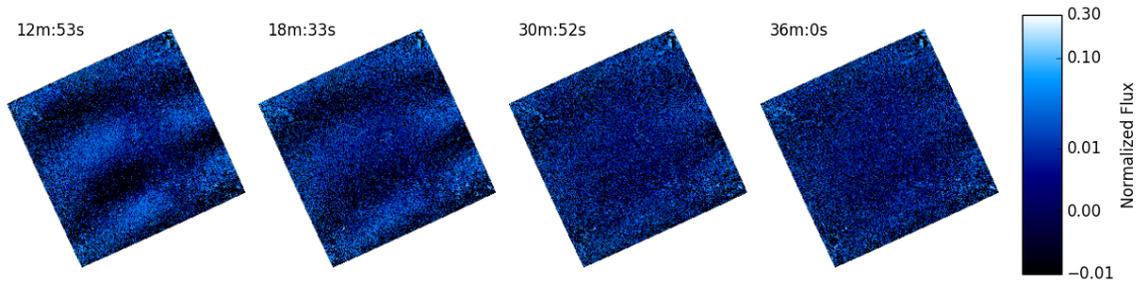


Figure 2.3: A time series of normalized difference images from four polarimetry datacubes taken from the same observation set (60s exposures), immediately following a spectroscopic observation of the same target. Persistence manifests as the thick vertical bar structure seen clearly in the first two images and diminishes to negligible levels by the fourth frame. The time labels indicate the time between the end of the last spectroscopic observation and the depicted polarimetry frame.

method improves the polarized intensity contrast starting at $\sim 0.3''$, where the data are limited by photon and read noise in the BOX extraction. While the improvement can be as small as 5% at 0.5 arcseconds, it reaches over 20% at larger separations. As a result, weighted PSF extraction is now the default option in the current version of the DRP and is used exclusively in the remaining sections of this work. Establishing the relative contributions of different noise sources using the PSF extraction technique has been left for future work.

2.4.1 Persistence

Persistence, also known as latency, occurs in detectors when electrons previously freed by incident light, get trapped in the detector crystal lattice. When the detector is read out, the trapped electrons are indistinguishable from newly released electrons and cause a bias in the raw image. These electrons are not expelled when the detector is reset, but instead remain trapped and decay as a function of time. Persistence in the GPI H2RG can be as high as $21 \text{ e}^-/\text{s}$; calculated by taking a 60s exposure immediately after a saturating exposure.

In GPI observations persistence is most apparent in polarimetry observations, due to the differential nature of the measurement. Though the fractional value of the persistent signal relative to a lenslet spot's counts may be small, because the polarization measurement relies on the difference in signal between the two spots, the bias can have a significant effect. This effect is strongest when polarimetric observations are taken immediately following a spectroscopic observation. Persistent spectra from the previous observation can coincide with the locations of one or both of the two polarization spots (e.g. Figure 2.2). The result is a polarization bias that depends on the strength of the persistence and the exact alignment of the spectrum on the two polarization spots. The relative alignment of the spectrum and the polarization spots changes across the detector resulting in a spatially dependent polarization bias (Figure 2.3). Because the strength of the polarization bias changes with time it can not be fully removed during the double differencing procedure and can masquerade as polarized emission.

Some detectors have shown a persistence which is a non-linear function of the total light accumulated during an integration, making the effects of persistence worse for brighter stars (Long et al. 2012). In our experience in the GPIES campaign we have found that back to back spectroscopic/polarimetric observations for the brightest stars ($I < 5$) can result in significant amounts of persistence lasting over

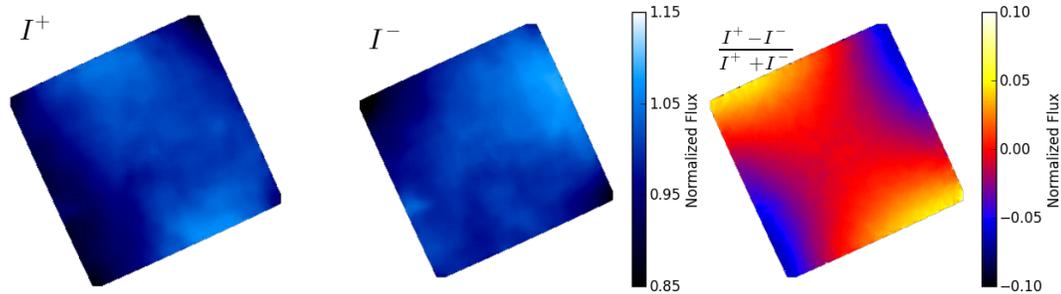


Figure 2.4: Images of the two orthogonal polarization slices (left, center) and their normalized difference (right) from a low-pass filtered GCAL flat field polarization datacube. The quadrupole pattern seen in the normalized difference is present in both GCAL flat field data, as well as on-sky observations. Each individual slices has been normalized by its median value.

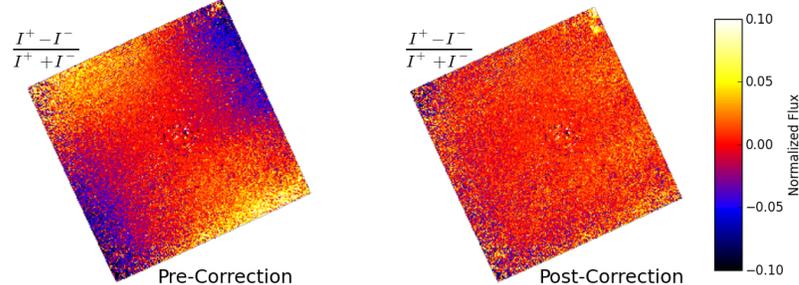


Figure 2.5: The normalized difference image of a polarization datacube from the HD 118666 dataset before (left) and after division by GCAL low spatial frequency flat field shown in Figure 2.4 (right). Dividing by the flat field successfully removes the polarization bias. Both images are shown with the same color stretch.

20 minutes after the polarimetric observations have begun (Figure 2.3). Detailed characterization of persistence is difficult because of the complicated physics involved in understanding detector crystalline structure. Our efforts in understanding the effects of persistence including detailed modeling and methods to mitigate its effects are ongoing.

2.5 Flat Fielding

GPI’s polarization datacubes are subject to a polarization bias signal that is proportional to the incident signal and changes across the field of view (Figure 2.4 and Figure 2.5). The flat field resembles a quadrupole pattern where the strength of the bias is strongest in the four corners, ranging from about -5% to 5% . This signal appears in polarization observation of both the Gemini GCAL calibration unit and on-sky targets. The bias does not change with HWP position, indicating that its origin is downstream of the HWP in the optical train.

This signal can be compensated for by dividing each polarization datacube by a “low spatial frequency polarization flat field” datacube, a spatially filtered polarization flat field datacube (Figure 2.4). Polarization flat field datacubes are created from observations of the GCAL Quartz-Halogen lamp, and are assembled into polarization datacubes in the same manner as standard polarization observations.

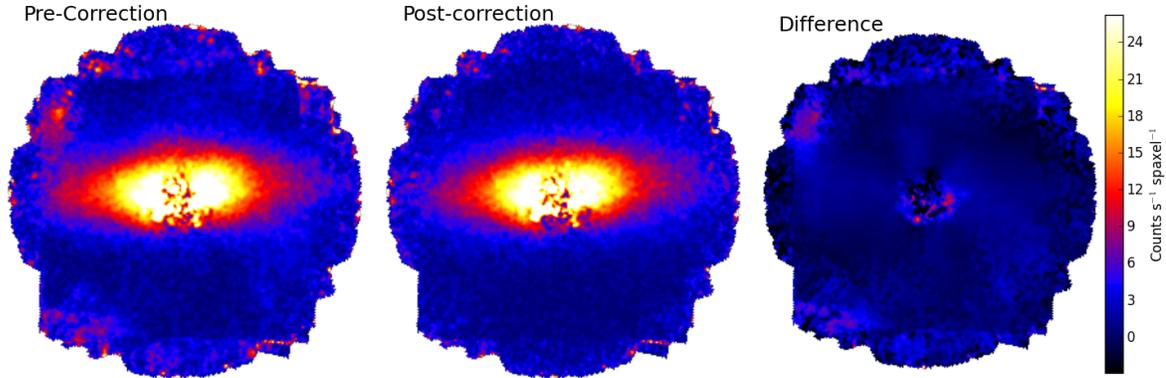


Figure 2.6: Linear polarized intensity images of the β Pictoris disk before (left) and after (center) flat field correction using a low spatial frequency polarized flat field. The disk images have been rotated so that the outer disk, with a position angle of 29° is horizontal, as in Millar-Blanchaer et al. (2015) Millar-Blanchaer et al. (2015), where this data was first presented. The difference between the left and the center panel can be seen on the right. The correction affects mostly the outer regions, but in this case the inner regions have mild corrections (relative to the local flux levels) as well.

Each 2D slice of the cube (i.e. one orthogonal polarization slice) is then normalized by dividing by its mean. The flat field datacube is then filtered spatially using a Fourier transformed based-filter. This has been implemented in the DRP as a standard calibration recipe template called “Create Low Spatial Frequency Polarized Flat-Field” and is available in DRP versions 1.4 and higher. This flat can be divided from a polarization datacube using a primitive called “Divide by Low Spatial Freq. Polarized Flat Field”. Figure 2.5 displays a polarization datacube before and after dividing by a low spatial frequency polarized flat-field. We have found that dividing by a polarization datacube that has not been smoothed results in a significant number of bad lenslets being artificially injected into the datacube.

In combined Stokes datacubes the flat-field bias can manifest as added noise at larger separations. For example, Figure 2.6 displays a linear polarized intensity image of the β Pic debris disk, originally presented by Millar-Blanchaer et al. (2015), before and after flat field correction. Significant noise can be seen at the edge of the field. When each polarization datacube is divided by a low spatial frequency flat field before being combined into a Stokes cube the noise is largely eliminated. Because this signal is a multiplicative function of the input intensity, it cannot be fully compensated for by the double differencing algorithm, which compensates for static bias offsets in each lenslet. However, for short sequences (e.g. 10-minute) where the PSF is relatively stable, the double differencing will compensate for most of the signal. For example, in the HD 118666 observations, the flat-field correction provides negligible improvement over the double differencing algorithm alone. Nonetheless, the flat field correction is now included as a standard primitive in the DRP recipes handling polarization datacubes.

2.6 Instrumental Polarization

Polarization induced by optics upstream of GPI’s HWP, or instrumental polarization, can masquerade as astrophysical signal since it also modulates with the rotation of the HWP. In reduced GPI polarization images instrumental polarization manifests as a polarization signal that is proportional to the residual total intensity of the stellar PSF at any given location (Figure 2.7) and whose position angle is constant

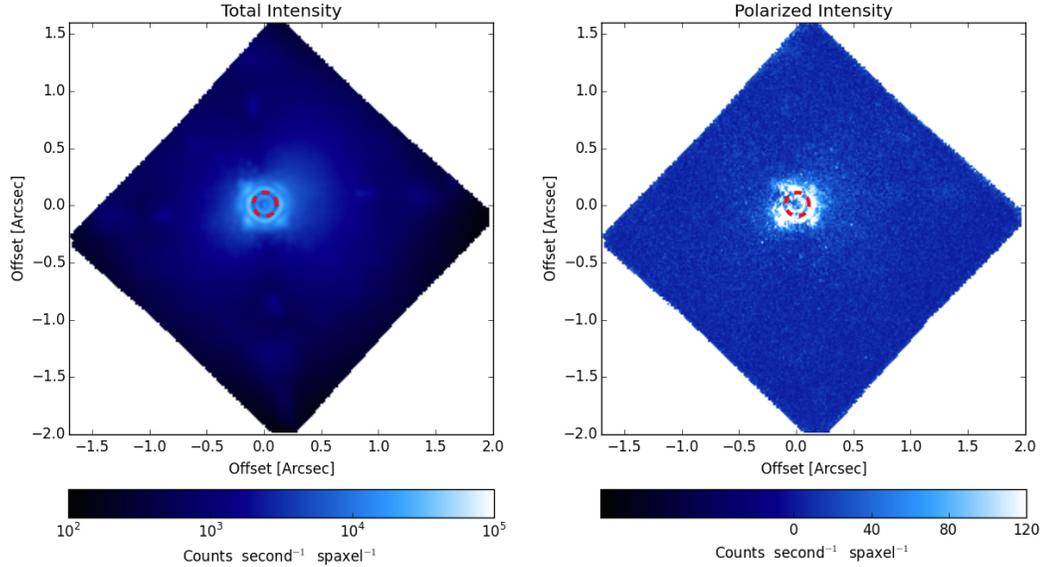


Figure 2.7: GPI H -band total intensity and linear polarized intensity images of HD 118666. The linear polarized intensity is predominantly due to instrumental polarization and its strength is proportional to the total intensity at a given location. The red dashed circle in each image denotes the angular extent of the H -band focal plane mask.

across the frame. If a Stokes cube is converted to the radial convention, then the instrumental polarization manifests as a quadrupole pattern due to the constant position angle (e.g. Figure 2.8). The following two subsections describe our efforts to characterize GPI’s instrumental polarization at multiple wavelengths and the current method used to subtract it from polarization datacubes.

2.6.1 Unpolarized Standard Stars

In this section we present an analysis of GPI observations of several unpolarized standard stars, which we use to characterize GPI’s instrumental polarization. This work builds upon the initial work carried out by Wiktorowicz et al. (2014a), who leveraged a large amount of field rotation in coronagraphic observations of β Pic to separate astrophysical signal from instrumental polarization (which keeps a constant position angle over time) to measure the H -band instrumental polarization to be $0.4354 \pm 0.0075\%$. Here we present measurements of the instrumental polarization in all 5 of GPI’s broadband filters (Y , J , H , $K1$ and $K2$).

We observed three unpolarized standard stars, HD 82386, HD 99171, and HD 210918, in GPI’s direct mode (i.e. without a focal plane mask) as polarimetric calibrators during several of GPI’s commissioning runs in 2014 (Table 2.1). To avoid saturating the detector within the minimum exposure time, the AO loop was set to open. In all exposures the tip/tilt loop was open, except for the last eight images of HD 210918. Each observation sequence consisted of 4 images where the HWP was rotated by 22.5° in between each exposure. We reduced the raw data to polarization datacubes with the GPI DRP, using the methods described in §2.2. Because we were observing in direct mode, the AO loop was open during these observations, and the star position and the PSF shape varied from exposure to exposure. We therefore opted to carry out our analysis on the total intensity in each individual polarization slice, rather than

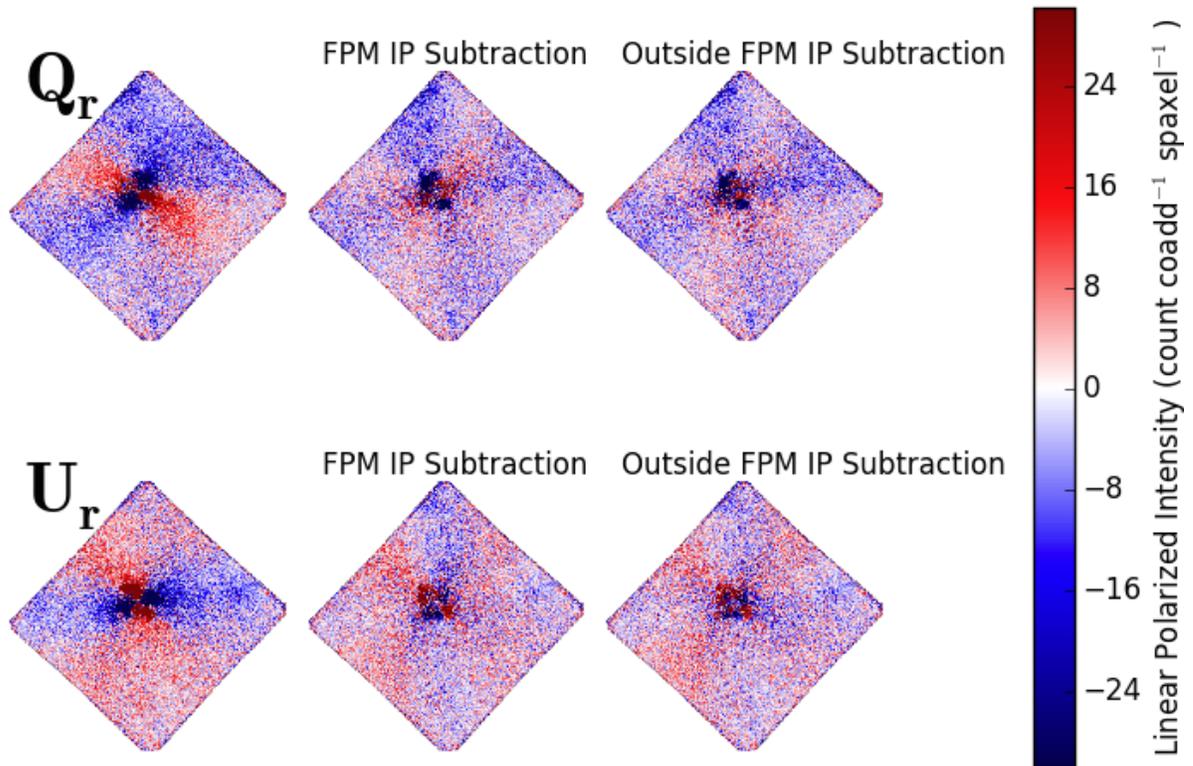


Figure 2.8: Q_r and U_r images of the GPI HD 118666 H -band observations. The first column of images shows the reduced cubes with no instrumental polarization (IP) subtraction. The instrumental polarization appears as a quadrupole pattern in both the Q_r and U_r images, offset 45° from each other. The second and third column display the results of subtracting the instrumental polarization as measured from behind the coronagraph FPM and just outside the FPM, respectively. When subtracting instrumental polarization measured from outside the FPM, slight gains are made in the inner regions (most readily seen in the Q_r image). A faint residual octopole can be seen in all of the IP-subtracted Q_r and U_r images. Understanding and compensating for this residual is a work in progress.

Table 2.1. Unpolarized Standard Stars Observed in Direct Mode as Polarimetric Calibrators

Unpolarized Standard	UT Date	Filter Sequence	# of HWP Sequences per Filter	Exp. Time (s)
HD 82386	2014-03-21	<i>H, J, Y, K1</i>	2	12
HD 82386	2014-03-25	<i>K1, H, J, Y</i>	2	12
HD 99171	2014-05-12	<i>H, J, Y, K1, K2, H</i>	1	30
HD 210918	2014-09-10	<i>H</i>	4	15

complete the analysis pixel-by-pixel as is normally done for GPI data. This procedure avoids comparing pixels with greatly different signal to noise ratios as the PSF position and shape changes between images. For each datacube we summed the total flux in each slice and placed it into a single pixel in that slice. To calculate the Stokes vector associated with each waveplate sequence, the four polarization datacubes (8 single-value pixels) were combined using the Combine Polarization Sequence primitive in the DRP with the default settings. For observing configurations where multiple waveplate sequences were taken in a row we calculated a Stokes vector for each waveplate sequence.

The linear polarized fraction from each Stokes vector can be seen in Figure 2.9 as a function of wavelength. Among all the observations, there is only one waveplate sequence taken in the *K2*-band. The polarization datacube images of this *K2* data set show a Morié pattern, which is a known artifact from the data pipeline procedure caused by a misalignment between the polarization spot locations in the polcal and the data. Unfortunately, simply offsetting the calibration solution or using the standard flexure correction procedure could not mitigate this problem. We therefore consider that this measurement may be spoiled.

All of the stars observed in our sample are unpolarized standard stars and we consider the detected polarization signals due to instrumental polarization. The percentage (%) mean polarization fractions for *Y, J, H,* and *K1* are 0.40 ± 0.27 , 0.62 ± 0.37 , 0.56 ± 0.17 , and 0.76 ± 0.59 respectively. The uncertainties represent the sample standard deviations. We find our *H*-band observation to be consistent with the measurement made by Wiktorowicz et al. (2014a) using β Pic. However, all of our observations are shorter sequences and have less field rotation than the observations of β Pic and as a result the errors are larger. Our results indicate that the instrumental polarization appears to be wavelength-independent to within our errors. This is consistent with the analysis presented by Wiktorowicz et al. (2014a) who use their *H*-band measurement to calibrate GPI *J, H* and *K1* observations of the polarized stars HD 77581 and HD 78344. Using this calibration they successfully fit a Serkowski law to a combination of their GPI measurements and previous visible light measurements.

We note that noise in the measurements of *Q* and *U* can introduce a bias when calculating the mean polarization fraction due to the squared *Q* and *U* terms in the calculation of *P*. Even when both *Q* and *U* measurements have zero means, the means of Q^2 and U^2 will not be zero as long as the sample standard deviations of *Q* and *U* are not zero. This effect will introduce a positive bias in the mean of *P*, with the magnitude of the bias dependent on the sample standard deviations of *Q* and *U*. We take a numerical approach to estimate this bias. For each band, we draw two large random samples: one from a normalized Gaussian with the mean of zero and σ of the sample standard deviation of Q/I , and the other with the same Gaussian except for having the σ being the sample standard deviation of U/I . We then use those values to calculate *P*. By taking the average of *P*, we estimate the bias in the polarization fraction in percentage (%) for *Y, J, H,* and *K1* to be 0.29, 0.30, 0.37, and 0.56, respectively. All of these

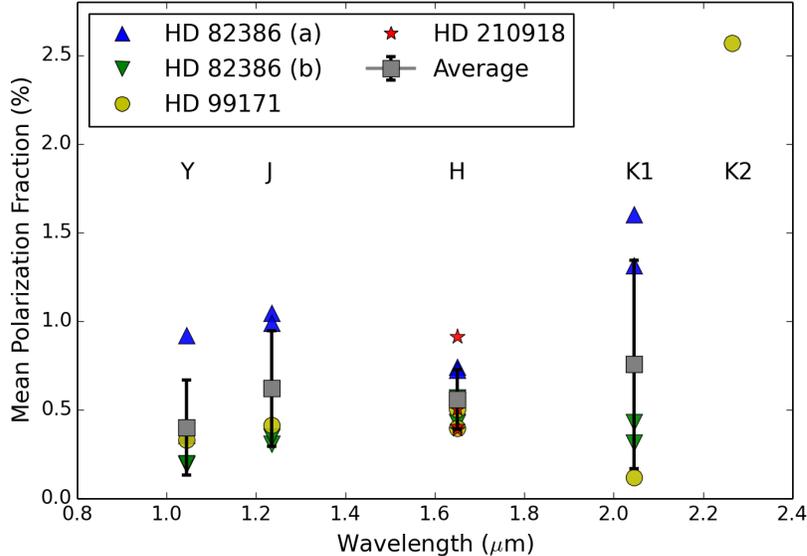


Figure 2.9: Polarization fraction measurements of unpolarized stars expressed as a percent. The different targets are denoted with different shapes and different colors. HD 82386 (a) and (b) data sets were taken 4 days apart. The $K2$ measurement may have been corrupted due to a misaligned polarization calibration solution. The mean polarization fractions for Y , J , H , and $K1$ across all observations are shown as grey squares, with error bars representing the sample standard deviations. The instrumental polarization appears to be wavelength-independent within errors, with an average of $\sim 0.6\%$.

biases are lower than the P measurements of the unpolarized standard stars, showing the presence of the instrumental polarization over this bias. Due to the varying PSF during the open loop observations we consider the variation in the sample’s Q/I and U/I to represent an upper limit in the uncertainty in Q and U and so these bias estimates are considered to be strong upper limits.

2.6.2 Subtracting Apparent Stellar Polarization from Polarization Datacubes

In a standard coronagraphic observing sequence the instrumental polarization can be subtracted in each polarization datacube by measuring the apparent stellar polarization. The apparent stellar polarization is measured as the mean normalized difference in a polarization datacube at the location of the focal plane mask (see Figure 2.7), and contains contributions from the instrumental polarization, interstellar polarization and possibly polarized scattered light on angular scales less than GPI’s diffraction limit. We expect that for most of GPI’s targets the instrumental polarization will be the dominant term. The normalized difference is defined as the difference of flux between the two orthogonal polarization slices divided by the total flux in both slices. Any light in this area will be light that has diffracted around the FPM and should be almost entirely due to the light of the star (a similar effect as an Arago or Poisson spot). The polarized flux caused by the apparent stellar polarization at a given spatial location can then be estimated by scaling the measured fractional polarization behind the coronagraph by the total flux at that location. This signal can then be subtracted out from each datacube individually. This method has been implemented in the most recent release of the GPI DRP as a primitive called “Subtract Mean Stellar Polarization” and has been used in a number of recent GPI publications (e.g. Millar-Blanchaer

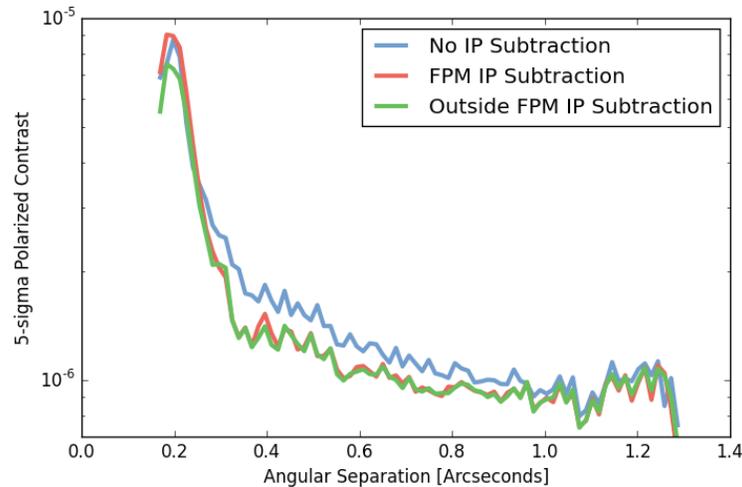


Figure 2.10: Linear polarized intensity 5-sigma contrast for the observations of HD 118666 as a function of angular separation using different instrumental polarization subtraction techniques. For this dataset, when the instrumental polarization (IP) is estimated using intensity measured behind the FPM, gains in contrast are seen between ~ 0.3 and 1.0 arcseconds. Further contrast gains can be achieved at smaller separations when the instrumental polarization is measured using intensity from just outside the FPM. Note that when there is significant astrophysical flux at small angular separations, measuring the instrumental polarization outside of the FPM runs the risk of subtracting real polarized signal from the datacubes.

et al. 2015; Kalas et al. 2015; Draper et al. 2016).

To demonstrate the effects of this procedure we applied it to the GPI commissioning observations of the unpolarized standard star HD 118666 discussed in §2.4. Figure 2.10 displays the improvement in linear polarized intensity contrast in the final combined Stokes cube when subtracting the instrumental polarization. Noticeable improvements are realized between ~ 0.3 and 1 arcseconds. Figure 2.8 displays Q_r and U_r images (where the effects of apparent stellar polarization are most apparent) of HD 118666, before and after subtraction. Though gains are made in the inner regions, it appears that we are still limited by systematics rather than random noise inside of $0.25''$.

In some cases (e.g. for very faint stars or very short exposure times) the amount of flux behind the FPM may be extremely low, resulting in a poor S/N estimate for the apparent stellar polarization. In this case, it may instead be estimated using the light just outside of the coronagraph, where the stellar flux is the highest (Figure 2.8). The increased flux in this area results in a higher S/N measurement. However, caution must be exercised when using this region to measure instrumental polarization; if there is a highly polarized source near the edge of the FPM, the assumption that the measured fractional polarization is due solely to instrumental/stellar polarization may break down. The “Subtract Mean Stellar Polarization” gives the user the ability to choose from which area they wish to measure the instrumental polarization, with the default being behind the FPM.

The general strategy of subtracting the instrumental polarization from each polarization datacube individually has several advantages. First, if the host star exhibits some level of polarization, we can use this method to measure and subtract the stellar polarization that may dilute any measurement of polarized circumstellar material. However, stellar polarization is typically due to polarization from interstellar dust grains, and the magnitude of the signal increases with distance from the earth (Fosalba

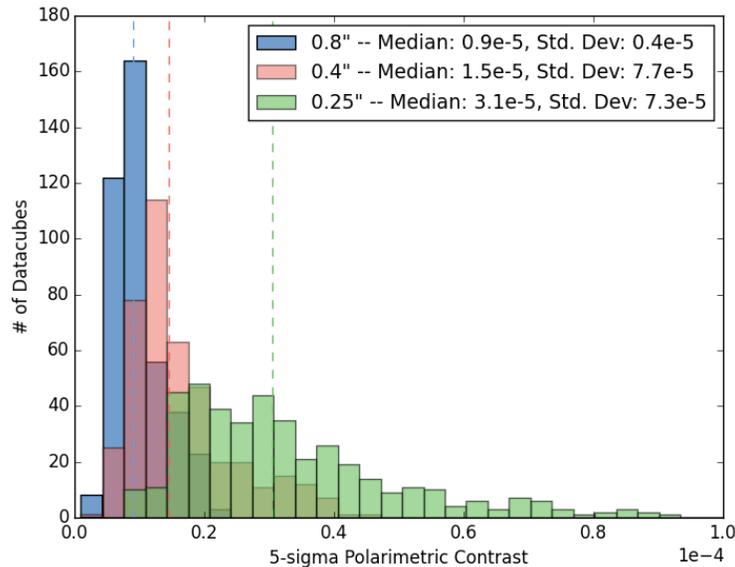


Figure 2.11: A histogram of polarized intensity contrasts from single polarization datacubes obtained throughout the GPIES survey, measured at separations of 0.25'' (green), 0.4'' (red) and 0.8'' (blue). The dotted lines represent the median of each distribution. The sample has been culled of any targets that have detected debris disks.

et al. 2002). The majority of targets that are appropriate for observation with GPI are nearby and should have negligible interstellar polarization. A second advantage is that this method is robust against a changing instrumental polarization, that may vary with time or telescope elevation. In fact, by recording the fractional polarization measured in each frame we are able to monitor and track any changes in the instrumental polarization over time. The analysis of this data is ongoing and will be published at a later time.

2.7 GPI Exoplanet Survey Contrasts

The GPIES campaign is a multi-year Gemini South program with the goal of discovering and characterizing directly imaged exoplanets around young nearby stars using GPI's spectroscopy mode. The campaign also includes a debris disk component with the goal of imaging and characterizing debris disks using GPI's polarimetry mode. Debris disk observations are split into two categories: a shorter snapshot sequence, whose purpose is detecting disks previously unseen in scattered light; and a deeper observation sequence, to obtain higher S/N data for detailed disk characterization. The exoplanet search will target a total of 600 stars in spectroscopy mode and those with a known infrared excess are observed as a polarimetric snapshot immediately following the spectroscopic observations. In total there are roughly 60 targets that will receive a snapshot observation.

A typical polarimetry snapshot is between 8 and 16 1-minute observations, with the HWP rotating between each. At the beginning of the campaign it was standard to obtain 8 1-minute observations for a snapshot. However, because polarimetric snapshots are obtained immediately after a spectroscopic sequence, persistence can significantly impact the depth of the observation sequence (§2.4.1). As a result,

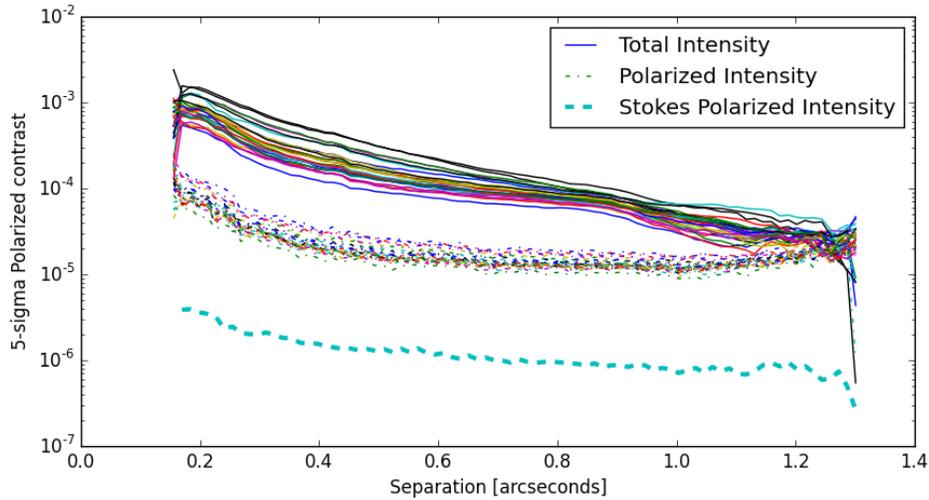


Figure 2.12: Contrasts from the GPIES observations of HD 10472. Displayed at the top are the total intensity and linear polarized intensity contrasts from each polarization datacube for the entire observation set (26 1.5-minute exposures). The bottom line displays the linear polarized intensity contrast of the combined Stokes datacube.

the length of a snapshot was first increased to 12 frames and as of Jan 2016 the standard snapshot time was further increased to 16 frames. Disks detected in a snapshot, or those previously resolved in scattered light are observed with a deep polarimetric sequence, typically 40 1-minute observations. In practice, telescope tracking errors and/or the opening of the adaptive optics control loops can lead to one or more frames being unusable. Thus, the exact number of frames used when forming a Stokes datacube can vary, but is nonetheless on the order of 16 and 40, for snapshot and deep observations, respectively.

As of 2016 May 19, GPIES has observed 24 polarimetric snapshots and 18 deep sequences, amounting to 861 polarization datacubes and 42 Stokes datacubes. Figure 2.11 displays histograms of the polarized intensity contrasts at angular separations of $0.25''$, $0.4''$ and $0.8''$ of all the polarization datacubes without detected disks. The distributions at $0.4''$ and $0.8''$ show strong peaks near a contrast of $1e-5$, with sharp drops to smaller values and small tails that trail off to higher values. This likely indicates that we are reaching the photo/read out noise boundary at these separations for most of our observations, as suggested in §2.4. On the other hand, the distribution for $0.25''$ is much broader. At these separations we believe we are limited by instrumental polarization and our ability to remove it, though future tests will confirm this. Note that in this plot the datacubes have neither been cleaned with the double differencing process nor had instrumental polarization subtracted, because the contrast is typically measured before the cubes are combined in any way. Thus, the values shown here can be considered upper limits.

