UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

HIGH CONTRAST IMAGING OF PLANET-FORMING DISKS WITH EXTREME ADAPTIVE OPTICS SYSTEMS

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

KELLEN D. LAWSON
Norman, Oklahoma
2022

Downloaded from SHAREOK. Persistent Handle URL: https://hdl.handle.net/11244/335850

HIGH CONTRAST IMAGING OF PLANET-FORMING DISKS WITH EXTREME ADAPTIVE OPTICS SYSTEMS

A DISSERTATION APPROVED FOR THE
HOMER L. DODGE DEPARTMENT OF PHYSICS AND ASTRONOMY

BY THE COMMITTEE CONSISTING OF

Dr. John Wisniewski, Chair

Dr. Edward Baron

Dr. Thayne Currie

Dr. Nathan Kaib

Dr. Kerry Magruder

Dr. Kieran Mullen

© Copyright by KELLEN D. LAWSON 2022
All Rights Reserved.

To my wife, Jess, who ceaselessly inspired and supported me in pursuit of this aspiration.

And to the many others — a community of family, friends, & mentors — who helped along the way.

Acknowledgements

This work was conducted in collaboration with an extensive list of co-authors whose invaluable insights and other contributions were critical for its completion. Please refer to the publications noted at the start of Chapters 2–5 for details.

This research is based on data collected at Subaru Telescope, which is operated by the National Astronomical Observatory of Japan. We are honored and grateful for the opportunity of observing the Universe from Maunakea, which has cultural, historical and natural significance in Hawaii.

We wish to acknowledge the critical importance of the current and recent Subaru telescope operators, daycrew, computer support, and office staff employees. Their expertise, ingenuity, and dedication is indispensable to the continued successful operation of Subaru.

This research was enabled, in no small part, by the efforts of countless software developers who labor tirelessly and selflessly to provide and maintain tools for their communities. This includes the developers of the following tools: Matplotlib (Hunter, 2007; Caswell et al., 2021), NumPy (Harris et al., 2020), SciPy (Virtanen et al., 2020), Astropy (Astropy Collaboration et al., 2013, 2018), CuPy (Okuta et al., 2017), LMFIT (Newville et al., 2014), diskmap (Stolker et al., 2016), the CHARIS Data Reduction Pipeline (Brandt et al., 2017), and the CHARIS Data Processing Pipeline (Currie et al., 2020a).

The results of Chapter 5 were enabled by funding from Sigma Xi Grants in Aid of Research.

Table of Contents

Abstract viii

1 Introduction 1
- 1.1 Studying Exoplanets and Circumstellar Disks with High-contrast Imaging . . . 1
- 1.2 Imaging Planet-forming Disks in the Era of Extreme Adaptive Optics . . . 6

2 Integral Field Spectroscopy of the HD 15115 Debris Disk System 10
- 2.1 Introduction . . . 10
- 2.2 Data . . . 12
  - 2.2.1 Observations . . . 12
  - 2.2.2 CHARIS Data Reduction . . . 12
  - 2.2.3 Results . . . 13
- 2.3 Disk Morphology . . . 18
- 2.4 Modeling the Debris Disk of HD 15115 . . . 19
  - 2.4.1 Disk Forward Modeling . . . 19
  - 2.4.2 Model Debris Disks . . . 21
  - 2.4.3 Model Optimization Using Differential Evolution . . . 22
  - 2.4.4 Model Results . . . 25
  - 2.4.5 Modeling Discussion . . . 31
  - 2.4.6 Comparison with Recent Studies . . . 33
- 2.5 Disk Surface Brightness and Color . . . 34
  - 2.5.1 Disk Photometry . . . 34
  - 2.5.2 Surface Brightness Results . . . 40
- 2.6 Limits on Planets . . . 42
- 2.7 Conclusions and Future Work . . . 43

Appendix . . . 45
- 2.A Scattering Phase Function Comparison . . . 45
- 2.B Python implementation of differential evolution . . . 45

3 Integral Field Spectroscopy of the HD 36546 Debris Disk System 49
- 3.1 Introduction . . . 49
- 3.2 Data . . . 50
  - 3.2.1 Observations . . . 50
  - 3.2.2 CHARIS Data Reduction . . . 51
  - 3.2.3 Results . . . 52
- 3.3 Modeling the Debris Disk of HD 36546 . . . 54
  - 3.3.1 Disk Forward Modeling . . . 54
  - 3.3.2 Model Results . . . 58

3.3.3 Modeling Discussion . . . 58
3.4 Disk Spine Trace and Surface Brightness . . . 63
3.4.1 Disk Spine . . . 63
3.4.2 Disk Surface Brightness . . . 64
3.4.3 Surface Brightness Power-Law . . . 67
3.4.4 Surface Brightness Discussion . . . 69
3.5 Limits on Planets . . . 72
3.6 Conclusions and Future Work . . . 73
Appendix . . . 76
3.A Additional Model Results . . . 76
3.B RDI-KLIP Reduction using Radial Profile Subtraction . . . 76

4 Integral Field Spectropolarimetry of Planet-forming Disks 80
4.1 Introduction . . . 80
4.2 CHARIS Spectropolarimetry Mode . . . 82
4.2.1 Half-wave Plate Cycles . . . 82
4.2.2 Instrument Configurations . . . 82
4.2.3 Calibration Data . . . 83
4.2.4 Observations . . . 83
4.3 CHARIS PDI Data Reduction and Processing . . . 84
4.3.1 Cube Extraction . . . 84
4.3.2 Preprocessing . . . 84
4.3.3 Half-Wave Plate Cycle Matching . . . 86
4.3.4 Polarimetric Differential Imaging . . . 86
4.4 Preliminary CHARIS Spectropolarimetry Results . . . 89
4.4.1 AB Aurigae . . . 90
4.4.2 TW Hydrae . . . 90
4.5 CHARIS PDI Disk Forward Modeling . . . 92
4.6 Summary and Conclusions . . . 93

5 Constrained Reference Star Differential Imaging: Enabling High Fidelity Imagery of Highly Structured Circumstellar Disks 96
5.1 Introduction . . . 96
5.2 Disk Signal Attenuation in Reference Star Differential Imaging . . . 98
5.3 Mitigating RDI Oversubtraction: Constrained RDI . . . 99
5.3.1 Polarized Intensity Constrained RDI . . . 100
5.3.2 Model Constrained RDI . . . 101
5.3.3 Optimization of Constrained RDI . . . 102
5.4 Throughput Assessment using Simulated Data . . . 102
5.4.1 PCRDI Throughput . . . 102
5.4.2 MCRDI Throughput . . . 104
5.5 Application of PCRDI to IFS Observations of AB Aurigae . . . 106
5.6 Broader Applications . . . 109

5.7 Conclusions . . . 110
Appendix . . . 112
5.A Forward Modeling for PCRDI Products . . . 112
5.B Validation of Constrained RDI Using Forward Modeling . . . 112
5.C Supporting Software and Run-time . . . 113
5.C.1 Run-time . . . 113

6 Conclusions **115**
6.1 Chapter Summaries . . . 115
6.2 Future Work . . . 116
6.2.1 Applying Constrained RDI to Study Disks and Search for Embedded Planets . . . 116
6.2.2 Meta-parameter Optimization for PSF Subtraction . . . 117

Abstract

Circumstellar disks serve as benchmark systems to study how and where exoplanets form. With the advent of ground-based extreme adaptive optics (AO) facilities, we have entered a revolutionary era for disk studies. It is now feasible to simultaneously resolve the structures within disks that could be caused by planets and directly identify the young companions that may induce them. The use of integral field spectrographs (IFSs) — which provide images at an array of wavelengths — enhances this capability by helping to characterize disk material and any identified companions.

Using data from Subaru's extreme AO system, SCExAO, paired with the CHARIS IFS, we present near-infrared (NIR) observations of two systems hosting debris disks – gas-poor, dusty disks at the end stages of planetary system formation. For HD 15115, we introduce the use of differential evolution (DE) for disk model optimization — which finds a strong solution while evaluating many fewer models than common alternatives. From modeling, we find no evidence to support the presence of a (possibly planet-driven) misaligned inner-ring as was suggested previously. For the HD 36546 system, we show its disk has either an extremely shallow radial density profile or a "back-swept wing profile"; the latter case, combined with the blue NIR color of the disk, could suggest interaction with an unseen planet.

CHARIS recently added a new and unique integral field spectropolarimetry mode which produces high-contrast polarimetry at the same array of wavelengths as its standard IFS mode. We describe the software module which we designed to process data from this observing mode, present preliminary results for two planet-forming disk systems, and provide a summary of investigations enabled by this mode – including tests on embedded companions.

Finally, we introduce a new technique for removing starlight in imagery of circumstellar disk systems, "Constrained Reference Star Differential Imaging (RDI)", which uses prior knowledge to more accurately model (and remove) the starlight in a data sequence. Using simulated and on-sky data, we demonstrate the ability of this technique to effectively eliminate the variable loss of circumstellar light that normally occurs as starlight is removed. This enables spatial and spectral analysis of circumstellar disks in unprecedented detail and provides a much clearer path to the robust identification and validation of exoplanets embedded within them.

CHAPTER 1

Introduction

1.1. Studying Exoplanets and Circumstellar Disks with High-contrast Imaging

To date, nearly 5000 exoplanets have been confirmed – with the overwhelming majority being identified using "indirect" detection methods. The most prolific of these are the transit method, in which planets passing in front of a star are deduced from periodic dips in the star's brightness (e.g., Borucki et al., 2010), or the radial velocity (RV) method, in which planets are identified from periodic changes in a star's velocity induced by the planet's gravitational pull (e.g., Joergens & Neuhäuser, 2003).

Direct imaging of exoplanets is possible, but requires sophisticated high-contrast imaging techniques and instrumentation to separate the faint circumstellar signal (CSS) from the much brighter pattern of diffracted light from the parent star (the stellar point spread function, or PSF). Though challenging, direct imaging provides an important complement for indirect detection techniques in that it is an effective strategy for probing the occurrence of planets at orbital separations and system ages where indirect detection is weakest – namely, large separations and young ages. While large separation planets require longer baselines of observation for RV and transit detection, they are easier to detect with direct imaging by virtue of a more favorable ratio of star-to-planet light further from the star. Similarly, while the increase in stellar activity for young stars introduces significant barriers for indirect detection (e.g., Paulson et al., 2004), young planets will be brighter than older planets (since planets cool and become dimmer over time) and thus are generally easier to detect using direct imaging. As a result, while direct imaging has contributed only $\sim 1\%$ (55 out of 4884) of all exoplanet discoveries, it has contributed $\sim 71\%$ (50 out of 70) of planets with semi-major axes larger than 10 au and $\sim 65\%$ (36 out of 55) of those with ages younger than 100 Myr [1](see Figure 1.1).

High-contrast imaging can also be used to study circumstellar disks — the disks of dust and/or gas surrounding a star which provide the material for planet formation. Circumstellar disks around newly formed stars are called "protoplanetary disks": an evolutionary phase which typically persists for a few Myr and is marked by the significant presence of gas (Williams & Cieza, 2011). Notably, this phase is thought to coincide with the epoch of giant planet formation. "Transitional disks" (or "transition disks") are a class of protoplanetary disks which mark the end of the protoplanetary

[1]Based on data from the NASA Exoplanet Archive (https://exoplanetarchive.ipac.caltech.edu/; queried 10 Jan 2022), including all entries labeled "Published Confirmed".