In Figure 2.12, we display contrast curves from observations of HD 10472, as an example of a typical GPIES deep polarization observing sequence. HD 10472 was observed on 2015 December 12, as part of the GPIES campaign and the observation set consist of 26 1.5-minute exposures. The figure simultaneously displays the total intensity contrast and polarized intensity contrast of the polarization datacubes, as well as the linear polarized intensity contrast of the final Stokes cube. The median total intensity contrasts of the polarization datacubes at $0.25''$, $0.4''$ and $0.8''$ are $4.2e-4$, $1.7e-4$ and $0.8e-4$, respectively. The median linear polarized intensity contrasts of the polarization datacubes at $0.25''$, $0.4''$ and $0.8''$ are $0.44e-4$, $0.19e-4$ and $0.12e-4$, respectively, providing an improvement of roughly a factor of

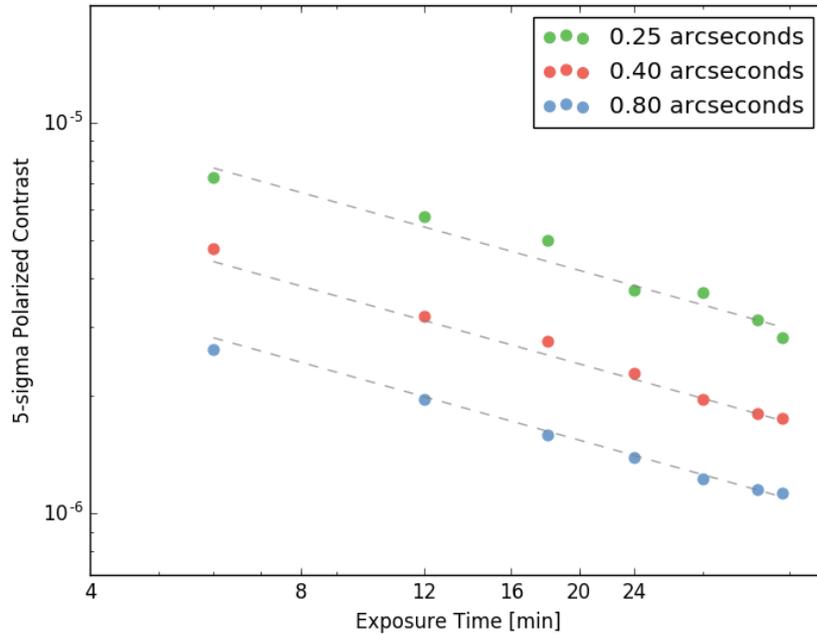


Figure 2.13: Polarized contrast at 0.25", 0.40" and 0.80" as a function of total exposure time for the HD 10472 dataset. The polarized contrast at all three separations appears to decrease as a function of $\sqrt{\text{Exposure Time}}$ (grey dashed lines).

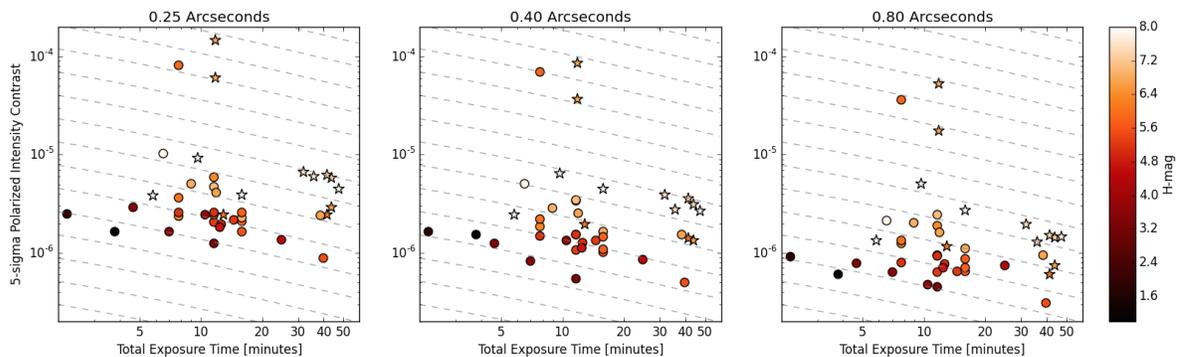


Figure 2.14: Polarized intensity contrasts at 0.25", 0.4" and 0.8" from all the observations throughout the GPIES survey as a function of number of exposures and H-magnitude. Each exposure consists of a 60-s observation, so the x-axis can be considered a proxy for exposure time. Observations with detected disks are marked with a star symbol, and non-detections are marked with circles. The grey dashed lines indicate contrasts decreasing as $\sqrt{\text{Exposure Time}}$.

10 at 0.25'' and 0.4'', and a factor of 6.5 at 0.8''. The linear polarized intensity contrasts of the final Stokes cube at 0.25'', 0.4'' and 0.8'' are 0.024e-4, 0.015e-4 and 0.010e-4, respectively. Thus, we gain factors of roughly 175, 110 and 80 between the total intensity and final linear polarized intensity at 0.25'', 0.4'' and 0.8'', respectively. In this observation set, we find that polarized contrast at all three separations decreases as a function of $\sqrt{\text{Exposure Time}}$ (Figure 2.13).

Final contrasts for all the Stokes cubes produced so far in GPIES (both snapshot and deep observations) can be seen in Figure 2.14 as a function of exposure time and H-magnitude. The plot includes observations both with non-detections and with detected disks. Because the disk flux likely increases the measured contrast, contrasts associated with detected disks can be considered upper limits on the true sensitivity of the observations. In general, deeper contrasts are achieved for brighter targets, and targets with the same H-magnitude appear to gain in contrast roughly as $\sqrt{\text{Exposure Time}}$ (with a few exceptions).

2.8 Conclusions

The Gemini Planet Imager has now been on-sky for over 2.5 years and is producing exciting results through the GPIES campaign, the debris disk Large and Long Program and through Gemini queue observations. Over this time we have developed new data analysis techniques that we have implemented as part of the publicly available GPI DRP.

Three techniques in particular have allowed us to reduce systematics and improve upon the contrast in GPI's Polarimetry mode. First, by using a weighted PSF extraction in the assembly of polarization datacubes from raw data we were able to lower the photon noise/read noise floors of the polarization datacubes, which improves the contrast between 0.3'' and the edge of the field. Second, by applying a polarized flat-field we can reduce systematics near the edge of the field. Third, by measuring and subtracting the apparent stellar polarization we can subtract the instrumental polarization, with contrast benefits from the inner working angle to about 1''. All of these improvements have been included in the GPI DRP and are used in standard reduction recipes of GPIES data. In addition we perform a multi-wavelength analysis of GPI's instrumental polarization, that indicates that the instrumental polarization appears to be roughly wavelength independent.

The combined datasets of the GPIES campaign demonstrate that the polarized contrast can be improved with increased exposure time and the achievable contrast depends on the brightness of the source. An examination of histograms of the contrast of polarization datacubes and the $\sqrt{\text{Exposure Time}}$ dependence of the contrast of a deep polarization sequence indicate that sensitivity at 0.4'' and 0.8'' is likely dominated by photon/read noise (with small contributions from instrumental polarization). We believe to be dominated by residual instrumental polarization or other polarization systematics at separations smaller than 0.25''. Nonetheless, we find that the polarized contrast decreases as a function of $\sqrt{\text{Exposure Time}}$. The best contrasts achieved so far as part of the campaign at 0.25'', 0.4'' and 0.8'' are 9e-07, 5e-07 and 3e-07 respectively, obtained with 40 60s frames on a star with a 5.5 H magnitude. These measurements and others presented throughout this work can be used as baseline estimates of GPI's polarimetry mode's performance when planning future observations.

A full characterization of the different noise contributions and systematics when using the weighted PSF subtraction technique has been left for future work. Ongoing work to further improve our sensitivity includes the development of methods to subtract persistence in raw data and to subtract instrumental

polarization at small inner working angles.

Chapter 3

β Pictoris' Inner Disk in Polarized Light and New Orbital Parameters for β Pictoris *b*

A version of this work has appeared in the Astrophysical Journal as "Beta Pictoris' Inner Disk in Polarized Light and New Orbital Parameters for Beta Pictoris b" by Millar-Blanchaer et al. 2015.

3.1 Chapter Summary

We present *H*-band observations of β Pic with the Gemini Planet Imager's (GPI's) polarimetry mode that reveal the debris disk between $\sim 0.3''$ (6 AU) and $\sim 1.7''$ (33 AU), while simultaneously detecting β Pic *b*. The polarized disk image was fit with a dust density model combined with a Henyey-Greenstein scattering phase function. The best fit model indicates a disk inclined to the line of sight ($\phi = 85.27^{\circ+0.26}_{-0.19}$) with a position angle $\theta_{PA} = 30.35^{\circ+0.29}_{-0.28}$ (slightly offset from the main outer disk, $\theta_{PA} \approx 29^{\circ}$), that extends from an inner disk radius of $23.6^{+0.9}_{-0.6}$ AU to well outside GPI's field of view. In addition, we present an updated orbit for β Pic *b* based on new astrometric measurements taken in GPI's spectroscopic mode spanning 14 months. The planet has a semi-major axis of $a = 9.2^{+1.5}_{-0.4}$ AU, with an eccentricity $e \leq 0.26$. The position angle of the ascending node is $\Omega = 31.75^{\circ} \pm 0.15$, offset from both the outer main disk and the inner disk seen in the GPI image. The orbital fit constrains the stellar mass of β Pic to $1.61 \pm 0.05 M_{\odot}$. Dynamical sculpting by β Pic *b* cannot easily account for the following three aspects of the inferred disk properties: 1) the modeled inner radius of the disk is farther out than expected if caused by β Pic *b*; 2) the mutual inclination of the inner disk and β Pic *b* is $\sim 4^{\circ}$, when it is expected to be closer to zero; and 3) the aspect ratio of the disk ($h_0 = 0.137^{+0.005}_{-0.006}$) is larger than expected from interactions with β Pic *b* or self-stirring by the disk's parent bodies. The inferred disk properties can be further tested with future multi-color polarimetric observations of the disk at similar radii, which will allow for more sophisticated disk models that include more realistic dust grains models.

3.2 Introduction

The dynamical interactions between exoplanets and their local debris disks provide a unique window into the understanding of planetary system architectures and evolution. In this regard, the β Pic system is important as it is one of the rare cases where both a planet and a debris disk have been directly imaged.

The β Pic system first garnered interest after Smith & Terrile (1984) followed up a prominent Infrared Astronomical Satellite (IRAS) infrared excess detection (Aumann 1985) and imaged an edge-on circumstellar disk in dust scattered light. Since then, many observational and theoretical studies have helped to paint a picture of a dynamically active system that contains a fast-spinning directly imaged ~ 10 - $12 M_J$ planet (Lagrange et al. 2009; Snellen et al. 2014; Chilcote et al. 2015), an asymmetric debris disk (Lagage & Pantin 1994; Kalas & Jewitt 1995), infalling small bodies (Beust & Morbidelli 1996; Kiefer et al. 2014), multiple planetesimal belts (Okamoto et al. 2004; Wahhaj et al. 2003), a carbon-rich gas disk (Roberge et al. 2006) and a circling gas cloud that may indicate a recent collision between planetesimals (Dent et al. 2014). In this chapter we examine the nature of the dynamical relationship between the planet, β Pic *b*, and the debris disk using polarimetric imaging and modeling of the innermost region of the disk.

The overall structure of the disk—a depleted inner region, an extended outer region, and an apparent warp—has been well-established in the literature. Smith & Terrile (1984) originally used optical depth arguments to infer that β Pic's disk must be depleted of grains interior to a radius of ~ 30 AU. Burrows et al. (1995) used HST/WFPC2 to image the disk in optical scattered light and described qualitatively a vertical warp in the midplane structure somewhere between $1.5''$ and $10''$ radius. The first quantitative measurements of the midplane warp were derived from ground-based adaptive optics (AO) observations in the near infrared (NIR; Mouillet et al. 1997). In these data, the peak height of the warp is at $\sim 3''$ radius, ~ 58 AU assuming heliocentric distance of 19.44 pc (van Leeuwen 2007), and corresponds to 3° deviation from the position angle (PA) of the midplane measured beyond ~ 100 AU.

Two geometrical interpretations of the apparent warp have been proposed. The first is that we are observing a single disk warped by forcing from a planet on an inclined orbit. Using numerical models and semi-analytic arguments, Mouillet et al. (1997) demonstrated that a planet inclined by 3° – 5° to a hypothetical disk can replicate the observed structure via a secular perturbation. The inferred mass of the planet depends on when the planet's orbit was perturbed out of coplanarity, because in this paradigm the warp propagates radially outward on million year timescales. Augereau et al. (2001) applied this model to explain several other observational features of the disk such as the larger scale asymmetries.

Alternatively, the structure could be composed of two disks, with symmetric linear morphologies, superimposed on the sky plane. Two disks would appear to create a warp in the midplane of the primary disk because of a $\sim 3^\circ$ difference in position angle. Based on high angular resolution optical data obtained with HST/STIS that clearly showed the warp component, Heap et al. (2000) postulated that the sky plane contains “two disks 5° apart.” This interpretation is also favored in subsequent studies based on multi-color HST/ACS/HRC observations of β Pic's disk (Golimowski et al. 2006). More detailed analytic modeling of these data are consistent with two disks with a relative position angle on the sky of $3.2 \pm 1.3^\circ$ (Ahmic et al. 2009). Ahmic et al. (2009) also find that the fainter inclined disk has a line of sight inclination $6.0 \pm 1.0^\circ$, whereas the brighter, primary disk is consistent with being exactly edge-on. More recently Apai et al. (2015) presented a re-reduction of the early HST/STIS observations, coupled with newer observations obtained 15 years apart. They find that these observations are consistent with the two-disk interpretation, but they also examine a scenario where β Pic *b* is perturbing the disk.

Table 3.1. Summary of GPI Observations of β Pic

Date	Observing Mode	Exposure Time (s)	Parallactic Rotation ($^{\circ}$)	Seeing ($''$)	β Pic <i>b</i>	
					Separation (mas)	PA ($^{\circ}$)
2013-11-16	K1-Spec.	1789	26	1.09	430.3 ± 3.2	212.31 ± 0.44
2013-11-16	K2-Spec.	1253	18	0.93	426.0 ± 3.0	212.84 ± 0.42
2013-11-18 ^a	H-Spec.	2446	32	0.68	428.1 ± 2.7	212.22 ± 0.39
2013-12-10	H-Spec.	1312	38	0.77	418.8 ± 3.6	212.64 ± 0.53
2013-12-10 ^b	J-Spec.	1597	18	0.70	419.1 ± 6.2	212.16 ± 0.81
2013-12-11	H-Spec.	556	64	0.46	419.2 ± 5.1	212.26 ± 0.72
2013-12-12	H-Pol.	2851	91	0.43	426.6 ± 7.0	211.80 ± 0.68
2014-03-23	K1-Spec.	1133	47	0.47	412.5 ± 2.7	212.08 ± 0.41
2014-11-08	H-Spec.	2147	25	0.77	362.94 ± 4.1	212.17 ± 0.65
2015-01-24	H-Spec.	716	5	0.85	347.73 ± 4.7	212.17 ± 0.65

^aThese observations were published by Macintosh et al. (2014), but have been rereduced here to maintain homogeneity across the datasets

^bThese observations were published by Bonnefoy et al. (2014), but have been rereduced here to maintain homogeneity across the datasets

The perturbing planet scenario requires a planet with a mass, semi-major axis, and mutual inclination with respect to the flat outer disk sufficient to create the warp. Lagrange et al. (2009) discovered β Pic *b*, a planet with a mass and separation appropriate for creating the warp; with additional astrometric measurements, its orbit was constrained to $a \sim 9$ AU, $i \sim 89^{\circ}$ and $e \sim 0.1$ (Chauvin et al. 2012; Macintosh et al. 2014). If the planet is secularly perturbing the disk, we expect it to be in the same plane as the inner disk and misaligned from the flat outer disk (though it may appear to be aligned in projection). One technical challenge is that the planet location, the inner warp and the outer disk have been measured on different angular scales and are detected using different observing strategies. Therefore, systematic errors in the position angle calibrations between different data sets lead to uncertainty in the relative orientations of these three structures. For example, Currie et al. (2011) reported that the planet’s orbit is misaligned with the inner disk, but Lagrange et al. (2012b) noted that they are consistent with alignment when all sources of error are accounted for.

Lagrange et al. (2012b) attempted to solve these problems by constructing observations where a single instrument is used to simultaneously detect both the planet and the disk. The results show β Pic *b* positioned 2-4 $^{\circ}$ above the southwest disk midplane at the 2010 epoch of observation (“above” means north of the SW midplane or at a larger PA than the SW midplane). Therefore, β Pic *b*’s orbit is not coplanar with the main, flat, outer disk. Instead the position above the main midplane in the SW is in the direction of the warped component. This projected misalignment is consistent with the necessary mutual inclination between the planet’s orbit and the main, flat, outer disk.

A different technical challenge is imaging the disk along the minor axis direction very close to the star in order to establish small inclinations away from edge-on (Kalas & Jewitt 1995). A small inclination away from edge-on (85-89 $^{\circ}$) is difficult to ascertain at large separations because the sharpness of the disk midplane in projection (i.e., the shape of a cut perpendicular to the midplane) is a combination of the disk scale height and the small inclination to the line of sight. Closer to the star, however, the small inclination combined with an asymmetric scattering phase function tends to shift the isophotes so that the disk does not exactly intersect the star. For example, if the disk midplane appears to pass “above” the star, then that is taken as evidence that the disk comes out of the sky plane above the star, and enhanced forward scattering leads to the apparent misalignment between the midplane and the star.

For β Pic, Milli et al. (2014) discovered that the disk midplane traces a line that lies above the star. They inferred a ~ 86 – 89° disk inclination from modeling the data, using a Henyey-Greenstein phase function with $g = 0.36$. One significant issue with inferring the line-of-sight inclination from their L' dataset is that the $3.8 \mu\text{m}$ morphology of the disk within 10 AU is a combination of scattered light and thermal emission. Therefore the very warm dust near the star contributes to the detected flux within $0.5''$. Milli et al. (2014) concluded that shorter wavelength observations, that are less-contaminated by thermal emission, are necessary to disentangle the geometry of the system within $0.5''$ radius.

The technique used to image β Pic b relies on angular differential imaging (ADI; Marois et al. 2006) to achieve sub-arcsecond inner working angles (Lagrange et al. 2012b; Milli et al. 2014; Nielsen et al. 2014). For ground-based observations, this technique typically provides more effective point spread function (PSF) subtraction than using PSF reference stars images, which are subject to the time variability of the AO-corrected PSFs. However, when applied to extended objects—such as circumstellar disks—ADI often causes significant self-subtraction (e.g. Milli et al. 2012), impacting the accuracy of derived model disk parameters. These effects can be mitigated with forward modeling (e.g. Esposito et al. 2014), but self-subtraction can be largely avoided through polarimetric differential imaging (PDI; Kuhn et al. 2001). PDI takes advantage of the fact that scattered light is inherently polarized while stellar radiation is not, to subtract the unpolarized stellar PSF, revealing the polarized disk underneath.

In this chapter we present polarimetric observations of β Pic's debris disk at $1.6 \mu\text{m}$ (H -band), taken with the Gemini Planet Imager. The data simultaneously reveal the debris disk in polarized light and β Pic b in unpolarized light. These observations provide a unique perspective on the vertical extent of the disk at small angular separations, where ADI self-subtraction is typically the most severe. In addition, we present new astrometric measurements of the companion β Pic b taken with GPI's spectroscopy mode, which we use to provide an updated orbital fit.

In §3.3 we provide a description of the observations and data reduction steps for both polarimetry and spectral mode data. We describe our analysis of the disk image in §3.4, and our orbit fitting in §3.5. In §3.6, we discuss our interpretation of the two fits, both individually and in the context of the disk-planet interaction. We present our conclusions in §3.7.

3.3 Observations and Data Reduction

The Gemini Planet Imager (GPI) is a recently commissioned NIR instrument on the Gemini South telescope, designed specifically for the direct imaging of exoplanets and circumstellar disks (Macintosh et al. 2014). The optical path combines high-order adaptive optics (Poyneer et al. 2014), with an apodized pupil Lyot coronagraph (Soummer et al. 2011) that feeds an integral field spectrograph (IFS; Larkin et al. 2014). The coronagraph system masks out the central star, while simultaneously suppressing diffraction caused by the telescope and its support structure. Within the IFS, GPI's focal plane is sampled by a lenslet array at a spatial scale of $14.13 \text{ mas/lenslet}$ (see §3.5.1) over a $\sim 2.8'' \times 2.8''$ square field of view. The light from each lenslet is passed through either a spectral prism, to allow for low resolution ($R \sim 45$) integral field spectroscopy, or a Wollaston prism, for broadband integral field polarimetry. During observations, Gemini's Cassegrain rotator is turned off to allow the sky to rotate while the orientation of the PSF remains static with respect to the instrument.

The complexity of the instrument results in an intricate path from raw data to a fully processed datacube, requiring many calibrations and transformations to obtain a final data product. As a result,

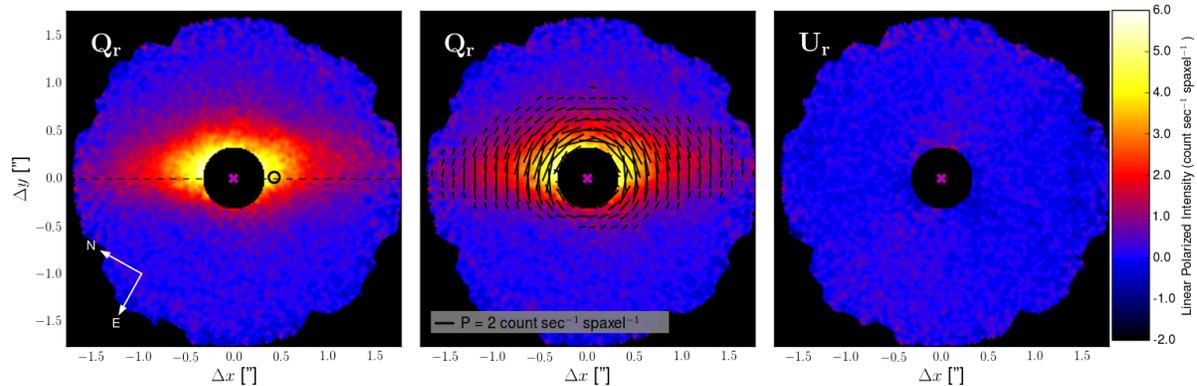


Figure 3.1: *Left*: The β Pic debris disk in polarized intensity, Q_r , from observations with GPI’s polarimetry mode. The disk image has been rotated so that the midplane of the outer disk (P.A.= 29.1° ; black dashed horizontal line) is horizontal. The star’s location and β Pic b ’s location are marked by the magenta \times and black circle, respectively. *Center*: The Q_r image with polarization vectors overlotted. Though the centrosymmetric nature of the polarization is captured in the transformation to the radial stokes images, the vectors serve to emphasize this behavior. *Right*: The radial polarized intensity, U_r from the same datacube as the image on the left, shown at the same color scale. For circumstellar material, the flux is expected to be solely in the Q_r image, thus this images provides an estimate of the noise in the disk image.

the GPI Data Reduction Pipeline (DRP) has been designed as a dedicated software application for processing GPI data. A full description of the GPI DRP can be found in Perrin et al. (2014c), Maire et al. (2010) and references therein. A walkthrough of the data reduction process for GPI polarimetry data can be found in Perrin et al. (2015). Below, we provide a brief summary of the observations and relevant data reduction steps. All the data described herein were reduced using the GPI Pipeline version 1.2 or later¹.

3.3.1 Polarimetry Mode Observations

Polarimetric observations of β Pic were carried out on 2013 December 12 UT, while performing a series of AO performance and optimization tests (Table 3.1). β Pic was observed for a total of forty-nine 60 s frames, during which the field rotated in parallactic angle by $\sim 91^\circ$. Between each image the half waveplate (HWP) modulator was rotated by 22.5° . For 25 frames, GPI’s two Sterling cycle cryocoolers (Chilcote et al. 2012) were disabled to minimize vibration² in the telescope and instrument to improve the AO performance. The average seeing throughout the observations, as measured on the external Gemini DIMM, was $0.43''$ and β Pic b can easily be seen in the majority of raw detector images, even before extraction into datacubes. Each raw data frame was dark-subtracted, corrected for bad-pixels and then ‘destriped’ to remove correlated noise in the raw image caused in part by vibration induced by the cryocoolers (Ingraham et al. 2014).

In GPI’s polarimetry mode (pol. mode), a Wollaston prism splits the light from each lenslet into two spots of orthogonal polarization states on the detector. Flexure effects within the instrument cause these lenslet PSF spots to move from their predetermined locations on the detector, typically by a fraction of a pixel. For each frame, the PSF offset was determined using a cross-correlation between the raw frame

¹<http://planetimager.org/datapipeline>

²Since the time of these observations, the level of vibration caused by the cryocoolers has been significantly mitigated through the use of a new controller, which drives the two coolers 180° out of phase (Hartung et al. 2014). The vibration caused by the two coolers interferes destructively and the overall effect is significantly damped.

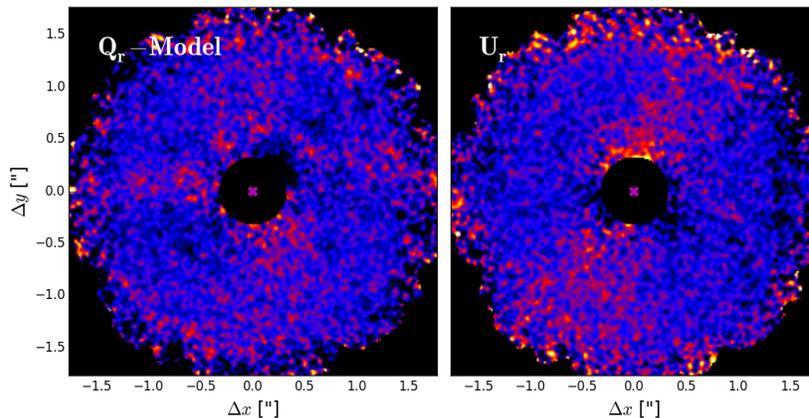


Figure 3.2: *Left*: The residuals of the Q_r image minus the disk model, shown at a stretch to emphasize their structure. The residual structure in the NW-SE direction is likely due to imperfect subtraction of the instrumental polarization, but may also represent the difference between the true scattering properties of the grains and the Henyey-Greenstein function used in the model. *Right*: The U_r image, shown at the same stretch as the residuals. The star’s location is marked with a magenta \times .

and a set of lenslet PSF models measured using a Gemini Facility Calibration Unit (GCAL) calibration frame. The overall method is described in Draper et al. (2014) using high-resolution microlenslet PSFs. Here we use a Gaussian PSF model, which is less computationally intensive, but provides similar results. The raw frames were then reduced to polarization datacubes (where the third dimension carries the two orthogonal polarizations) using a weighted PSF extraction centered on the flexure-corrected location of each of the lenslets’ two spots (see Perrin et al. 2015).

Each cube was divided by a reduced GCAL flat field image, smoothed using a low pass filter. The flat field corrects simultaneously for throughput across the field and a spatially varying polarization signal. In theory, this polarization signal should be removed during the double differencing procedure later in the pipeline; however, we have found empirically that this polarization signal is best divided out of each cube individually.

To determine the position of the occulter-obscured star, a Radon-transform-based algorithm (Pueyo et al. 2015) was used to measure the position of the elongated satellite spots (Wang et al. 2014). Knowledge of the obscured star’s location is critical when combining multiple datacubes that must be both registered and rotated. Each datacube was then corrected for distortion across the field of view (Konopacky et al. 2014). The datacubes were then corrected for any non-common path biases between the two polarization spots using the double differencing correction described by Perrin et al. (2015), before being smoothed with a 2-pixel FWHM Gaussian profile.

Instrumental polarization, due to optics upstream of the waveplate, converts unpolarized light from the stellar PSF into measurable polarization that, if left uncalibrated, can mimic an astrophysical signal. This signal was removed from each difference cube individually, first by measuring the average fractional polarization (i.e., the difference of the two orthogonal polarization slices divided by their sum) inside of the occulting mask, where the flux is due solely to star light diffracting around the mask. We assume that this fractional polarization signal is due to polarization of unpolarized stellar flux by the instrument and telescope. For each lenslet, the fractional instrumental polarization was then multiplied by the

total intensity at that location, and then subtracted off in a similar manner to the double differencing correction. Using this method we find the instrumental polarization to be $\sim 0.5\%$, a similar level to that reported by Wiktorowicz et al. (2014b) using the same dataset.

The difference cubes were then shifted to place the obscured star at the center of the frame and then rotated to place North along the y -axis and East along the x -axis. All of the polarization datacubes were then combined using singular value decomposition matrix inversion to obtain a three dimensional Stokes cube, as described in Perrin et al. (2015). Non-ideal retardance in GPI's HWP makes GPI weakly sensitive to circular polarization, Stokes V . Measurements of the circular polarization of an astrophysical source would require knowledge of the HWP's retardance well beyond the current level of calibration. Therefore, in almost all cases the Stokes V cube slice should be completely disregarded.

The Stokes datacube was then transformed to 'radial' Stokes parameters: $(I, Q, U, V) \rightarrow (I, Q_r, U_r, V)$ (Schmid et al. 2006). Under this convention, each pixel in the Q_r image contains all the linear polarized flux that is aligned perpendicular or parallel to the vector connecting that pixel to the central star. A positive Q_r value indicates a perpendicular alignment and a negative Q_r indicates a parallel alignment. Note that this sign convention is opposite that used in Schmid et al. (2006), where positive values of Q_r correspond to a parallel alignment. The U_r image holds the flux that is aligned $\pm 45^\circ$ to the same vector. For an optically thin circumstellar disk, the polarization is expected to be perpendicular and all the flux is expected to be positive in the Q_r image. The U_r image should contain no polarized flux from the disk and can be treated as a noise map for the Q_r image. The final reduced disk image can be seen in Figure 3.1.

3.3.2 Spectral Mode Observations

Observations of β Pic in spectroscopic mode (spec. mode) were carried out during four separate GPI commissioning runs, as well as during an ongoing astrometric monitoring program scheduled during regular general observing time. In total, we present ten individual sets of observations over eight unique epochs (Table 3.1). Two of the observation sets have been previously published: the H -band dataset from 2013 November 18 (Macintosh et al. 2014) and the J -band dataset from 2013 December 10 (Bonnetfoy et al. 2014). Here we have re-reduced all the data in a consistent manner in an effort to reduce systematic biases and maximize the homogeneity of the dataset. As with the polarization mode observations, those observations that were taken during the instrument's commissioning were carried out during AO optimization tests and therefore have a range of exposure times and filter combinations.

All datasets were reduced with standard recipes provided by the GPI DRP. Raw data frames were dark subtracted and destriped for microphonics in the same manner as the polarimetry observations. A short exposure arc lamp image was taken contemporaneously with each science observation to measure the offsets of the lenslet spectra due to flexure within the IFS. The mean shift was calculated for a subset of lenslets across the field of view relative to a high SNR arc lamp image taken at zenith via a Levenberg-Marquardt least-squares minimization algorithm (Wolff et al. 2014).

The raw detector image was then transformed into a spectral datacube, using a box extraction method. For observations obtained with the $K1$ and $K2$ filters, thermal sky observations were taken immediately before or after the observation sequence. Sky background cubes were created in the same manner described above and subtracted from science datacubes. Finally, all cubes were corrected for distortion (Konopacky et al. 2014).

Each data-set was PSF subtracted using the methods outlined in Pueyo et al. (2015). To minimize

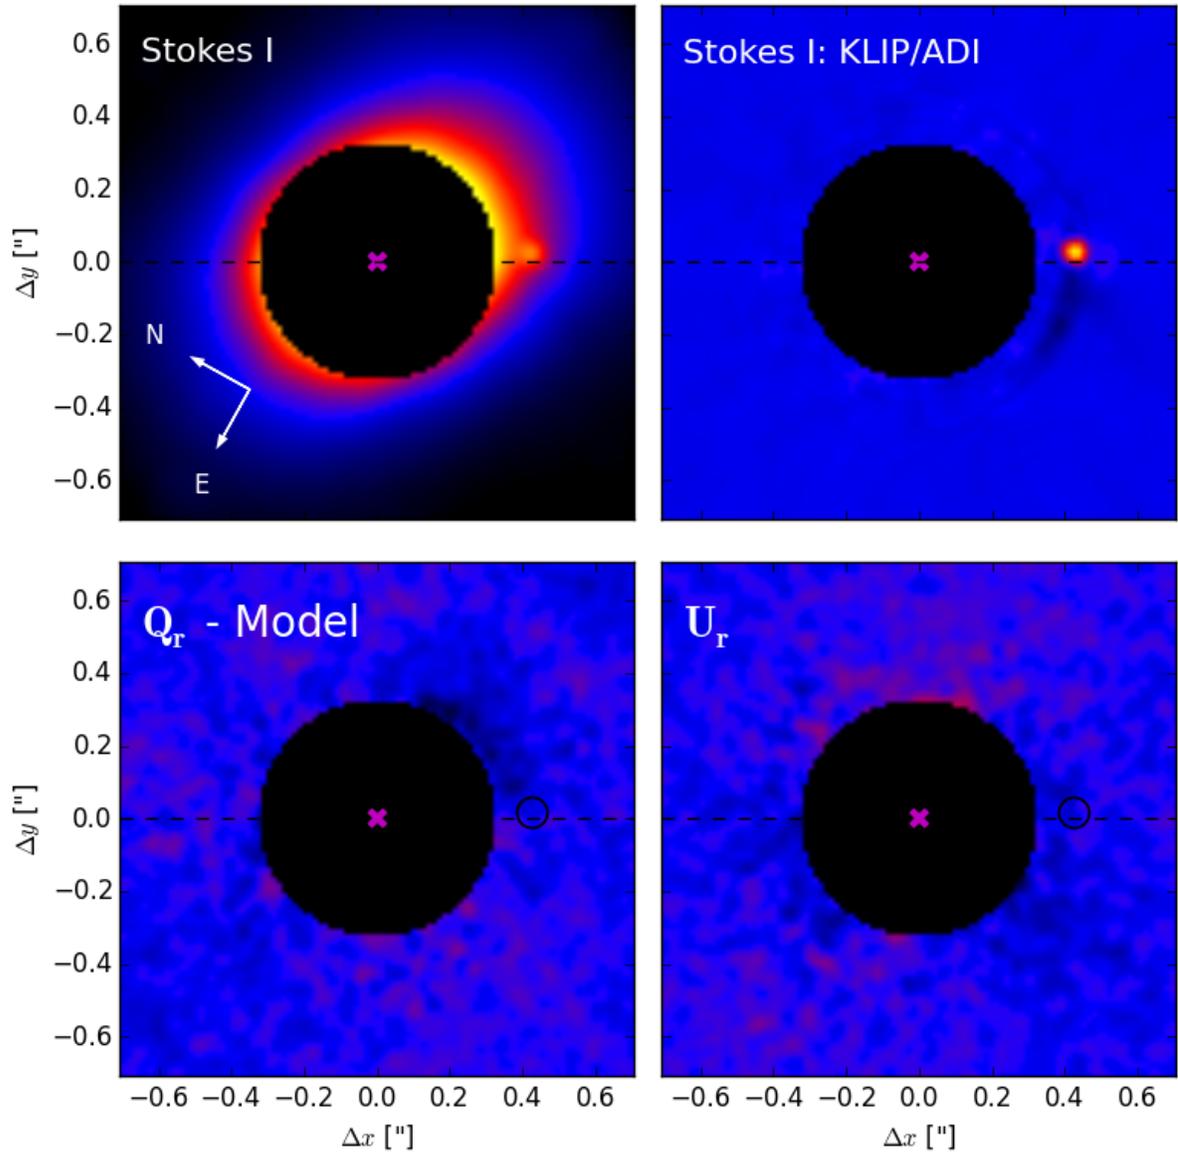


Figure 3.3: *Top Left:* The Stokes I image from the polarimetry observations, without PSF subtraction. The dashed line indicates the position angle of the outer disk. The planet can be seen at a separation of $\sim 0.4''$ just above the horizontal line, to the SW from the central star. *Top Right:* The Stokes I image after applying KLIP/ADI PSF subtraction. The planet is recovered at a very high SNR. *Bottom Left:* The polarized intensity image, Q_r , after disk model subtraction. The black circle indicates the location of the planet in the Stokes I images. *Bottom Right:* The radial Stokes image (same as in Figure 3.1), U_r . No point source polarization signal is detected for β Pic b in either Stokes Q_r or U_r . All images have been rotated so that the outer disk's PA is horizontal (dashed black line). In all four images the star's location is marked with a magenta \times .

systematic biases, the ensemble of datasets was treated as uniformly as possible. The main steps of this data reduction process include: high-pass filtering, to remove the remaining PSF halo; wavelength-to-wavelength and cube-to-cube image registration, to correct for atmospheric differential refraction and sub-pixel stellar motion across the observing sequence; subtracting the speckles using the KLIP principal component analysis algorithm (Soummer et al. 2012) on each wavelength slice in each cube; rotation to align the north angle of each image; and co-adding the resulting cubes in time.

For the epochs in which β Pic b was observed on consecutive nights, relative alignment was tested using both the cross correlation method described in Pueyo et al. (2015) and the absolute stellar locations based on the satellite spot centroids derived using the GPI DRP. For these epochs we found better consistency in the location of β Pic b using the DRP centroids, which we then chose to adopt for all datasets.

The KLIP algorithm was implemented using both spectral differential imaging (SDI; Marois et al. 2000) and ADI, building for each slice a PSF library that takes advantage of the radial and azimuthal speckle diversity (in wavelength and in PA, respectively). Due to the relative brightness of β Pic b with respect to the neighboring speckles we limited the exploration of KLIP parameter space to two zone geometries and two exclusion criteria (1 and 1.5 PSF FWHM) for each dataset. For each slice, the 30 PSFs that were the most correlated in the region where β Pic b is located were used for PSF subtraction, except for the J-band data which required 50 PSFs for satisfactory subtraction.

To determine the optimal number of principal components to use for each dataset, we examined both the evolution of the extracted spectrum and the astrometric stability as a function of wavelength as we increased the number of components. This latter test helps us to rule out the pathological cases for which either a residual speckle (i.e. insufficiently aggressive PSF subtraction) or self-subtraction (i.e. over-aggressive PSF subtraction) bias planet centroid estimates. We checked for potential biases by comparing astrometric positions measured when using only a high-pass filter with those measured when applying KLIP. Finally, we checked for self-consistency by injecting six synthetic point sources at the separation of β Pic b and at different position angles. Based on these tests we concluded that the astrometric measurements do not feature systematics either introduced by residual speckles or biases associated with KLIP above the uncertainty levels reported in §3.5.1.

3.4 Disk Results

The debris disk is recovered in polarized light from $\sim 1.7''$ (32 AU), to an inner working angle of $\sim 0.32''$ (6.4 AU); see Figure 3.1. While GPI's H -band focal plane mask extends to a radius of $\sim 0.12''$, uncorrected instrumental polarization and other noise sources dominate over the disk emission at separations smaller than $\sim 0.32''$.

A comparison of the Q_r and U_r images indicate that the disk is detected at a high signal-to-noise ratio out to the edge of the GPI field. Overplotting linear polarization vectors indicates that the emission is perpendicularly polarized, as expected. This property is captured in the transformation to radial Stokes parameters, but we have included the vectors in Figure 3.1 for additional clarity.

Morphologically, the disk appears vertically offset from the midplane of the outer disk in the NW direction, indicative of a slight inclination relative to the line of sight. This is consistent with previous models of the disk at similar angular separations (e.g. Milli et al. 2014).

The U_r image shows low level structure in the form of a dipole-like pattern with positive emission in

Table 3.2. Model Parameters and Prior Distributions for the β Pic Disk Model

Parameter	Symbol	Range	Prior Distribution
Inner Radius	R_1	1 – 100 AU	Uniform in $\text{Log}(R_1)$
Outer Radius	R_2	$R_1 - 500$ AU	Uniform in $\text{Log}(R_2)$
Density Power Law Index	β	0.5 – 4	Uniform in β
Scale height aspect ratio	h_0	0.01 – 2	Uniform in $\text{Log}(h_0)$
HG asymmetry parameter	g	0 – 1	Uniform in g
Line of Sight Inclination	ϕ	$80^\circ - 90^\circ$	Uniform in $\cos \phi$
Position Angle	θ_{PA}	$25^\circ - 35^\circ$	Uniform in θ_{PA}

the E-W direction. Figure 3.2 displays the U_r image with a color scale that emphasizes this structure. In the radial Stokes basis, this is the pattern produced by a constant linear polarization across the field, which could be associated with residual instrumental polarization that was not successfully subtracted during the data reduction process. Since the level of these residuals is much lower than the disk emission, we defer improvement of our instrumental polarization subtraction procedure for future work.

The disk is not detected in total intensity (Stokes I ; Figure 3.3), where images are dominated by the residual uncorrected PSF of the star itself. Due to both the extended nature of the disk at these angular scales and frame-to-frame variation of the PSF (compounded by the AO tests carried out during the observing sequence), ADI PSF subtraction has proven unsuccessful. Without an unbiased total intensity image of the disk, characterization of the polarization fraction remains out of reach at present. As a result, we opt to model only the polarized intensity.

3.4.1 Disk Modeling

The principal objective of our disk modeling is to retrieve basic geometric properties of the disk. The modeling approach adopted here is to combine a simple recipe for the 3D dust density distribution with a parametric model of the polarized scattering phase function and then fit to the data using the parallel-tempered sampler from the **emcee** Markov chain Monte Carlo (MCMC) package (Foreman-Mackey et al. 2013). Parallel tempering uses walkers at different ‘temperatures’ to broadly sample the posterior distributions and is therefore a useful strategy when the likelihood surface is complex.

For a disk seen in edge-on projection, the radial dust density distribution becomes degenerate with the dust scattering properties. This degeneracy is typically broken with the use of physical grain models, which describe scattering properties (including polarization) as a function of wavelength. In practice, observations are fit to grain models either using simultaneous polarization and total intensity information (e.g. Graham et al. 2007), or multicolor images (Golimowski et al. 2006). With only single wavelength polarized intensity images available, we instead use the Henyey-Greenstein (HG) scattering function (Henyey & Greenstein 1941) to describe the scattering efficiency as a function of scattering angle. The shape of the HG scattering function is a function of only one parameter, the expectation value of the cosine of the scattering angle, $g = \langle \cos \theta \rangle$, and thus provides a useful tool to approximate grain scattering when using physical models is impractical. The applicability of the HG scattering function to the modeling of our polarized intensity images is discussed in §3.6.1.

Our dust density model, expressed in stellocentric coordinates, $\eta(r, z)$, follows a power law between an inner radius, R_1 , and an outer radius, R_2 , and has a Gaussian vertical profile with RMS height $h_0 r$

and constant aspect ratio, h_0 :

$$\eta(r, z) \propto \left(\frac{r}{R_1}\right)^\beta \exp\left[-\frac{1}{2}\left(\frac{z}{h_0 r}\right)^2\right]$$

where r is the distance from the star, z is the height above the disk midplane and β is the power law index of the dust density. Inside R_1 and outside R_2 the dust density is zero. The dust density distribution is combined with the Henyey-Greenstein function, $H(\theta, g)$, to generate a scattered light image of the disk as seen in 2D projection from the observer's frame, where the intensity for a given pixel (x', y') is calculated as the integral along the line-of-sight direction \hat{z}' :

$$I(x', y') = I_0 + \int_{z'=-R_2}^{R_2} \frac{N_0}{r^2} \eta_\phi(r, z) H(\theta, g) dz'.$$

Here, $\eta_\phi(r, z)$, represents the dust density distribution, but tilted with a disk inclination, ϕ , relative to the observer's line of sight. The scattering angle $\theta = \theta(x', y', z')$ is a function of position.

The $1/r^2$ term accounts for the diminishing stellar flux as a function of distance from the star. The disk's position angle, θ_{PA} , is implemented as a coordinate transformation between the stellocentric coordinates and the projected observer's coordinates. The constant, I_0 , and the flux normalization, N_0 , have been included to account for any possible biases and the conversion between model flux and detector counts, respectively. In summary, our model has a total of nine free parameters: $R_1, R_2, \beta, h_0, g, \phi, \theta_{PA}, N_0, I_0$.

Within the current model there exists a degeneracy between forward scattering ($g > 0$) with an inclination of $\phi < 90^\circ$ and backwards scattering ($g < 0$) with an inclination of $\phi > 90^\circ$. In an effort to conserve computation time we chose to assume forward scattering and place a prior constraint on the scattering parameter, $g > 0$, which is consistent with the model of Milli et al. (2014). A summary of the model parameters and their prior distributions can be found in Table 3.2.

We fit the model to the GPI disk image using the parallel-tempering sampler from the **emcee** package. In the H -band, the FWHM of Gemini's diffraction limit is $0.043''$, equal to about three GPI pixels. We therefore apply a 3×3 pixel binning to both the Q_r and U_r images before fitting. This improves the noise statistics and speeds up the execution time of the MCMC fit, without sacrificing spatial information. At each angular separation in the Q_r image, the errors were estimated as the standard deviation of a 3 pixel-wide annulus centered at that separation in the U_r image. The error estimates therefore contained photon noise, read noise and the unsubtracted instrumental polarization.

The MCMC sampler was run for 2500 steps with 2 temperatures, 128 walkers and burn-in of 500 steps. Additional temperature chains were not employed because of the additional computational cost incurred and the lack of evidence that the Markov chain sampler was selecting only local islands of high probability. One strength of using ensemble sampling over other types of sampling for MCMC fitting is that large speed-ups are possible via parallel-processing. On a 32-core (2.3 GHz) computer the entire MCMC run took nearly five days to complete.

After the run, the maximum auto-correlation across all parameters was found to be 85 steps, indicating that the chains should have reached equilibrium (Foreman-Mackey et al. 2013 recommend $\sim \mathcal{O}(10)$ autocorrelation times for convergence). In addition, the chains were examined by eye and appeared to have reached steady-state by the end of the burn-in phase. The posterior distributions (Figure 3.4) were estimated from the zero temperature walkers, using only one of out every 85 steps to ensure statistical

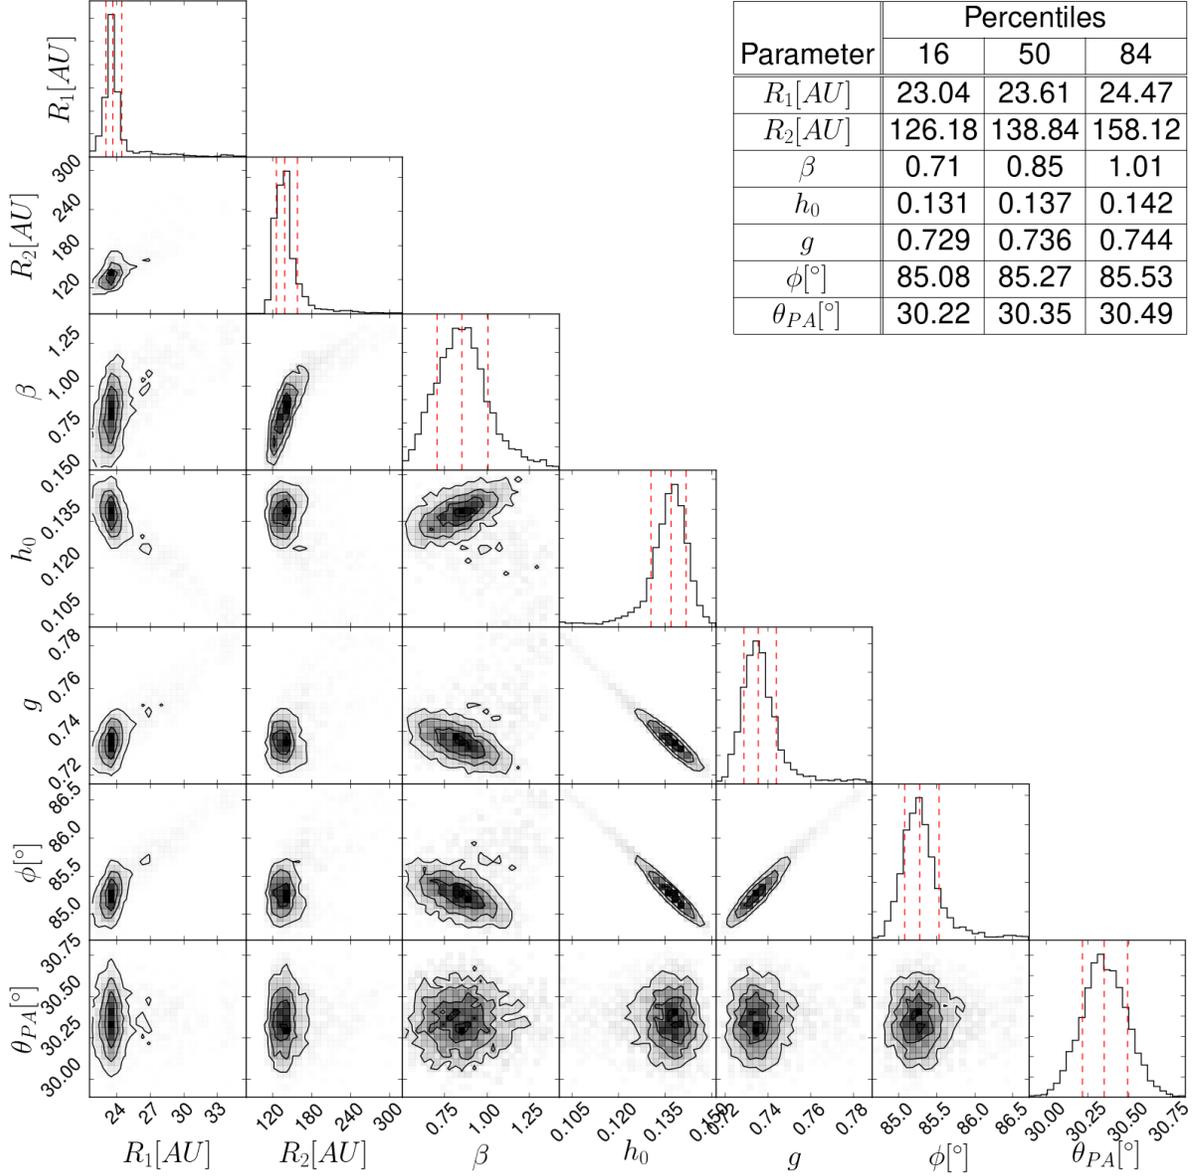


Figure 3.4: The posterior distributions of the model parameters from MCMC disk model fitting to the Q_r disk image. The diagonal histograms show the posterior distributions of each parameter marginalized across all the other parameters. In each plot, the dashed lines indicate the 16%, 50% and 84% percentiles. The off-diagonal plots display the joint probability distributions with contour levels at the same percentiles. The normalization term, N_0 , and the constant offset term, I_0 , have been excluded from this plot because they convey no relevant astrophysical information. *Inset table:* The 16%, 50% and 84% percentiles from the marginalized distributions.

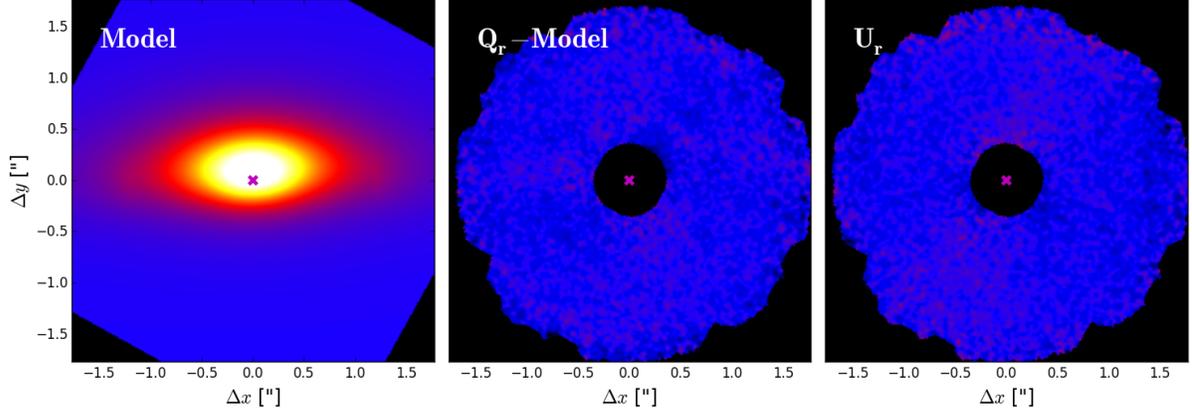


Figure 3.5: *Left*: The disk model generated with the median values from the marginalized posterior distributions (as found in Figure 3.4). The inner edge of the disk is at a projected separation of $1.2''$, but contributes negligible light to the observed surface brightness. *Center*: The residuals of the Q_r image minus the disk model. The level of the residuals is very similar to the U_r image. *Right*: The U_r image, reproduced here as a point of comparison to the residual image. In all images the star's location is marked with a magenta \times . The images have been rotated so that the outer disk is horizontal and all are displayed at the same colour scale as Figure 3.1.

independence of the surviving samples. The expected covariance between the inclination ϕ and g (Kalas & Jewitt 1995) is reproduced. Degeneracies are also found between ϕ and g , h_0 and g , h_0 and β , h_0 and ϕ , β and g , and β and ϕ , but in all cases, the parameters appear to be well constrained.

The 16%, 50% and 84% percentiles for each parameter are displayed in a table in the upper right corner of Figure 3.4. Marginalized across all parameters we find a disk inclined relative to the line of sight by $\phi = 85.27^\circ_{-0.19}^{+0.26}$, with an inner radius of $R_1 = 23.6_{-0.6}^{+0.9}$ AU, an outer radius of $R_2 = 139_{-11}^{+19}$ AU and an aspect ratio of $h_0 = 0.137_{-0.006}^{+0.005}$. The position angle of the disk is $\theta_{PA} = 30.35^\circ_{-0.28}^{+0.29}$, where the error include GPI's systematic error in position angle ($\sim 0.2^\circ$). Note that the systematic uncertainty in the position angle is not reflected in Figure 3.4. The dust is well fit by forward scattering grains, with a scattering asymmetry parameter of $g = 0.736_{-0.007}^{+0.008}$. These results are further discussed in §3.6.1.

Figure 3.5 displays the best fit model and the residuals of the Q_r image minus the model. The best fit model was generated using the median value of each parameter in the marginalized posterior distribution. We find that the highest likelihood disk model successfully reproduces the GPI data. When examined at a different color scale, the residuals image displays similar low-level structure as that of the U_r image (Figure 3.2). The structure in the NW-SE direction is likely the Q_r counterpart of the residual instrumental polarization that's seen in the U_r image. A possible alternative explanation is that the structure could be due to a mismatch between the true scattering properties of the dust and the Henyey-Greenstein scattering function at small angular separations. A second structure can be seen along the disk midplane to the NE of the star. This asymmetric brightness feature is possibly due to an overdensity of dust, that would increase the scattering at that location. Indeed, the β Pic disk is known to have multiple brightness asymmetries (Apai et al. 2015 provide a good summary). However, the feature is detected at similar brightness levels as the residual instrumental polarization and may yet be an uncharacterized artifact of the polarimetry reduction. Deeper observations of the disk will be required to distinguish between a true brightness asymmetry and instrumental effects.

3.5 Planet Results

3.5.1 Astrometry in Spectroscopy Mode

We describe here in broad terms our astrometric measurements and estimation of uncertainties, without delving into the details of each individual dataset. For each epoch, the entirety of the dataset is combined to estimate the planet's position relative to β Pic. The errors on this relative position are a combination of the error on the star's position, the planet's position, GPI's pixel scale and the accuracy to which we know GPI's orientation relative to true North.

For each dataset, the stellar position was calculated using two methods. The wavelength slices of each datacube were first registered using the relative alignment procedure described in §3.3.2 and then collapsed into a broadband image. A Radon transform was then performed on the radially elongated satellite spots to find the stellar position (as in Pueyo et al. 2015). The stellar position was also estimated using the geometrical mean of the satellite spot locations provided by the GPI DRP. Most *H*-band datasets show agreement between two methods at the 0.05 pixel level, with the exception of the 2013 Dec. 17 commissioning sequence, during which extensive AO performance tests were being carried out. For *K*-band datasets the difference between the two methods is no more than 0.05 pixels and for the *J*-band data-set it is 0.2 pixels. We found greater consistency in the relative location of β Pic *b* between observations obtained on consecutive nights when using the Radon method, and therefore chose to adopt the centroids measured with the Radon method for all measurements. For each dataset, we considered the difference between the two methods as our estimate for the uncertainty on stellar position.

The location of β Pic *b* (in detector coordinates) was estimated at each wavelength channel, at each epoch and in each filter using the modified matched filter described in Pueyo et al. (2015). The uncertainty in β Pic *b*'s location was estimated as the scatter in the position of the planet as a function of wavelength and number of principal components. We found the uncertainty to range from 0.05 pixels, for the datasets with significant field rotation and where the planet was at larger separations, up to 0.15 pixels, for the later epochs where the planet is significantly closer to the stellar host and SDI is less effective.

We estimated GPI's pixel scale using the methods described in Konopacky et al. (2014) by combining all the data presented therein with four new observations of Theta 1 Ori B, taken between September 2014 and January 2015. We find an updated pixel scale value of 14.13 ± 0.01 mas/lenslet. Konopacky et al. (2014) measured a PA offset of $-1.00 \pm 0.03^\circ$ during GPI commissioning. Subsequently, version 1.2 of the GPI DRP was updated to incorporate that 1° offset and correct for it automatically. Using the new measurements of Theta 1 Ori B, we find a residual PA offset of $-0.11 \pm 0.25^\circ$.

Based on the measured location of β Pic *b* and its parent star in detector coordinates we calculated the separation and position angle at each epoch. The separation was converted to milliarcseconds using the new platescale estimate and the PA was adjusted by -0.11° . The separation and position angle from each measurement can be found in Table 3.1. Uncertainties on these quantities were combined with the errors on the star position and planet position to yield the errors presented in the table.

3.5.2 Astrometry in Polarimetry Mode

β Pic *b* is detected in the Stokes *I* image as a point source superimposed on the extended PSF halo (Figure 3.3). After applying PSF subtraction using a python implementation of KLIP/ADI (Wang et al.

2015) to the image, the planet is recovered at extremely high SNR. The planet’s position in the Stokes I image was estimated using the **StarFinder** IDL package (Diolaiti et al. 2000), which requires the user to input a PSF model for precision astrometry. In GPI’s polarimetry mode the entire bandpass is seen by each frame and therefore the satellite spots are elongated and cannot be used as a PSF reference, as they are in spectroscopy mode. Instead, we used a GPI PSF generated with AO simulation software (Poyneer & Macintosh 2006).

To estimate astrometric errors we used **StarFinder** to measure β Pic b ’s location in the total intensity image from each of the 49 polarization data cubes. The RMS difference between the planet location in the individual cubes and the Stokes I image was taken to be the error in the planet location. The error on the location of the star is estimated from the RMS scatter of the measured star’s position across the set of cubes. This error tracks the motion of the star behind the coronagraph between frames, which we expect to be larger than the errors on the star’s position determined by the Radon transform, and therefore likely overestimates the errors.

The position of β Pic b in the polarimetry mode observations can be found in Table 3.1. As with the spectroscopy mode data, the errors represent a combination of the errors on the star’s and planet’s positions, GPI’s pixel scale and GPI’s PA offset on the sky.

3.5.3 Orbit fitting

Using the ten newly obtained astrometric points (nine from spec. mode and one from pol. mode), combined with the datasets presented by Chauvin et al. (2012) and Nielsen et al. (2014), we fit for the six Keplerian orbital elements of β Pic b plus the total mass of the system using the parallel-tempered sampler from **emcee** (Foreman-Mackey et al. 2014). While astrometric datapoints have been published in other papers, in an effort to minimize systematics between datasets, we limited ourselves to only these two large datasets where considerable effort has been made to calibrate astrometric errors. The fitting code was previously used in Kalas et al. (2013), Macintosh et al. (2014), and Pueyo et al. (2015). We also fit the radial velocity measurement of the planet from Snellen et al. (2014), which allows us to constrain the line-of-sight orbital direction and break the degeneracy between the locations of the ascending and descending node.

The model fits seven parameters: the semi-major axis, a ; the epoch of periastron, τ ; the argument of periastron, ω ; the position angle of the ascending node, Ω ; the inclination, i ; the eccentricity, e ; and the total mass of the system, M_T . Our orbital frame of reference followed the binary star sign convention used in Green (1985). Under this convention the ascending node is defined as the location in the orbit where the planet crosses the plane of the sky (centered on the star), moving southward in projection. Note that this is 180° different from the convention used in Chauvin et al. (2012). The projected position angle of the ascending node on the sky is defined from North, increasing to the East. The argument of the periastron is defined as the angle between the ascending node and the location of the periastron in the orbit, with ω increasing from the ascending node. The epoch of periastron, τ , is defined in units of orbital period, from October 10, 1995 (Julian date 2450000.5). A summary of the orbital parameters and their prior distributions can be found in Table 3.3.

The MCMC sampler was run for 10,000 steps with 10 temperatures and 256 walkers after a “burn-in” of 2000 steps. After the run, the maximum auto-correlation across all parameters was found to be 25 steps, indicating that the chains should have reached equilibrium. The posterior distributions (Figure 3.6) were estimated using the zero temperature walkers, using only one of out every 25 steps.

Table 3.3. Orbit Fit Parameters and Prior Distributions for β Pic b

Parameter	Symbol	Range	Prior Distribution
Semi-major axis	a	4 – 40 AU	Uniform in $\text{Log}(a)$
Epoch of Periastron	τ	–1.0 – 1.0	Uniform in τ
Argument of Periastron	ω	$-2\pi - 2\pi$ rad	Uniform in ω
Position Angle of the Ascending Node	Ω	$25^\circ - 85^\circ$	Uniform in Ω
Inclination	i	$81^\circ - 99^\circ$	Uniform in $\cos i$
Eccentricity	e	0.00001 – 0.99	Uniform in e
Total Mass	M_T	0 – 3 M_\odot	Uniform in M_\odot

In Figure 3.6, the epoch of periastron was wrapped around to be only positive between 0 and 1 and the arguments of the periastron was wrapped around to range from 0 to 360° . A random selection of 500 accepted orbits are plotted on top of the astrometric and radial velocity datapoints in Figures 3.7 and 3.8, respectively. While the orbital fits are generally consistent with most of the astrometric datapoints, the majority of the orbital solutions fall more than $1\text{-}\sigma$ from the measured radial velocity.

We find that the planet has a semi-major axis of $9.25_{-0.45}^{+1.46}$ AU, an inclination of $\sim 89^\circ 01 \pm 0^\circ 30$ and an ascending node at a position angle of $31^\circ 75 \pm 0^\circ 15$. We take the median of the marginalized posterior distribution to be the best estimator of each parameter’s value, and the 68% confidence values as the errors. Following this convention, the eccentricity of the orbit is found to be $e = 0.07_{-0.05}^{+0.11}$. However, the eccentricity is a positive definite quantity and typical estimators (e.g., the mean and median) will overestimate the true eccentricity when it is small ($e < 0.1$). When considering eccentricities of radial velocity planets, Zakamska et al. (2011) consider several different estimators and find that for small eccentricities the mode of the distribution is the least biased. The mode of our distribution falls in the smallest eccentricity bin indicating an eccentricity very close to zero. Therefore, it is perhaps more appropriate to quote the upper limit on the eccentricity, $e < 0.26$ (95% confidence).

For orbits with higher eccentricities ($e > 0.1$), the epoch and argument of periastron have strong peaks at ~ 0.5 periods and $\sim 170^\circ$, respectively. At lower eccentricities these two parameters remain degenerate, with a large range of acceptable values. Overall, the marginalized distributions reveal that these parameters are still relatively unconstrained. The ensemble of accepted orbits at the end of the run have a reduced χ^2 of $1.55_{-0.05}^{+0.09}$.

Marginalized across all parameters, the total mass of the system is $1.61 \pm 0.5 M_\odot$. At $\sim 11 M_{Jup}$, β Pic b contributes less than 1% to the total mass, giving β Pic itself a mass of $M_{\beta Pic} = 1.60 \pm 0.5 M_\odot$. This falls slightly below the range estimated by Crifo et al. (1997) ($M \sim 1.7 - 1.8 M_\odot$) and just within the range of Blondel & Djie (2006) ($1.65 M_\odot < M < 1.87 M_\odot$), who both use evolutionary models and the HR diagram to date β Pic. This estimate provides a slightly smaller value than that presented in Nielsen et al. (2014), though it is still consistent within their errors ($1.76_{-0.17}^{+0.18} M_\odot$)

By combining the semi-major axis and stellar mass values of each walker at each accepted iteration, we are able to create a posterior distribution for the orbital period, from which we derive that $P = 22.4_{-1.5}^{+5.3}$ yr. The large upper limit is due to the extended tail in the semi-major axis distribution.

3.5.4 Planet Polarization

Giant exoplanets may have polarized emission in the NIR either due to rotationally induced oblateness (Marley & Sengupta 2011) or asymmetries in cloud cover (de Kok et al. 2011). For β Pic b , the recently measured rotational period of ~ 8 hours would induce a polarization signature due to oblateness of less

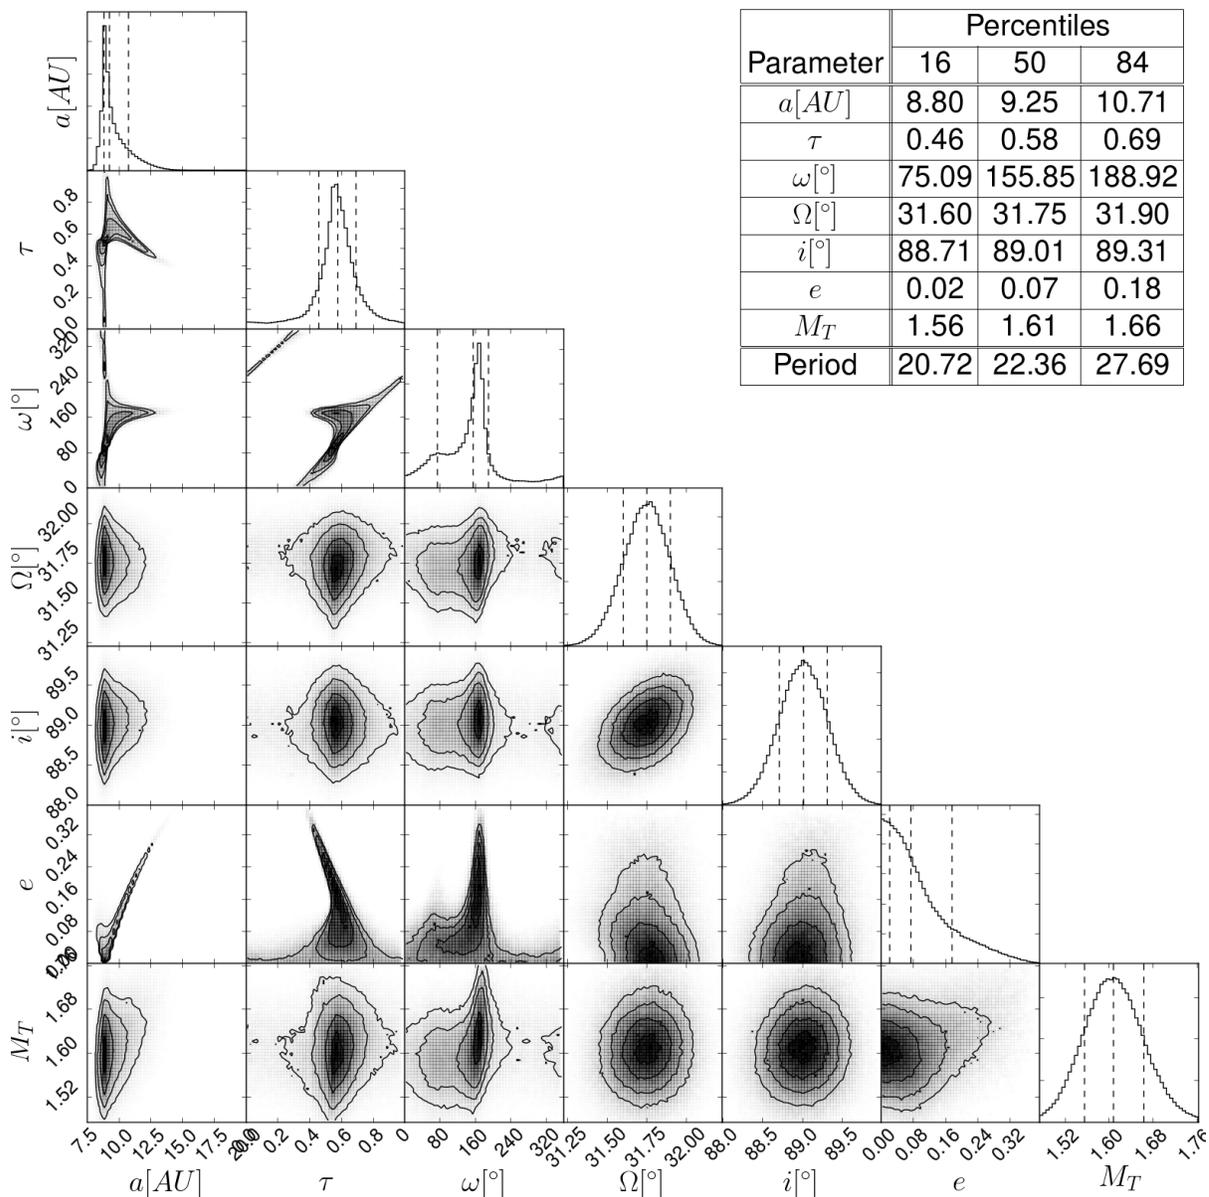


Figure 3.6: The posterior distributions of the model parameters from the MCMC orbit model fitting to the astrometry data points of β Pic b . The diagonal histograms show the posterior distributions of each parameter marginalized across all the other parameters. In each plot, the red dashed lines indicate the 16%, 50% and 84% percentiles. The off-diagonal plots display the joint probability distributions with contour levels at the same percentiles. *Inset table:* The 16%, 50% and 84% percentiles from the marginalized distributions.