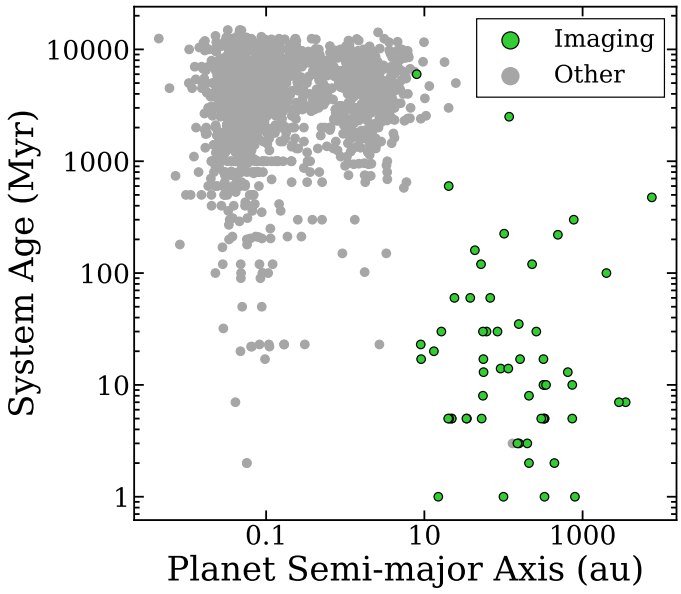

Figure 1.1: Semi-major axis (au) versus system age (Myr) for all confirmed exoplanets. Planets detected with direct imaging (green markers) can be seen predominantly occupying the large separation - low age regime. In contrast, those detected with other methods (grey markers) tend to have smaller separations and older ages. Based on data from the NASA Exoplanet Archive (https://exoplanetarchive.ipac.caltech.edu/; queried 10 Jan 2022), including all entries labeled "Published Confirmed".

phase. Transitional disks are delineated by the appearance of radial gaps or cavities in their dust distributions, which may be evident from analysis of spectral energy distributions (SEDs) or from resolved imagery (Francis & van der Marel, 2020). The presence of these structures within protoplanetary disks, as well as others like spiral arms or warping, has been linked to both unseen planetary-mass drivers and non-planetary phenomena (e.g., gravitational instability; Dong et al., 2018). Observing disk systems in this stage, when planets might be directly recovered alongside such features, is critical toward reaching a better understanding of the role that disks play in the evolution of planets (and vice versa). Even if planets are not identified among these features, continuous study enables diagnosis of their likely driver (i.e., planetary or non-planetary; Dong et al., 2018), and thus serves to vet targets for future surveys having improved sensitivity.

The final stage of circumstellar disks are gas-depleted, dust-dominated "debris disks", which have been identified around stars spanning ages from as young as 3–10 Myr (HD 36546; Currie et al., 2017) to as old as ~ 8 Gyr (τ Ceti; Hughes et al., 2018). Given the relatively short timescale for dust dissipation ($\lesssim 10$ Myr Kenyon & Bromley, 2002), the dust within these disks is thought to be continuously replenished by collisions between planetesimals. As such, observations of debris disks provide a window into the recent dynamical history of these systems. Much like other disks, observing features within debris disks (such as gaps or asymmetries) can provide clues regarding

the presence and location of yet-unseen exoplanets (e.g., Lagrange et al., 2009).

Ground-based direct imaging of exoplanets and circumstellar disks is enabled by three key technologies. The first of these, adaptive optics (AO), corrects for the distortion of incoming light (e.g., from turbulence in Earth's atmosphere) to provide a substantially sharper and more stable stellar PSF (Beckers, 1993, see Figure 1.2). The second, coronagraphy (Lyot, 1939), uses one of a number of distinct coronagraph designs to reject the bright central light from a star, facilitating the recovery of other nearby light sources (Guyon, 2009). Even with the use of AO and coronagraphy, the light of circumstellar sources is still typically obscured beneath the starlight that remains. The final capability needed for direct imaging of exoplanets and disks is a method for isolating and removing this starlight. This is typically achieved using one or more "differential imaging" techniques in which observations are made in such a way that the pattern of diffracted starlight can be disentangled from that of any circumstellar light in post-processing. Among these techniques are:

- Angular differential imaging (ADI; Marois et al., 2006), in which the field of view is allowed to rotate as the telescope tracks the target throughout a sequence of exposures. As this occurs, the diffraction pattern (which results predominantly from diffraction within the optics) remains fixed while any circumstellar signal appears to rotate about the image center. In the simplest version of ADI data reduction, the stellar PSF can be isolated by simply median-combining the frames to eliminate the "moving" circumstellar signal.

- Spectral differential imaging (SDI; Sparks & Ford, 2002), in which observations are made at two or more wavelengths. Since the size of the diffraction pattern scales with wavelength, the stellar PSF can be isolated by rescaling the images by wavelength to produce a sequence in which the stellar PSFs at different wavelengths are aligned, but in which circumstellar features are then radially displaced. As with ADI, in the simplest case, the stellar PSF could then be isolated by median combining these rescaled frames.

- Reference star differential imaging (RDI), in which images of a star with no circumstellar signal are used to isolate the stellar PSF of the instrument.

Since the stellar PSF changes on the timescales of individual exposures, the resulting stellar PSF model in each case is imperfect. The simplest methods for modeling the PSF with these data tend to perform poorly at small separations from the star (where the PSF changes more rapidly) – leaving significant residual starlight that inhibits recovery of CSS. More advanced techniques, such as Locally Optimized Combination of Images (LOCI; Lafrenière et al., 2007) and Karhunen-Loève Image Projection (KLIP; Soummer et al., 2012), better model starlight by combining multiple frames to minimize residuals with a target frame, but can also cause a significant fraction of CSS

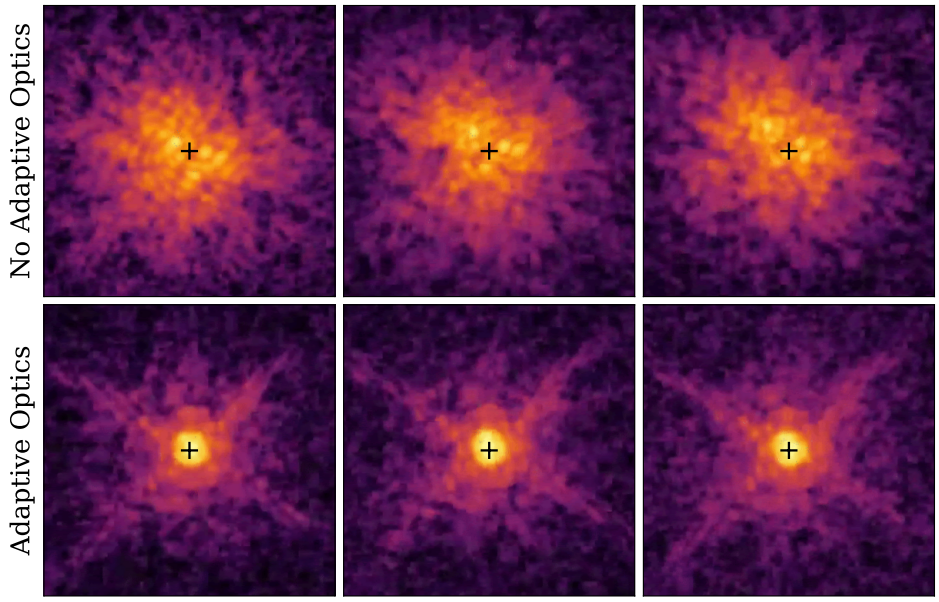

Figure 1.2: A comparison of three successive images of a star without the use of adaptive optics (top row) and with adaptive optics (bottom row) using the Subaru Telescope's first generation 'AO188' system. Adaptive optics provides a PSF that is both sharper and more stable over time — easing recovery of faint circumstellar sources such as exoplanets and disks.

to be erroneously removed (or "attenuated"). For ADI and SDI, this signal loss results from both "self-subtraction", in which circumstellar signal is included in the PSF model (e.g., by insufficient rotation between exposures for ADI), and "over-subtraction", in which the presence of the CSS in the target frame results in an over-bright PSF model (Pueyo, 2016). In the case of RDI, CSS attenuation is necessarily the result of over-subtraction alone (see Figure 1.3). If not addressed, this signal loss can bias measurements of companion position and brightness, alter the assumed orientation and morphology of disks, cause disk structure to be mistaken for planets, and even entirely erase exoplanets, disks, and disk features (e.g., Milli et al., 2012; Currie et al., 2019).

Commonly, attenuation of disk and planet light is addressed through "forward modeling" (e.g., Wang et al., 2015; Currie et al., 2015a; Pueyo, 2016; Mazoyer et al., 2020). In forward-modeling, the CSS attenuation caused by a particular PSF subtraction procedure is induced on a synthetic planet or disk image. By comparing the known input signal with the attenuated output signal, throughput for a particular circumstellar source can be determined. While we can determine the attenuated result for a given input image, this *cannot* be directly inverted (that is, we cannot analytically determine the input for a given attenuated result). Since any unresolved companions ("point-sources") can be assumed to manifest with the same shape (that of the instrumental PSF), throughput estimates can be made for a candidate anywhere within the field of view for a relatively small computational cost. For circumstellar disks, assessing throughput is much more challenging. The fraction of disk

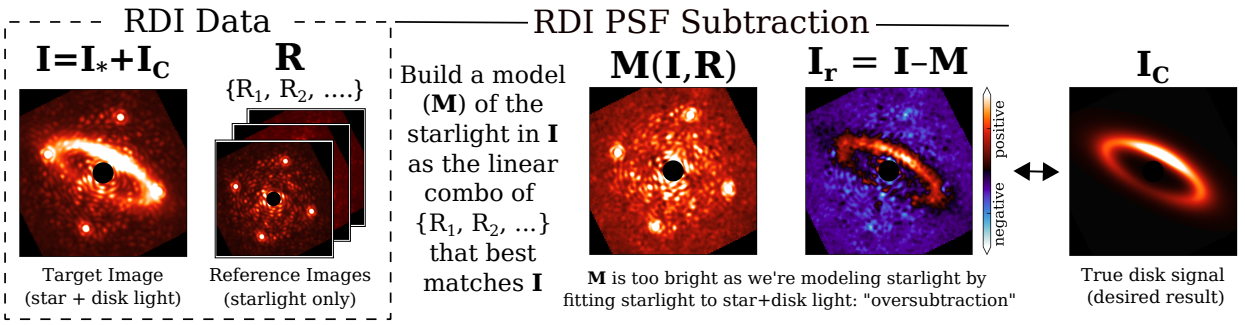

Figure 1.3: Typical RDI PSF-subtraction for a target image, I, containing both stellar signal, I_*, and circumstellar signal, I_C. Using a sequence of reference images, $\mathbf{R} = \{R_1, R_2, \ldots\}$, a PSF model, M, is constructed by identifying the linear combination of reference images that will minimize the residuals with the target image. As a result of the disk signal present in the target image, the PSF model is too bright. Subtracting this model from the reference image thus leads to "over-subtraction" and a significantly attenuated residual, I_r.

signal lost at any given position will depend not only on the disk flux and stellar flux at that location – but also at every other location considered. To account for this, forward modeling must be carried out on an array of disk models spanning plausible morphologies and orientations to identify the attenuated models that reasonably match the attenuated data (Currie et al., 2015a; Pueyo, 2016; Goebel et al., 2018; Boccaletti et al., 2019).

Another differential imaging technique – polarimetric differential imaging (PDI) – enables the recovery of polarized intensity (PI) and provides an important complement for the total intensity (I) data of other differential imaging techniques by producing *unattenuated* imagery of polarized sources (Kuhn et al., 2001). Disk signal, which is largely comprised of starlight scattered toward us by material in the disk, is typically strongly polarized ($\sim 25-50\%$; Perrin et al., 2009; Uyama et al., 2017). However, as the emission from young planets is unpolarized (to a reasonable approximation), planets generally cannot be detected in PI alone.

Using these techniques, early direct imaging surveys contributed considerably to our understanding of exoplanet populations and of how planets form and evolve. For example, the Strategic Exploration of Exoplanets and Disks with Subaru (SEEDS) survey (Tamura, 2009) conducted high-contrast H-band (1.6 μm) observations of nearly 400 nearby stars with the 8.2m Subaru Telescope's High-Contrast Coronographic Imager for Adaptive Optics (HiCIAO) instrument (Suzuki et al., 2010) and adaptive optics system (AO188; Hayano et al., 2008). From this, they reported detection limits for the complete sample, along with the detection of "*dozens*" of circumstellar disks with sub-structure (Tamura 2016; see Figure 1.4; Janson et al. 2013; Uyama et al. 2017), two planetary-mass (or near-planetary mass) companions (κ And b, GJ 504 b; Carson et al., 2013; Kuzuhara et al., 2013), and a number of additional sub-stellar (likely brown dwarf) companions.