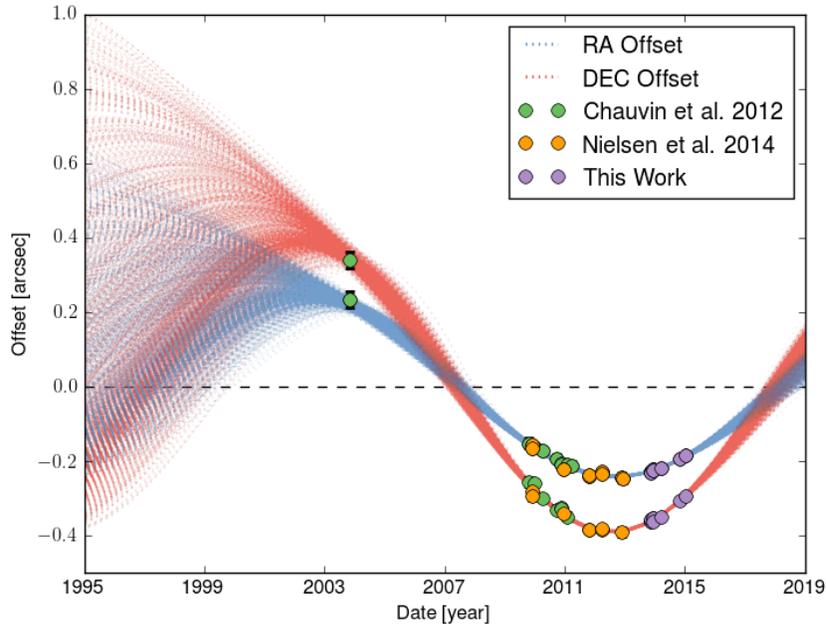


Figure 3.7: The RA (blue) and DEC (red) offsets of β Pic b from β Pic for a random selection of 100 accepted orbits (dotted lines) from the MCMC run. The 29 data points used in the fit are overplotted, with the colors indicating their source. Error bars on the datapoints are smaller than the markers, except for the 2003 measurement from Chauvin et al. (2012).

than 0.1% (below GPI's current sensitivity limit, Wiktorowicz et al. 2014b). Therefore, any detected polarization signal would be indicative of cloudy structure.

To estimate β Pic b 's polarization, we first created a disk-free linear polarized intensity image by combining the model-subtracted Q_r image with the U_r image ($P = \sqrt{Q_r^2 + U_r^2}$). The total polarized flux at the location of β Pic b , within an aperture of radius $1.22\lambda/D$, was then compared to the flux of 38 independent apertures at the same angular separation. We find that β Pic b 's polarized flux is 0.5σ from the mean flux of the independent apertures, consistent with zero linear polarization signal from the planet (see Figure 3.3). While this measurement does not provide any evidence for cloud structure, it does not exclude the possibility either; the magnitude of cloud-induced polarization depends on many factors, including the atmospheric temperature and pressure profile, the composition, the nature of the inhomogeneities, rotation, and viewing angle. The PSF variability during the observations makes accurate recovery of the total intensity of the planet difficult, and thus we leave the characterization of an upper limit on the planet's polarization fraction for future work.

3.6 Discussion

3.6.1 The Debris Disk

With GPI we probe the projected disk between $0.3''$ and $1.5''$ at high spatial resolution. The work presented here has two advantages over previous attempts to model the disk at similar angular separations. First, the polarized intensity images provide a unique view of the disk, in particular the vertical extent is free of any biases associated with ADI PSF subtraction. Second, the MCMC fitting allows us to fully

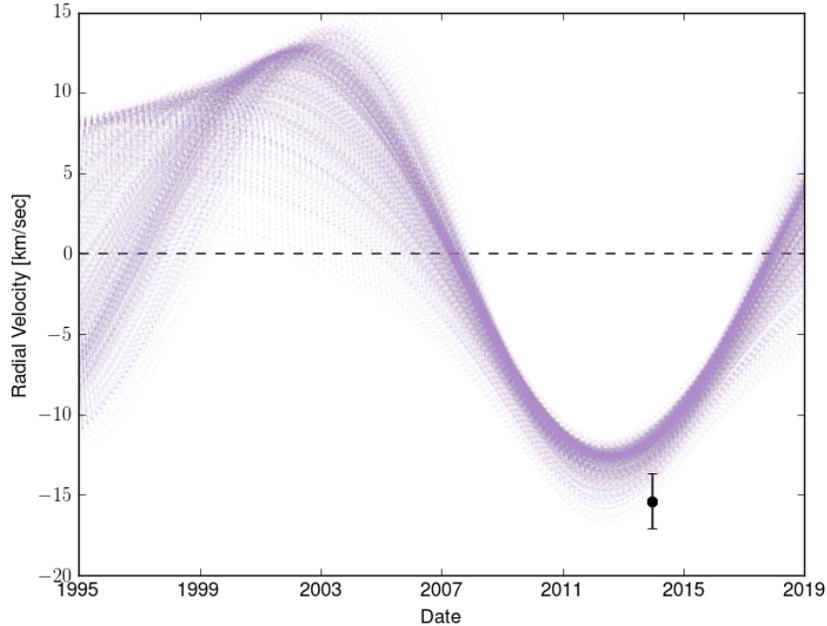


Figure 3.8: The radial velocity of β Pic b from a random selection of 100 accepted orbits (purple dotted lines) from the MCMC run and the single measurement from Snellen et al. (2014).

explore the multi-dimensional parameter-space and place quantitative confidence intervals on the model parameters.

MCMC fitting requires evaluation of the likelihood function for each set of parameters that is examined. The cost of fitting depends on the computational expense of evaluating the model and the dimensionality of the model parameter space. For that reason we have limited our exploration to optically thin scattering, an analytic recipe for the phase function, and a simple model of the dust distribution. We do not consider multi-component disks (as modeled for the outer disk, e.g., Ahmic et al. 2009) and we assume that the disk aspect ratio is constant. Regardless of these simplifications, we find that this model provides an excellent fit to our polarized image.

The Henyey-Greenstein scattering function is often used to model the total intensity scattering efficiency of dust grains, but has not been used extensively for polarized intensity. This is at least partially due to the fact that in most circumstances where polarized intensity is measured, total intensity is obtained as well, allowing for more sophisticated modeling of the dust scattering. In addition, the scattering efficiency of polarized intensity of small spherical particles approaches zero at very small scattering angles, a feature that is not captured by the HG function. While the exact shape of the HG function cannot fully reproduce the polarized scattering efficiency function for physical models, a quick informal survey of possible grain models indicates that our best fit $g = \langle \cos \theta \rangle \approx 0.7$ can be reproduced by Mie scattering particles with a radius of $\sim 1 \mu\text{m}$ and an index of refraction of $m = 1.033 - 0.01i$, similar to the porous, icy grains inferred by Graham et al. (2007) for AU Mic. However, as previously mentioned, a true characterization of the physical grain scattering properties will require either an unbiased total intensity image, or polarized intensity images at other wavelengths. We leave the characterization of the dust properties of the inner disk for future work.

The observations of Milli et al. (2014) have a field of view ($0.4'' - 3.8''$) that overlaps with our disk

the forward scattering nature of the dust grains means that the inner edge itself contributes minimally to the observed surface brightness at its projected separation. This serves to emphasize the critical role of dust scattering when interpreting the surface brightness as a function of radius; a smooth surface density by itself does not necessarily exclude features in the radial dust profile.

Note that our model has been defined with a sharp cut-off inside the inner radius, and caution should be used when interpreting the value. There may be dust inside of the inner radius with a lower surface density. For example, the true dust density inside the inner radius may have a slowly decreasing inner power-law, such as those considered in Milli et al. (2014).

Imaging and spectroscopic studies in the mid-IR have probed similar regions of the debris disk at wavelengths where contrast between the stellar flux and the dust (thermal) emission is more favorable than in the optical and NIR. Okamoto et al. (2004) found spectroscopic evidence for dust belts at 6, 16, and 30 AU. Wahhaj et al. (2003) fit a series of four dust belts to deconvolved 18 μm images and found their best fit radii to be 14, 28, 52, and 82 AU. With the exception of the belts close to ~ 15 AU, all of these belts are either well outside or at the very edge of our field of view. The 6 AU belt seen by Okamoto et al. (2004) is below our inner working angle. We note that the Okamoto et al. (2004) belt at ~ 16 AU only occurs on the NE side, at roughly the same location as the tentative brightness asymmetry seen in our disk model residuals. We see no evidence of the other belts in projection, but we model the disk with a continuous dust distribution and therefore may not be sensitive to dust at their locations. Mid-IR imaging by Weinberger et al. (2003) indicate emission within 20 AU that is significantly offset in position angle from the main outer disk. In our disk image we see no indication of this offset.

Previous studies of β Pic's debris disk in polarized scattered light have been carried out both in the optical (Gledhill et al. 1991; Wolstencroft et al. 1995) and the NIR (Tamura et al. 2006). These observations image the disk at separations of $15'' - 30''$ and $2.6'' - 6.4''$, in the optical and NIR, respectively. At these angular separations the total intensity observations are not limited by the PSF halo and when combined with the polarized images, polarization fraction can be used to model the dust grains (Voshchinnikov & Krügel 1999; Krivova et al. 2000). Tamura et al. (2006) combine the optical measurements with their K -band data and find that the observations could be explained by scattering from fluffy aggregates made up of sub-micron dust grains. Unfortunately, the lack of total intensity images and a non-overlapping field of view make a direct comparison between our observations and this past work difficult.

3.6.2 β Pic b

In general, our orbit fit is consistent with those previously published (e.g. Chauvin et al. 2012; Macintosh et al. 2014; Nielsen et al. 2014), but the longer temporal baseline and increased astrometric precision significantly tighten the constraints on the orbital parameters. In particular, we find that the position angle of the ascending node of the planet lies in between the main outer disk and warp feature, consistent with Nielsen et al. (2014).

At first glance, the errors on our orbital elements appear comparable to those in Macintosh et al. (2014). However, our fit includes the total mass of the system as an additional free parameter. Nielsen et al. (2014) modeled the system's total mass as a free parameter in their orbital fit and found that accounting for the uncertainty in the system's total mass resulted in larger uncertainties in the planet's orbital elements. In particular, they find that with a floating system mass the eccentricity distribution has a long tail that peaks at high eccentricities. Due to a degeneracy between semi-major axis and

eccentricity, this stretched the semi-major axis distribution to higher values as well. In Figure 3.6, we find that the eccentricity is now significantly better constrained ($e < 0.26$), and while the degeneracy remains, the semi-major axis is constrained to be < 10.7 AU with 84% confidence.

For each orbit defining our posterior distribution, we calculate the epoch of closest approach and find that it will fall between 2017 November 20, and 2018 April 4 with 68% confidence. With our derived inclination of $i = 89.01 \pm 0.36^\circ$, the updated transit probability is $\sim 0.06\%$, assuming that the planet will transit if the inclination is within 0.05° from 90° . This is a reduction by a factor of ~ 50 from the estimate in Macintosh et al. (2014), who found $i = 90.7 \pm 0.7$.

Even though the likelihood of a planet transit is small, it is still possible that dust particles orbiting within the planet's Hill sphere ($R_{Hill} \approx 1$ AU) will transit. Indeed the transit of a ring system surrounding an exoplanet was recently detected around J1407 (Mamajek et al. 2012). In the outer solar system, satellites around the giant planets have stable orbits within a Hill sphere about the planet out to $\sim 0.5 R_{Hill}$ when in prograde orbits and $\sim 0.7 R_{Hill}$ in retrograde orbits (Shen & Tremaine 2008). For β Pic b , assuming a planetary mass of $11 M_J$, a semimajor axis of 9.25 AU, a circular orbit and a stellar mass of $1.61 M_\odot$, we calculate a Hill radius of ~ 1.2 AU. Thus, stable orbits within the Hill sphere will transit if the planet's inclination is within 3.8° and 5.3° of edge-on, for prograde and retrograde orbits, respectively. Our new constraints on the inclination indicate that these orbits will almost certainly transit. However, the true transit probability will depend not only on the exact location of the dust, but also its orientation relative to the observer. For example, dust that fills the stable orbits and is orbiting face-on relative to the observer will transit, but if it is orbiting edge-on it will not.

The presence of infalling comets (a.k.a. falling evaporating bodies, or FEBs) has been previously inferred by redshifted absorption features in β Pic's spectrum (Beust & Morbidelli 1996). Thébaud & Beust (2001) suggested that a massive ($M \geq M_{Jup}$) planet within ~ 20 AU on a slightly eccentric orbit ($e \gtrsim 0.05 - 1$), could be responsible for imparting highly elliptical orbits on bodies within a 3:1 or 4:1 resonance, that then plunge towards the star. In this scenario the argument of the periastron of the planet is restricted to a value of $\omega' = -70 \pm 20^\circ$ from the line of sight. Using our definitions, the equivalent requirement is $\omega = 200 \pm 20^\circ$. Our results neither confirm nor rule out the infalling comet scenario. While the marginalized distribution of the argument of periastron allows for a broad range of values, if the orbit is indeed eccentric, then ω peaks strongly around 170° , just outside of the acceptable values for this scenario. Thébaud & Beust (2001) find that if the eccentricity of the massive perturber (here assumed to be β Pic b) is as large as $e \approx 0.1$ then the infalling comets most likely originate in the 3:1 resonance, which occurs between 18 AU and 22 AU based on our 68% confidence range for β Pic b . The inner edge of the dust in our scattered light images falls at $R_1 = 23.61^{+0.86}_{-0.57}$ AU, outside of range of values for the 3:1 resonance. However, as noted above, our inner radius is sharply defined, and there may still be material inside. For smaller perturber eccentricities the infalling comets originate in the 4:1 resonance, which occurs between 22.2 AU and 26.5 AU. Our disk model does not constrain whether there is an excess of bodies librating in the 4:1 resonance.

3.6.3 The Disk-Planet Interaction

A planet on an inclined orbit is thought to be responsible for the warp feature in the region of the disk outside our field of view at ~ 80 AU (Mouillet et al. 1997; Augereau et al. 2001). The directly-imaged planet β Pic b (Lagrange et al. 2009) has a mass, semi-major axis, and inclination consistent with producing the warp (e.g., Dawson et al. 2011). The updated position angle of β Pic b 's ascending node,

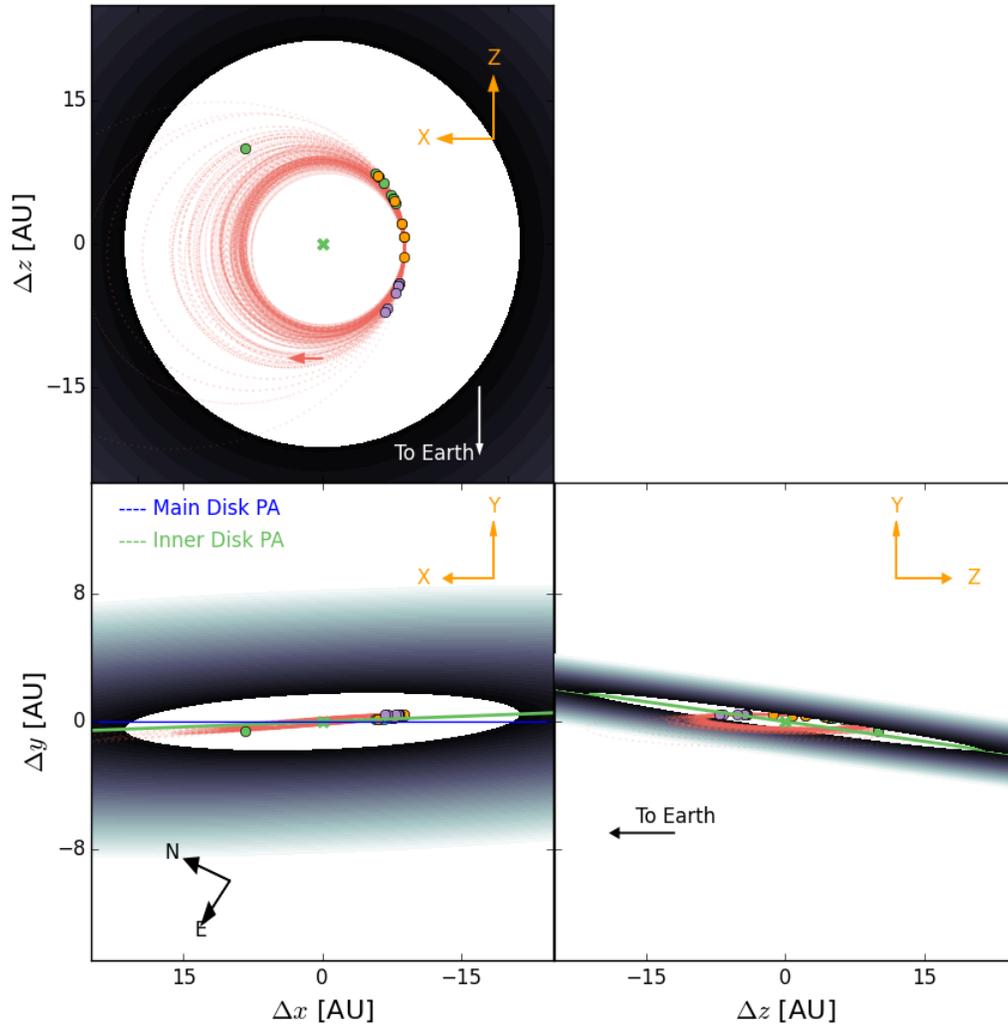


Figure 3.10: Three orthogonal projections of the system: the plane of the sky (*bottom left*), a top down view (*top left*), and a side view (*bottom right*). The system has been rotated in three dimensions so that the midplane of outer disk is horizontal (blue line) in the *bottom left* plot. Each image includes the best fit disk midplane (greyscale decreasing as r^{-1}), a random selection of 100 accepted orbits (dotted red lines) and the location of β Pic b according to a likely orbit at the same epochs as the measurements of Chauvin et al. (2012), Nielsen et al. (2014) and those included in this work (purple markers). *Bottom left:* The positions of the planet emulate the direct imaging astrometry points. The green line indicates the position angle of our inner disk model. *Top left:* The orbital inclination of the planet and the mismatch between the planet and main disk’s position angles on the sky result in the top panel being slightly offset from the planet’s orbital plane. The RV measurement of the planet (Snellen et al. 2014) breaks the degeneracy in the orbital direction and allows us to calculate the line of sight (Z) coordinate for each epoch. The red arrow indicates the direction of motion in the orbit. *Bottom right:* The green line displays the inclination of the inner disk relative to the observer’s point of view. In all panels, the green x indicates the location of the star.

$\Omega = 31.75 \pm 0.15^\circ$, is offset by 2.65° with respect to the flat outer disk ($29.1 \pm 0.1^\circ$; Apai et al. 2015), consistent with producing a warp tilted by 5° counter-clockwise with respect to the flat outer disk. (As illustrated by Apai et al. 2015, Fig. 21, our azimuthal viewing angle of the warped disk affects the degree to which the inner disk and the planet's orbit appear aligned with the flat outer disk in projection and also affects the projected height of the warp.) Although the planet's updated orbit remains consistent with sculpting the *outer disk*, several features of the *inner* disk that we measured here are unexpected solely from sculpting by β Pic *b* (Figure 3.10).

First, the inner edge of the disk is at 23.6 AU, about twelve Hill radii from the planet. We performed a simulation using `mercury6` (Chambers 1999) of a planet with orbital parameters set to the median values in Figure 3 embedded in a disk of test particles initially spanning 10 to 40 AU. On the timescale of hundreds to thousands of orbits, the planet clears out the disk to ~ 15 AU, with the inner edge persisting at that location over the 20 Myr stellar lifetime (Binks & Jeffries 2014; Mamajek & Bell 2014). An inner edge at ~ 15 AU is in agreement with simulations by Rodigas et al. (2014); cf. their Table 2.

We have not explored disk models with gradual inner edges (e.g., Milli et al. 2014), so there may be material between 15 and 23.6 AU with a lower surface density, or planet bodies that are less collisionally active. If the disk inside 23.6 AU is truly cleared out, an undetected low-mass planet in between β Pic *b* and the disk's inner edge could be responsible; we find that a planet could exist on a stable orbit in that region.

Second, we expect the inner disk to be centered on β Pic *b*'s orbital plane. Given a warp located at ~ 80 AU, the width of a secular cycle (i.e., the difference in semi-major axis for which the planetesimals are 2π out of phase in their oscillations about the planet's orbital plane) is only about 1 AU at a radius of 40 AU and the timescale of a secular cycle is about forty times shorter than at the location of the warp. Therefore, close to the planet, a sufficient number of secular cycles should have passed that the parent bodies' free inclination vectors are randomized about the forced inclination from the planet. Under certain conditions, we found that our simulation could produce a parent bodies sky plane inclinations distribution with peaks at $\sim i_{\text{planet,sky}} \pm i_{\text{planet,outerdisk}}$ (one of which could correspond to $\sim 85^\circ$), where $i_{\text{planet,sky}}$ is the line of sight inclination of the planet and $i_{\text{planet,outerdisk}}$ is the mutual inclination between the planet and the outer disk. However, we expect that even in these circumstances the measured disk midplane would be aligned with the planet's orbital plane. Moreover, damping by collisions, small bodies, or residual gas—provided that it occurs on a timescale shorter than half a secular timescale—reduces the free inclination, decreasing the disk scale height but keeping it centered about the planet's orbital plane.

Instead, the average plane of the inner disk appears mutually inclined with respect to the planet's orbit. If the polarized intensity images were dominated by scattered light from the outer disk, a mutual inclination with respect to the planet could be consistent, (depending on the semi-major axis of the dominant dust; see Figure 1 from Dawson et al. 2011), but in the current disk model the observed light is dominated by a close-in disk. Contribution from another planet to the forced plane of the disk is a possibility but the available parameter space for an additional planet that tilts the disk toward us, yet is too low mass to escape detection, is quite limited. In the future, we plan to explore a wider range of dust-scattering models to ensure that this result (a disk mutually inclined to the planet's orbit at ~ 25 AU) is not dependent on the assumed dust properties.

Finally, the scale height of the disk appears larger than expected from stirring by β Pic *b* or self-stirring of the parent bodies. In the absence of damping, the total thickness of the disk would be $\sim 2i_p$,

corresponding to a scale height aspect ratio of about 0.06 for a planet inclined by 3.6° with respect to the primordial plane. Self-stirring to the escape velocity of 10 km planetesimals would contribute only about 0.001 to the aspect ratio; self-stirring to the escape velocity of Plutos would be required. In practice, we do not expect most parent bodies participating in the collisional cascade to be stirred to random velocities of the largest bodies (e.g., Pan & Schlichting 2012); their steady-state random velocities depend on the balance between stirring, damping by smaller bodies and each other, and radiation forces. The scale height is also significantly larger than observed further out in the disk (Ahmic et al. 2009)—even at 50 AU (Milli et al. 2014). The robustness of the scale height to the dust scattering model should be explored further; for example, a significant contribution from polarized back scattering (not modeled here) could result in a smaller inferred scale height. If the current inferred large scale height in the very inner disk is robust, a sub-detection planet located between β Pic b and the inner edge of the disk and mutually inclined with respect to β Pic b is a possible explanation.

Nesvold & Kuchner (2015) recently simulated the dynamical and collisional behavior of β Pic's planetesimals and dust grains using SMACK (Nesvold et al. 2013), which models planetesimals across a range of sizes using super particles. They find that collisional damping is not important in shaping the morphology of the disk. Their detailed model also does not predict the surprising observational features discovered here: they find the planet clears a gap only out to 14.5 AU and that the disk is centered about the planet's orbital plane (see their Figure 3). They find that some planetesimals in the inner disk are scattered by each other or the planet to inclinations larger than $2i_p$, increasing the thickness of the inner disk by about 50%, not enough to account for the ($\sim 200\%$ larger) observed scale height.

3.7 Conclusion

We have presented new images of the β Pic debris disk in polarized light that reach angular separations previously inaccessible to both space and ground-based telescopes. The use of PDI as a means of PSF subtraction circumvents the need for ADI PSF subtraction which can cause self-subtraction, especially in vertically extended disks like that of β Pic at the angular separations explored by GPI. The disk image was modeled with a radial power-law dust distribution combined with a Henyey-Greenstein scattering function. The disk model indicates an inclined disk at a position angle on the sky between the main outer disk and the warped feature with an inner edge at ~ 23 AU.

The conclusions about the geometry of the disk are based on the assumption that a Henyey-Greenstein scattering phase function can accurately represent the true scattering properties of the constituent dust grains. Future imaging studies, such as multi-color polarimetry at similar angular separations, will allow for the use of more sophisticated dust grain models that will be able to further examine the inner part of the disk and to test our results.

In addition, we presented ten new astrometric measurements of the planet β Pic b , which we combine with previous measurements to fit an orbital solution. The solution improves upon those previously published by tightening the constraints on the Keplerian orbital elements, particularly the inclination and position angle of the ascending node. We leave the total mass of the system as a free parameter, allowing us to constrain the stellar mass of β Pic to within 5%.

When considered together, the disk model and the orbital fit indicate that the dynamics of the inner edge of the disk are not consistent with sculpting by the planet β Pic b alone. This could be explained by an as-of-yet undetected planet in-between the known planet and the inner edge of the disk. Under this

scenario the less massive, further out planet would dynamically influence the inner regions of disk, while the more massive β Pic b would have a greater affect at larger radii, causing the well know warp. If there is in fact another planet at this location, this will have significant consequences for our understanding of the planet formation history and dynamical evolution of this system.

Chapter 4

Imaging an 80 AU Radius Dust Ring Around the F5V Star HD 157587

A version of this work was submitted to the Astrophysical Journal on June 10, 2016.

4.1 Chapter Summary

We present *H*-band near-infrared polarimetric imaging observations of the F5V star HD 157587 obtained with the Gemini Planet Imager (GPI) that reveal the debris disk as a bright ring structure at a separation of $\sim 80\text{--}100$ AU. The new GPI data complement recent HST/STIS observations that show the disk extending out to over 500 AU. The GPI image displays a strong asymmetry along the projected minor axis as well as a fainter asymmetry along the projected major axis. We associate the minor and major axis asymmetries with polarized forward scattering and a possible stellocentric offset, respectively. To constrain the disk geometry we fit two separate disk models to the polarized image, each using a different scattering phase function. Both models favor a disk inclination of $\sim 70^\circ$ and a 1.5 ± 0.6 AU stellar offset in the plane of the sky along the projected major axis of the disk. We find that the stellar offset in the disk plane, perpendicular to the projected major axis is degenerate with the form of the scattering phase function and remains poorly constrained. The disk is not recovered in total intensity due in part to strong adaptive optics residuals, but we recover three point sources. Considering the system's proximity to the galactic plane and the point sources' positions relative to the disk, we consider it likely that they are background objects and unrelated to the disk's offset from the star.

4.2 Introduction

Circumstellar debris disks, composed of planetesimals and dust, are remnants of the planet formation process. Therefore, their study can provide insights into the planet formation and evolution history of the systems in which they reside. The dust grain composition of a disk traces grain growth and erosion, and, if spatially resolved, disk morphology can provide evidence of dynamical interactions with nearby planets. Such an interaction can manifest as a warp (e.g. Beta Pic; Burrows et al. 1995; Mouillet et al. 1997), a stellocentric offset (e.g. HR 4796A; Wyatt et al. 1999; Telesco et al. 2000) or a sharp radial profile at the inner edge of a dust ring (e.g. Fomalhaut; Kalas et al. 2005; Quillen 2006).

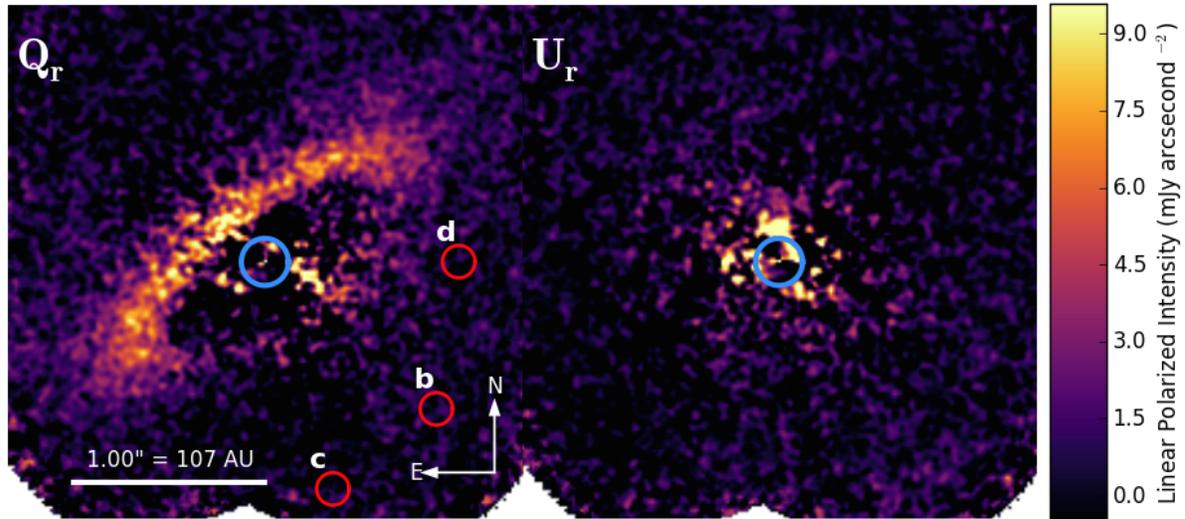


Figure 4.1: *Left*: GPI H -band radial polarized intensity image (Q_r) of the HD 157587 debris disk. The image has been smoothed with a Gaussian kernel ($\sigma = 1$ pixel). The blue circle indicates the size of the central focal plane mask ($0''.12$ radius). The red circles and associated letters designate point sources seen in total intensity in Figure 4.4 (§4.4.2). *Right*: The U_r image of HD 157587. For single scattering from circumstellar material we expect no contributions to the U_r image. Thus, this image can be used as a noise map for the Q_r image. The image appears to be largely free of correlated structure, except at small inner working angles. Both images have been cropped from the full GPI field of view to display only the inner $2''.6 \times 2''.6$ region. No polarized emission was seen outside the cropped region.

Debris disks are imaged via their thermal emission in infrared or millimeter wavebands, which typically traces the location of millimeter sized bodies, or via scattered light in the visible and near-infrared (NIR), which is more sensitive to micron-sized dust. Observations of debris disks in scattered light are typically able to resolve finer spatial scales than longer wavelength observations (though ALMA’s spatial resolution is now competitive), but are challenging due to the extreme contrast ratios between the faint dust-scattered light and the bright host stars. Instrumental point-spread functions (PSFs) extend the stellar emission out to angular separations where debris disks are found, obscuring the scattered light from the dust. For ground-based observations this problem is compounded by the atmosphere, which scatters light from the PSF out to farther separations.

The Gemini Planet Imager (GPI; Macintosh et al. 2014) is an instrument on the Gemini South 8-m telescope that has been designed specifically to mitigate these challenges. It employs a high-order adaptive optics (AO) system, combined with an apodized-pupil Lyot coronagraph and an integral-field spectrograph, to image exoplanets and debris disks at angular separations down to $\sim 0''.1$. The GPI Exoplanet Survey (GPIES) is a long-term Gemini South program targeting 600 nearby stars with the goal of discovering and characterizing young Jovian exoplanets. A secondary goal of the survey is to image and characterize debris disks. Stars with previously resolved debris disks and survey stars that exhibit infrared excesses are observed using GPI’s polarimetry mode. The polarimetry mode is implemented as a rotatable half-wave plate (HWP) modulator and a Wollaston prism analyser. This mode has been designed to take advantage of the inherent polarization of light scattered off circumstellar dust grains, to further suppress the unpolarized starlight and reveal the disk beneath (Perrin et al. 2015).

Here we present GPIES observations of the debris disk around HD 157587, an F5V star with an infrared excess $L_{IR}/L_{star} = 7.9 \times 10^{-4}$, (McDonald et al. 2012) at a distance of 107.4 pc (van Leeuwen

2007). HST/STIS coronagraphic imaging (GO-12998; PI Padgett) first revealed the dust scattered light extending to $>7''$ radius, with a morphology resembling a fan (such as for HD 15745; Kalas et al. 2007), where the straight edge of the fan lies along the southwestern side of the nebulosity (Padgett & Stapelfeldt 2015). The inner working angle of these data corresponds to a projected separation of ~ 100 AU. Our new scattered light images, obtained as part of the GPIES campaign, detect the structure of the circumstellar dust in the projected 30 - 130 AU radial region.