Similarly, the Gemini/NICI Planet-Finding Campaign (Liu et al., 2010) observed more than 200 stars and reported detection limits and the discovery of four brown-dwarf companions (Biller et al., 2014). Using space-based coronagraphy from the Space Telescope Imaging Spectrograph (STIS) instrument aboard the Hubble Space Telescope (HST), Schneider et al. (2014) searched for exoplanets and exoplanet-signposts within 11 nearby disk systems, reporting a broad array of intriguing disk features and providing measurements in complementary optical wavelengths for use by future studies.

Perhaps most notably, the detection by these high-contrast imaging surveys (and smaller contemporary campaigns) of numerous 'super-Jovian' mass companions at wide separations, such as κ And b ($r \sim 60$ au, $M \sim 13\ M_J$), challenged prior understanding regarding the formation of sub-stellar companions. Formation by the conventionally-favored core-accretion mechanism is not expected to produce such massive planets at these separations (Chabrier et al., 2014). While they might be explained by formation at smaller separations and subsequent migration to their observed locations, this possibility is difficult to reconcile with the young ages and lack of significantly eccentric orbits for a number of these (and more difficult still in the case of multi-planet systems such as HR 8799; Marois et al. 2010; Pueyo et al. 2015).

Though transformationally successful, these surveys were limited by a) typically conducting observations at a single wavelength (which limited the characterization of detected exoplanets and disks), and b) the performance of first-generation AO systems and coronagraphs, which restricted detections to beyond $\sim 0\farcs2$ (corresponding to ~ 30 au, roughly the orbital separation of Neptune, for the nearby Taurus-Auriga star forming region).

1.2. Imaging Planet-forming Disks in the Era of Extreme Adaptive Optics

Recently, "extreme" AO systems for ground-based observatories, such as Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE; Beuzit et al., 2019), Gemini Planet Imager (GPI; Macintosh et al., 2015), and Subaru Coronagraphic Extreme Adaptive Optics (SCExAO; Jovanovic et al., 2015), have entered use. While near-infrared (NIR) imaging with first-generation AO systems was limited to recovery of planets and disks beyond $\sim 0\farcs2$, extreme AO systems (combined with state-of-the-art coronagraphs) can push to separations as small as $\sim 0\farcs05$[2] (Jovanovic et al., 2015). This has enabled ground-based imaging of exoplanets and disk features at orbital separations comparable to the orbits of the outer planets in our solar system. Notably, extreme AO also provides the precision necessary to simultaneously image protoplanetary disk features alongside the planets thought to be inducing them ("embedded" planets). In combination with integral field spectrographs (IFSs), which provide imaging data at multiple wavelengths simultaneously, we now have the instrumentation to conduct incredibly detailed spatial and spectral analysis of these systems.

[2]For an inner working angle (IWA) of $1\lambda/D$ in H-band ($\lambda = 1.6\ \mu$m).

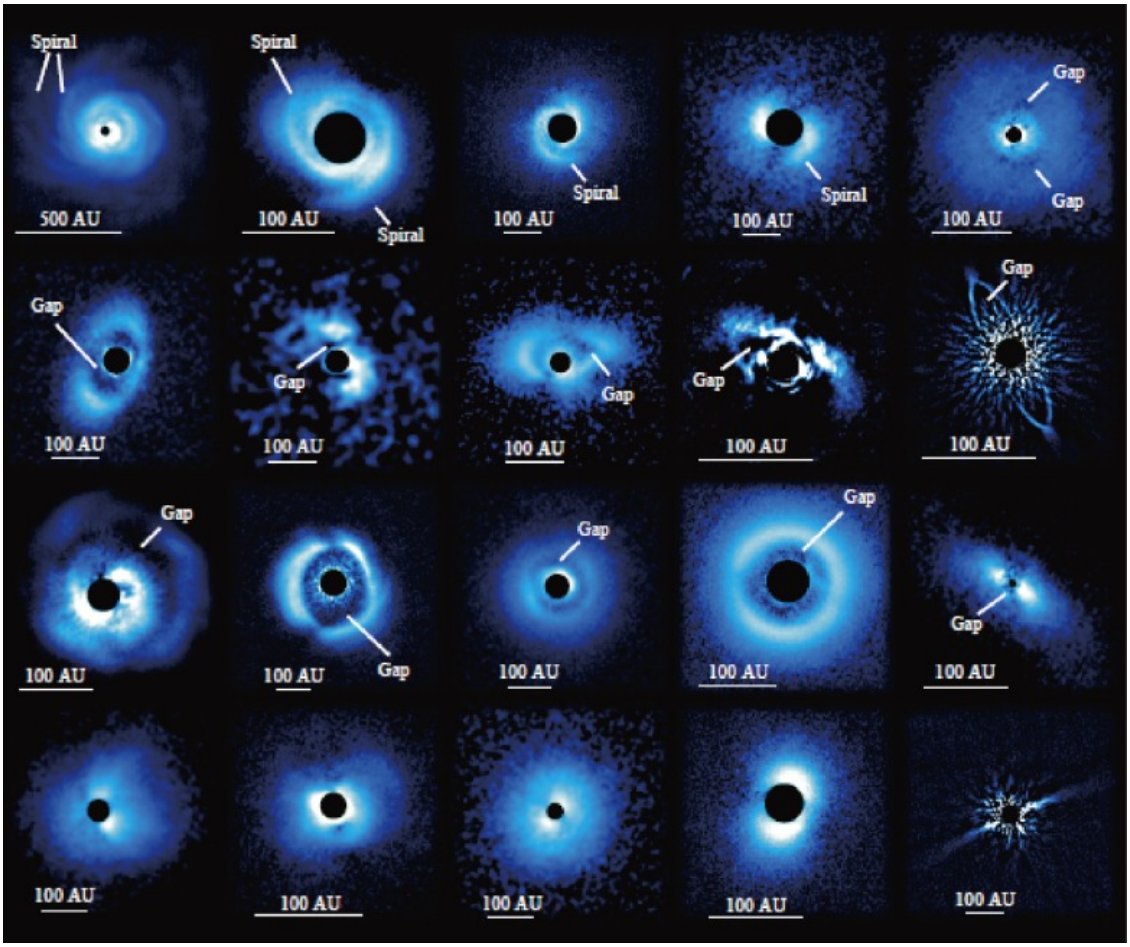

Figure 1.4: Reproduced from Tamura (2016), a selection of disk images from the SEEDS survey showing compelling substructure in polarized intensity (compiled by J. Hashimoto). **First row** (from left to right): AB Aur, SAO 206462, MWC 758, LkHa 330, TW Hya. **Second row**: PDS 70, SZ 91, WLY 2-48, LkCa 15, HR 4796A. **Third row**: AB Aur (close-up), HD 142527, HD 169142, RX J1604.3-2130A, GM Aur. **Fourth row**: RY Tau, SR 21, MWC 480, UX Tau A, HIP 79977.

A particularly noteworthy and striking example of the early success of extreme AO systems was the detection of PDS 70b – the first conclusive detection of a planet within a protoplanetary disk (Keppler et al., 2018, see Figure 1.5).

A hindrance for the direct imaging capabilities of modern (and near-future) observatories results from the aforementioned loss of circumstellar signal during post-processing – which is not only spatially variable, but also spectrally variable. Although these facilities can reach the contrasts and spatial resolution necessary to recover embedded exoplanets, the chaotic environments of young circumstellar disks (see, e.g., Figure 1.6) can manifest small-scale disk features that are easily mistaken for planets based on morphology (even without signal loss). While multi-wavelength or IFS observations might enable spectrophotometric differentiation of planets from fine disk structure, variable CSS attenuation inhibits the extraction of reliable spectral information. Similarly,

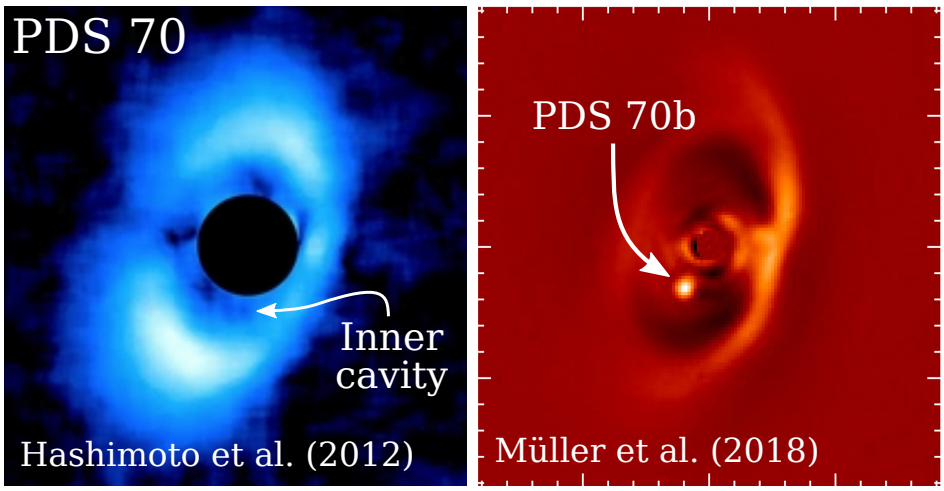

Figure 1.5: **Left** – The PDS 70 disk in polarized intensity (H-band; Subaru AO188/HiCIAO); altered from Hashimoto et al. (2012). Based on the inner cavity in this data, the presence of one or more gas giant planets within the disk was suggested (Hashimoto et al., 2012; Dong et al., 2012). **Right** – Subsequently, using VLT/SPHERE observations, Keppler et al. (2018) reported the discovery of a gas giant planet, PDS 70b, within the cavity; altered from Müller et al. (2018).

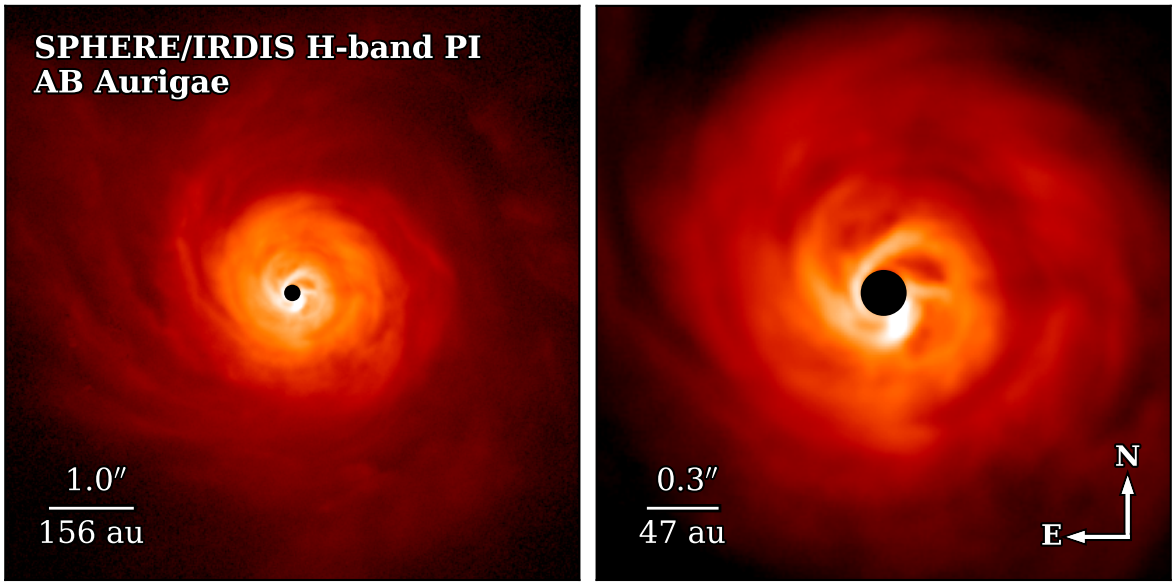

Figure 1.6: Extreme adaptive optics polarized intensity imagery of the AB Aurigae protoplanetary disk system using VLT's SPHERE/IRDIS (originally reported in Boccaletti et al. 2020). This system, presented in a $7'' \times 7''$ field of view (FOV) on the left and a $2\farcs5 \times 2\farcs5$ FOV on the right, shows complex disk structure at both large and small scales. Both images use a logarithmic color scale. Systems like AB Aurigae are perhaps the most likely to hold evidence of ongoing planet formation, but are also the systems where it is the most challenging to verify.

though comparison of total intensity and polarized intensity images should highlight self-luminous embedded planets by their diminished fractional polarization ($\mathrm{F_{pol}} = \mathrm{PI/I}$) relative to disk material,

the attenuation present in total intensity products again complicates this. The ability to mitigate this signal loss, or to quantify it in a tractable manner, is therefore of significant importance for realizing the exquisite capabilities of modern observatories.