4.3 Observations and Data Reduction

We observed HD 157587 with GPI’s polarimetry mode in the H -band on 2015 August 28 UT. The observations consisted of twenty-eight 90 s frames, with the HWP position angle cycling between $0^\circ 0$, $22^\circ 5$, $45^\circ 0$ and $67^\circ 5$. Throughout the sequence the field rotated by a total of 46° . The average airmass was 1.02 and the seeing as measured by the Gemini Differential Image Motion Monitor and Multi-Aperture Scintillation Sensor was $0.61''$ and $0.63''$, respectively. The AO system telemetry reported a post-correction wavefront rms error of 216 ± 20 nm across the sequence.

The data were reduced using the GPI data reduction pipeline version 1.3 (Maire et al. 2012b; Perrin et al. 2014c). The raw data were dark subtracted, cleaned of correlated detector noise, bad pixel corrected, flexure corrected and then combined into a polarization datacube (where the third dimension holds two orthogonal polarization states). Each datacube was divided by a polarized flat field and corrected for non-common path errors via a double differencing algorithm (Perrin et al. 2015). The instrumental polarization was assumed to be equivalent to the apparent stellar polarization in each polarization datacube, measured as the mean normalized difference of pixels with separations between 7 and 13 pixels from the star’s location. The star’s position behind the focal plane mask was determined from the satellite spots using a radon-transform-based algorithm; Wang et al. 2014. The estimated instrumental polarization was then subtracted from each pixel, scaled by the pixel’s total intensity (Millar-Blanchaer et al. 2015). The region selected to measure the instrumental polarization was just outside of the coronagraph edge where the residual PSF flux, and hence the flux from instrumental polarization, is maximized. We assume that this area is devoid of any significant polarized structure and that any measured difference between the two polarization states is due to the instrumental polarization.

The datacubes were corrected for geometric distortion, smoothed with a Gaussian kernel ($\sigma = 1$ pixel) and then combined into a Stokes datacube by solving a set of equations that describe the linear polarization states measured in each of the individual exposures given the waveplate and sky rotation angles (Perrin et al. 2015). The Stokes datacube was subsequently converted to the radial Stokes convention ($[I, Q, U, V] \rightarrow [I, Q_r, U_r, V]$; Schmid et al. 2006). The sign convention is such that a positive Q_r corresponds to a polarized intensity whose vectors are oriented perpendicular to a line connecting a given pixel to the central star and negative values are parallel to the line. Under this convention (and the assumption of low optical depth) the U_r image should contain no disk flux and will only contain noise. Thus the Q_r image should contain all of the disk polarized intensity as positive values. Finally, the flux of the four satellite spots were measured and flux calibration was carried out as described in Hung et al. (2015). The final Q_r and U_r images can be seen in Figure 4.1.

The polarization datacubes were also processed separately using the pyKLIP (Wang et al. 2015) implementation of the Karhunen-Loève Image Projection (KLIP) algorithm (Soummer et al. 2012) to attempt to recover the disk in total intensity and search for point sources (§4.4.2).

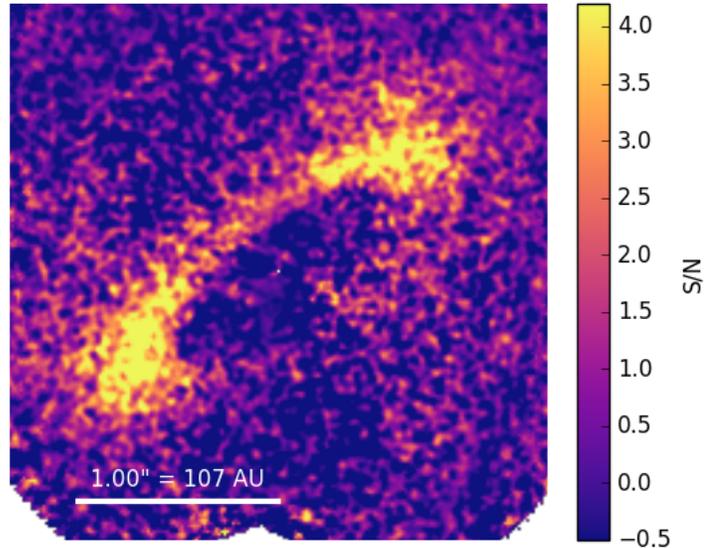


Figure 4.2: Signal to noise map of the GPI H -band Q_r image. The ansae are clearly detected here with a S/N greater than 5, while the region near the minor axis is closer to a S/N of 3. In this map the projected minor axis (in the NE direction) appears at a low S/N even though it is the brightest region in the Q_r image. This is due to the elevated noise in the inner regions of the U_r image that overlap with the edge of the disk.

4.4 Results

4.4.1 Polarized Intensity Image

The Q_r image displays an inclined, ring-like structure with a strong brightness asymmetry in the NE-SW direction (the projected minor axis). The inner edge of the ring has projected semi-major and semi-minor axes of $\sim 0''.65$ and $\sim 0''.2$, respectively. The region interior to the inner edge of the ring appears to be cleared of any scattering material. However, the residual systematics at smaller separations in the Q_r and U_r images do not exclude additional dust at these smaller radii. Outside of the ring the surface brightness decreases quickly and reaches the noise floor within our field of view (FOV), which extends to a radius of $1''.7$ along the semi-major axis of the the ring.

A comparison between the Q_r and U_r images indicates that the ring detection is robust and that the morphology is not due to instrumental effects. The U_r image appears to be dominated by uncorrected systematics interior to $\sim 0''.275$, a region which intersects the ring near its minor axis. Outside of $\sim 0''.275$, there appear to be no coherent structures in the U_r image.

The strong NE-SW asymmetry seen in the polarized images is reminiscent of the asymmetries seen in other disks recently imaged in polarized light by GPI, for example: HR 4796A, Perrin et al. (2015); HD 106906, Kalas et al. (2015); HD 131835, Hung et al. (2015); and HD 61005 (Esposito et al., submitted). In all of these disks, this asymmetry is interpreted as the disk being tilted such that the brighter side is closer to the observer and the observed brightness asymmetry is mostly due to strong forward scattering in the polarized scattering phase function. Indeed, a recent analysis of Cassini observations (albeit *total intensity visible light* observations) of Saturn's G and D rings indicate that collisionally generated dust is expected to be strongly forward scattering (Hedman & Stark 2015).

In addition to the NE-SW asymmetry, the Q_r image also displays a mild brightness asymmetry between the SE and NW sides of the disk, visible as two main features: a) The SE ansa appears brighter

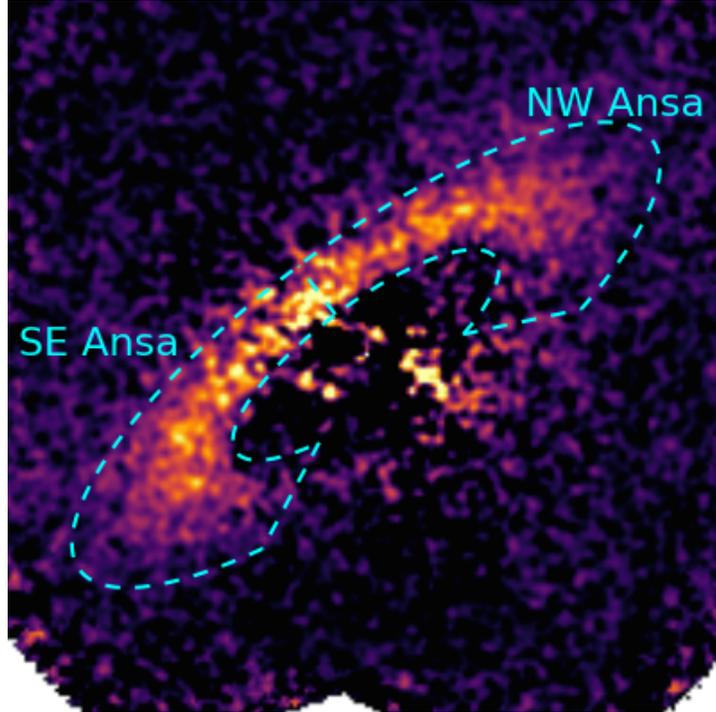


Figure 4.3: The regions used to calculate the magnitude of the brightness asymmetry between the two ansae overplotted on top of the Q_r image from Figure 4.1.

and reveals more of the backside of the disk than the NW ansa, causing the ansa to appear hook-like, and b) the SE side of the disk appears brighter along the bright NE edge of the disk, about the NE semi-minor axis. These features are confirmed in a signal-to-noise ratio (S/N) map (Figure 4.2), created by dividing the Q_r image at each point by the standard deviation of an annulus in the U_r image at the same angular separation.

To estimate the magnitude of this asymmetry we created a custom shaped aperture for each ansa (Figure 4.3). The two apertures are mirror images of each other, with the axis of symmetry coincident with the projected minor axis (a position angle of 37° , §4.5). By summing the flux in each aperture, we find the brightness ratio between the SE ansa and the NW ansa to be 1.15 ± 0.02 , where the errors at each pixel are calculated in the same manner as when creating the S/N map.

This brightness asymmetry may be explained by a stellar offset, which in turn may be caused by a perturbing planet in an eccentric orbit that imparts a forced eccentricity to the dust’s parent bodies (e.g. Wyatt et al. 1999). For small eccentricities, the morphology of the disk remains axisymmetric (to first order), but the host star is no longer located at the geometric center of the disk. A brightness asymmetry can then be seen as the result of one side of the disk being closer to the star than the other and receiving increased stellar irradiation. This offset also warms the closer dust grains, an effect known as pericenter glow which can be observed in thermal emission (Wyatt et al. 1999).

4.4.2 Total Intensity Image

Each individual polarization datacube was summed across its two polarization channels to create a total intensity datacube. The entire set was then processed with pyKLIP using a large range of Karhunen-

Table 4.1. Properties of the Candidate Point Sources Around HD 157587

Label	S/N	Separation	Position Angle	<i>H</i> -band Flux Ratio ^a	Radial Separation ^b
b	6.6	1''180 ± 0''002	228°9 ± 0°2	(3.2 ± 0.8) × 10 ⁻⁶	364 AU
c	4.2	1''248 ± 0''005	195°9 ± 0°2	(2.7 ± 0.7) × 10 ⁻⁶	380 AU
d	2.8	1''002 ± 0''004	269°3 ± 0°3	(1.9 ± 0.7) × 10 ⁻⁶	210 AU

Loève (KL) modes, exclusion criteria and optimization regions. No disk emission was detected in any reduction. This is due in part to prominent stellar residuals resulting from imperfect AO correction caused by strong winds throughout the observation set. These winds held a roughly constant position angle during the observations, which caused the AO residuals to rotate relative to the instrument frame, mimicking the rotation of an astrophysical source. Aggressive PSF subtraction is able to suppress nearly all of this signal, but also suppresses any of the disk emission. In addition to the wind residuals, any angular differential imaging-based method will be subject to self-subtraction for such an azimuthally extended disk, compounding the difficulties in detecting the disk in total intensity.

Although no dust-scattered light was detected, we recovered three point sources (Figure 4.4) in the PSF-subtracted data. Their measured properties are summarized in Table 4.1. While the faintest of the three sources (source *d*) is recovered at less than 3-sigma significance, we find that it is stable as a function of KL modes and appears as a point source in both our most conservative (i.e., with low number of KL modes) and aggressive (i.e., with a high number of KL modes) reductions, which does not hold true for other low significance point source candidates in the data. Thus, we consider all three point sources to be real.

The flux and position of the point sources were calculated using a Gaussian matched filter. The flux S/N was determined by comparing the flux of the point sources with the noise at the same radial separation. Because the point sources lie outside of the region with strong wind residuals, we used a parallelogram-shaped region to mask out the wind residuals when estimating the noise. To correct for algorithm throughput and to characterize the position uncertainties, artificial point sources of known brightness and position were injected into the data at similar separations but at different azimuthal positions with respect to the point sources *b, c* and *d*, avoiding the region with strong wind residuals. Algorithm throughput was estimated by measuring the flux of the artificial point sources after PSF subtraction. The scatter in the position and flux of the artificial planets were used as the uncertainties on the position and flux of the point sources, respectively. To obtain the total error in the astrometry, we use the reported plate scale and North angle from De Rosa et al. (2015b) and add the uncertainties in quadrature. For our flux conversion, we used the flux of the satellite spots to convert the flux of the point sources to contrast units, using the standard GPI calibrations for the flux ratio of the satellite stars relative to the central PSF (Wang et al. 2014). The scatter in the satellite spot fluxes was used as the uncertainty in the flux conversion factor.

While it is possible that one or more of these three point sources is associated with HD 157587, we note that the star’s projected position on the sky is near the galactic plane ($[l, b] = [6^\circ 0, 9^\circ 4]$) and it is likely that most, if not all, of these sources are background objects. We further discuss the potential relationship of these point sources to the debris disk in §4.6.

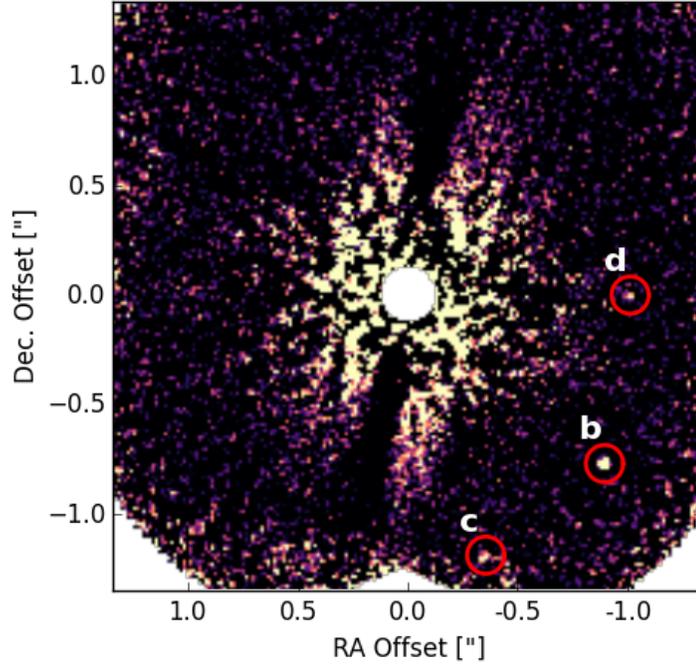


Figure 4.4: The PSF-subtracted total intensity (Stokes I) image of HD 157587 at H -band. No disk is detected, in part because of strong AO residuals caused by winds. Residuals due to the winds can be seen as nearly vertical dark streaks to the north and south of the obscured star. The three point source described in Table 4.1 are marked by the red circles.

4.5 Disk Modeling

To recover basic geometric properties of the disk, we modeled the Q_r image using two modified versions of the disk model presented in Millar-Blanchaer et al. (2015). The original model describes the three-dimensional dust density as a radial power law centered on the host star with a Gaussian height profile and constant aspect ratio (the ratio of the disk scale height to the radial separation). Optically thin (single) scattering is assumed and a Henyey-Greenstein (HG) function is used as a polarized scattering phase function. A disk image is calculated by combining the dust density profile with the scattering phase function and integrating along the line of sight. The model includes nine free parameters: inner radius, R_1 ; outer radius, R_2 ; power law index for the radial dust distribution, β ; disk aspect ratio, h_0 ; average scattering cosine, g ; inclination, ϕ ; position angle on the sky, θ_{PA} ; a flux normalization factor N_0 and a constant offset term applied to the entire image I_0 .

Motivated by the SE-NW asymmetries seen in Figure 4.1, we adapted the disk model to allow for a stellar offset in the plane of the disk. This was implemented as two new free parameters: ΔX_1 , an offset along the projected semi-major axis of the disk (i.e. parallel to the position angle in the plane of the sky; positive offset is in the SE direction), and ΔX_2 , an offset in the disk plane perpendicular to ΔX_1 (positive offset is towards the backside of the disk). The true orientation of ΔX_2 relative to the observer is parallel to the inclination vector ϕ and will have components both along the projected semi-minor axis of the disk and along the observer's line of sight. For example, if the disk is edge-on ($\phi = 90^\circ$), then ΔX_2 is parallel to the line of sight and if the disk is face-on, then ΔX_2 will be in the plane of the sky.

We considered two different options for the polarization scattering phase function. We first used the

HG function, as in the original model, and we also considered a model where the polarization scattering phase function is described by a Rayleigh scattering function multiplied by an HG function. This second model was considered because the Rayleigh scattering function produces a peak in the polarization fraction at 90° scattering angles and therefore places increased importance on the ansae, where the hook-feature is seen. In this second model the HG function describes the scattering phase function of the (undetected) total intensity.

A similar strategy was used by Graham et al. (2007), who simultaneously fit total intensity and polarization fraction observations of the AU Mic debris disk with the same combination of an HG function and Rayleigh scattering function. One difference between their analysis and ours is that their fit includes the maximum polarization fraction, p_{max} , as a free parameter. In our case, with no total intensity detection, we fit only the polarized intensity, not the polarization fraction, and p_{max} is folded into our flux normalization term, N_0 .

Following Millar-Blanchaer et al. (2015), we fit the two disk models to the image using the affine invariant Markov Chain Monte Carlo (MCMC) ensemble sampling package `emcee` (Foreman-Mackey et al. 2013). The package uses an ensemble of “walkers” each with its own MCMC chain, where each step of each walker depends on the state of the other walkers. Before fitting, we first applied 3×3 pixel binning to both the Q_r image and the U_r image. At H -band, the diffraction limit of Gemini is about 3 GPI pixels. Therefore, by binning the data, we can improve S/N without any significant loss of spatial information. This has the added benefit that each binned pixel can be considered independent and correlated errors between the binned pixels are minimized. The error estimates for each pixel were calculated as in §4.3, by taking the standard deviation of concentric annuli in the (binned) U_r image. Preliminary fitting runs revealed that there was no evidence for an offset term, I_0 and as a result we opted to drop the offset as a free parameter.

For the first model (HG-only), the fitting procedure was run with 240 walkers and a burn-in stage of 500 iterations followed by a full run of 1300 iterations. The burn-in serves to initialize the state of the walkers and only walker positions from the full run are used to calculate the posterior. The fitting code was executed on the Edison supercomputer at the National Energy Research Scientific Computing Center (NERSC), employing one process for each walker. For each iteration, a new model was generated at the same resolution as a full GPI frame, smoothed with a Gaussian filter ($\sigma = 1$ pixel) and then binned 3×3 , to replicate the final steps of the data reduction.

After the run, the maximum autocorrelation across all parameters was found to be 65 iterations, indicating that the chains had iterated longer than the recommended $O(10 \times t_{autocorrelation})$ for convergence (see Foreman-Mackey et al. 2013). As a second check, the ensemble chains were examined by eye after the burn-in stage and appeared to have converged. The entire procedure took 46 hours and 30 minutes to complete using 240 Intel Ivy-Bridge 2.4-GHz cores (10 nodes each with 24 cores).

We fit the data to the second model (HG + Rayleigh), using the same `emcee` setup, but only 1000 iterations were completed due to time constraints on the NERSC Edison supercomputer. For this run the autocorrelation time was found to be 91, again, indicating that the chains should have converged.

Marginalized posterior distributions for each parameter of each of the two models were obtained by sampling the MCMC chains of each walker at intervals equal to one autocorrelation time. A summary of the results of the fitting procedure can be found in Table 4.2, which reports the best-fit value for each parameter as the median of its posterior, and the errors are taken to be the 68% confidence intervals. The normalization factor N_0 is considered a nuisance parameter and is not included in the table. The

Table 4.2. Best-Fit Disk Model Parameters for the HD 157587 Disk

Parameter	Symbol	Units	Prior Range	Best-Fit Values	
				HG Model	HG + Rayleigh Model
Inner radius	R_1	AU	40 – 120	81 ± 2	78.9 ± 0.8
Outer radius	R_2	AU	$R_1 - 1000$	216 ± 16	211_{-15}^{+21}
Density power law index	β	...	0.1 – 4	2.23 ± 0.15	2.2 ± 0.2
Scale height aspect ratio	h_0	...	0.0001 – 0.5	0.079 ± 0.005	0.084 ± 0.006
HG asymmetry parameter	g	...	0 – 1	0.285 ± 0.012	0.65 ± 0.03
Line of sight inclination	ϕ	$^\circ$	45 – 85	72.2 ± 0.4	$68.3_{-0.8}^{+0.7}$
Position angle	θ_{PA}	$^\circ$	105 – 145 $^\circ$	127.0 ± 0.3	127.1 ± 0.3
Sky plane offset	ΔX_1	AU	$-R_1 - R_1$	1.6 ± 0.6	1.4 ± 0.6
Line of sight offset	ΔX_2	AU	$-R_1 - R_1$	2.1 ± 1.6	$-5.7_{-2.2}^{+2.1}$
Chi-Squared	χ^2	2648	2670
Reduced Chi-Squared	χ^2_{red}	0.89	0.90

posteriors for each parameter all appeared to be single peaked and approximately Gaussian, with the exception of the outer radius, R_2 that displayed a slightly elongated tail towards higher values. Slight degeneracies were found between ϕ and ΔX_2 . However, all parameters appeared to be well constrained in the marginalized distributions. Images of the best-fit models for both the HG fit and the HG + Rayleigh fit can be seen in Figure 4.5, alongside an image of the residuals of each model subtracted from the Q_r image.

The best-fit HG (HG + Rayleigh) model reveals disk inclination of 72.2 ± 0.4 ($68.3_{-0.8}^{+0.7}$), with a relatively steep radial density power law index of 2.25 ± 0.15 (2.2 ± 0.2), that extends from an inner radius of 81 ± 2 AU (78.9 ± 0.8 AU) to an outer radius of 216 ± 16 AU (211_{-15}^{+21} AU). The best-fit inclinations are roughly consistent with that determined by measuring the aspect ratio in the STIS disk image ($\sim 75^\circ$; Padgett et al., in prep). The outer radius is constrained by the level of forward scattering seen in the Q_r image towards the NE, which quickly reaches the noise floor. Therefore, this fit parameter is governed by the sensitivity of these observations, and represents a lower limit on the true disk outer radius. Indeed, the disk is seen to extend to a radius greater than 500 AU in the HST/STIS images. However, it is possible that the HST/STIS observations probe a different dust population.

The results of the disk fitting also indicate that the disk is offset from the star the along the projected semi-major axis, ΔX_1 , by 1.6 ± 0.6 AU (1.4 ± 0.6 AU). This feature is consistent between both models, with both fits finding an offset with a significance slightly below 3σ . The direction of this offset is consistent with the direction of the brightness asymmetries seen in Figure 4.1 and it causes a faint brightness asymmetry in both model images in Figure 4.5.

Visually, there are two major differences between the two models. First, there is an apparent deficit of light near the projected semi-minor axis in the HG + Rayleigh model which is due to the low polarization fraction at small (and large) scattering angles of the Rayleigh function. The deficit is in an area that coincides with a region of reduced S/N in Figure 4.2 and this area was therefore down-weighted in the MCMC fitting relative to the ansae. In Figure 4.1, the brightness appears to increase towards the NE, with a maximum near the semi-minor axis and no apparent deficit. This may be a real feature of the disk, but could also be possibly due to residual systematics in the image. Indeed, there is a strong residual along the SW semi-minor axis. Comparing the residuals between the two models in Figure 4.5, there appears to be additional flux at the location of the deficit in the HG + Rayleigh model. However, the nearly identical reduced- χ^2 value indicate that both model fit the data equally well.

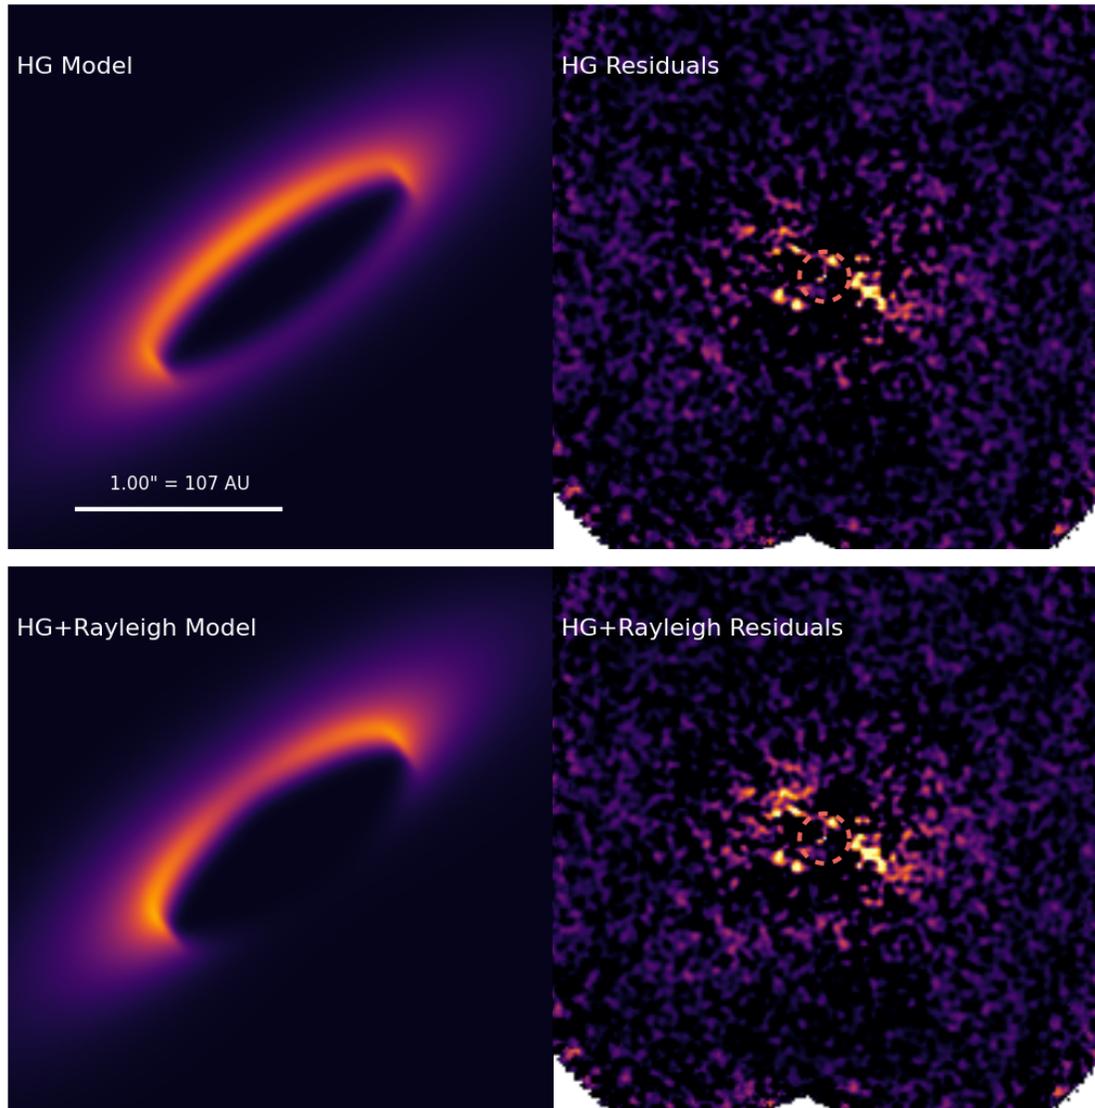


Figure 4.5: *Left:* Images of the best-fit H -band, Q_r models described in Table 4.2. *Right:* The residuals of each model subtracted off the Q_r image. The top row displays the disk image and residuals for model which uses only an HG function to describe the polarized scattering phase function. The bottom row displays the disk image and residuals for the model that uses an HG function combined with Rayleigh scattering. The images are displayed at the same colour scale and orientation as Figure 4.1.

The second major difference between the two models is the brightness of the backside of the disk, which is barely visible in the HG + Rayleigh model but appears in the image of the HG model. A close examination of the HG + Rayleigh model residuals image, reveals faint residuals at the location of the backside in the HG model that is not apparent in the HG-only model residuals. These residuals may indicate a very low S/N detection of the backside of the disk in the Q_r image.

Between the two models, the ΔX_2 offset varies in both magnitude and sign. In the HG-only best-fit model ΔX_2 is positive, indicating that the star is offset away from the observer along the inclination vector, relative to the disk center. The HG + Rayleigh model is best fit by a negative offset, where the star is offset towards the observer. This discrepancy illuminates a degeneracy between this offset and the exact form of the polarized scattering function. We therefore consider this parameter to remain poorly constrained. Future studies that are able to constrain the grain scattering properties or image the backside of the disk will be able to elucidate this remaining unknown.

4.6 Discussion

The dust seen in scattered light images of debris disks is thought to have originated in disks or belts of planetesimals, where collisional cascades grind km-sized bodies down to micron-sized dust (see Wyatt 2008, and references therein). In order to initiate these cascades, the constituent planetesimals must be dynamically stirred such that their eccentricities reach a high enough level (on the order of 10^{-3} to 10^{-2}) to allow their orbits to cross and for their collisions to be destructive. Stirring mechanisms include: self-stirring, where objects on the order of 2000 km located inside the belts induce the cascade from within (e.g. Kenyon & Bromley 2004); planet-induced stirring, where a nearby planet excites the disk (e.g. Mustill & Wyatt 2009b); or dynamical interaction with a passing star (e.g. Kenyon & Bromley 2002). With the exception of the stellar fly-by scenario, the strength of all of these mechanisms should diminish with stellar age. As a given system reaches a steady-state configuration it cools dynamically and the collision rate slows. Therefore, the scattered light luminosity should dim with age as the small grains are removed via radiation pressure, Poynting–Robertson drag and/or stellar winds, and can no longer be replenished through collisions. Note that for disks of all ages, large collisions or other transient events may cause a temporary increase in dust production and create a short-term increase in disk brightness (e.g. Wyatt et al. 2007).

The only published age estimates of HD 157587 are presented by Feltzing et al. (2001) and Casagrande et al. (2011). Feltzing et al. (2001) fit metallicity, effective temperature and absolute magnitude to the outputs of a rapid stellar evolution algorithm (Hurley et al. 2000) based on the evolutionary tracks produced by Pols et al. (1998). They estimate HD 157587 to have an age of 2.2 ± 0.5 Gyr. Casagrande et al. (2011) use a Bayesian analysis to fit effective temperature, metallicity and Johnson V magnitude to the Padova evolutionary tracks (Bertelli et al. 2008; 2009) and the BASTI isochrones (Pietrinferni et al. 2004; 2006; 2009). They find an age of $3.0_{-1.5}^{+1.7}$ Gyr and $3.0_{-1.4}^{+1.0}$ Gyr at 95% confidence for the Padova and BASTI models, respectively.

Such advanced ages suggest that the star has evolved along the HR diagram away from the zero age main sequence (ZAMS). Alternatively, this offset from the ZAMS may be indicative of a pre-main sequence star moving towards the ZAMS, rather than away from it. However, considering the timescales of evolution along the pre-main sequence tracks (~ 12 Myr for an F5 star, Siess et al. 2000) compared to main sequence tracks (on the order Gyrs), it is more likely that the star has been found as it moves away from the ZAMS.

Conversely, applying the star’s proper motion, radial velocity and parallax to the BANYAN II webtool (Gagné et al. 2014; Malo et al. 2013)¹ indicates a 91% probability that the star is a young (< 1 Gyr) field star, and a 9% probability that the star is an old field star. Additionally, with velocities of $[U, V, W]$

¹<http://www.astro.umontreal.ca/gagne/banyanII.php>

$= [-7, -17, -8]$ km/s (Holmberg et al. 2009), we find that HD 157587 could be kinematically associated with several relatively young moving groups. For example, in Chereul et al. (1999) the Pleiades stream 2-5 has $[U, V, W] = [-12.0 \pm 5.3, -21.6 \pm 4.7, -5.3 \pm 5.9]$ km/s, whereas the Centaurus-Lupus stream 2-12 has $[U, V, W] = [-12.4 \pm 6.1, -16.5 \pm 4.6, -7.4 \pm 3.1]$ km/s. In Asiain et al. (1999), the Pleiades moving group B2 has $[U, V, W] = [-10.7 \pm 5.3, -18.8 \pm 3.7, -5.6 \pm 2.2]$ km/s, whereas members of Lower Centaurus Crux have $[U, V, W] = [-6.8 \pm 4.7, -18.5 \pm 6.5, -6.4 \pm 1.7]$ km/s. In all of these cases the age distributions determined by various methods lie in the range $10^7 - 10^9$ years. Thus, the stellar kinematics possibly support a relatively young age for HD 157587, while its photometry supports a much older age.

If HD 157587 is truly 2.5 – 3 Gyr old, then its disk’s optical/NIR emission would be unusually bright when compared to the overall population of imaged debris disks (for a good summary see Figure 1 from Choquet et al. 2016). In fact, the disk around HD 157587 would represent one of only five debris disks imaged in scattered light with ages greater than one Gyr. The other four old disks seen in scattered light are HD 207129 (Krist et al. 2010), HD 202629 (Krist et al. 2012), HD 53143 (Kalas et al. 2006) and HD 10647 (q¹ Eri; Stapelfeldt et al. 2007). With a spectral type of F5V, HD 157587 is the earliest type star of the five, making it potentially the earliest type star older than one Gyr with a debris disk seen in scattered light.

Of the other four old scattered-light debris disks (i.e. > 1 Gyr), HD 10647 is the most similar to HD 157587. HD 10647 is an F8V star with a debris disk at a radial distance of ~ 85 AU (Liseau et al. 2010). The disk has been imaged in scattered light with HST/ACS (Stapelfeldt et al. 2007) and in the infrared with Herschel/PACS (Liseau et al. 2010). As seen here for HD 157587, the HST images show evidence of a brightness asymmetry between the two ansae of the disk. Interestingly, HD 10647 is known to host a Jupiter mass planet at 2 AU. However, this planet is at too great a distance from the disk to have significant dynamical influence on the disk. Liseau et al. (2008) posit that the disk asymmetry suggests the presence of a second planet at larger distances.

In our images of the HD 157587 disk we see a slight brightness asymmetry between the SE and NW ansae of the disk that we have modeled as being caused by a stellocentric offset. If the offset is due to a perturbations from a substellar companion, it is highly unlikely that it is one of the three point sources imaged in total intensity. Using the COND evolutionary models (Baraffe et al. 2003) and assuming an age of 3 Gyr, we find their luminosities (flux ratios) correspond to brown dwarf masses between 30 and 40 M_{Jup} . If these three objects reside within the plane of the disk, their deprojected separations correspond to stellocentric radii between 210 AU and 380 AU, assuming circular orbits. Considering the smooth radial extent of the disk seen in the STIS images, we deem such an alignment unlikely, as the presence of such massive companions would cause significant disturbances to the disk morphology. Though it is possible that the expulsion timescale for the STIS grains are short enough hide any such feature.

Future imaging observations that reveal the relative motion of these three sources are required to understand their true relationship to HD 157587, if any. However, it is worth noting that the likelihood of a field star flyby is rare - less than 1% in 100 Myr of a 500 AU approach (Kenyon & Bromley 2002) - and considering HD 157587’s proximity to the galactic plane, these three objects are likely background sources. If HD 157587 is a younger star, as suggested by the stellar kinematics, it is possible that a lower mass perturber at smaller angular separations is still bright enough in thermal emission to be detected via direct imaging.

4.7 Conclusion

Using GPI we have imaged the dust ring around HD 157587 in H -band polarized intensity. The image reveals an inclined disk that appears to be cleared of material inside of a projected major axis of ~ 80 AU. The FOV of our observations overlaps with the inner regions of previous STIS images of the disk, and our analysis returns a similar disk inclination to that derived with the STIS data. The disk has a strong polarized brightness asymmetry in the NE-SW direction, where we interpreted the bright side of the disk to be tilted towards the observer. A similar brightness asymmetry has been seen in polarized observations of a number of other recently imaged disks, suggestive of similar grain compositions, size distributions and/or dust grain morphologies. Future detailed studies of these disks' dust composition that include multicolour observations or polarization fraction measurements will be able to further explore the similarities and differences of their dust grain populations.

A second, weaker, brightness asymmetry is seen between the two ansae that could be due to a stellocentric offset in the plane of the sky. To test this hypothesis we used Bayesian MCMC methods to fit the polarized disk image to two disk models, one that used a HG polarized scattering phase function and one that combined a HG function with Rayleigh scattering phase function. Both models reveal an offset dust disk with an inner radius of 80 AU and an inclination of about 70° . The center of the disk is found to be offset approximately 1.5 AU from the star's location in the plane of the sky and both models reproduce the brightness asymmetry between the two ansae. This offset could be confirmed with longer wavelength imaging using ALMA, which would trace thermal emission and therefore have less of a dependence on the scattering properties of the grains.

In general the two model fits return similar disk properties, with the exception of ΔX_2 , the offset in the disk plane. We find that the form of the polarized scattered phase function is degenerate with the magnitude and direction of this offset and without further information on the form of the scattering phase function this value will remain poorly constrained.

The total intensity observations are dominated by stellar residuals at the location of the disk in the polarized intensity image and no disk was recovered. However, three point sources were recovered. Considering HD 157587's proximity to the galactic plane and the positions of the point sources relative to the disk, we consider these point sources to be background objects. Nonetheless, follow up observations are required to confirm this proposition.

The currently published ages of the system that rely on stellar evolutionary tracts indicate an age well over 1 Gyr old. However, such an evolved age is at odds with the stellar kinematics. The stellocentric offset, suggest that this system has a complicated dynamical history and may harbour one or more unseen planets. This notion is reinforced by the similarities between HD 157587's stellar properties and disk morphology, and those of the RV planet host HD 10647. If the stellocentric offset is due to perturbations by one or more planets, further detailed study of the system's debris disk will be required to thoroughly characterize the system; the system's advanced age would make it ill-suited for direct imaging planet searches and radial velocity measurements would require prohibitively long time baselines. On the other hand, if the disk is younger than 1 Gyr, as implied by the stellar kinematics, then it presents itself as a prime target for deeper direct imaging observations which may be able to image the disk's perturber.

This work was supported in part by NASA's NEXSS program, grant number NNX15AD95G. Portions of this work were performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Portions of this were also carried out with the support of NSF grants AST-1413718 and AST-1518332, and NASA grant NNX15AC89G. This

research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Chapter 5

Polarization Gratings for Visible and Near-infrared Astronomy

This work has been published in the proceedings of the 2014 SPIE Astronomical Telescopes + Instrumentation conference as “Polarization gratings for visible and near-infrared astronomy” by Millar-Blanchaer et al. 2014.

5.1 Chapter Summary

We report on the development of polarization gratings that can be used for polarimetry and/or high throughput broadband spectroscopy in astronomy. Polarization gratings are able to overcome fundamental limitations on the diffraction efficiency of conventional gratings to provide near 100 % diffraction efficiency over a broad bandwidth. The broad spectral coverage of these devices will be useful for observations of gamma-ray bursts and supernovae of unknown redshift, where spectral features may fall over a range of wavelengths. As a spectropolarimeter a polarization grating would be ideal, for example, for the study of dusts and hazes, whose polarimetric properties vary with wavelength. We present the results of a series of laboratory measurements of the diffraction efficiency and modulation efficiency of a prototype grating designed for operation from 500 to 900 nm. We find that the grating is able to achieve greater than 90 % diffraction efficiency from 500 to 850 nm and modulate incident circular polarized light with an efficiency of $\sim 99\%$. Our future plans include on-sky testing at a small local telescope, with an eventual goal of incorporating a polarization grating into the design of a microshutter array-based multi-object visible/NIR spectrograph for a 10 m class facility.

5.2 Introduction

A polarization grating is a type of diffraction grating that can be designed to diffract incident light with high efficiency into only the $m=+1$ and/or $m=-1$ orders (Crawford et al. 2005; Escuti et al. 2006; Provenzano et al. 2006; Packham et al. 2010). The amount of light diffracted into each order depends on the circular polarization state of the incoming beam, so that they act simultaneously as polarizing beam splitters and spectrally dispersive gratings. In contrast to conventional surface relief gratings, which affect the phases of orthogonal polarizations equally, polarization gratings (PGs) are able to control

the two phases separately. As a result, PGs can readily overcome the fundamental limitations on the diffraction efficiency of ruled gratings (Wyrowski 1991). Conventional gratings may be blazed to achieve relatively high diffraction efficiencies, but only over a small wavelength range. In these cases diffraction efficiency must often be traded for a spectral coverage. PGs, on the other hand, are able to achieve near perfect diffraction efficiency over a near unity bandwidth, $\Delta\lambda/\lambda_0 \approx 1$ (Oh & Escuti 2007; Escuti & Oh 2010). This highly desirable property makes PGs relevant as dispersing elements not only for polarimetric applications, but for normal spectroscopy as well.

The development of these devices has been spearheaded by one of the co-authors of this work at North Carolina State University, where the applications to date have mostly been in display, telecommunications, beam steering, and remote sensing fields. PGs have been previously considered by Packham et al. (2010) for a mid infrared spectropolarimeter for SOFIA and in the near future a PG is expected to be installed as an instrument upgrade to MMT-POL (Packham et al. 2012). Whereas these two efforts have mostly involved using PGs in the IR, our work focuses on using PGs in the visible and near infrared (NIR).

Our eventual goal is to develop a PG that operates from 600 to 1600 nm that can be included in a spectropolarimeter on a large telescope. As a first step we have designed and fabricated a set of PGs that have been optimized for near unity diffraction efficiency between 500 and 900 nm. In this chapter we present the results of diffraction efficiency and modulation efficiency measurements carried out at the University of Toronto on one of the gratings. Following further lab tests the grating will be included in the design of a simple visible spectropolarimeter for a small local telescope.

In §5.3 we summarize the basic properties of polarization gratings and provide a description of our visible PGs. A description of our lab test set-up and our laboratory measurements can be found in §5.4. In §5.5 we consider science targets for PGs in the visible and NIR regimes and explore several possible instrument architectures. Our future plans are discussed in §5.6.

5.3 Polarization Gratings

The PGs developed at NCSU are based on liquid crystal polymer thin film coatings, which take advantage of the Pancharatnam-Berry effect to induce a phase delay on incident light. A simple PG can be made with uniaxially birefringent liquid crystals that rotate in one direction along the grating surface. By designing the thin film to induce a half-wave of retardance for a central wavelength, it can be shown (e.g. McManamon et al. 2009) that a change in orientation from 0 to π in liquid crystal orientation can impart a phase delay of 0 to 2π to on-axis light. By regularly repeating a 180° rotation across the grating plane, one can construct a PG that follows standard diffraction theory, where a 180° rotation defines the grating period. In particular, the periodic pattern causes PGs to diffract light as a function of wavelength following the grating equation. The 2π phase delay has the effect of sending light into only the $m=+1$ and/or $m=-1$ orders at a very high efficiency.

The birefringence of the liquid crystals creates a sensitivity to the circular polarization state of the incident light, which as a result sends purely left-hand circular or right-hand circular polarized light exclusively to the $m=+1$ or $m=-1$ order, respectively. Unpolarized and linearly polarized light is diffracted equally between the two orders. Regardless of the incident polarization, the total intensity can be recovered by summing the $m=+1$ and $m=-1$ orders.

This design is able to deliver near unity diffraction efficiency (when you sum the $m=+1$ and $m=-1$

Table 5.1. Design Parameters for the Polarization Gratings Tested at the University of Toronto.

Parameter	Value
Central Wavelength, λ_0	700 nm
Diffraction Efficiency Range	500 – 900 nm
Grating Period	5 μm
Grating Diameter	20 mm

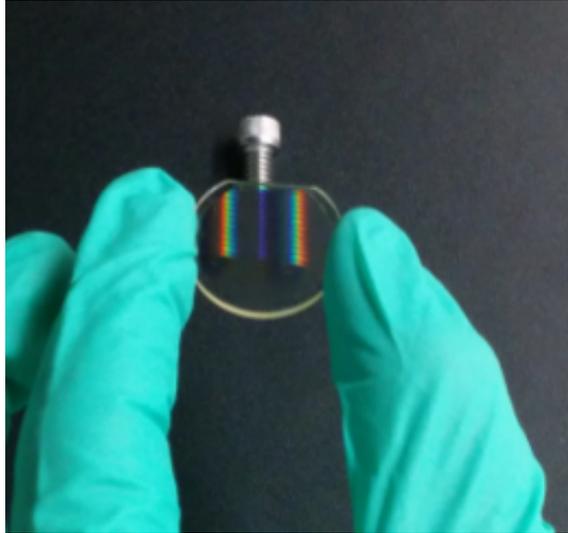


Figure 5.1: A photo of a visible PG designed for testing at the University of Toronto. Where the PG is placed above the screw, images appear in the $m=+1$ and $m=-1$ orders. At bluer wavelengths the efficiency is not as high and the zeroth order can be seen at the center of the grating.

orders), but only at the central wavelength, where the thickness of the thin film induces a half-wave of retardance. The wavelength coverage can be extended by sandwiching together two polarization gratings, each with a chiral nematic twist of opposite handedness induced in the liquid crystals along the direction of the optical axis. For a two layer grating the diffraction efficiency can be shown to be:

$$\eta_{\pm 1} = A \left(\frac{1 \mp V}{2} \right) (\cos^2 X + \Phi^2 \text{sinc}^2 X) \quad (5.1)$$

where V is Stokes V , $X = \sqrt{\Phi^2 + \Gamma^2}$, $\Gamma = \pi \Delta n d / \lambda$, and $A = 2\Gamma \text{sinc} X$ (see Oh & Escuti 2007, for a derivation). These equations are governed by the thin film properties: Δn is the birefringence of the liquid crystals, d is the thickness of the film and Φ is the total rotation of the liquid crystals along the optical axis of each layer. By optimizing these three parameters, the grating can be designed to have a diffraction efficiency above 99 % over a bandwidth $\Delta\lambda/\lambda_0 \sim 0.57$. In reality manufacturing defects can result in slightly lower performance, but very high diffraction efficiency is still achievable at some wavelengths (see §5.4 for lab results). The sensitivity to circular polarization can be seen explicitly in this equation, where the difference of the two orders is directly proportional to Stokes V .

The bandwidth coverage can be extended even further to $\Delta\lambda/\lambda_0 \sim 1$ by using three layers of liquid crystal, each with their own twist (Komanduri et al. 2013). Our work concerns this three layer design, where the extended wavelength coverage will be useful for extending the wavelength coverage into the near infrared.

With the eventual goal of producing a grating that can operate in both the visible and NIR in mind, we decided to first work with a three layer prototype grating designed only for the visible regime. Narrowband PGs have been produced for beam steering applications in the NIR (Oh et al. 2010; Kim et al. 2011), but broadband PGs covering 600 to 1600 nm require additional design work. Visible PGs, on the other hand, have already been produced at NCSU (Oh & Escuti 2008) and the subtleties of their fabrication are well understood. In addition, working in visible wavelengths is simpler from an instrument testing point of view: you can use relatively cheap silicon detectors and no cryogenics are needed.

A set of four prototype polarization gratings with identical design parameters have been designed and fabricated for near unity throughput from 500 to 900 nm (see Table 5.1). The clear aperture is 18-20 mm and the grating period is 5 μm . These gratings will be used for laboratory characterization (see §5.4) and in future astronomical instruments for broadband spectroscopy and spectropolarimetry (see §5.6). In this paper we focus on the analysis of just one of the four gratings, pictured in Figure 5.1.

5.4 Laboratory Characterization

The main goal of these lab tests is to characterize the diffraction efficiency and the modulation efficiency of the prototype gratings. Here we define the diffraction efficiency as the fraction of the total flux leaving the PG, that gets dispersed into the $m=+1$ and $m=-1$ orders at a given wavelength. The modulation efficiency is the amplitude of the flux difference between the $m=+1$ and $m=-1$ orders normalized by their sum as the incident polarization is modulated from $V = +1$ to $V = -1$. The normalized flux difference is equal to the fractional circular polarization seen by the PG. A large modulation efficiency is required for the measurements of faint polarization signals in noisy conditions. With these two measurements in mind, the lab set-up has been designed to relay light from a broadband source through a set of optional polarizing optics to the polarization grating, and then focus the $m=-1$, $m=0$ and $m=+1$ orders onto a detector.

5.4.1 Lab Set-up

The lab set-up is depicted in Figure 5.2. Light from a tungsten lightbulb enters the fiber via a fiber collimation package (not pictured). The collimation package helps to focus the light onto the fiber tip and greatly increases the collecting area of the fiber. Using the fiber allows us to separate our light source from the rest of the optical system, significantly reducing the scattered light in the system and the output of the fiber acts as a point source on the focal plane. The light exits the fiber and is then formed into a ~ 20 mm diameter beam by a collimating lens, after which it passes through a series of removable filters and polarizing optics. The filters include both Neutral Density filters, to control the total intensity reaching the detector, and bandpass interference filters with central wavelengths spaced every 50 nm from 400 to 900 nm (10 nm FWHM).

A Thorlabs LPVIS100-MP2 thin-film nanoparticle linear polarizer (LP) can be placed into the beam to generate linearly polarized light, which is then converted to circular polarized light by the Thorlabs AQWP10M-980 Achromatic Quarter Waveplate. Both components were mounted in manual rotation stages, where 0° corresponds to the polarization axis and the fast axis of the LP and QWP, respectively, being vertical in the lab reference frame. The combination of these two components allows for the full exploration of the surface of the Poincaré sphere. Assuming ideal polarizing components, with the LP

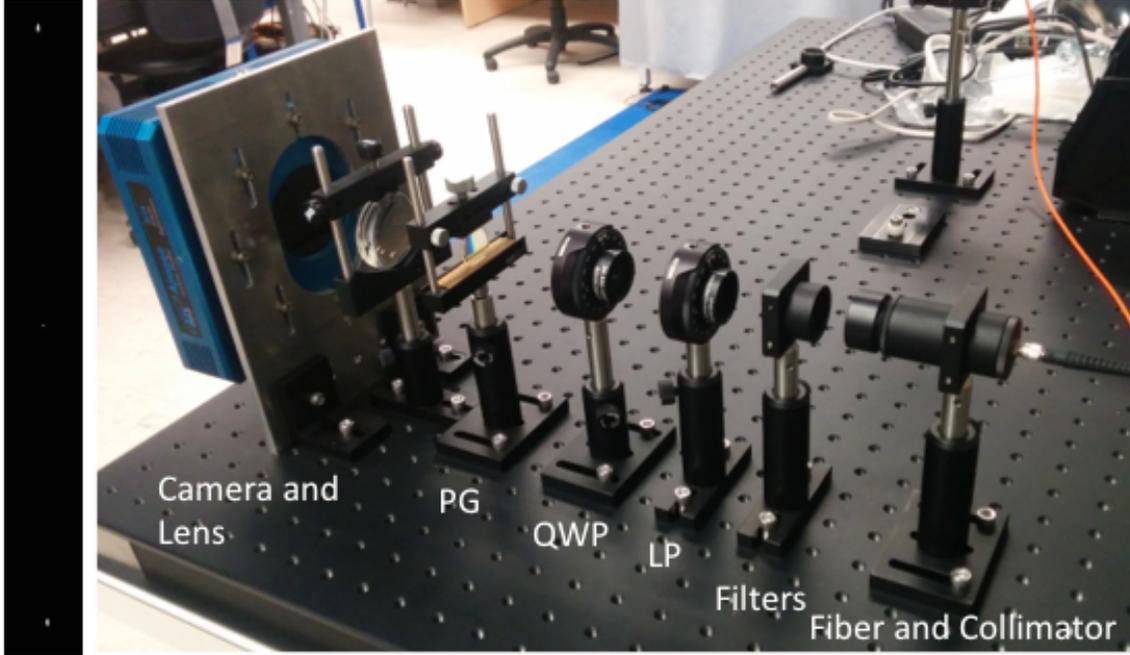


Figure 5.2: *Left* - An example of the raw detector image from our lab set-up with the 900 nm filter in the beam. The first two orders, which are slightly extended due to the 10 nm FWHM bandwidth of the interference filters, can be seen at the top and bottom sides of the image, with the zeroth order in the center. *Right* - Our laboratory test set-up, from right to left (following the path of the light): optical fiber, collimating lens, filter set, LP, QWP, PG, camera lens, Apogee F16M camera. The set-up is able to produce circular polarized light from Stokes $V = +1$ to $V = -1$, and focus all three orders on the detector across the visible regime.

fixed at 0° the Stokes V reaching the PG is equal to $\sin(2\psi)$, where ψ is the position angle of the QWP read off the rotation stage. In reality, the QWP only delivers a quarter-wave retardance from 690 to 1200 nm, and so the above equation is only valid in that range. Thus we focus our polarization analysis on the wavelengths from 700 to 900 nm, where the set-up is able to generate Stokes V from $V = +1$ to $V = -1$.

After the LP and QWP, the collimated light beam passes through the PG and is focused on to the detector by an F/1.5 camera lens. The detector is an Apogee F16M camera, a 4096x4096 pixel array with 9 μm pixels. The large surface area of the detector (3.68 x 3.68 cm) allows both of the first diffraction orders and the zeroth order to be measured simultaneously across our entire wavelength range (see Figure 5.2).

5.4.2 Lab Measurements

Diffraction efficiency measurements were made with one bandpass filter in place at a time and both the LP and QWP out of the beam. The diffraction efficiency at each wavelength, $\eta_{\Sigma\pm 1}$, was calculated as:

$$\eta_{\Sigma\pm 1}(\lambda) = \frac{\sum_{m=\pm 1} F_m(\lambda)}{\sum_{m=-1,0,1} F_m(\lambda)} \quad (5.2)$$

where $F_m(\lambda)$ was the flux measured in the m th order when the bandpass filter of central wavelength λ was in the beam. For these calculations the light diffracted into higher orders is assumed to be negligible

compared to the zero order leakage, which has yet to be experimentally confirmed here, but has been the case for previous visible PGs (Oh & Escuti 2008). The overall transmittance of the grating (i.e. F_{out}/F_{in}) will be a combination of the reflectance, absorption and diffraction efficiency. The thin-film grating itself has negligible absorption and the absorption of the substrate will depend on the choice of glass. The reflectance can be controlled with an anti-reflection coating. In our case we expect both the absorption and reflectance to be very low and so the transmittance and diffraction efficiency should be very similar.

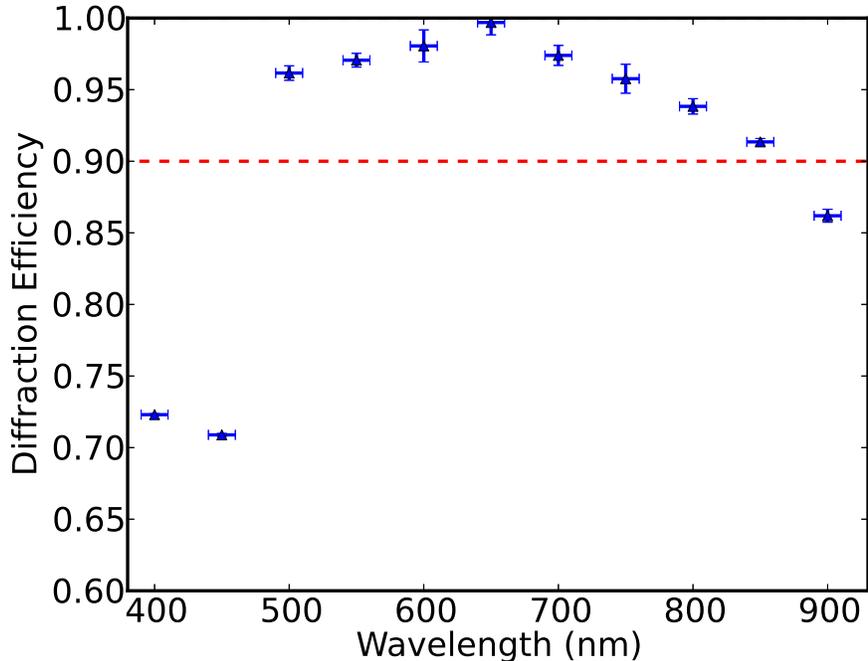


Figure 5.3: The diffraction efficiency of the visible PG. The red dashed line indicates 90 % diffraction efficiency. The error bars in the wavelength direction represent the 10 nm FWHM of the bandpass filters. A peak efficiency of 99.7 % is achieved at 650 nm.

The measured diffraction efficiency from 400 to 900 nm can be seen in Figure 5.3. The camera lens was found to create a small amount of coma in the spots produced in the $m=+1$ and $m=-1$ orders. We measured the flux in each spot using aperture photometry, with apertures large enough to contain the central spots, but not necessarily the entire coma. As a consequence, the flux measurements are likely underestimated for the two first order spots, which in turn would underestimate the diffraction efficiency calculation. However, the effect is likely small since most of the flux is concentrated in the central spot. The error in the flux measurement of each spot was estimated by measuring the background noise outside of the apertures (avoiding potential contamination from the extended coma) and considering the photon noise of the measured signal. These estimates were then propagated through Equation 5.2 to obtain error estimates on the diffraction efficiency.

We find that the PG is able to produce a diffraction efficiency of greater than 90 % between 500 and 850 nm, with a peak of 99.7 % at 650 nm. At 900 nm the diffraction efficiency was measured to be 86.2 %. These results indicate that the polarization grating can achieve high diffraction efficiency across a large bandwidth, as expected.

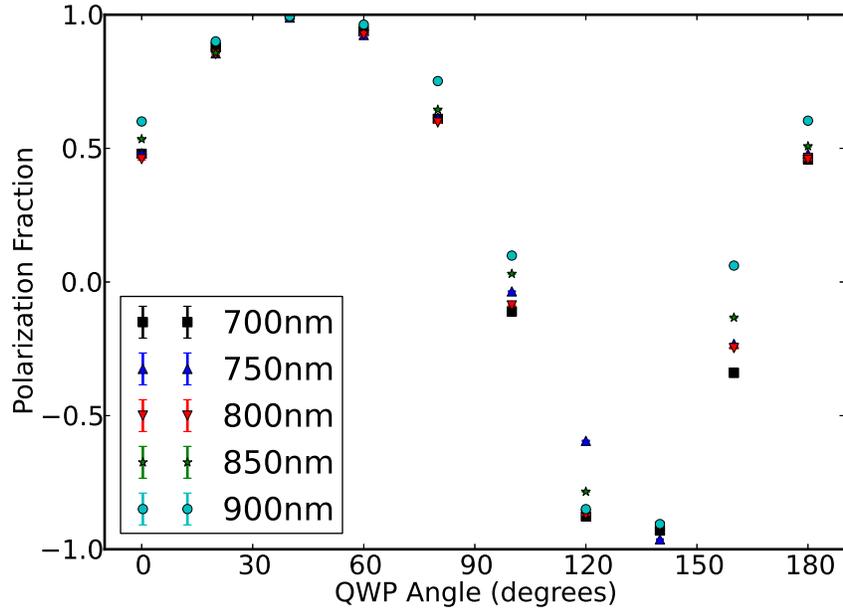


Figure 5.4: The modulation efficiency of the PG from 700 to 900 nm, measured every 20° . A peak modulation efficiency of 99.9 % is reached at 700 nm.

The fractional circular polarization measured by the PG, P_{frac} , has been calculated as:

$$P_{frac} = \frac{F_1(\psi) - F_{-1}(\psi)}{F_1(\psi) + F_{-1}(\psi)} \quad (5.3)$$

where $F_m(\psi)$ was the flux measured in the m th order when the QWP was at an angle of ψ . The measurements of the circular polarization fraction with the QWP at position angles from 0° to 180° can be found in Figure 5.4. The errors were estimated by calculating the error on each spot's flux measurement and the propagated through Equation 5.3. From Equation 5.1, the circular polarized fraction is expected to follow the incident circular polarization state, i.e. Stokes $V = \sin(2\psi)$. A sinusoidal pattern is recovered at all wavelengths, with high levels of modulation across the entire range. The shape of the polarization fraction curve in Figure 5.4 appears to be slightly distorted from a perfect sinusoid: between 0° and 90° the shape appears wider than expected and between 90° and 180° it is narrower. This discrepancy is possibly due to either a measurement error or misaligned optics, as previous PG studies have reproduced the expected curve with much higher fidelity (Escuti et al. 2006).

The modulation efficiency could be measured either directly by measuring the maximum and minimum polarization fractions, which should occur at 45° and 135° , or by fitting a sine curve to the data. Because of the poor fit to a sine curve, and the lack of appropriate measurements neither of these methods are possible. Instead we consider the modulation efficiency to be the maximum absolute value measured, which should provide a lower limit on the actual modulation efficiency. The polarization fraction reaches a peak polarized fraction of 99.9 % at 700nm, and all wavelengths peak above 98 %. Thus, regardless of the exact shape of the curve in Figure 5.4, the fact that modulation efficiencies of $\sim 99\%$ are reached indicates the PG is indeed an efficient polarizing beam splitter.

5.5 Scientific Application of Polarization Gratings

The broad high-throughput spectral coverage of PGs will be very useful for conducting observations of diverse types of astronomical objects, especially those objects that can benefit significantly from broadband medium-resolution spectroscopic information. Any objects of unknown nature, e.g., previously-unknown objects discovered in a new survey or in a serendipitous manner, or of unknown redshift naturally belong to this group. For these types of objects, coarse broadband information is more urgent and can pave the way to more detailed observations to understand their nature and redshift. Due to the recent rapid development in the detector technology, now it is possible to have a large format (e.g., 2k x 2k pixels) detector array that has high quantum efficiencies in a broad wavelength range in the visible-NIR regime such as the substrate-removed HgCdTe Hawaii II RG (H2RG) array that has 60 % quantum efficiency not only in the NIR but also over a significant portion of the visible waveband. We are particularly interested in the PG-based visible-NIR spectroscopy simultaneously covering the 0.6 to 1.6 μm range at $R \sim 1,000$ spectral resolving power using an H2RG for observations of gamma-ray bursts and/or supernovae of high-redshift or unknown redshift. Polarization gratings will also be very useful in measuring the polarization state of astronomical objects, such as dust and haze in planetary systems and exoplanet atmospheres whose polarimetric properties vary with wavelength (Stam et al. 2004).

5.5.1 Polarization Grating-Based Spectrograph Designs

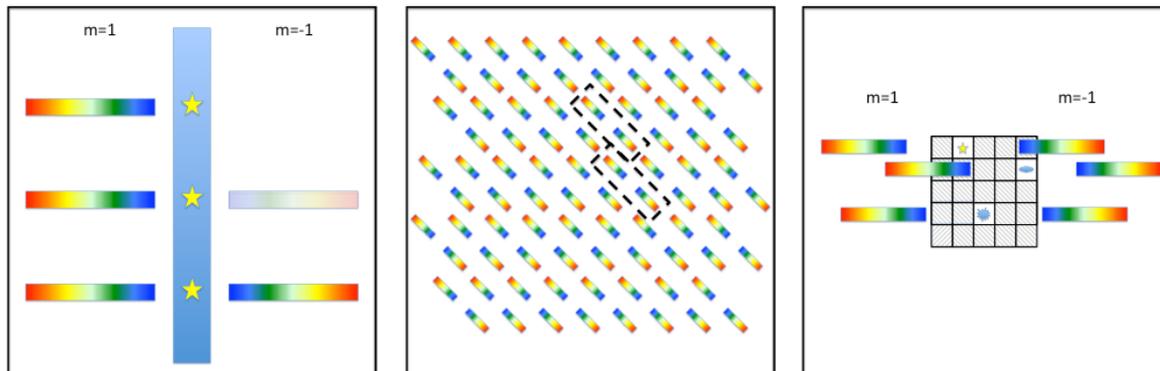


Figure 5.5: *Left* - An example of three different spectra dispersed by a PG in a long slit spectrograph. The top source is completely polarized and so only the $m=+1$ order is visible. The middle source is partially polarized and so the $m=-1$ order is visible, but with a lower intensity. The bottom source is randomly polarized and so both $m=+1$ and $m=-1$ orders are visible with equal intensity. *Center* - The focal plane of a lenslet array based integral field spectrograph equipped with a PG. When viewing an extended source each lenslet produces a pair of spectra (dashed black line), which allows for spectropolarimetry at each spatial location. *Right* - The focal plane of a microshutter array-based spectropolarimeter using a PG (Moon et al. 2014). A microshutter array picks off sources from the FOV which each produce a pair of spectra on the focal plane.

Since PGs function as diffraction gratings, one could imagine including them in most of the standard spectrograph architectures seen in modern instruments. A simple slit-based dual-channel spectropolarimeter could be constructed with a quarterwave-plate (QWP) as a modulator and a PG as an analyzer. The design would follow the basic design of a spectrograph, but with the inclusion of a QWP in the collimated beam, to convert incident linear polarized light to circular polarized light. Figure 5.5 depicts the focal plane arrangement for a slit-based spectropolarimeter as well as two other possible

design schemes of a PG-based spectrograph: a lenslet-based integral field spectrograph and a multi-object spectrograph. Each of the three implementations could be used both as a normal spectrograph, by summing the two dispersion orders, or as a spectropolarimeter, by taking the difference of the two orders at different QWP rotation angles.

5.6 Future Plans and Conclusions

The unexpected shape of Figure 5.3 curve clearly shows that a deeper understanding of the lab measurements is required. While this is likely due to measurement error, if it isn't then the behaviour must be fully understood before including the PG in an instrument. In addition we wish to characterize the modulation behaviour of the PG from 500 to 700 nm to fully understand the polarimetric behaviour of the PG.

Following the lab tests we plan to design and build a proof-of-concept visible slit-based spectropolarimeter with the prototype grating to be tested on the 16-inch University of Toronto campus telescope. We will use such an instrument to measure the polarization of highly polarized stars and of polarized solar system objects. Assuming satisfactory on-sky results we plan to bring the visible spectropolarimeter to a larger telescope. In conjunction with these efforts, we plan to design and test a PG that extends the wavelength coverage into the NIR: with high diffraction efficiency from 600 to 1600 nm. This grating will then be included in a microshutter array-based multi-object spectrograph currently under development at the University of Toronto (Moon et al. 2014).

In conclusion, polarization gratings are a unique new technology that have many potential applications in astronomy. Their broadband diffraction efficiency is unmatched by other gratings and their polarimetric capabilities allow for broadband spectropolarimetry. Our laboratory tests indicate that high diffraction efficiencies ($> 90\%$) are achievable across large bandwidths, with extremely high efficiencies ($> 99\%$) at some wavelengths. Modulation efficiencies on the order of $\sim 99\%$ mean that the PGs are very effective polarizing beam splitters. In the very near future these devices will be on-sky providing many new and exciting scientific developments.

Chapter 6

On-sky Observational Tests of a Polarization Grating for Visible Astronomy

A version of this work titled “On-sky Tests of a Polarization Grating for Visible Astronomy” by Millar-Blanchaer et al. has been submitted to be published in the proceedings of the 2016 SPIE Astronomical Telescopes + Instrumentation conference.

6.1 Chapter Summary

Polarization gratings (PGs) are a type of diffraction grating based on birefringent liquid crystal polymers that simultaneously act as a polarizing beam splitter and as a spectral dispersive element. PGs are also capable of providing high diffraction efficiency ($> 90\%$) over a very broad wavelength range. These properties make PGs ideal for spectropolarimetry and/or high throughput, broad wavelength observations for a range of astronomical objects. Here we report on the design and the results of on-sky observational testing of a prototype spectropolarimeter instrument that employs a PG optimized for operation from 500 to 900 nm. The prototype was mounted on a 16-inch telescope at the University of Toronto, where we carried out observations of the polarized twilight sky, a polarized standard star and two spectroscopic standard stars. Using these observations we demonstrate the PG’s capability to measure linear polarization fraction and position angle as well as recover spectra from astronomical objects.

6.2 Introduction

A PG is a thin-film optical device capable of simultaneously diffracting light according to its wavelength, like a conventional surface relief grating, and of acting as a polarizing beam splitter (Crawford et al. 2005; Escuti et al. 2006; Provenzano et al. 2006; Packham et al. 2010). PGs are fabricated using a uniaxial birefringent liquid-crystal polymer (LCP), where the birefringent axis rotates along one spatial direction. A rotation of 180° is equivalent to the period in a surface relief grating. The rotating birefringence induces

phase delay between incident left hand circular (LHC) and right hand circular (RHC) polarized light, causing the light to diffract into the $m=-1$ and $m=+1$ orders for LHC and RHC polarizations states, respectively. Thus, incident light can be sorted by both wavelength and polarization state. By inducing a rotation of the LCPs in the film along the optical axis, PGs can be designed such that when the two orders are summed their combined diffraction efficiency can reach near unity over bandwidths up to $\Delta\lambda/\lambda \approx 1$. These features make PGs ideal for high-throughput, broad wavelength spectroscopy and/or spectropolarimetry.