Herein, we present efforts in pursuit of this goal. In Chapters 2 and 3, we introduce the use of differential evolution for disk forward modeling optimization in application to observations of two debris disk systems, HD 15115 and HD 36546, using the Subaru Telescope's extreme AO system, SCExAO, paired with the Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS) NIR IFS. This technique identifies strong disk models in a much more tenable timeframe and enables our analysis of the composition and morphology of these disks. Then, in Chapter 4, we introduce the new integral field spectropolarimetric imaging mode for CHARIS and summarize ways that it can be used to advance studies of planet-forming disk systems. Finally, in Chapter 5, we describe and demonstrate a novel technique for RDI-based PSF subtraction of circumstellar disk systems which is both straightforward to implement and capable of effectively eliminating CSS attenuation.

CHAPTER 2

Integral Field Spectroscopy of the HD 15115 Debris Disk System[1]

2.1. Introduction

Gas-poor, dusty debris disks around stars are key laboratories for studying planetary system structure and the late stages of their formation (Wyatt, 2008; Hughes et al., 2018). Scattered light imagery of debris disks around young stars clarifies the disks' structures and can identify the signatures of sculpting planets and in-situ formation and erosion of icy Kuiper belt objects (Kalas et al., 2005; Kenyon & Bromley, 2008). Further, these studies enable analysis of the composition and scattering properties of the material within the disks, potentially providing reference points for the evolution of the Kuiper belt (Currie et al., 2015b). High-contrast imaging produced using so-called "extreme adaptive-optics" (exAO) facilities (e.g. SPHERE (Beuzit et al., 2019), GPI (Macintosh et al., 2015), and SCExAO (Jovanovic et al., 2015; Lozi et al., 2018; Currie et al., 2019)) provides the opportunity to study these systems to smaller inner working angles than was possible with conventional AO. This enables the assessment of the disks' scattering phase functions at previously inaccessible angles, as well as placing more significant constraints on the presence of embedded planets.

The debris disk around HD 15115, an F2V star at a distance of 49.0 ± 0.1 pc (Gaia Collaboration et al., 2018) and an estimated age of $\sim 10 - 100$ Myr (e.g. Moór et al., 2006; Rhee et al., 2007; Gagné et al., 2018), could be a particularly good target for studying planetary system structure and the results of the initial formation stages. Discovery optical scattered-light imagery from the Hubble Space Telescope (HST) Advanced Camera for Surveys (ACS) reveals a highly inclined disk with an "extreme" east-west length asymmetry in the HST/ACS F606W bandpass ($\lambda_{pivot} = 5886$ Å, FWHM $= 2325$ Å), resolving the eastern extent out to $\sim 7''$ while the western extent is recovered to the edge of the field of view at $12\farcs38$ (Kalas et al., 2007). Their measurements of the disk's surface brightness on either side show an approximately symmetric brightness at $2\farcs0$, with the west becoming brighter than the east at larger separations ($\Delta m \sim 1$ at $6''$). Follow-up H-band Keck/NIRC2 adaptive-optics (AO) imagery resolved the disk at $\rho \sim 1'' - 3\farcs3$ and revealed a blue (F606W-H) color on both sides and a brightness asymmetry beyond $2''$. The combination of the reported color and its highly inclined orientation led to HD 15115's disk being informally referred to as "the Blue Needle". Follow-up observations from ground-based AO and space expanded the wavelength range over which HD 15115's disk is resolved and further clarified its properties

[1]This chapter is reproduced from Lawson et al. (2020) with permission.

at separations beyond 1″. HST/NICMOS 1.1 μm data revealed evidence of more complicated color gradients and a wavelength dependence for the disk's asymmetry and an angular separation dependence for its colors (Debes et al., 2008).

Subsequent studies found evidence of a bow-like shape in the disk at $\sim$ 1″–2″ (Rodigas et al., 2012; Mazoyer et al., 2014; Sai et al., 2015), consistent with a ring-like disk at $\sim$ 90 au. Using archival near-infrared imaging from the Gemini Observatory, Mazoyer et al. (2014) conclude that, while the system's ring is asymmetrical in brightness, the geometry of the ring itself is symmetric about the parent star. Schneider et al. (2014) reported HST Space Telescope Imaging Spectrograph (STIS) data which significantly improved upon the visible light photometry and morphology of the disk. These data showed that the bowing and asymmetry in visible wavelengths continue down to 0.″4. Additionally, this HST/STIS imaging revealed a previously unseen morphological bifrucation on the east side of the outer disk Schneider et al. (2014).

Recent results suggest the possible existence of multiple debris ring components. Engler et al. (2019) reported the first extreme AO observations of HD 15115, consisting of VLT/SPHERE total intensity data in J and H band, and polarized intensity data in J band. They recover the disk over stellocentric separations of $\rho \sim 1.''0 - 5.''5$ and suggest, from peaks in their polarized intensity profiles, the possibility of a distinct non-coplanar inner disk having a fiducial radius of $\sim 1.''3$. Attempting to investigate this using their total intensity imagery, they are unable to reveal conclusive evidence regarding the disk's geometry. MacGregor et al. (2019) reported 1.3 mm Atacama Large Millimeter/submillimeter Array (ALMA) observations of the system with a synthesized beam size of $0.''58 \times 0.''55$. From these observations, evidence exists for either a distinct inner ring, with radius $\sim 0.''95$, or a significant gap in the canonical disk at a separation of $1.''2$. Notably, however, they report a lack of evidence in their data to support the misalignment of the inner disk hypothesized by Engler et al. (2019). Additionally, MacGregor et al. (2019) found an absence of the east-west brightness asymmetry typically reported in previous NIR and optical imagery (e.g. Kalas et al. 2007; Mazoyer et al. 2014). They suggest that the large-grain dust population probed by ALMA was unaffected by the mechanism responsible for the asymmetry reported by other studies over similar separations. As perturbations from planetary mass companions are often used to explain disk asymmetries and more complex, multi-ringed disk geometries (e.g. MacGregor et al., 2019), the details of these occurrences in the HD 15115 disk are significant. To better clarify the presence or absence of additional ring components and brightness asymmetries over a wide wavelength range, high-contrast imaging data interior to 1″, matching the coverage of STIS data from Schneider et al. (2014), are needed.

In this work, we report new near-infrared scattered light imagery of the HD 15115 system using the Subaru Coronagraphic Extreme Adaptive Optics (SCExAO) system and the Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS) integral field spectrograph in broadband

(spanning near-infrared J, H, and K bands, $1.13 - 2.39$ μm) mode (Groff et al., 2016). This imagery provides a view of the disk to separations a factor of $\sim 3 - 5$ smaller than previous recent studies ($\rho \sim 0\farcs2$). We conduct analysis of the disk's color in NIR and optical wavelengths by combining CHARIS IFS data with prior HST STIS imagery. Through both spine tracing and forward modeling, we investigate the details of the system's geometry and offer constraints for the presence of additional rings or planet companions within CHARIS's $2'' \times 2''$ field of view.

2.2. Data

2.2.1. Observations

HD 15115 was observed on 2017 August 30 and 2017 September 07 using the Subaru Telescope's SCExAO paired with the CHARIS integral field spectrograph operating in low-resolution (R $\sim$ 20), broadband (1.13–2.39 μm) mode, and utilizing SCExAO's Lyot coronagraph with 217 mas diameter occulting spot. CHARIS has a nominal pixel scale of $0\farcs0164$ pixel^{-1}, which has been revised to $0\farcs0162$ pixel^{-1} (Currie et al., 2018). Both sets of data were collected in *angular differential imaging* mode (ADI; Marois et al. 2006), achieving total parallactic angle rotations of ΔPA $= 76°$, and $56°$ with total integration times of $t_{int} = 81$ and 55 minutes respectively. Each set is made up of 80 individual exposures, with August 30 images having exposure times of 60.48 seconds and September 07 images having exposure times of 41.3 seconds. Sky frames were obtained for both data sets.

For the September 7 data, the conditions were good, with the "slow" (long coherence time) seeing having a full-width at half-maximum (FWHM) in V band of $\theta_{\rm V} \sim 0\farcs5$. SCExAO's real-time telemetry data estimated H-band Strehl ratios of $\sim 80\%$. The conditions for the August 30 data were comparable. No telemetry data were recorded for the August observation, but the point spread function (PSF) quality appeared slightly superior by-eye.

Additionally, we make use of HST/STIS analysis-quality (AQ) imagery (STIS/50CCD, $2000 - 10500$ Å, $\lambda_{pivot} = 5752$ Å), originally reported and analyzed in Schneider et al. (2014), to better explore the colors of HD 15115's disk (see Section 2.5).

2.2.2. CHARIS Data Reduction

CHARIS data were extracted from raw CHARIS reads using the CHARIS Data Reduction Pipeline (Brandt et al., 2017). Extracted data take the form of image cubes with dimensions $(N_\lambda, N_x, N_y) = (22, 201, 201)$ (i.e. 201×201 pixel images for each of 22 wavelength channels). Subsequent basic image processing – e.g. sky subtraction, image registration, spectrophotometric calibration – was carried out as in Currie et al. (2011, 2018).

PSF subtraction was performed by application of both the *Karhunen-Loève Image Projection* (KLIP; Soummer et al. 2012) and the *Adaptive, Locally Optimized Combination of Images* (A-LOCI;

Currie et al. 2012a, 2015a) algorithms independently.

We performed PSF subtraction with settings geared towards the detection of a) the HD 15115 debris disk and b) companions plausibly responsible for sculpting the disk. Table 2.1 lists our parameter choices for each reduction, with the motivations for these choices summarized hereafter.

Disk Detection – The HD 15115 debris disk is oriented nearly edge-on in the plane of the sky (e.g. Kalas et al., 2007; Mazoyer et al., 2014). To detect the disk, we performed PSF subtraction exploiting ADI only, not SDI. For A-LOCI, tuning the geometry of the optimization and subtraction regions – i.e. their relative azimuthal and radial widths – is essential toward the recovery of disk flux. Combining the minimum rotation gap with azimuthally elongated regions allowed LOCI coefficients to be computed with less perturbation by the radially extended disk flux while still producing a strong reconstruction of the speckle noise. For KLIP, we performed PSF subtraction in full annuli. To limit self-subtraction of the disk, we imposed a minimum rotation gap of $\delta \sim 1.25 - 1.50\ \lambda/D$ when selecting suitable reference frames for both A-LOCI and KLIP reductions.

Companion Detection – To achieve deeper contrast limits needed to detect faint planets, we used a combination of ADI and then SDI (on the ADI residuals) using A-LOCI following Currie et al. (2018). For the ADI component, optimization regions were a factor of 20 smaller (50 PSF footprints) and the rotation gap was reduced to δ = 0.5–0.75. For a given section of the science image, up to the 50 most correlated sections from the reference image library were used to build a reference PSF (with the number of available images depending on the portion of the 80 exposures satisfying the minimum rotation gap requirement for the section). For the SDI component, the optimization zone covers an annular region with the same depth $\Delta\ r_{\rm sub}$ as the subtraction zone but the smaller annular wedge-shaped subtraction zone is masked.

2.2.3. Results

Both PSF subtraction techniques yield strong detections of the disk to $\rho \sim 0\farcs15 - 0\farcs25$ in CHARIS data (Figures 2.1, 2.2), improving upon the $0\farcs4$ angular separation achieved with optical HST/STIS data Schneider et al. (2014). CHARIS data mark a substantial improvement over previous ground-based, near-IR scattered light imaging of the disk, with conventional AO data limited to $\rho \gtrsim 1''$ and extreme AO imaging from Engler et al. (2019) detecting the disk exterior to $\rho \sim 0\farcs75$ - $1''$. This improvement is owed in part to the significant field rotation achieved, with $\Delta{\rm PA} = 76^\circ$ and $\Delta{\rm PA} = 56^\circ$ for our two sets of observations, versus e.g. $\Delta{\rm PA} = 23^\circ$ for the data from Engler et al. (2019).