Previously we reported on lab tests of a PG designed for high diffraction efficiency between 500 and 900 nm (Millar-Blanchaer et al. 2014). Our tests confirmed that the grating is able to provide a diffraction efficiency over 90 % between 500 and 850 nm, while maintaining it over 86 % at 900 nm. We also showed that the grating is able to modulate incident circular polarized light with a > 99 % modulation efficiency.

In this chapter we present the instrumental design and the results from on-sky observational testing of a prototype spectropolarimeter that we have developed employing the same polarization grating used in our lab tests. The on-sky test observations were carried out on a Boller & Chivens 16-inch telescope at the University of Toronto (UofT), and our test targets included the polarized twilight sky, a polarized standard star and two spectroscopic standards. As shown below, the promising results from these test observations serve to validate the on-sky capabilities of PGs and provide a glimpse of their future potential. In §6.3 we describe the prototype design and the specifications of the UofT telescope. We detail our data reduction steps in §6.4. We present the results of our on-sky test in §6.5, followed by some concluding remarks in §6.6.

6.3 Polarization Grating Instrument Design

The ultimate goal of our test observations was to demonstrate both the spectroscopic and spectropolarimetric capabilities of PGs on-sky. The prototype spectropolarimeter (hereafter PSP) that we developed for the test observations is a simple slit-based dual channel spectropolarimeter. The polarimeter aspect of the instrument is implemented using a rotatable quarter-wave plate (QWP) as a modulator and the PG as an analyzer. The combination of the QWP and the PG make the PSP sensitive to linear polarization (Stokes Q and U). PGs are naturally sensitive to LHC and RHC polarized light (Stokes V), which could be measured if the QWP were removed from the optical train. However, we chose to focus on the measurement of linear polarized light.

The optical design of the PSP was driven by the desire to achieve the maximum possible resolving power using the 16-inch UofT telescope, taking into consideration the $\sim 2.5''$ seeing in Toronto and using an available Apogee Alta U16M 4096×4096 $9 \mu\text{m}$ -pixel CCD camera. The PSP optical system consists of a slit, a collimating lens, a rotatable QWP, the PG, a camera lens and the Alta U16M detector (Figures 6.1 and 6.2), with many parts being off-the-shelf components. In this system, the 0th order spot from the grating falls at the center of the detector and the $m = \pm 1$ orders extend in opposite directions. A spectrum can be obtained with a single frame by adding together the two orders, while a polarized spectrum can be obtained by combining four frames with the QWP changing its position angle by 45° between each observation (§6.4).

We conducted the test observations with the PSP on the UofT 16-inch telescope, an equatorial mounted Cassegrain telescope with an $F/18$ output beam. The PSP was mounted on the rear of the

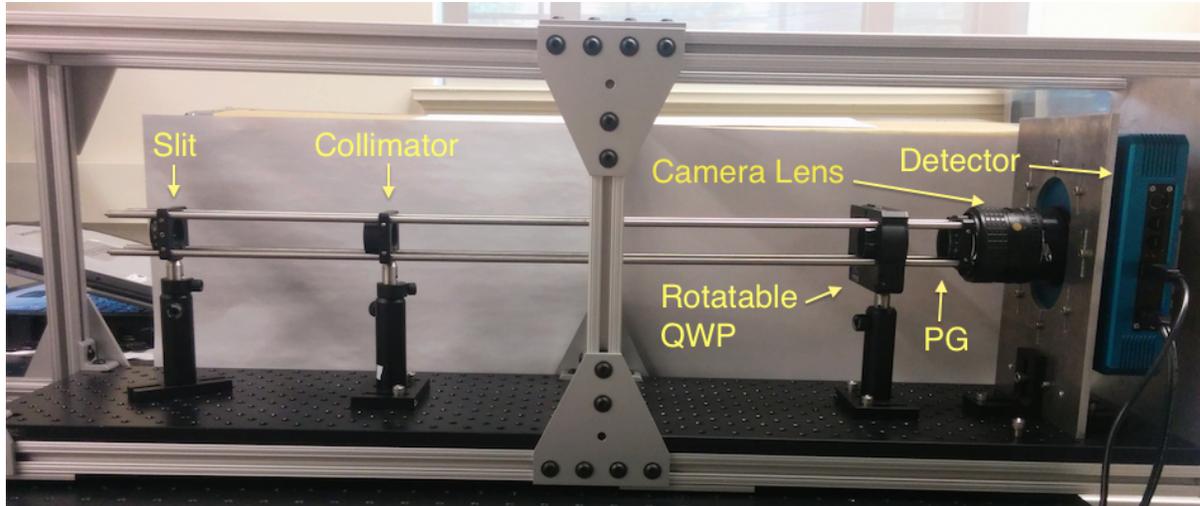


Figure 6.1: A photo of the PSP instrument taken in the laboratory. Light from the telescope enters through a mounting plate (not pictured) to the left of the slit and proceeds through the components to the detector at the far right. The instrument consists of (from left to right): a slit, a collimating lens, a rotatable QWP, a PG, a camera lens and a detector.

Table 6.1. Telescope and PSP Design Parameters

Parameter	Value
Telescope Diameter	404.6 mm (16 inches)
Telescope Focal Ratio	F/18
Telescope Focal Length	7 315.2 mm
Slit Width	0.1 mm (2.8'')
PG Period	5 μ m
PSP Collimator Focal Length	300 mm
PSP beam diameter	16.7 mm
PSP Camera Focal Length	100 mm
PSP Spatial Plate Scale	0.85''/pixel
PSP Single-Pixel Wavelength Coverage	\sim 0.445 nm
Maximum Wavelength on the Detector	\sim 900 nm

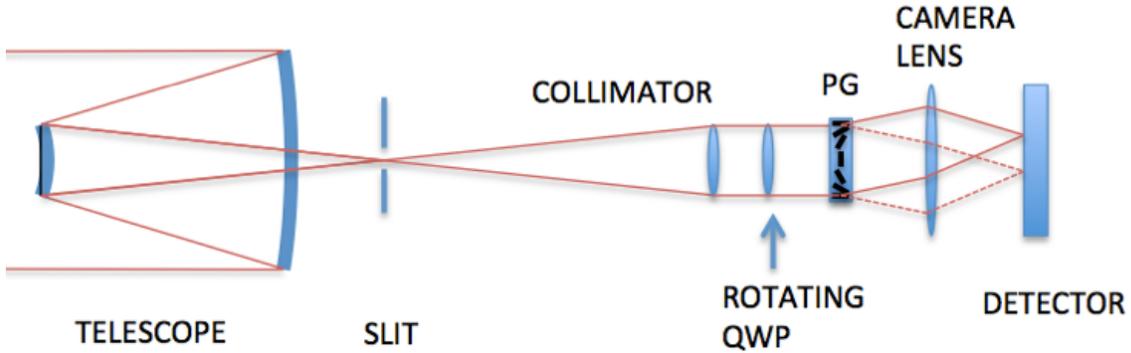


Figure 6.2: A cartoon schematic of the optical layout of the PSP (not to scale).

telescope at the Cassegrain focus (Figure 6.3). At the telescope focus is a $100\ \mu\text{m}$ width slit, which is equivalent to the angular size of $\sim 2.8''$. Light passing through the slit is collimated by a lens of 300 mm focal length to provide a beam size of 16.7 mm, which nearly fully illuminates the 18 mm diameter of the PG. Before reaching the PG, light passes through a QWP mounted in a motorized rotation stage. The QWP is a Thorlabs AQWP10M-980 Quartz-MgF₂ Achromatic Quarter-wave Plate optimized for operation between 690 and 1200 nm. The QWP was mounted with a $10.0^\circ \pm 0.5^\circ$ rotational offset relative to the instrument’s vertical axis (i.e. the North-South axis). After the QWP, light passes through the PG, where it is sorted into the $m = \pm 1$ orders according to its polarization while at the same time being dispersed by wavelength. A slight rotational misalignment of the PG with the vertical axis of the instrument causes the spectra to appear slightly rotated in the raw data, an effect that is compensated for during data processing. After the PG, light then enters the camera lens, a 100 mm, $F/2.8$ Nikon lens, mounted on the Apogee camera. The combination of collimator and camera focal lengths were chosen to maximize the wavelength coverage (i.e. the longest wavelength on the detector for both orders), while sampling the slit with ~ 3 pixels. The design results in a wavelength of ~ 900 nm reaching the top and bottom edges of the detector for the $m=+1$ and $m=-1$ orders, respectively, and a plate scale of $0.85''$ per pixel. At that plate scale, the $2.8''$ slit covers slightly more than 3 pixels and the spectral resolving power between 500 and 900 nm ranges from 300 to 543. The key design parameters of the optical design can be found in Table 6.1.

In practice we found that the telescope was unable to keep a point source centered in the slit even when using the secondary telescope in Figure 6.3 as a guider; as a result, we removed the slit from the optical train for all observations except those of the twilight sky (§6.5). This could also possibly be due to a changing gravity vector throughout the observations, which could introduce flexure-related misalignment. The true spectral resolving power of a given observation will then depend on the spatial extent of the point spread function (PSF), which can be measured from the 0th order spot.

6.4 Data Reduction

The raw PSP data for a single observation consists of a PSF (or rectangular slit for twilight observations) centered on the detector and the two first-order spectra extending vertically (in the spectral direction) to the top and bottom edges of the detector. To save on read-out time, only a 200×4096 pixel subframe



Figure 6.3: A photo of the PSP mounted on the 16-inch UofT telescope.

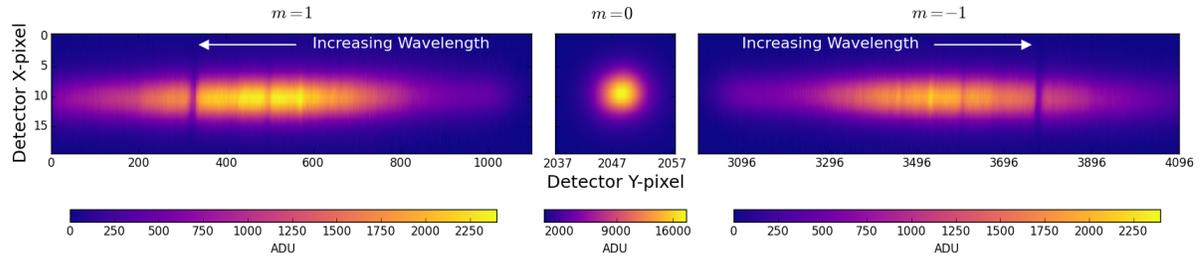


Figure 6.4: An example of a single raw data frame from the PSP, showing the 0th order PSF (center) and two first order spectra (left and right) in a 180 s observation of the polarized standard star HD 183143. The image has been rotated 90° from the detector frame of reference to display the spectrum horizontally. Note that for display purposes the horizontal stretch of the two spectra is different from that of the central PSF. Each order represents an orthogonal linear polarization state. The wavelength range covered by the two spectra is 420–900 nm. The dark bands near the detector y-pixels (on the horizontal axis here) 300 and 3800 correspond to a wavelength of ~ 765 nm and are likely due to Fraunhofer A-band O_2 atmospheric absorption. The bright feature seen near y-pixels 600 and 3500 corresponds to a wavelength of ~ 650 nm and likely represents H- α emission from the star. Indeed, HD 183143 is known to have strong H- α emission (e.g. Chentsov et al. 2003).

containing the spectra is read out for each observation. For each raw frame, two spectra can be extracted, one from each order. The two spectra can then be summed to obtain a total intensity spectra or they can be combined with other observations to obtain a polarized spectrum (§6.4.1). Here we summarize the steps to extract the two spectra from the raw data.

A raw data frame is first dark-subtracted using a dark frame of equal exposure time. The rotational offset of the PG from the vertical axis is then determined empirically for each observation by fitting a straight line to the pixel locations of the peak flux in each horizontal row of the two spectra. Using the observations of HD 183143 (§6.5) we have measured this rotational offset to be $0.61^\circ \pm 0.01^\circ$. The spectra are then vertically aligned by rotating the image about the 0th order PSF, whose center is determined by fitting a 2D Gaussian to the PSF. Figure 6.4 shows an example spectrum after the rotation. The image has been cropped to only show the regions with stellar flux. For observations of the twilight sky, the slit is left in place. In this case, rather than measure the PG rotation offset from the data we instead rotate the spectra by 0.61° .

The flux at each wavelength is measured by calculating the flux in each detector row. For each row, the background level is estimated as the mean flux value in two 10 pixels-wide regions on either side of the spectrum, each centered 50 pixels away from the center of the spectrum (i.e., at a distance of slightly over 10 times the PSF FWHM). The background is subtracted and then the flux in each row is calculated using a weighted sum of the pixel values within the row. The weights are equal to each pixel’s relative contribution to the horizontal component of the PSF, which is assumed to be a 1D Gaussian with a width equal to the horizontal width of the best-fit 2D Gaussian to the 0th order PSF. An error estimate is calculated for each row using photon noise and read noise errors and scaled by the same PSF weights used to extract the flux values. Throughout this work we determine each row’s corresponding wavelength using the grating equation, assuming a grating period of $5 \mu\text{m}$ and the nominal focal length of the camera. Figure 6.5 displays an example of the typical spectra extracted from a single observation with the PSP.

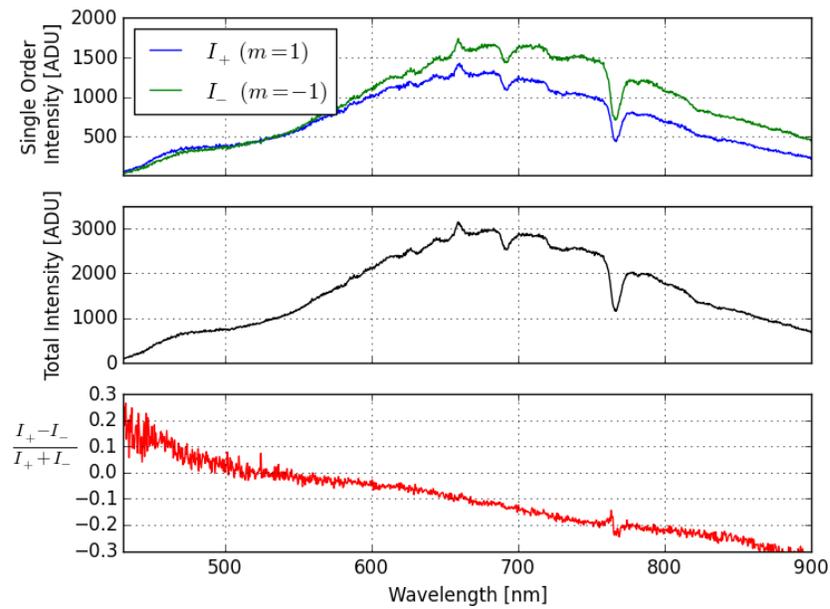


Figure 6.5: The single order intensities (I_+ and I_- ; top), the total intensity (center) and the normalized difference $((I_+ - I_-)/(I_+ + I_-)$; bottom) of a single PSP observation of HD 183143 (§6.5). The decreasing trend in the normalized difference is a constant bias across all frames and is likely due to a slight misalignment of the PG. Because it does not modulate with the rotation of the QWP, it does not register as a polarization signal when combining multiple frames. The small feature in the normalized difference near 775 nm is likely due to a small wavelength offset between the two channels and could be corrected with a more rigorous wavelength calibration.

6.4.1 Linear Polarization Extraction

Linear polarization can be measured by combining four or more observations, with the QWP rotating by 45° between consecutive frames (e.g. $[0^\circ, 45^\circ, 90^\circ, 135^\circ]$). To estimate the input Stokes vector ($S_{in} = [I_{in}, Q_{in}, U_{in}, V_{in}]$) we use a Mueller matrix formalism that describes the PSP's expected behaviour when exposed to polarized light. We note that our sensitivity to circular polarization (Stokes V) is extremely weak and only due to the wave plate's departure from ideal retardance. As mentioned above, PGs are naturally sensitive to circular polarization, but here our primary focus is on linear polarization. Stokes V could be measured in a single observation with the QWP removed by differencing the two spectra.

The Stokes vector reaching the detector in each grating order ($S_{det}^\pm(\lambda)$), can be found as a function of the incident Stokes vector $S_{in}(\lambda)$:

$$S_{det}^\pm(\theta, \lambda) = M_{PG}^\pm \cdot M_{QWP}(\theta, \lambda) \cdot S_{in}(\lambda), \quad (6.1)$$

where M_{PG}^\pm and $M_{QWP}(\theta, \lambda)$ are the Mueller matrices for the PG and QWP, respectively. The ' \pm ' superscript in M_{PG}^\pm represents the Mueller matrix for an ideal PG for each of the two orders:

$$M_{PG}^\pm = \begin{bmatrix} 0.5 & 0.0 & 0.0 & \pm 0.5 \\ 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 \\ -0.5 & 0.0 & 0.0 & \mp 0.5 \end{bmatrix}.$$

The Mueller matrix for the QWP, $M_{QWP}(\theta, \lambda)$, is a standard waveplate Mueller matrix with retardance $\delta(\lambda)$ (expressed as a fraction of a wave) and position angle θ . Thorlabs provides typical retardance values for their QWPs between 690 and 1200 nm. To assess the polarization behaviour of the PG at shorter wavelengths we estimate the QWP's retardance by fitting the manufacturer-provided retardance values to a model waveplate consisting of a quartz plate and a MgF_2 plate. The free parameters in the fit are the thicknesses of the two plates, and the quartz and MgF_2 birefringences were taken from Shields & Ellis (1956) and Dodge (1984), respectively. The best-fit retardance and the manufacturer-provided retardance for the Thorlabs QWP can be seen in Figure 6.6. Measurements of the polarization outside of the ideal wavelength range should be interpreted only as approximate values until a proper calibration of the QWP's retardance is carried out.

Each PSP observation with QWP position angle θ makes a measurement of Stokes $I_{det}^+(\theta, \lambda)$ and $I_{det}^-(\theta, \lambda)$. The expected $I_{det}^\pm(\theta, \lambda)$ can be found by expanding Equation 6.1:

$$I_{det}^\pm(\theta, \lambda) = 0.5I_{in}(\lambda) \pm 0.5[-Q_{in}(\lambda) \sin 2\theta \sin 2\pi\delta(\lambda) + U_{in}(\lambda) \cos 2\theta \sin 2\pi\delta(\lambda) + V_{in}(\lambda) \cos 2\pi\delta(\lambda)]. \quad (6.2)$$

For an ideal quarter-wave plate $\delta(\lambda) = 0.25$ and Equation 6.2 simplifies to:

$$I_{det}^\pm(\theta, \lambda) = 0.5I_{in}(\lambda) \pm 0.5[-Q_{in}(\lambda) \sin 2\theta + U_{in}(\lambda) \cos 2\theta]. \quad (6.3)$$

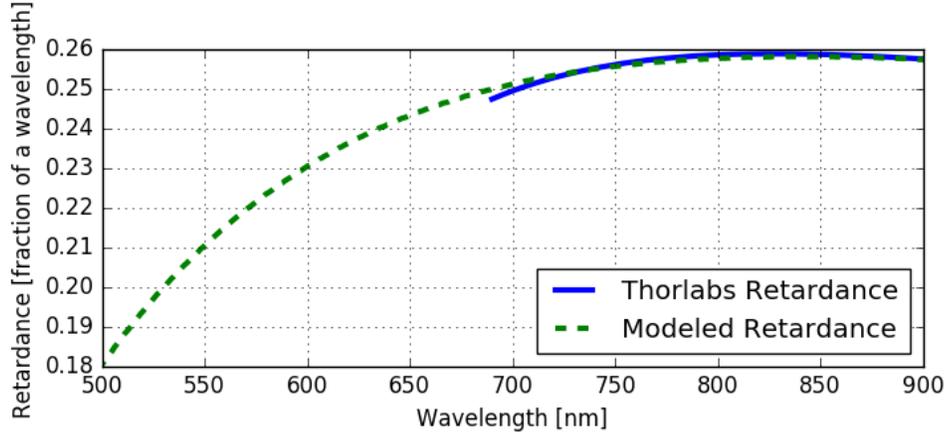


Figure 6.6: The retardance of the Thorlabs QWP as a function of wavelength. To estimate the retardance at wavelengths shorter than 690 nm, we model the QWP as a quartz plate combined with a MgF_2 plate. Our model differs slightly from the Thorlabs-provided values, and diverges from ideal behaviour (i.e. a retardance of 0.25 waves) by more than 10 % at wavelengths shorter than ~ 585 nm.

An ensemble of n observations with QWP position angles $[\theta_1, \dots, \theta_n]$ can then be described as:

$$\begin{bmatrix} I_{det}^+(\theta_1, \lambda) \\ I_{det}^-(\theta_1, \lambda) \\ \vdots \\ I_{det}^+(\theta_n, \lambda) \\ I_{det}^-(\theta_n, \lambda) \end{bmatrix} = M_{meas}(\lambda) \cdot S_{in}(\lambda) = 0.5 \begin{bmatrix} 1 & -\sin 2\theta_1 \sin 2\pi\delta(\lambda) & \cos 2\theta_1 \sin 2\pi\delta(\lambda) & \cos 2\pi\delta(\lambda) \\ 1 & \sin 2\theta_1 \sin 2\pi\delta(\lambda) & -\cos 2\theta_1 \sin 2\pi\delta(\lambda) & -\cos 2\pi\delta(\lambda) \\ \vdots & \vdots & \vdots & \vdots \\ 1 & -\sin 2\theta_n \sin 2\pi\delta(\lambda) & \cos 2\theta_n \sin 2\pi\delta(\lambda) & \cos 2\pi\delta(\lambda) \\ 1 & \sin 2\theta_n \sin 2\pi\delta(\lambda) & -\cos 2\theta_n \sin 2\pi\delta(\lambda) & -\cos 2\pi\delta(\lambda) \end{bmatrix} \cdot \begin{bmatrix} I_{in}(\lambda) \\ Q_{in}(\lambda) \\ U_{in}(\lambda) \\ V_{in}(\lambda) \end{bmatrix}, \quad (6.4)$$

where $M_{meas}(\lambda)$ is called the measurement matrix. Included in θ_n is the 10.0° rotational offset of the QWP relative to the vertical axis. A typical sequence will have a minimum four QWP angles, $[\theta_1, \theta_2, \theta_3, \theta_4] = [0^\circ, 45^\circ, 90^\circ, 135^\circ]$. However, any set of four QWP positions that increase (or decrease) by intervals of 45° will enable sufficient modulation of Equation (6.2) to measure Stokes Q and U . We have therefore left Equation (6.4) in its most general form. The input Stokes vector $S_{in}(\lambda)$ can then be estimated by inverting the measurement matrix and multiplying by the measured I_{det}^\pm vector. Because both the QWP retardance and the source polarization may vary with wavelength, this calculation is performed independently at each wavelength. The covariance matrix of S_{in} can be estimated from the covariance matrix of the measured fluxes $C_{I_{det}^\pm}$:

$$C_{S_{in}} = M^+ C_{I_{det}^\pm} (M^+)^T, \quad (6.5)$$

where M^+ and $(M^+)^T$ are the measurement matrix's pseudo-inverse and its transpose, respectively. Figure 6.8 displays an example of the polarized spectrum derived from observations of HD 183143, a polarized standard star (discussed further in §6.5). The linear polarized intensity P and the position angle (PA) of the linear polarization can be calculated as $P = \sqrt{Q^2 + U^2}$ and $PA = 0.5 \arctan(U/Q)$, respectively.

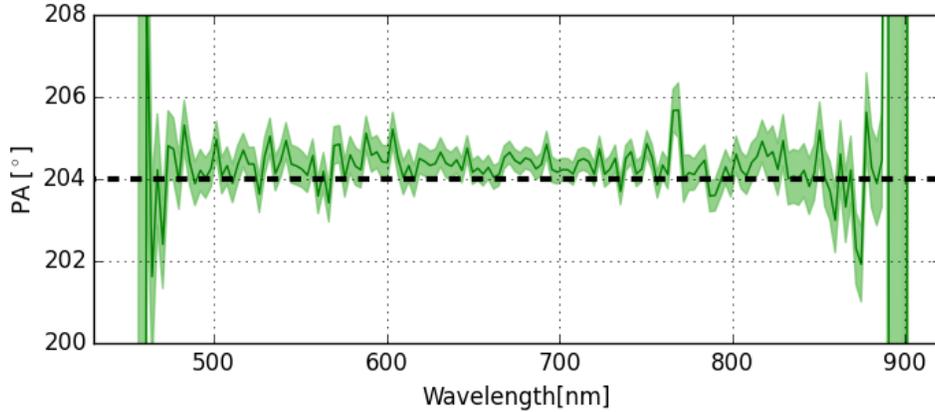


Figure 6.7: The measured position angle of the twilight polarization as a function of wavelength. The expected PA for pure Rayleigh scattering is denoted with the black dashed line. Between 500 nm and 900 nm our measurement is systematically higher by $\sim 0.3^\circ$. This could be due an uncalibrated rotation of the instrument or possibly indicates a departure from our assumption of pure Rayleigh scattering.

6.5 Observational Test Results

Our on-sky tests consist of observations of the polarized twilight sky to test our data reduction routines and verify our ability to measure the PA of linear polarized light; observations of the polarized standard star HD 183143 to measure polarization fraction and position angle as a function of wavelength; and observations of two spectroscopic standard stars to test the PSP’s ability to recover spectra. Table 6.2 summarizes the on-sky observations, which were all carried out on 2016 April 27 Universal Time. Because of telescope tracking errors, keeping a target centered on the slit for multiple exposures in a row became prohibitively difficult and, as a result, the slit was removed for all observations except those of the twilight sky.

6.5.1 Twilight Sky Observations

The twilight sky is a bright, strongly polarized source, predominantly because of Rayleigh scattering in the upper atmosphere. Other minor contributors include aerosol scattering, multiple scattering effects, and secondary illumination from light scattered off clouds, water, and land (e.g. Harrington et al. 2011). Under the assumption of pure single Rayleigh scattering, the polarization angle will be offset 90° from the azimuthal position of the sun at the time of the observations. Thus it serves as a useful target with which to test our ability to recover PA.

Twilight sky observations were obtained during the beginning of the night on 2016 April 27, with the focal plane slit in place. Since the twilight sky fully illuminates the slit, we performed a sum along the spatial direction for each spectral row on the detector, rather than using a weighted sum to extract the two spectra from each observation. The four sets of twilight spectra were then combined, as described in §6.4.1 to obtain a Stokes vector and subsequently a PA measurement at each wavelength (Figure 6.7).

Our measurements recover the expected polarization between 475 and 850 nm, albeit with a systematic offset of roughly 0.3° . This offset could be due to a failure in our assumption of pure Rayleigh scattering, a misalignment within the instrument (e.g. the QWP angle is only known to 0.5°), or both.

Table 6.2. PSP Observations on 2016 April 27 UT

Object	Type	V mag.	Exp. Time (s)	# of Exposures	PSF Width (Pixels)	Airmass
Twilight Sky	PA Calibrator	N/A	60	4	N/A	1.00
HD 183143	Pol. Standard	6.86	180	4	2.9	1.28
HD 93521	Spec. Standard	7.03	240	4	1.9	1.05
HD 106252	Spec. Standard	7.36	120	9	2.7	1.32

Table 6.3. Polarization Fraction and Position Angle Measurements of HD 183143

Band	Measured		Literature Values (Dolan & Tapia 1986)	
	P/I [%]	PA [°]	P/I [%]	PA [°]
V	5.0 ± 0.8	176.5 ± 3.3	6.08	179.8
R	5.7 ± 0.3	178.6 ± 1.6	5.74	179.9
I	5.1 ± 0.5	179.2 ± 2.1	5.08	179.8

Note. — The errors on the measurements represent the standard deviation across the band.

6.5.2 Polarized Standard - HD 183143

In order to assess our ability to measure polarization fraction we observed the polarized standard star HD 183143, which has a relatively high polarization fraction of $\sim 6\%$ (Heiles 2000). The star was observed for a standard four QWP-position sequence (Table 6.2). The two spectra were first extracted from each raw observation and then binned by three pixels (which is roughly equivalent to the Gaussian σ measured from the central PSF) before being combined to obtain a Stokes vector at each wavelength. The binning reduces the noise without significantly degrading the spectral resolution. The measured Stokes I , Q and U , alongside the polarization fraction and position angle are shown in Figure 6.8. Polarization fraction measurements are a positive definite quantity and are biased by the errors in the Q and U measurements. To estimate the true polarization fraction the measured polarization fraction must be debiased (e.g. Simmons & Stewart 1985; Vidal et al. 2016). To first order the true polarization fraction can be estimated as $P_{frac,true} = \sqrt{P_{frac,measured}^2 - \sigma_{P_{frac}}^2}$. Using the errors in Q and U , we estimate $\sigma_{P_{frac}}$ and find the bias to be on the order of 1% of the fractional polarization signal. Note that fractional polarization displayed in Figure 6.8 has not been debiased. However, because the estimate of the bias is small relative to our signal it does not significantly affect our conclusions.

Comparing our measurements to the V , R and I broadband measurements presented by Dolan & Tapia (1986), we find that we recover both the polarization fraction and the spectral dependence in the R and I bands (i.e. between 600 and 900 nm). Our measurements at wavelengths shorter than 600 nm underestimate the polarization fraction relative to theirs. However, at these wavelengths our measurements rely on modeled QWP retardance values and, consequently, are less reliable than those at longer wavelengths. Table 6.3 contains our band-averaged linear polarized fraction and PA for the V , R and I bands. We also provide the values measured by Dolan & Tapia (1986), who do not quote specific errors for their measurements, but suggest that they achieve $\sigma_{P/I} \approx 0.02\%$ and $\sigma_{PA} = 0.1^\circ$. We find that our measurements in the R and I bands are consistent with theirs, but that our V band

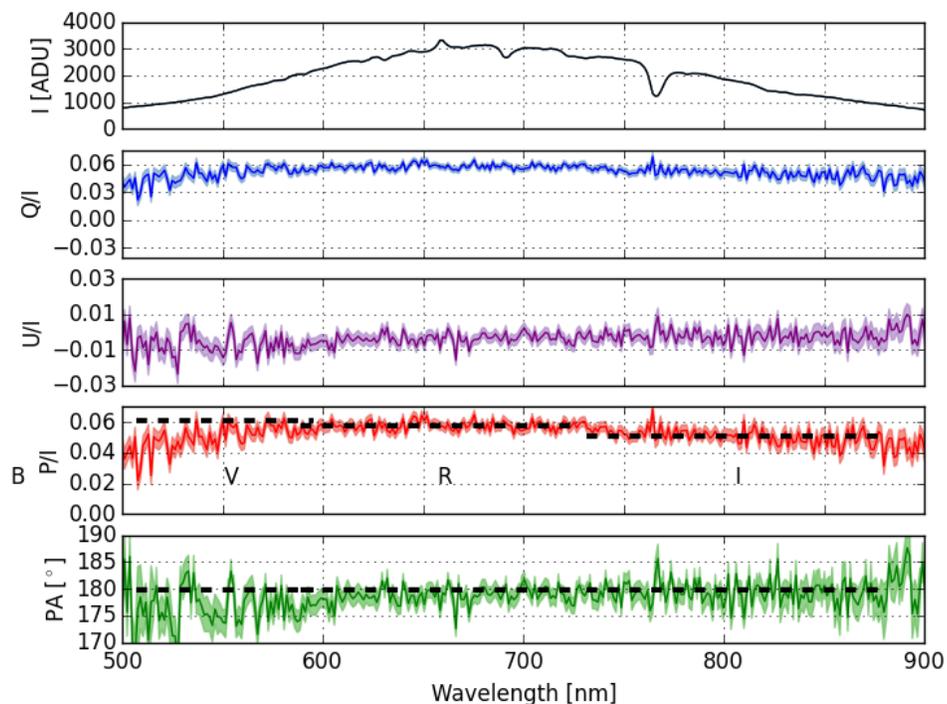


Figure 6.8: The polarized spectrum of HD 183143. The image displays the total intensity spectrum (Stokes I ; top), the normalized Stokes Q spectrum (Q/I ; center-top), the normalized Stokes U spectrum (U/I ; center), the polarization fraction (P/I ; center-bottom) and the PA as a function of wavelength (bottom). The shaded regions in each plot represent the $1 - \sigma$ errors and the thick lines are the measurement. The black dotted lines display the broadband polarized fraction and position angle measurements of HD 183143 presented by Dolan & Tapia (1986). The total intensity spectrum shown here is the raw measured spectrum and has not yet been calibrated for spectral throughput. Note that the polarization fraction has not been debiased. However, the errors estimated for the Q and U signal result in a bias of less than 1% of the signal and therefore do not significantly affect our conclusions.

measurements are more than 1-sigma below theirs. The poor performance at shorter wavelengths is likely due to the unknown (and non-ideal) retardance of the QWP at shorter wavelengths.

6.5.3 Spectrophotometric Calibration

A flux-calibrated spectrum can be obtained for a given observation by multiplying the measured spectrum by the spectral response of the instrument, which incorporates the wavelength-varying throughput of the instrument as well as atmospheric throughput. Here we estimate the spectral response of the PSP by dividing the measured spectrum of a spectroscopic standard by its calibrated spectrum obtained from the CALSPEC database¹. This method has the advantages that it is simple and can be used to correct for atmospheric absorption, but is not robust to changing atmospheric conditions or changes in airmass.

We calibrate the spectral response of the PSP using observations of the spectroscopic standard star HD 93521 (Table 6.2). The two spectra were extracted from each raw data file and then summed to obtain the total intensity spectrum. Throughout the observations the total flux of the observations varied frame by frame. To compensate for this, each spectrum (except the first) was multiplied by a scaling factor that brought the total integrated flux from 500 nm to 900 nm to the same value as the first spectrum of the set. After scaling, the mean spectrum was determined by averaging all 4 observations. To compare the CALSPEC spectrum to our measure spectrum, the CALSPEC spectrum was first interpolated to the wavelength values of the PSP spectrum, and then smoothed by a Gaussian with a FWHM equal to that of the observations. Finally, the spectral response was calculated by dividing the interpolated and smoothed CALSPEC spectrum by the measured spectrum.