The quality of the detection varies from J band, where the disk detection is contaminated by residual speckles, to H and K band where the images are free of strong residuals exterior to $0\farcs25$ (Figure 2.3). The detection in the broadband images is strongest in the Aug 30 data (especially for the A-LOCI reduction), with a signal-to-noise per resolution element (SNRE) along the trace of the

Table 2.1. PSF Subtraction Algorithm Settings

Data	Param. Tuning	A-LOCI				KLIP			
		g	N_A	$\delta_{\rm FWHM}$	$\Delta r_{\rm sub}$	$N_{\rm PCA}$	$N_{\rm zones}$	$\delta_{\rm FWHM}$	$\Delta r_{\rm sub}$
Aug 30	disk	0.1	1000	1.50	4	5	1	1.25	4
Sep 07	disk	0.1	1000	1.50	4	5	1	1.25	4
Aug 30[a]	companion	1.0	50	0.5, 1	10				
Sep 07[a]	companion	1.0	50	0.75, 1	10				

Note. — Algorithm settings for A-LOCI and KLIP utilized for PSF subtraction of each of the three data sets. 'g' refers to the aspect ratio of the optimization regions, with $g < 1$ producing azimuthally elongated sections and $g > 1$ producing radially elongated sections. 'N_A' refers to the area of optimization regions in units of PSF cores. '$\delta_{\rm FWHM}$' indicates the minimum rotation gap in units of PSF FWHM (for both A-LOCI and KLIP). '$\Delta r_{\rm sub}$' gives the radial size of subtraction regions in units of pixels (for both A-LOCI and KLIP). '$N_{\rm PCA}$' indicates the number of principal components utilized in construction of the model PSF. '$N_{\rm zones}$' is the number of subsections into which each KLIP optimization annulus was divided (with a value of 1 corresponding to full annuli). In all A-LOCI reductions, we also truncated the covariance matrix to zero out (normalized) singular values smaller than 1.25×10^{-6} and constructed a reference PSF from only the 50 most correlated images. a) The two entries for the rotation gap refer to the ADI rotation gap and the SDI radial movement gap.

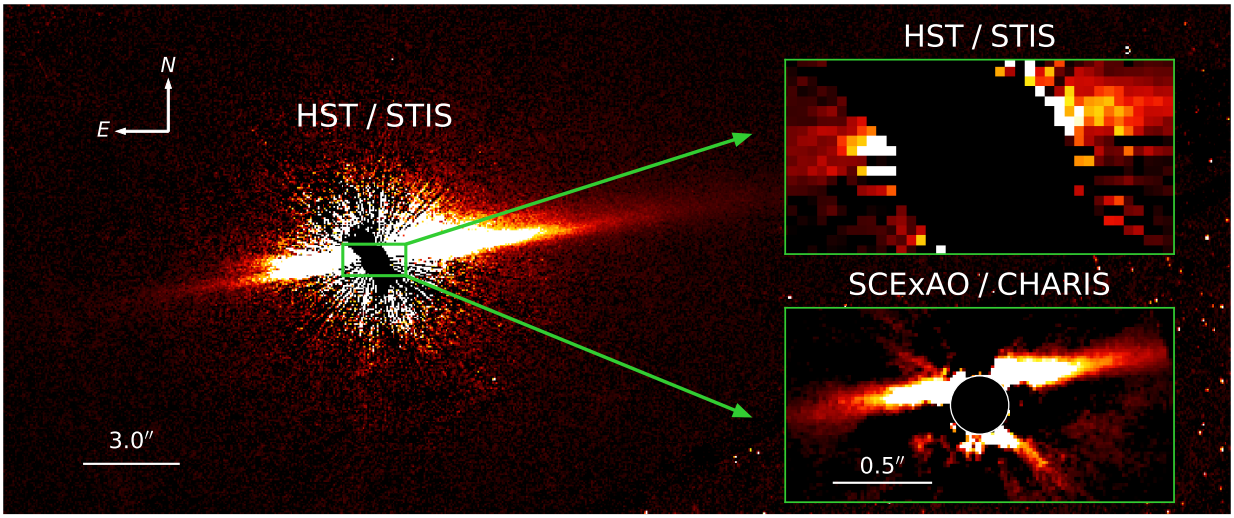

Figure 2.1: HST/STIS imagery of HD 15115 (originally reported in Schneider et al. (2014)) with lower and upper inset images corresponding to CHARIS imagery and the same HST/STIS imagery scaled to the CHARIS field of view. CHARIS imagery presented is the average of the results for A-LOCI reductions of August 30 and September 07 data using settings for disk detection (see Table 2.1). The CHARIS image's central mask has a radius of 0.″15. The STIS, STIS inset, and CHARIS images have the same orientation and are depicted with linear display stretches spanning $0 - 0.02$ $mJy\ arcsec^{-2}$, $0 - 4.0\ mJy\ arcsec^{-2}$, and $0 - 5.21\ mJy\ arcsec^{-2}$ respectively.

disk of $\sim$ 5–7 for most regions exterior to 0.″25[2] (Figure 2.4). In H and K bands, the disk detection is strong, generally achieving SNRE $\sim 3 - 5$ along the disk, and peaking around 5.6 to the west; the J band detection of the disk is considerably weaker (though still definitive), with SNRE $\sim 2 - 3$ over the same regions and peaking around 4 in the west (see Figure 2.5).

CHARIS imagery reveals a strongly asymmetrically scattering disk whose maximum intensity is unambiguously offset (with a projected semi-minor axis of $\sim$ 0.″12) from the system's major axis throughout the $\sim 2'' \times 2''$ field of view. This indicates a view of the system entirely inside the bow-like feature originally described by Rodigas et al. (2012). For the assumption of preferential forward scattering (Hughes et al., 2018), the brighter ($\sim$ northern) side of the disk observed clearly in our data would be presumed as the near side. Signal to the west and slightly south of the center in Figures 2.2 & 2.4 (annotated in the latter) may constitute marginal detections of the disk's dimmer (presumably far) side, which has been recovered in previous ground-based imagery (e.g. Mazoyer et al. 2014; Engler et al. 2019).

[2]In the SNR calculation, we use a software mask to reduce the amount of disk signal included in the noise estimation. This approach increases the finite-element correction penalty (Mawet et al., 2014), so the gain in SNR is small.

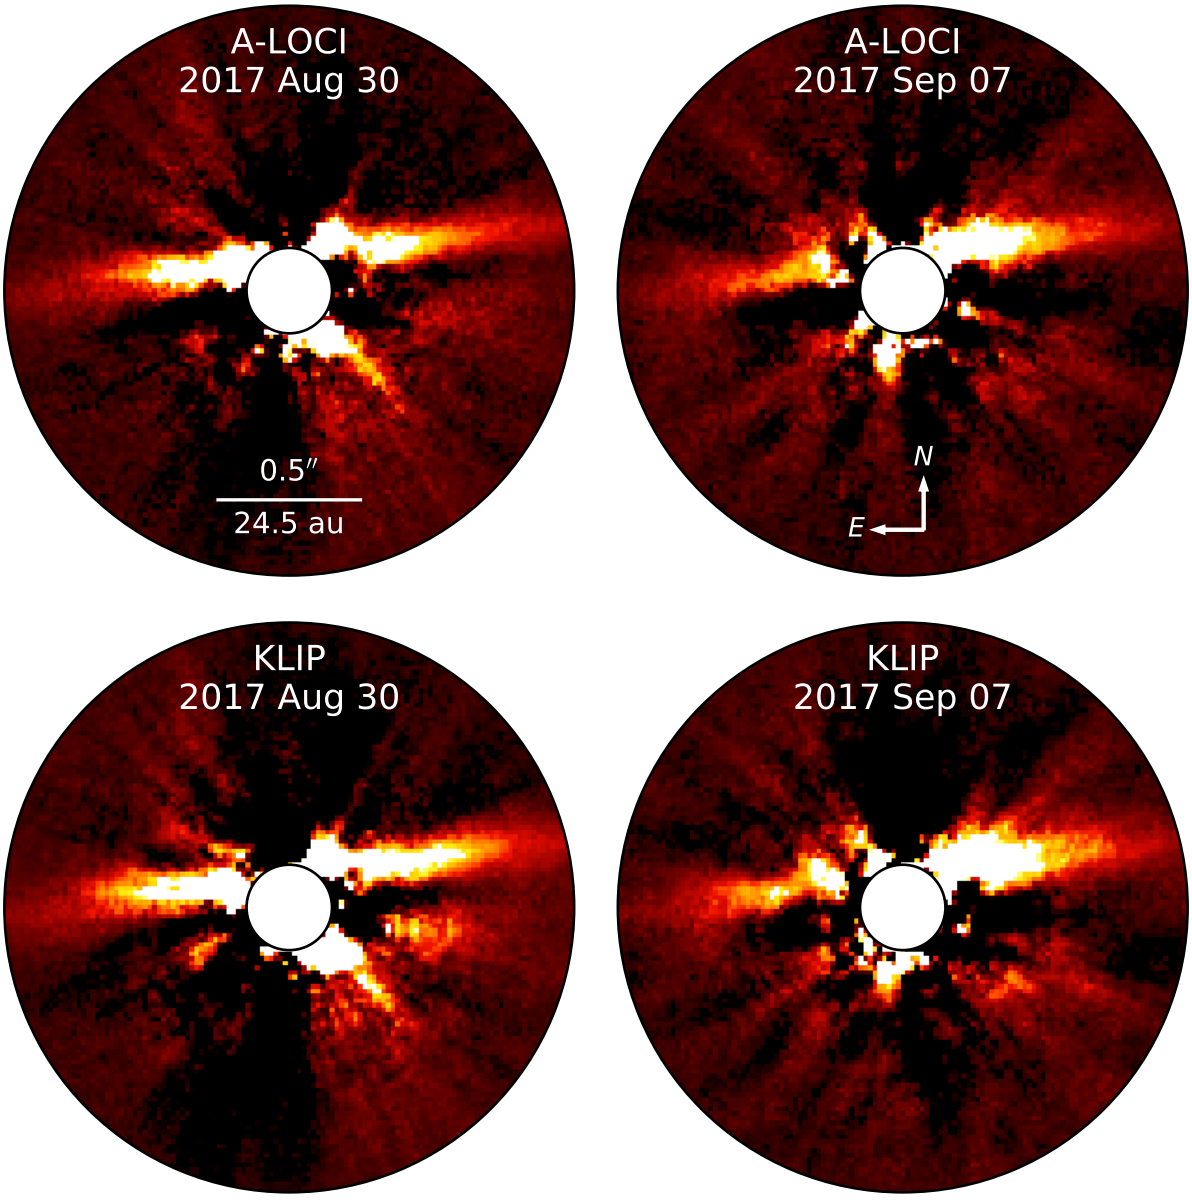

Figure 2.2: Wavelength-collapsed results for August 30 and September 7 observations following PSF subtraction using either A-LOCI or KLIP techniques with settings for disk detection (see Section 2.2.2). The central mask in each subplot has a radius of 0″.15, and the image has a linear display stretch spanning $-0.61 - 6.10\ mJy\ arcsec^{-2}$. In August 30 products, a plainly visible flux enhancement appears just beyond the inner software mask to the southwest, but is not evident in September 07 data. This feature is likely residual speckle noise (likewise for the similar feature to the southeast in September 07 imagery).

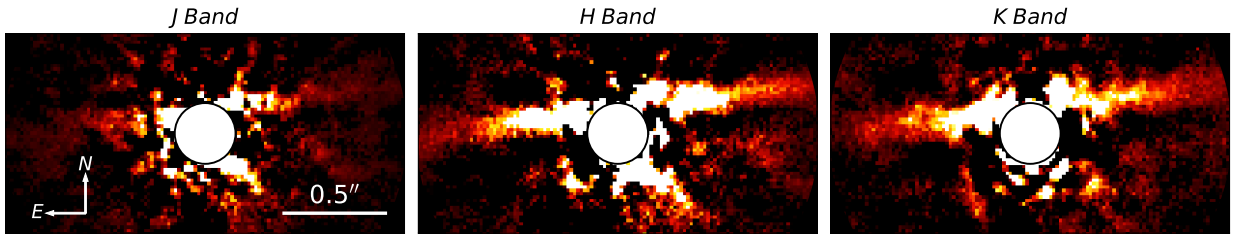

Figure 2.3: A-LOCI PSF subtracted imagery for August 30 using settings for disk detection, with wavelength channels combined to produce images comparable to J (channels $1-5$, $1.16-1.33$ μm), H (channels $8-14$, $1.47-1.80$ μm) and K (channels $16-21$, $1.93-2.29$ μm) bands. The central mask in each subplot has a radius of $0\farcs15$. Images are displayed with linear stretches spanning $0-24.13$ $mJy\ arcsec^{-2}$ (J-band) or $0-4.39$ $mJy\ arcsec^{-2}$ (H-band and K-band).

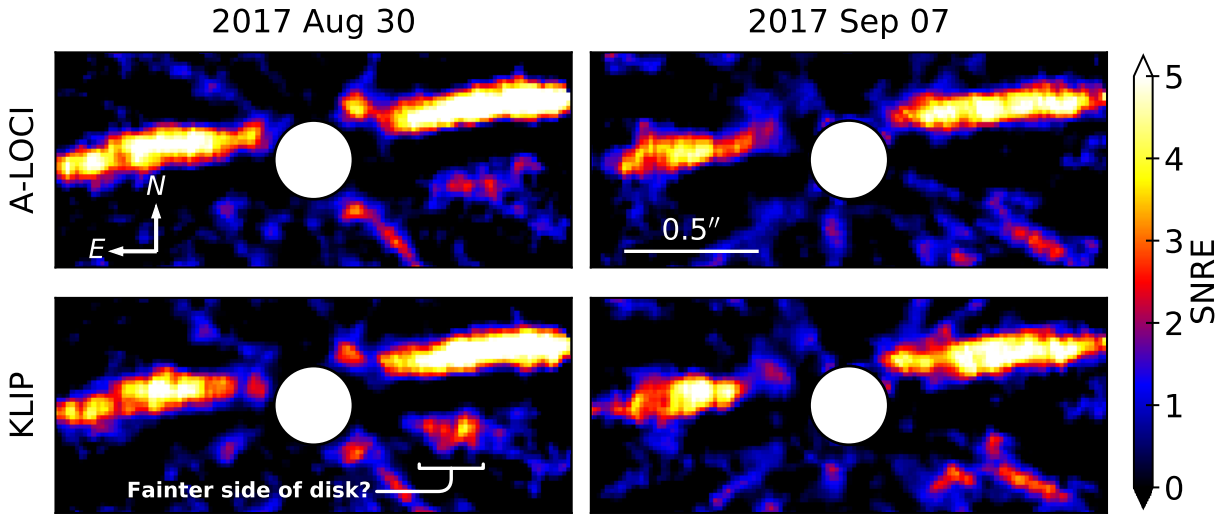

Figure 2.4: Maps of Signal-to-Noise per resolution element for August 30 and September 7 data following PSF subtraction using A-LOCI or KLIP techniques with settings for disk detection (see Section 2.2.2). The central mask in each subplot has a radius of $0\farcs15$. Possible signal from the disk's fainter side is indicated in the subplot for the KLIP reduction of August 30 data, peaking at a SNRE of ~ 4. This feature is also visible in the A-LOCI reduction of the same data, albeit at a lower SNRE (~ 2.5).

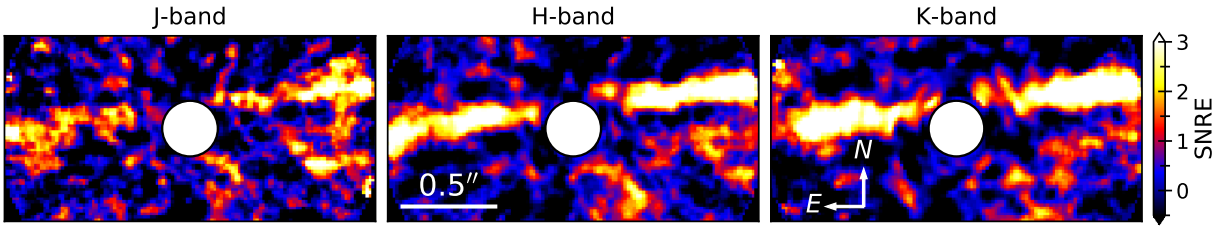

Figure 2.5: Maps of Signal-to-Noise per resolution element for J, H and K bands from A-LOCI PSF subtraction of August 30 data (see Section 2.2.2). The central mask in each subplot has a radius of 0.″15.

2.3. Disk Morphology

To make estimates of the disk's geometric parameters, we seek to identify the position of peak brightness along the disk's bright ($\sim$ northern) edge for each of the wavelength-collapsed (CHARIS broadband) final images resulting from both A-LOCI and KLIP reductions of the August 30 and September 07 data (4 images in total). We begin by rotating the image based on the PA reported in Engler et al. (2019) (PA = $278.^{\circ}9 \pm 0.^{\circ}1$, which is then adjusted to account for the CHARIS PA correction discussed below) to orient the disk's major axis along the x-axis[3]. For each x-direction integer pixel position, we identify the approximate y position of the brightness peak by taking the median of the locations of the brightest 5% of pixels in that column. To determine the precise peak position, we then fit a Lorentzian profile to the array of flux values within 15 pixels ($\sim$0.″25) of the approximate peak, seeking the y position of the profile's peak (and taking the standard error from this fit as the uncertainty). For these fits, each flux value is weighted by the inverse of the corresponding noise levels from noise maps (described in Section 2.4.1). The y-position of the spine as a function of x-position is computed this way for each of the four images being analyzed. The spine profiles for the four images are then combined by taking the weighted average of the four values at each x-position as the nominal average spine y-position, with uncertainty corresponding to the standard error for a weighted average[4]. The resulting average spine profile is then fit with an

[3]The stated PA is assumed here only for the purpose of this initial rotation, which is carried out to simplify fitting of the spine across the narrowest part of the disk, where the peak will appear sharpest. This additionally allows meaningful measurements of the disk's projected FWHM for later use (see Section 2.5). Rotating the images based on the 1-σ upper and lower limits from Engler et al. (2019) instead ($279.^{\circ}0$ and $278.^{\circ}8$ respectively) changes the eventual PA measurement from this procedure by $\sim 0.^{\circ}01$, likely largely as a result of the rotation interpolation itself.

[4]For the jth image's i^{th} x-axis spine position, call the corresponding y-axis spine position y_{ij} with associated fit y-position uncertainty σ_{ij}. The standard error for the corresponding weighted average, μ_i, with weight $w_{ij} = 1 \,/\, \sigma_{ij}^2$, is then:

$$\sigma_i' = \sqrt{\frac{\sum_{j=1}^{4}(y_{ij} - \mu_i)^2 \cdot w_{ij}}{\sum_{j=1}^{4} w_{ij}}}$$

ellipse (which is centered on the star), described by projected semi-major axis (a), nominal position angle (PA_0, which accounts for the initial image rotation applied previously) and inclination (i). Our best fit (over the region $0\farcs25 \leq \rho \leq 1\farcs0$) is achieved for an ellipse with parameters: $\mathrm{a} = 80.1 \pm 3.3$ au, $\mathrm{PA}_0 = 277\fdg82 \pm 0\fdg05$, and $\mathrm{i} = 85\fdg76 \pm 0\fdg22$. The nominal position angle measurement is then corrected for the CHARIS PA offset of $-2\fdg20 \pm 0\fdg27$ [5] (Currie et al., 2018). This results in a final measurement of $\mathrm{PA} = 280\fdg02 \pm 0\fdg27$. The average spine positions and the best fit ellipse are visualized in Figure 2.6.

The position angle measured by this methodology falls above recent measurements made by MacGregor et al. (2019) ($\mathrm{PA} = 278 \pm 1°$) and Engler et al. (2019) ($\mathrm{PA} = 278\fdg9 \pm 0\fdg1$). The measured inclination is consistent with the values of both works (MacGregor et al. (2019) measured $\mathrm{i} = 86\fdg3 \pm 0\fdg4$ and Engler et al. (2019) measured $\mathrm{i} = 85\fdg8 \pm 0\fdg7$). Evaluating χ^2_ν for the spine parameters from Engler et al. (2019) with our data suggests that the difference we measure is significant, with these parameters giving $\chi^2_\nu = 36.6$ versus our best-fit of $\chi^2_\nu = 1.1$. Given that our measurements of the disk are made in the region of $\rho \sim 0\farcs2$–$1\farcs0$, it is possible that we are measuring overlapping signal of the canonical outer ring and the inner ring proposed in both works. If the disk profile observed is the result of an architecture featuring a distinct non-coplanar inner ring (a possibility suggested by Engler et al. 2019), we should expect the fit values to be skewed somewhere between those of the inner and outer component. The difference in measured PA could also be explained by an imprecise calibration for either instrument (or both). However, the CHARIS PA calibration utilized was performed using data collected only a day after our September observations (Currie et al., 2018): it should provide a reasonable assessment of the PA calibration for our data. A reevaluation of the CHARIS north PA and pixel scale using additional data obtained at additional epochs reaffirm these results (T. Currie 2020, in prep.).

Attenuation of disk flux during PSF subtraction can also have an effect on the measured position of the spine, and thus on the derived parameters as well. However, we note that carrying out the aforementioned measurements on attenuated models with known PA from our forward-modeling procedure (see Section 2.4.1) indicates that this effect is small. e.g. for our best-fit one ring model with true $\mathrm{PA} = 279\fdg8$, we measure $\mathrm{PA} = 279\fdg7 \pm 0\fdg3$.

2.4. Modeling the Debris Disk of HD 15115

2.4.1. Disk Forward Modeling

We implement a strategy of forward-modeling synthetic disks, as described in Currie et al. (2018, 2019), to investigate the details of HD 15115's debris disk. In this approach, coefficients (for A-LOCI) or Karhunen-Loève modes (for KLIP) retained from the science data reduction are applied

[5] Hereafter, any values of PA presented (e.g. in the case of disk modeling in Section 2.4.5) are already corrected for this PA offset.

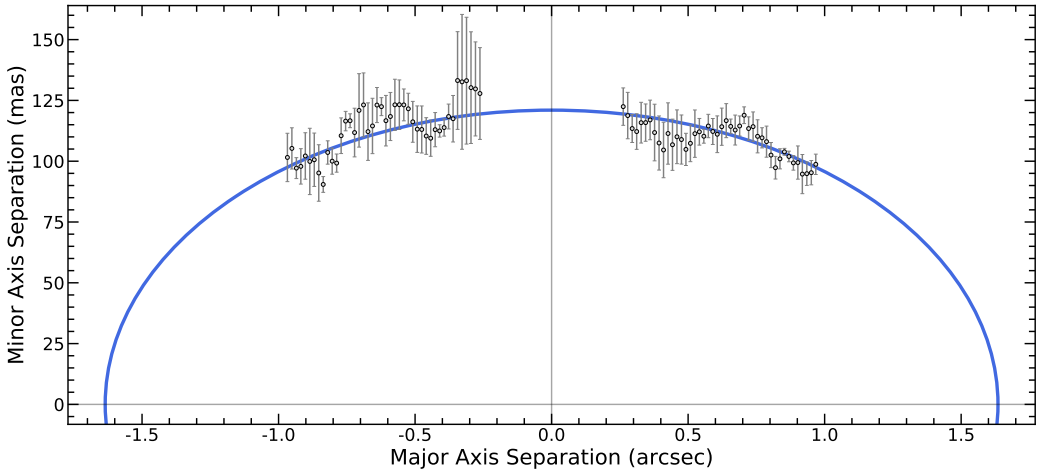

Figure 2.6: Stellocentric separation of disk spine fits along the major and minor axes. The spine positions are used for ellipse fitting and evaluating disk geometry. Points with errors correspond to the weighted average position from all four utilized CHARIS broadband images (see Section 2.3). The blue arc depicts the best-fit ellipse solution for the data shown, and provides the disk major and minor axes against which the spine positions are plotted here.

to image cubes containing only the signal of a model disk which has been rotated to reproduce the array of observed position angles and convolved with the instrumental point-spread function.