We validated this calibration procedure by applying the spectral response to observations of the spectroscopic standard HD 106252. The mean total intensity spectrum of all the HD 106252 observations was calculated in the same fashion as for HD 93521. We then multiplied the measured spectrum by the spectral response derived from HD 93521, in order to obtain a flux calibrated spectrum. The spectrum of HD 106252 can be seen alongside the CALSPEC spectrum in Figure 6.9. For comparison, the CALSPEC spectrum of HD 106252 has been smoothed with a Gaussian with a FWHM equal to that of the observations.

In general, the flux-calibrated spectrum matches the CALSPEC spectrum well. The PSP spectrum recovers nearly all of the larger-scale spectral features, including the deep absorption feature near 450 nm, as well as the overall spectral shape, except at the bounds of the wavelength coverage. There is a slight offset in wavelength between the two spectra, which we attribute to our simple wavelength assignment to each spectral bin. A large emission feature in the PSP spectrum appears near 765 nm that is not apparent in the CALSPEC spectrum. This feature coincides with a deep absorption band in the raw data (Figure 6.9) and is likely the $O_2 - A$ Fraunhofer absorption band. The feature seen in the calibrated spectrum could be due to either a slight offset in wavelength between the data and the spectral response, or to a different airmass between the two observations, which would lead to different levels of absorption.

6.6 Conclusion

Polarization gratings are an exciting new technology that have many applications in astronomical instrumentation. Because they act as a diffraction grating, they can be used in place of a conventional

¹<http://www.stsci.edu/hst/observatory/crds/calspec.html>

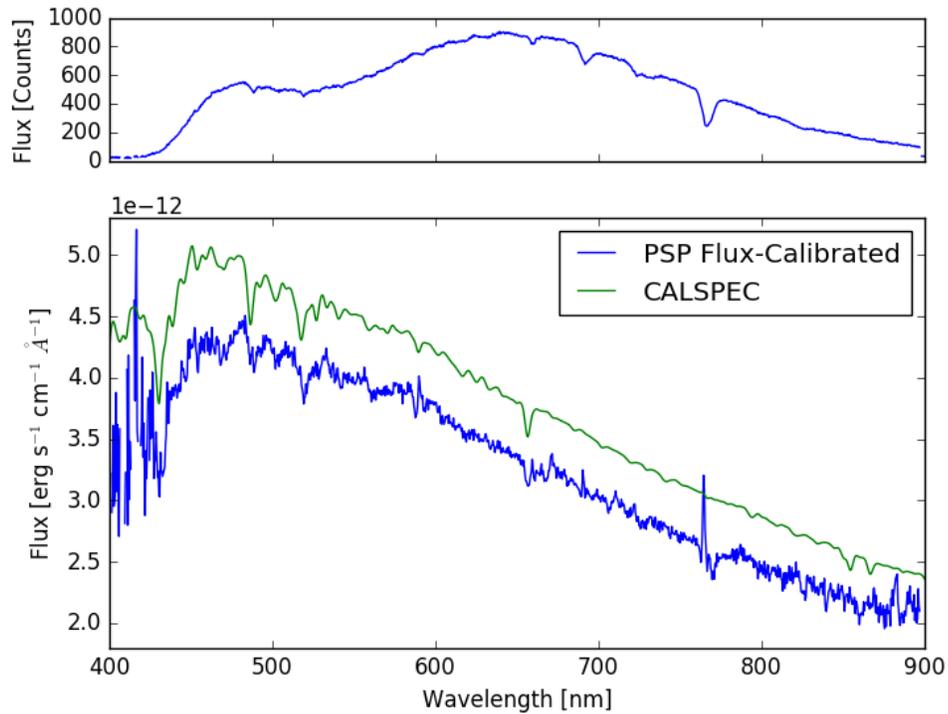


Figure 6.9: *Top*: The mean total intensity spectrum of HD 106252. *Bottom*: The flux-calibrated spectrum of HD 106252 as measured by the PSP. The CALSPEC spectrum has been overplotted, offset upwards by 0.5×10^{-12} erg s⁻¹ cm⁻¹ Å⁻¹ for display purposes.

ruled grating in an optical system to enable high-throughput broadband spectropolarimetry. Their high-throughput nature has applications across many domains of astronomical study, for instance, time-domain observations of faint sources. Their polarimetric capabilities can be used in fields ranging from the study of supernovae to the study of exoplanets and circumstellar disks.

We have used a PG optimized for visible wavelengths to make observations of both polarized sources and standard spectroscopic calibrators. Our observations of the twilight sky were able to recover linear polarized fraction. We were also able to measure the linear polarized fraction and position angle of a polarization standard star, where both measurements were consistent with literature values. Finally we were successful in recovering the visible spectrum of a spectroscopic standard star.

An important step in characterizing a new polarimeter is testing its response to unpolarized light. This characterization can reveal instrumental biases (such as instrumental polarization) and can serve as an absolute measure of performance. Unfortunately due to time constraints and difficult observing conditions in Toronto, this test was not completed.

Regardless, the observations presented here serve to demonstrate that PGs can be used both for total intensity spectroscopic observations as well as spectropolarimetric observations. Slight discrepancies between previous measurements and our own can be attributed to rudimentary calibration of the instrument or non-ideal components, and there are no indications of significant systematics being introduced by the PG. Eventually we hope to extend this work into the near-infrared with a PG designed for high throughput between 600 nm and 1600 nm.

Chapter 7

Conclusions

Polarimetry is a useful technique for the study of extrasolar planetary systems. It can be used to improve contrast in scattered-light imaging of debris disks and provide useful information about grain composition. Hopefully in the near future polarized thermal emission will be detected from the atmospheres of extrasolar planets, opening up a new realm of planet characterization. This thesis has worked towards taking advantage of the unique information afforded by the measurement of polarized light to further our knowledge-base and understanding of extrasolar planets. The following sections provide a brief summary of the work presented herein and a look to the future.

7.1 Thesis Summary

The work presented in this thesis consists of 5 major components, summarized below:

- The reduction of GPI polarimetry data involves a complicated set of procedures that requires careful thought in order to achieve the best results. By optimizing the performance of the instrument and our data analysis techniques we are able to achieve greater sensitivities and, consequently to image fainter debris disks (and perhaps one day, polarized planets). In Chapter 2 we presented several modifications to the data reduction process that have allowed us to improve the sensitivity of our polarized measurements. The modifications include a polarized flat-field correction, a weighted PSF scheme that improves polarization datacube assembly from the previous box aperture extraction, and the ability to measure and subtract the instrumental polarization. In addition we present H-band polarimetric contrast measurements using examples from the GPIES campaign. These performance measurements can serve as benchmarks for future observations in the evaluation of the performance of both GPI and future instruments.
- β Pic is an interesting system to study because many previous observations have revealed the system to be dynamically active and that it likely has a complicated past. As such, it was observed numerous times throughout GPI commissioning and also during an orbital monitoring campaign, in order to image both the disk in polarized intensity and to build up a time baseline of observations of the planet β Pic *b*. Our observations of the disk reached down to previously inaccessible inner working angles and revealed a slightly inclined disk with strong forward scattering dust grains. Our observations of β Pic *b* were able to provide an additional ~ 1.5 years of astrometric measurements

on top of previous measurements from other instruments, which we used to calculate an updated orbital fit. The combination of the disk model and the new orbital parameters allowed us to examine the dynamical relationship between the disk and the planet near the planet’s semi-major axis. We find that the disk morphology is incompatible with sculpting by β Pic *b* at the separations probed by GPI. One possible explanation is the presence of a second planet in the system, located between β Pic *b* and the inner edge of the disk, though this scenario requires further study.

- The HD 157587 system was imaged as part of the GPIES survey, following the first resolved images from HST. The GPI data is able to reach a smaller inner working angle than the HST data and reveals the disk to be a bright ring-like structure. We associate a strong brightness asymmetry along the projected minor axis with strong forward scattering, as has been recently seen in a number of other GPI disks. A fainter brightness asymmetry along the projected major axis is consistent with being a feature caused by a stellocentric offset of the geometric centre of the disk, possibly due by secular forcing by an unseen eccentric planet. Different age indicators for HD 157587 simultaneously indicate a stellar age above and below 1 Gyr. If this disk is truly older than 1 Gyr, the study of the debris disk may be the only method for constraining a potential perturbing planet’s characteristics.
- PGs are useful for astronomical observations because they can be used for high-bandwidth, high-throughput spectropolarimetric observations. Our laboratory tests confirmed that the gratings behave as expected, providing a > 90 % diffraction efficiency between 600 and 850 nm, and staying above 85 % out to 900 nm. In addition, we confirmed that they are able to modulate incident polarization with a modulation efficiency of ~ 99 %. We discussed several possible spectropolarimeter design setups that include polarization gratings as a dispersing mechanism and consider several science cases.
- In order to verify the performance of polarization gratings in an astronomical setting, we designed and built a prototype spectropolarimeter instrument that we mounted on the University of Toronto’s 16-inch campus telescope. Using this instrument we carried out observations of polarized sources and spectroscopic standard targets and demonstrated the spectropolarimetric and spectroscopic capabilities of polarization gratings on-sky. The results of our tests indicate that PGs are indeed a powerful tool and can make measurements of both astronomical spectra in both total intensity and polarized intensity.

7.2 Future work

7.2.1 GPI Data Analysis and Debris Disk Observations

While great progress has been made in reducing systematics in GPI data, there is still much work to be done. For example, the modelling and removal of persistence could significantly improve the efficiency of the ‘snapshot’ disk observations in the GPIES campaign. The effects limiting the contrast in the inner regions of GPI polarization datacubes are still poorly understood, but thought to be due to imperfect instrumental polarization subtraction. Improving the sensitivity at small inner working angles via improved instrumental polarization subtraction will allow GPI to be pushed to its limits, and open doors to even more exciting science.

For both the debris disks studied here, only the polarization intensity has been recovered; as a result, we have completed no analysis of the grain properties of the disks. It appears that measuring the total intensity, and, hence, the polarization fraction may be out of reach for these disks using current technology. However, different dust grain compositions are expected to have different behaviours as a function of wavelength in the NIR; thus, combining multi-band polarimetric observations may be able to overcome this impediment. With this in mind, we have obtained yet-to-be-analyzed J- and K1-band observations of the β Pic disk that we plan to use to constrain the dust properties in the inner regions of the disk.

Over the next few years debris disk observations will continue through both the GPIES campaign and the debris disk LLP. Both of these efforts are likely to provide important insights into the structure and formation of planetary systems. Once these programs are completed, a potential future path for GPI is to move to the Gemini North telescope to gain access to the stars in the northern hemisphere, a simple move thanks to the similar design of the two telescopes. A move to Gemini North would also provide the opportunity for instrument upgrades.

7.2.2 Polarization Gratings

PGs are now ready to be used on-sky in dedicated instruments on large telescopes. In the near future we expect to extend our efforts with these gratings into the NIR regime. One design we have considered is a grating optimized for high-throughput between 600 nm and 1600 nm. This grating would be useful for the spectropolarimetry of directly imaged planets, whose degree of polarizations are expected to peak in the J- and H-bands, and the broad wavelength coverage of such a grating will be useful for the quick measurement of redshift of distant objects.

The thin-film nature combined with the diffractive properties of PGs make them convenient for drop-in instrument upgrades, as they can be placed at the same location as current diffraction gratings. One potential instrument upgrade could be the placement of a polarization grating in the prism wheel of the NIRC2 camera on the 10-m Keck II telescope. Combined with the Keck AO system, this setup could be used to perform the first spectropolarimetric measurements of currently known directly imaged exoplanets that have already been imaged by NIRC2. Preliminary discussions with the Keck community have indicated significant interest in such an upgrade and we hope to advance this project in the near future with detailed design studies and, eventually, implementation.

Bibliography

- Ahmic, M., Croll, B., & Artymowicz, P. 2009, *ApJ*, 705, 529
- Apai, D., Schneider, G., Grady, C. A., Wyatt, M. C., Lagrange, A.-M., Kuchner, M. J., Stark, C. J., & Lubow, S. H. 2015, *ApJ*, 800, 136
- Asiain, R., Figueras, F., & Torra, J. 1999, *A&A*, 350, 434
- Augereau, J. C., Nelson, R. P., Lagrange, A. M., Papaloizou, J. C. B., & Mouillet, D. 2001, *A&A*, 370, 447
- Aumann, H. H. 1985, *PASP*, 97, 885
- Aumann, H. H., et al. 1984, *ApJ*, 278, L23
- Baraffe, I., Chabrier, G., Barman, T. S., Allard, F., & Hauschildt, P. H. 2003, *A&A*, 402, 701
- Benz, W. 2000, *Space Science Reviews*, 92, 279
- Berdyugina, S. V., Berdyugin, A. V., Fluri, D. M., & Piirola, V. 2008, *ApJ*, 673, L83
- . 2011, *ApJ*, 728, L6
- Bertelli, G., Girardi, L., Marigo, P., & Nasi, E. 2008, *A&A*, 484, 815
- Bertelli, G., Nasi, E., Girardi, L., & Marigo, P. 2009, *A&A*, 508, 355
- Beust, H., & Morbidelli, A. 1996, *Icarus*, 120, 358
- Beust, H., et al. 2014, *A&A*, 561, A43
- Beuzit, J.-L., et al. 2008, in *Proc. SPIE*, Vol. 7014, *Ground-based and Airborne Instrumentation for Astronomy II*, 701418
- Binks, A. S., & Jeffries, R. D. 2014, *MNRAS*, 438, L11
- Blondel, P. F. C., & Djie, H. R. E. T. A. 2006, *A&A*, 456, 1045
- Boccaletti, A., et al. 2015, *Nature*, 526, 230
- Bohren, C. F., & Huffman, D. R. 2004, *Absorption and scattering of light by small particles* (Weinheim, Germany: Wiley), iD: 8996758 (UTL catalogue ckey)
- Bonnefoy, M., et al. 2014, *A&A*, 567, L9

- Bott, K., Bailey, J., Kedziora-Chudczer, L., Cotton, D. V., Lucas, P. W., Marshall, J. P., & Hough, J. H. 2016, *MNRAS*, 459, L109
- Bouchy, F., et al. 2005, *A&A*, 444, L15
- Brauer, F., Dullemond, C. P., & Henning, T. 2008, *A&A*, 480, 859
- Burrows, C. J., Krist, J. E., Stapelfeldt, K. R., & WFPC2 Investigation Definition Team. 1995, in *Bulletin of the American Astronomical Society*, Vol. 27, American Astronomical Society Meeting Abstracts, 1329
- Canovas, H., Ménard, F., de Boer, J., Pinte, C., Avenhaus, H., & Schreiber, M. R. 2015, *A&A*, 582, L7
- Casagrande, L., Schönrich, R., Asplund, M., Cassisi, S., Ramírez, I., Meléndez, J., Bensby, T., & Feltzing, S. 2011, *A&A*, 530, A138
- Cassan, A., et al. 2012, *Nature*, 481, 167
- Chambers, J. E. 1999, *MNRAS*, 304, 793
- Chauvin, G., Lagrange, A.-M., Dumas, C., Zuckerman, B., Mouillet, D., Song, I., Beuzit, J.-L., & Lowrance, P. 2004, *A&A*, 425, L29
- Chauvin, G., et al. 2012, *A&A*, 542, A41
- Chentsov, E. L., Ermakov, S. V., Klochkova, V. G., Panchuk, V. E., Bjorkman, K. S., & Miroshnichenko, A. S. 2003, *A&A*, 397, 1035
- Chereul, E., Crézé, M., & Bienaymé, O. 1999, *A&AS*, 135, 5
- Chilcote, J., et al. 2015, *ApJ*, 798, L3
- Chilcote, J. K., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8446, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 8
- Choquet, É., et al. 2016, *ApJ*, 817, L2
- Crawford, G. P., Eakin, J. N., Radcliffe, M. D., Callan-Jones, A., & Pelcovits, R. A. 2005, *Journal of Applied Physics*, 98, 123102
- Crifo, F., Vidal-Madjar, A., Lallement, R., Ferlet, R., & Gerbaldi, M. 1997, *A&A*, 320, L29
- Currie, T., Thalmann, C., Matsumura, S., Madhusudhan, N., Burrows, A., & Kuchner, M. 2011, *ApJ*, 736, L33
- Dawson, R. I., Murray-Clay, R. A., & Fabrycky, D. C. 2011, *ApJ*, 743, L17
- de Kok, R. J., Stam, D. M., & Karalidi, T. 2011, *ApJ*, 741, 59
- De Rosa, R. J., et al. 2015a, *ApJ*, 814, L3
- . 2015b, *ApJ*, 814, L3
- Debes, J. H., Weinberger, A. J., & Kuchner, M. J. 2009, *ApJ*, 702, 318

- Dent, W. R. F., et al. 2014, *Science*, 343, 1490
- Diolaiti, E., Bendinelli, O., Bonaccini, D., Close, L., Currie, D., & Parmeggiani, G. 2000, *A&AS*, 147, 335
- Dodge, M. J. 1984, *Appl. Opt.*, 23, 1980
- Dohnanyi, J. S. 1969, *J. Geophys. Res.*, 74, 2531
- Dolan, J. F., & Tapia, S. 1986, *PASP*, 98, 792
- Dominik, C., Blum, J., Cuzzi, J. N., & Wurm, G. 2007, *Protostars and Planets V*, 783
- Draper, Z. H., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 4
- Draper, Z. H., et al. 2016, *ArXiv e-prints*
- Eiroa, C., et al. 2013, *A&A*, 555, A11
- Escuti, M., & Oh, C. 2010, Multi-layer achromatic liquid crystal polarization gratings and related fabrication methods, *uS Patent App. 12/596,176*
- Escuti, M. J., Oh, C., Snchez, C., Bastiaansen, C., & Broer, D. J. 2006, in , 630207–630207–11
- Esposito, T. M., Fitzgerald, M. P., Graham, J. R., & Kalas, P. 2014, *ApJ*, 780, 25
- Esposito, T. M., et al. 2016, *ArXiv e-prints*
- Feltzing, S., Holmberg, J., & Hurley, J. R. 2001, *A&A*, 377, 911
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
- Foreman-Mackey, D., Price-Whelan, A., Ryan, G., Emily, Smith, M., Barbary, K., Hogg, D. W., & Brewer, B. J. 2014, *triangle.py v0.1.1*
- Fosalba, P., Lazarian, A., Prunet, S., & Tauber, J. A. 2002, *The Astrophysical Journal*, 564, 762
- Gagné, J., Lafrenière, D., Doyon, R., Malo, L., & Artigau, É. 2014, *ApJ*, 783, 121
- Gledhill, T. M., Scarrott, S. M., & Wolstencroft, R. D. 1991, *MNRAS*, 252, 50P
- Golimowski, D. A., et al. 2006, *AJ*, 131, 3109
- Graham, J. R., Fitzgerald, M. P., Kalas, P., & Clampin, M. 2013, in *American Astronomical Society Meeting Abstracts*, Vol. 221, American Astronomical Society Meeting Abstracts #221, 324.03
- Graham, J. R., Kalas, P. G., & Matthews, B. C. 2007, *ApJ*, 654, 595
- Green, R. M. 1985, *Spherical astronomy* (Cambridge: Cambridge University Press), iD: 3140285 (UTL catalogue ckey); Includes index.
- Haisch, Jr., K. E., Lada, E. A., & Lada, C. J. 2001, *ApJ*, 553, L153

- Harrington, D. M., Kuhn, J. R., & Hall, S. 2011, *Publications of the Astronomical Society of the Pacific*, 123, 799
- Hartung, M., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9148, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 0
- Heap, S. R., Lindler, D. J., Lanz, T. M., Cornett, R. H., Hubeny, I., Maran, S. P., & Woodgate, B. 2000, *ApJ*, 539, 435
- Hedman, M. M., & Stark, C. C. 2015, *ApJ*, 811, 67
- Heiles, C. 2000, *The Astronomical Journal*, 119, 923
- Henry, L. G., & Greenstein, J. L. 1941, *ApJ*, 93, 70
- Hines, D. C., et al. 2007, *ApJ*, 671, L165
- Holmberg, J., Nordström, B., & Andersen, J. 2009, *A&A*, 501, 941
- Hung, L.-W., et al. 2015, *ApJ*, 815, L14
- Hung, L.-W., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, *MNRAS*, 315, 543
- Ingraham, P., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 7
- Jensen-Clem, R., et al. 2016, *ApJ*, 820, 111
- Johansen, A., Blum, J., Tanaka, H., Ormel, C., Bizzarro, M., & Rickman, H. 2014, *Protostars and Planets VI*, 547
- Johansen, A., Oishi, J. S., Mac Low, M.-M., Klahr, H., Henning, T., & Youdin, A. 2007, *Nature*, 448, 1022
- Kalas, P., Duchene, G., Fitzgerald, M. P., & Graham, J. R. 2007, *ApJ*, 671, L161
- Kalas, P., Graham, J. R., & Clampin, M. 2005, *Nature*, 435, 1067
- Kalas, P., Graham, J. R., Clampin, M. C., & Fitzgerald, M. P. 2006, *ApJ*, 637, L57
- Kalas, P., Graham, J. R., Fitzgerald, M. P., & Clampin, M. 2013, *ApJ*, 775, 56
- Kalas, P., & Jewitt, D. 1995, *AJ*, 110, 794
- Kalas, P. G., et al. 2015, *ApJ*, 814, 32
- Kennedy, G. M., & Wyatt, M. C. 2010, *MNRAS*, 405, 1253
- Kenyon, S. J., & Bromley, B. C. 2002, *AJ*, 123, 1757
- . 2004, *AJ*, 127, 513

—. 2008, *ApJS*, 179, 451

—. 2010, *ApJS*, 188, 242

Kiefer, F., Lecavelier des Etangs, A., Boissier, J., Vidal-Madjar, A., Beust, H., Lagrange, A.-M., Hébrard, G., & Ferlet, R. 2014, *Nature*, 514, 462

Kim, J., Oh, C., Serati, S., & Escuti, M. J. 2011, *Appl. Opt.*, 50, 2636

Kimura, H., Kolokolova, L., & Mann, I. 2006, *A&A*, 449, 1243

Komanduri, R. K., Lawler, K. F., & Escuti, M. J. 2013, *Opt. Express*, 21, 404

Konopacky, Q. M., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 84

Krist, J. E., Stapelfeldt, K. R., Bryden, G., & Plavchan, P. 2012, *AJ*, 144, 45

Krist, J. E., et al. 2010, *AJ*, 140, 1051

Krivova, N. A., Krivov, A. V., & Mann, I. 2000, *ApJ*, 539, 424

Kuhn, J. R., Potter, D., & Parise, B. 2001, *ApJ*, 553, L189

Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, 660, 770

Lagage, P. O., & Pantin, E. 1994, *Nature*, 369, 628

Lagrange, A.-M., et al. 2009, *A&A*, 493, L21

—. 2012a, *A&A*, 546, A38

—. 2012b, *A&A*, 542, A40

Langlois, M., et al. 2014, in *Proc. SPIE*, Vol. 9147, *Ground-based and Airborne Instrumentation for Astronomy V*, 91471R

Larkin, J. E., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 1

Lee, E. J., & Chiang, E. 2016, *ArXiv e-prints*

Liseau, R., et al. 2008, *A&A*, 480, L47

—. 2010, *A&A*, 518, L132

Long, K. S., Baggett, S. M., MacKenty, J. W., & Riess, A. G. 2012, in *Proc. SPIE*, Vol. 8442, *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*, 84421W

Lucas, P. W., Hough, J. H., Bailey, J. A., Tamura, M., Hirst, E., & Harrison, D. 2009, *MNRAS*, 393, 229

Macintosh, B., et al. 2014, *Proceedings of the National Academy of Science*, 111, 12661

—. 2015, *Science*, 350, 64

- Maire, J., et al. 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7735, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 31
- Maire, J., et al. 2012a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8451, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 3
- Maire, J., et al. 2012b, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8451, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 3
- Malo, L., Doyon, R., Lafrenière, D., Artigau, É., Gagné, J., Baron, F., & Riedel, A. 2013, *ApJ*, 762, 88
- Mamajek, E. E., & Bell, C. P. M. 2014, *MNRAS*, 445, 2169
- Mamajek, E. E., Quillen, A. C., Pecaut, M. J., Moolekamp, F., Scott, E. L., Kenworthy, M. A., Collier Cameron, A., & Parley, N. R. 2012, *AJ*, 143, 72
- Maness, H. L., et al. 2009, *ApJ*, 707, 1098
- Marley, M. S., & Sengupta, S. 2011, *MNRAS*, 417, 2874
- Marois, C., Doyon, R., Racine, R., & Nadeau, D. 2000, *PASP*, 112, 91
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, 641, 556
- Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafrenière, D., & Doyon, R. 2008, *Science*, 322, 1348
- Marshall, J. P., et al. 2014, *A&A*, 565, A15
- Matthews, B. C., & Kavelaars, J. 2016, *Space Sci. Rev.*
- Matthews, B. C., Krivov, A. V., Wyatt, M. C., Bryden, G., & Eiroa, C. 2014, *Protostars and Planets VI*, 521
- Mayor, M., & Queloz, D. 1995, *Nature*, 378, 355
- McDonald, I., Zijlstra, A. A., & Boyer, M. L. 2012, *MNRAS*, 427, 343
- McManamon, P. F., Bos, P. J., Escuti, M., Heikenfeld, J., Serati, S., Xie, H., & Watson, E. 2009, *Proceedings of the IEEE*, 97, 1078
- Millar-Blanchaer, M., Moon, D.-S., Graham, J. R., & Escuti, M. 2014, in , 91514I–91514I–8
- Millar-Blanchaer, M. A., et al. 2015, *ApJ*, 811, 18
- Milli, J., Mouillet, D., Lagrange, A.-M., Boccaletti, A., Mawet, D., Chauvin, G., & Bonnefoy, M. 2012, *A&A*, 545, A111
- Milli, J., et al. 2014, *A&A*, 566, A91
- Moon, D.-S., Sivanandam, S., Kutyrev, A. S., Jr., S. H. M., & Graham, J. R. 2014, in
- Moro-Martín, A., et al. 2015, *ApJ*, 801, 143

- Morzinski, K. M., et al. 2015, *ApJ*, 815, 108
- Mouillet, D., Larwood, J. D., Papaloizou, J. C. B., & Lagrange, A. M. 1997, *MNRAS*, 292, 896
- Mustill, A. J., & Wyatt, M. C. 2009a, *MNRAS*, 399, 1403
- . 2009b, *MNRAS*, 399, 1403
- Nesvold, E. R., & Kuchner, M. J. 2015, *ApJ*, 815, 61
- Nesvold, E. R., Kuchner, M. J., Rein, H., & Pan, M. 2013, *ApJ*, 777, 144
- Nielsen, E. L., et al. 2014, *ApJ*, 794, 158
- Oh, C., & Escuti, M. J. 2007, in , 668211–668211–12
- Oh, C., & Escuti, M. J. 2008, *Opt. Lett.*, 33, 2287
- Oh, C., Kim, J., Muth, J., Serati, S., & Escuti, M. 2010, *Photonics Technology Letters, IEEE*, 22, 200
- Okamoto, Y. K., et al. 2004, *Nature*, 431, 660
- Packham, C., Escuti, M., Ginn, J., Oh, C., Quijano, I., & Boreman, G. 2010, *PASP*, 122, 1471
- Packham, C., Jones, T. J., Warner, C., Krejny, M., Shenoy, D., Vonderharr, T., Lopez-Rodriguez, E., & DeWahl, K. 2012, in , 84463R–84463R–12
- Padgett, D., & Stapelfeldt, K. 2015, in *Proceedings of the International Astronomical Union, Vol. 10, Young Stars & Planets Near the Sun*, 175–178
- Pan, M., & Schlichting, H. E. 2012, *ApJ*, 747, 113
- Panić, O., et al. 2013, *MNRAS*, 435, 1037
- Perrin, M. D., Graham, J. R., & Lloyd, J. P. 2008, *PASP*, 120, 555
- Perrin, M. D., et al. 2014a, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 3
- Perrin, M. D., et al. 2014b, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 3
- Perrin, M. D., et al. 2014c, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 3
- Perrin, M. D., et al. 2015, *ApJ*, 799, 182
- Petigura, E. A., Howard, A. W., & Marcy, G. W. 2013, *Proceedings of the National Academy of Science*, 110, 19273
- Pietrinferni, A., Cassisi, S., Salaris, M., & Castelli, F. 2004, *ApJ*, 612, 168
- . 2006, *ApJ*, 642, 797
- Pietrinferni, A., Cassisi, S., Salaris, M., Percival, S., & Ferguson, J. W. 2009, *ApJ*, 697, 275

- Pols, O. R., Schröder, K.-P., Hurley, J. R., Tout, C. A., & Eggleton, P. P. 1998, *MNRAS*, 298, 525
- Poyneer, L. A., & Macintosh, B. A. 2006, *Optics Express*, 14, 7499
- Poyneer, L. A., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9148, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 0
- Provenzano, C., Pagliusi, P., & Cipparrone, G. 2006, *Applied Physics Letters*, 89,
- Pueyo, L., et al. 2015, *ApJ*, 803, 31
- Quillen, A. C. 2006, *MNRAS*, 372, L14
- Rapson, V. A., Kastner, J. H., Andrews, S. M., Hines, D. C., Macintosh, B., Millar-Blanchaer, M., & Tamura, M. 2015a, *ApJ*, 803, L10
- Rapson, V. A., Kastner, J. H., Millar-Blanchaer, M. A., & Dong, R. 2015b, *ApJ*, 815, L26
- Roberge, A., Feldman, P. D., Weinberger, A. J., Deleuil, M., & Bouret, J.-C. 2006, *Nature*, 441, 724
- Rodigas, T. J., Malhotra, R., & Hinz, P. M. 2014, *ApJ*, 780, 65
- Schmid, H. M., Joos, F., & Tschan, D. 2006, *A&A*, 452, 657
- Schneider, G., et al. 2014, *AJ*, 148, 59
- Shen, Y., & Tremaine, S. 2008, *AJ*, 136, 2453
- Shields, J. H., & Ellis, J. W. 1956, *J. Opt. Soc. Am.*, 46, 263
- Siegler, N., Muzerolle, J., Young, E. T., Rieke, G. H., Mamajek, E. E., Trilling, D. E., Gorlova, N., & Su, K. Y. L. 2007, *ApJ*, 654, 580
- Siess, L., Dufour, E., & Forestini, M. 2000, *A&A*, 358, 593
- Simmons, J. F. L., & Stewart, B. G. 1985, *A&A*, 142, 100
- Smith, B. A., & Terrile, R. J. 1984, *Science*, 226, 1421
- Snellen, I. A. G., Brandl, B. R., de Kok, R. J., Brogi, M., Birkby, J., & Schwarz, H. 2014, *Nature*, 509, 63
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJ*, 755, L28
- Soummer, R., Sivaramakrishnan, A., Pueyo, L., Macintosh, B., & Oppenheimer, B. R. 2011, *ApJ*, 729, 144
- Stam, D. M., Hovenier, J. W., & Waters, L. B. F. M. 2004, *A&A*, 428, 663
- Stapelfeldt, K., Krist, J., Bryden, G., & Chen, C. 2007, in *In the Spirit of Bernard Lyot: The Direct Detection of Planets and Circumstellar Disks in the 21st Century*
- Strubbe, L. E., & Chiang, E. I. 2006, *ApJ*, 648, 652
- Su, K. Y. L., et al. 2006, *ApJ*, 653, 675

- Tamura, M., Fukagawa, M., Kimura, H., Yamamoto, T., Suto, H., & Abe, L. 2006, *ApJ*, 641, 1172
- Telesco, C. M., et al. 2000, *ApJ*, 530, 329
- Thalmann, C., et al. 2008, in *Proc. SPIE*, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, 70143F
- Thébaud, P., & Beust, H. 2001, *A&A*, 376, 621
- Thébaud, P., & Wu, Y. 2008, *A&A*, 481, 713
- Thureau, N. D., et al. 2014, *MNRAS*, 445, 2558
- van Leeuwen, F. 2007, *A&A*, 474, 653
- Vidal, M., Leahy, J. P., & Dickinson, C. 2016, *MNRAS*, 461, 698
- Voshchinnikov, N. V., & Krügel, E. 1999, *A&A*, 352, 508
- Wahhaj, Z., Koerner, D. W., Ressler, M. E., Werner, M. W., Backman, D. E., & Sargent, A. I. 2003, *ApJ*, 584, L27
- Wang, J. J., Ruffio, J.-B., De Rosa, R. J., Aguilar, J., Wolff, S. G., & Pueyo, L. 2015, *pyKLIP: PSF Subtraction for Exoplanets and Disks*, Astrophysics Source Code Library
- Wang, J. J., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 55
- Weidenschilling, S. J. 1977, *MNRAS*, 180, 57
- Weinberger, A. J., Becklin, E. E., & Zuckerman, B. 2003, *ApJ*, 584, L33
- Wiktorowicz, S. J. 2009, *ApJ*, 696, 1116
- Wiktorowicz, S. J., et al. 2014a, in *Proc. SPIE*, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V, 914783
- Wiktorowicz, S. J., et al. 2014b, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 83
- Wolff, S. G., Perrin, M. D., Maire, J., Ingraham, P. J., Rantakyro, F. T., & Hibon, P. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 7
- Wolstencroft, R. D., Scarrott, S. M., & Gledhill, T. M. 1995, *Ap&SS*, 224, 395
- Wyatt, M. C. 2006, *ApJ*, 639, 1153
- . 2008, *ARA&A*, 46, 339
- Wyatt, M. C., Dermott, S. F., Telesco, C. M., Fisher, R. S., Grogan, K., Holmes, E. K., & Piña, R. K. 1999, *ApJ*, 527, 918

Wyatt, M. C., Smith, R., Su, K. Y. L., Rieke, G. H., Greaves, J. S., Beichman, C. A., & Bryden, G. 2007, *ApJ*, 663, 365

Wyatt, M. C., et al. 2012, *MNRAS*, 424, 1206

Wyrowski, F. 1991, *Opt. Lett.*, 16, 1915

Zakamska, N. L., Pan, M., & Ford, E. B. 2011, *MNRAS*, 410, 1895

Zhang, K., Blake, G. A., & Bergin, E. A. 2015, *ApJ*, 806, L7