For this procedure, we consider all three sources of flux annealing described by Pueyo (2016): over-subtraction (speckle noise being subtracted from disk signal), direct self-subtraction (resulting from inclusion of disk signal in basis vectors), and indirect self-subtraction (resulting from perturbation of basis vectors by disk signal). For our disk reductions, indirect self-subtraction is expected to be the smallest of the three terms. For our KLIP reductions, we retained a small number of KL modes compared to the total number of reference images available. For A-LOCI, we adopted a large optimization area ($N_{\rm A}$ = 1000 PSF cores) that is azimuthally elongated, in contrast to the nearly edge-on debris disk. For both reductions, we adopted a large rotation gap of δ = 1.25–1.5 PSF cores. Thus, oversubtraction and self-subtraction likely dominate over indirect self-subtraction.

Once processed this way, the result for a given model can be compared to the result for science data to assess the relative strength of the model. The procedure for calculating χ^2_ν that we implement is as previously described in Goebel et al. (2018), but is briefly summarized here. First, each value in the model image, f_{model}, and the science image, f_{obs}, is replaced with the sum of values within a FWHM-sized aperture. Following this, a finite element corrected noise map is computed from the science image as described in Currie et al. (2011), additionally utilizing a software mask as described above in Section 2.2.3. Then, the model image is rescaled to minimize the inverse-

variance-weighted residuals with the science image in a region of interest. The difference of the scaled model and the science image, weighted by the noise map, is squared to create a χ^2 map. This map is then binned to the size of the instrumental PSF. Finally, the reduced χ^2 metric is computed from this as $\chi^2_\nu = \chi^2/\nu = \nu^{-1}\sum_i^N (f_{i,obs} - f_{i,model})^2/\sigma_i^2$, where the degree of freedom, ν, is given by the difference between the number of bins in the binned optimization region, N, and the number of free parameters in the model, M (Thalmann et al., 2013).

The region of interest considered is a rectangular box of un-binned dimensions 120 pixels $\times$ 30 pixels ($\sim 2\farcs0 \times 0\farcs5$) centered on the star and oriented to fall along the disk's approximate major axis. The region interior to 6 pixels ($\sim 0\farcs1$) is excluded. For $\rho \lesssim 6$ pixels, not-a-number, or NaN, values begin to appear when no reference frames can meet the minimum rotation gap requirement during PSF subtraction. The region of interest described is overlaid as a white rectangle in Figures 2.9 and 2.10).

We delineate models which are acceptably consistent with our observations as in Thalmann et al. (2013), i.e. those having $\chi^2_\nu \leq \chi^2_{\nu,min} + \sqrt{2/\nu}$.

2.4.2. Model Debris Disks

The forward modeling procedure outlined above was applied to model debris disks generated using a version of the GRaTeR software (Augereau et al., 1999). The formalism and assumptions of the models are detailed in Augereau et al. (1999), but summarized briefly hereafter. The models assume an optically thin disk with a radial dust grain distribution described by a smooth combination of two power laws and with a vertical distribution described by an exponential function. For simplicity, and in line with the analysis of Engler et al. (2019), we set the vertical exponential distribution to be Gaussian in shape ($\gamma = 2$), and restrict the flaring of the disk to be linear ($\beta = 1$). To describe the angular distribution of scattered light, it is common to adopt the Henyey-Greenstein (HG) phase function (Henyey & Greenstein, 1941). The HG phase function is parameterized by a single variable, the asymmetry parameter g, defined as the average of the cosine of the scattering angle, weighted by the (normalised) phase function, over all directions. However, as noted in Hughes et al. (2018), this formalism is not physically motivated and may introduce misleading results. Moreover, a simple HG phase function fails at reproducing the surface brightness profile we observe for HD 15115's disk (see Appendix 2.A). Hong (1985) implements a linear combination of three HG phase functions to describe the observed angular distribution of scattered light for zodiacal dust. To limit model freedom and avoid non-physical solutions, we adopted the phase function of Hong (1985) with the asymmetry parameters and corresponding weights identified therein: $g_1 = 0.7$, $g_2 = -0.2$, and $g_3 = -0.81$, with weights $w_1 = 0.665$, $w_2 = 0.330$, and $w_3 = 0.005$. Though allowing the asymmetry parameters and weights to vary during exploration of disk models may improve the eventual result, it would also massively increase the size and complexity of the

parameter space. As early testing showed that the empirically derived parameters of Hong (1985) reproduced our observed disk surface brightness quite closely, we chose to adopt them as-is to allow a more thorough exploration of the remaining disk parameters. See Appendix 2.A for comparisons of scattering phase functions with our data, including simple HG phase functions using commonly reported asymmetry parameters for HD 15115's disk.

Model geometries investigated fall under two archetypes. The first is a single ring model, defined by 6 parameters:

1. R_0, the radius of peak grain density in au
2. α_{in}, the power law index describing the change in radial density interior to R_0
3. α_{out}, the power law index describing the change in radial density exterior to R_0
4. $\frac{H_0}{R_0}$, the ratio of disk scale height at R_0 to R_0
5. PA, the position angle of the disk in degrees
6. i, the inclination of the disk in degrees

The second is a two ring model, taken to be the linear superposition of two single ring models[6] plus an additional parameter: $F_{max,2}/F_{max,1}$, the ratio of the peak flux of ring 2 to that of ring 1. This results in a model described by 13 parameters (allowing inclination and position angle to differ between the inner and outer disks). To better explore the parameter space of the inner ring, we reduce these to 7 parameters by setting the well-studied outer ring's parameters to approximately match the ones identified by prior studies of the disk (e.g. Engler et al. 2019): $R_{0,1}$ = 96 au, $\alpha_{in,1}$ = 2, $\alpha_{out,1}$ = -3, $\frac{H_{0,1}}{R_{0,1}}$ = 0.03, PA_1 = 278.°9, and i_1 = 86.°0.

For the purpose of χ^2_ν calculation, the overall scaling factor applied to the model (see Section 2.4.1) is considered to be an additional free parameter for the model. This results in $M = 7$ for the one ring model and $M = 8$ for the two ring model.

2.4.3. Model Optimization Using Differential Evolution

Though HD 15115's disk has been studied extensively in the region beyond 1.″0, the small separations observed with CHARIS provide a look at the disk in the $0.''2 - 1.''0$ regime, where the parameters of the posited inner ring could potentially be studied in much greater detail. However, given the nearly edge-on orientation of the system, a parameter space with significant model degeneracies and multiple local minima is possible. A broad but detailed search of the parameter spaces

[6] e.g. by coadding the synthetic images for the individual models, as in Boccaletti et al. (2019). This assumes that the rings are sufficiently optically thin that single-scattering dominates over multiple-scattering.

outlined in Section 2.4.2 is necessary to offer a meaningful assessment of any such degeneracies and to ensure that a unique and globally optimum solution is identified. A grid search for the 7 parameter two ring model quickly reaches an intractably large size; e.g., a coarse grid examining only 5 values of each parameter would require 78125 models be propagated through the time consuming forward modeling procedure (typically $\sim$ minutes per model). While Markov-Chain Monte-Carlo (MCMC) techniques are commonly used for similar purposes (e.g MacGregor et al. 2019), they can become trapped in local minima and their results can be dependant on the initialization. Moreover, MCMC exploration typically requires a number of model evaluations that is effectively unapproachable for ADI forward modeling procedures comparable to ours (e.g. the MCMC procedure of MacGregor et al. 2019 evaluates $\sim 10^6$ models).

Instead, we make use of the differential evolution algorithm (DE, Storn & Price 1997) to explore possible solutions for the parameters of each model. DE requires no initial assumptions about a solution, beyond boundaries within which to explore, and is capable of efficiently probing large, correlated parameter spaces by quickly evolving a population of trial solutions away from regions that offer inferior solutions and allowing population members to move between local minima. Though DE has not seen widespread use in the optimization of scattered light disk models, it has been used elsewhere in the study of astrophysics with noteworthy efficacy, e.g. to explore optimal capture trajectories for Jovian orbiters by the European Space Agency (Labroquère et al., 2014), to identify edge-on galaxies with abundances of extraplanar dust (Shinn, 2018), or to search for flaring stars in sparsely sampled time-series data (Lawson et al., 2019). While DE does not enable the robust determination of parameter likelihood distributions in the same manner as MCMC, it is extremely effective at quickly reaching a global solution with little to no tuning of algorithm control variables or specific experience with the algorithm itself, and with relatively few function evaluations needed (e.g. see comparison benchmarks in Storn & Price 1997). Additionally, the relative simplicity of DE makes it trivial to add into existing frameworks (the C-style psuedo-code for the algorithm presented in Storn & Price 1997 requires only 19 lines of code). For groups currently exploring model parameters using grid searches: in all but the rarest cases, DE will tend to identify a superior final model while evaluating many fewer models overall. For groups that might be interested in adopting this technique, we include a simple Python implementation of differential evolution, whose procedure is outlined below, in Appendix 2.B.

In the differential evolution procedure, we initialize a population of N_{pop} random model parameter sets, restricted to fall between boundaries set for each parameter. For each type of model, N_{pop} is set to be 10 times the number of free parameters[7] (60 for the single ring model, and 70 for the two

[7] Storn & Price 1997 suggest 5–10 population members per free parameter, though larger values are often used in recent implementations as well, e.g. the default value of 15 per free parameter in the implementation from the Python package SciPy (Virtanen et al., 2020).

ring model). The initial model population is run through the forward modeling routine to evaluate the fitness of the models (χ^2_ν). Following this, a mutation (v_i) for each member of the population (x_i) is created by adding a scaled difference of the parameters of two distinct, random population members (x_j, x_k) to the parameters of the current best solution (x_{best}): $v_i = x_{best} + m(x_j - x_k)$ (this "strategy" is called "best/1/bin" by the notation of Storn & Price 1997). For our purposes, the value of the mutation constant, m, is randomly selected in the range [0.5, 1.0] for every generation[8]. From a given mutation, a trial replacement (u_i) is formed by setting $u_i = x_i$ and allowing probability $P = 0.7$[9] for each parameter value in u_i to be replaced by the corresponding value from v_i. The fitness of the set of N_{pop} trial replacements is then evaluated using the same forward modeling routine. Finally, each member of the population, x_i, is replaced by its corresponding trial vector, u_i, if the fitness of the trial vector is superior to that of the population member. The procedure of creating and evaluating trial replacements for the population is repeated until the population becomes stagnant or converges to a single solution.

Fitness is evaluated at each stage by computing the combined χ^2_ν for both A-LOCI reductions as:

$$\chi^2_\nu = \frac{\chi^2_1 + \chi^2_2}{N_1 + N_2 - M} = \frac{1}{\nu}\sum_{i=1}^{2}\chi^2_i \ ,$$

where subscripts 1 and 2 correspond to the values for A-LOCI reductions of August 30 and September 07 data. Once the DE procedure is completed, the subset of models meeting the threshold $\chi^2_\nu \leq \chi^2_{\nu,min} + \sqrt{2/\nu}$ are then propagated through the forward modeling procedure for the KLIP reductions. The final best model is taken from among these as the one which minimizes the combined χ^2_ν for all four reductions:

$$\chi^2_{\nu,\mathrm{tot}} = \frac{\sum_{i=1}^{4}\chi^2_i}{(\sum_{i=1}^{4}N_i) - M} = \frac{1}{\nu}\sum_{i=1}^{4}\chi^2_i \ ,$$

with $i = 3$ and $i = 4$ indicating the two KLIP reductions. While it may be preferable to evaluate each set of trial models for all four reductions, this procedure cuts the total model optimization time nearly in half by assuming that the overall best model will be contained within the 'acceptable' fitness bounds of the first two reductions. See column 2 of Table 2.2 and Table 2.3 for the bounds adopted for each parameter.

[8] Storn & Price 1997 introduce a single mutation constant and suggest a value of 0.5. The "dithered" mutation constant implemented here is adopted from the default setting in the SciPy module for Python, where the authors suggest that a dithered mutation constant will typically speed convergence substantially (Virtanen et al., 2020).

[9] A crossover probability, P, anywhere from 0.1 to 0.9 is recommended in Storn & Price 1997. P=0.7 is adopted from the default value in the SciPy module for Python.

2.4.4. Model Results

One Ring Models – The differential evolution procedure for single ring models yields an optimal combined fit for the two A-LOCI reductions of χ^2_ν = 1.145. A visualization of the full sample of models explored using differential evolution is provided in Figure 2.7. Some parameters (α_{in} and α_{out}) have converged to the boundaries. While this could indicate that the bounds are too restrictive, we note that the adopted boundaries include all values of these parameters explored by other recent studies of HD 15115 utilizing GRaTeR (e.g. Mazoyer et al. 2014 use models with $-6 \leq \alpha_{out} \leq -4$ and $2 \leq \alpha_{in} \leq 10$, and Engler et al. 2019 use models with $-8 \leq \alpha_{out} \leq -2$ and $2 \leq \alpha_{in} \leq 10$). This result is discussed in more detail in Section 2.4.5. However, we note here that these particular parameters are unimportant for our overarching conclusions.

From the 660 models evaluated, 13 resulted in acceptable values of χ^2_ν ($\chi^2_\nu \leq 1.188$, for $\nu = 1077$ with the two A-LOCI reductions included). After forward modeling this subset for the remaining reductions, the final best model results in $\chi^2_{\nu,\mathrm{tot}} = 1.166$ with a revised acceptable limit of $\chi^2_{\nu,\mathrm{tot}} \leq 1.196$ (for $\nu = 2161$ when all four reductions are included). The parameters for the best overall model and ranges of acceptable parameters are included in Table 2.2. It should *not* be assumed that an acceptable solution can be produced for any arbitrary combination of parameter values falling within the acceptable ranges. A given value included in the acceptable range indicates that the value, paired with specific values of the other parameters, produces a model meeting the given threshold for acceptability.

Two Ring Models – Following the initial DE procedure, the optimal fit for two ring models results in $\chi^2_\nu = 1.169$ for the two A-LOCI reductions. A visualization of the full sample of models explored using differential evolution is provided in Figure 2.8. The best solution at this stage has an inner ring radius that falls at the upper boundary of allowed values. We note, however, that the selected bounds include the best-fit inner rings of both MacGregor et al. (2019) ($\mathrm{R}_{0,2} \sim 48$ au) and Engler et al. (2019) ($\mathrm{R}_{0,2} = 64$ au).

From the 910 models evaluated here, 20 resulted in acceptable values of χ^2_ν ($\chi^2_\nu \leq 1.212$ with $\nu = 1076$). After KLIP forward modeling for this subset, the final best model resulted in $\chi^2_{\nu,\mathrm{tot}} = 1.151$ with a revised acceptable limit of $\chi^2_{\nu,\mathrm{tot}} \leq 1.181$ ($\nu = 2160$). Of the 20 models evaluated for all four reductions, 12 pass this revised threshold for acceptability. For clarity, we note that the best model from the initial DE/A-LOCI only reductions (the model whose parameters are indicated in Figure 2.8) is ultimately excluded by this final acceptability threshold. This model simply ends up being a worse explanation for the KLIP results than other acceptable models, which pushes its combined score up sufficiently to be eliminated (with $\chi^2_{\nu,\mathrm{tot}} = 1.183$). The parameters for the best overall model and ranges of acceptable parameters are included in Table 2.3.

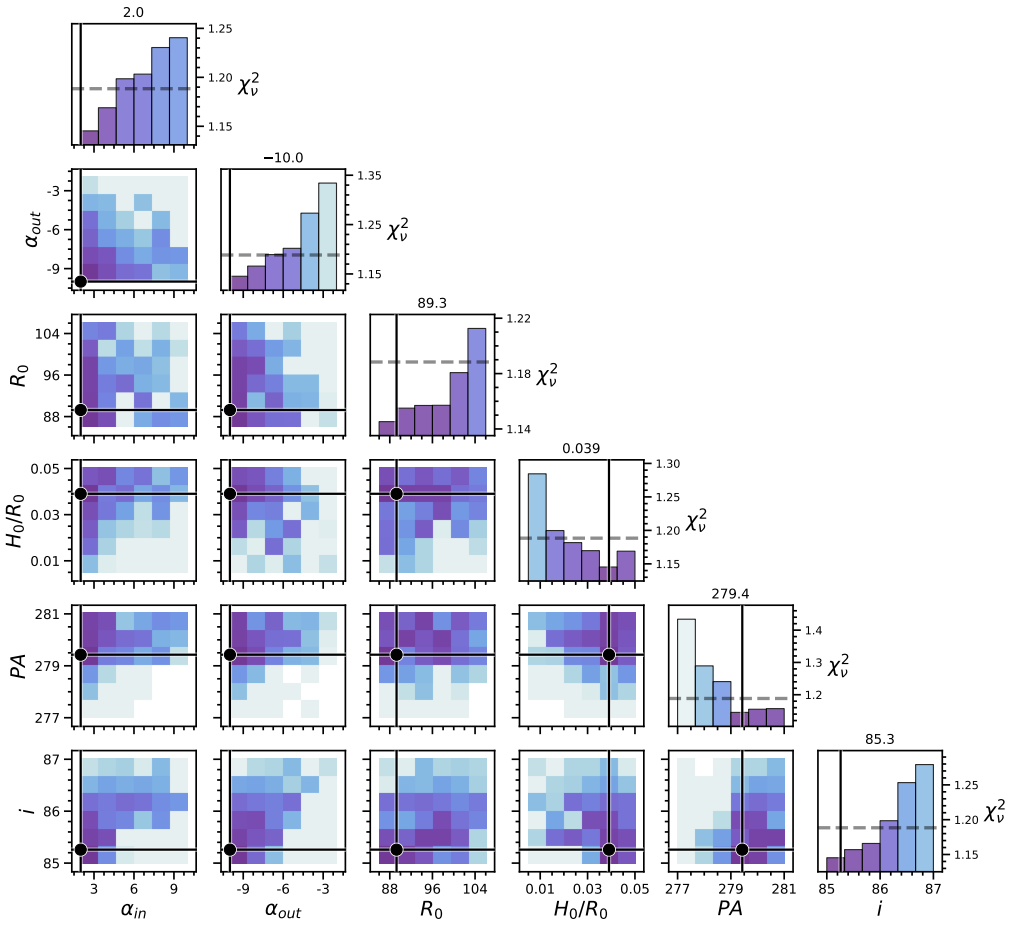

Figure 2.7: A variation of the corner plot for optimization of the single ring model to the A-LOCI reductions of HD 15115's August 30 2017 and September 07 2017 data. Each off-diagonal plot visualizes solutions as a function of two of the parameters, with each bin colored according to the quality of the best fit achieved with values of the two parameters in that range (and any values of the other parameters). Darker bins indicate smaller values of χ^2_ν, where χ^2_ν refers to the combined metric for the A-LOCI reductions of both data sets (see Section 2.4.3). Diagonal elements provide a one-dimensional view of each of the parameters, indicating the lowest χ^2_ν value (y-axis) achieved for the binned range of the given parameter (x-axis). For each parameter, the area within the bounds provided in Table 2.2 is divided into 6 equally sized bins. The best-fit solution values are indicated by black crosshairs, and their values given above the corresponding one-dimensional subplot. The threshold fitness for acceptable solutions with A-LOCI reductions ($\chi^2_\nu \leq \chi^2_{\nu,min} + \sqrt{2/\nu}$) is indicated by a horizontal grey dashed line in each diagonal plot; bars which end below this threshold line resulted in models meeting the acceptable fitness criteria for some values of the other parameters.

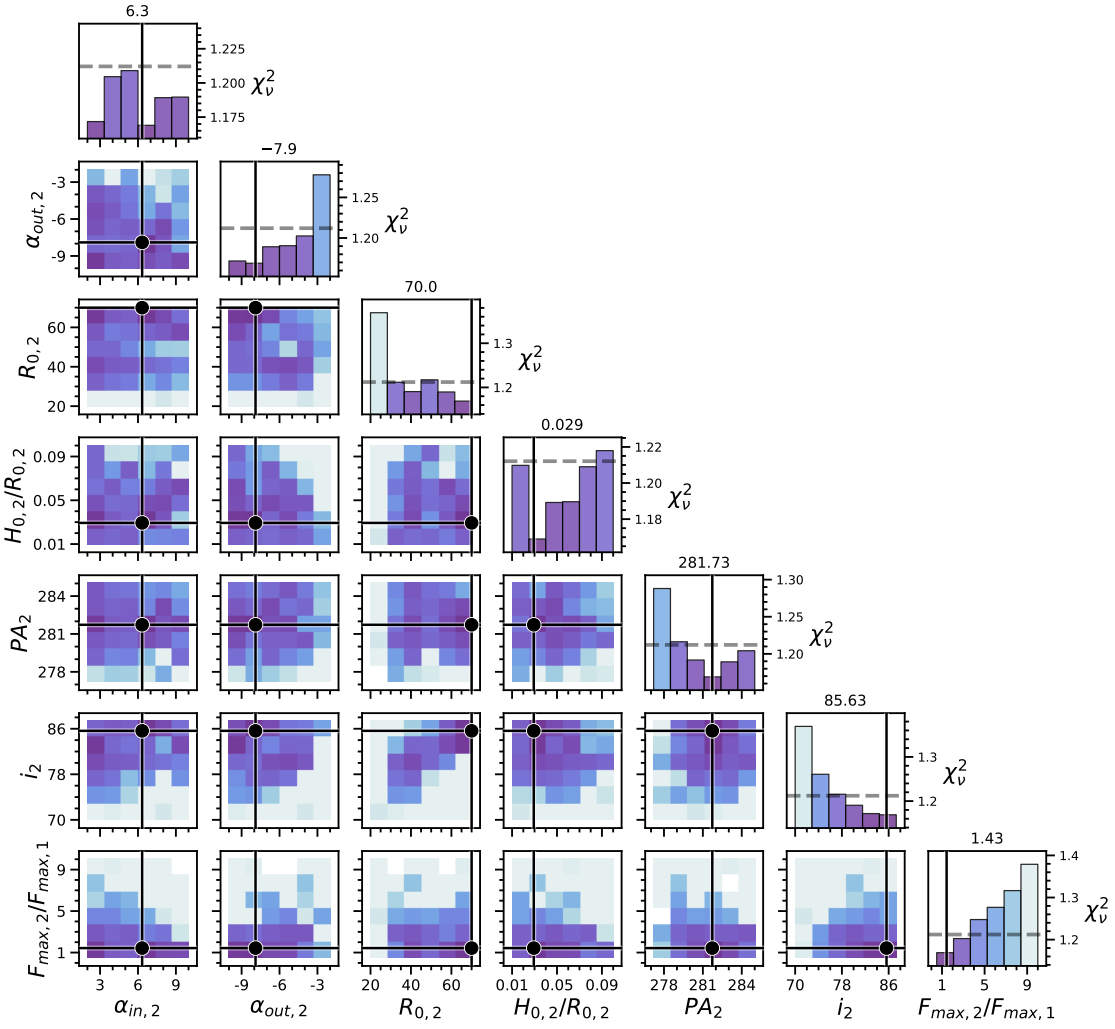

Figure 2.8: As Figure 2.7 but for optimization of the two ring model.

Table 2.2. One Ring Model Optimization Results

Parameter	Bounds	Best	Acceptable[a]
R_0 (au)	86.0 – 106.0	93.2	87.7 – 99.6
α_{in}	2.0 – 10.0	2.0	2.0 – 3.6
α_{out}	-10.0 – -2.0	-9.6	-10.0 – -7.5
H_0/R_0	0.01 – 0.05	0.05	0.04 – 0.05
PA (deg)	277.0 – 281.0	279.8	279.4 – 280.5
i (deg)	85.0 – 87.0	85.3	85.3 – 85.7
$\chi^2_{\nu,\mathrm{tot}}$	—	1.166	≤ 1.196

Note. — Optimization bounds, best fitting value, and the range of acceptable values for each varied parameter of the single ring model following propagation through A-LOCI and KLIP forward modeling procedures for August 30 and September 07 data. $\chi^2_{\nu,\mathrm{tot}}$ indicates the combined measure for all four reductions (see Section 2.4.3). $a)$ These ranges give the smallest and largest value of each parameter that resulted in an acceptable solution. Given the possibility of complicated correlations between parameters and the lack of perfect sampling, it cannot be stated conclusively that every value within these ranges can produce an acceptable solution.

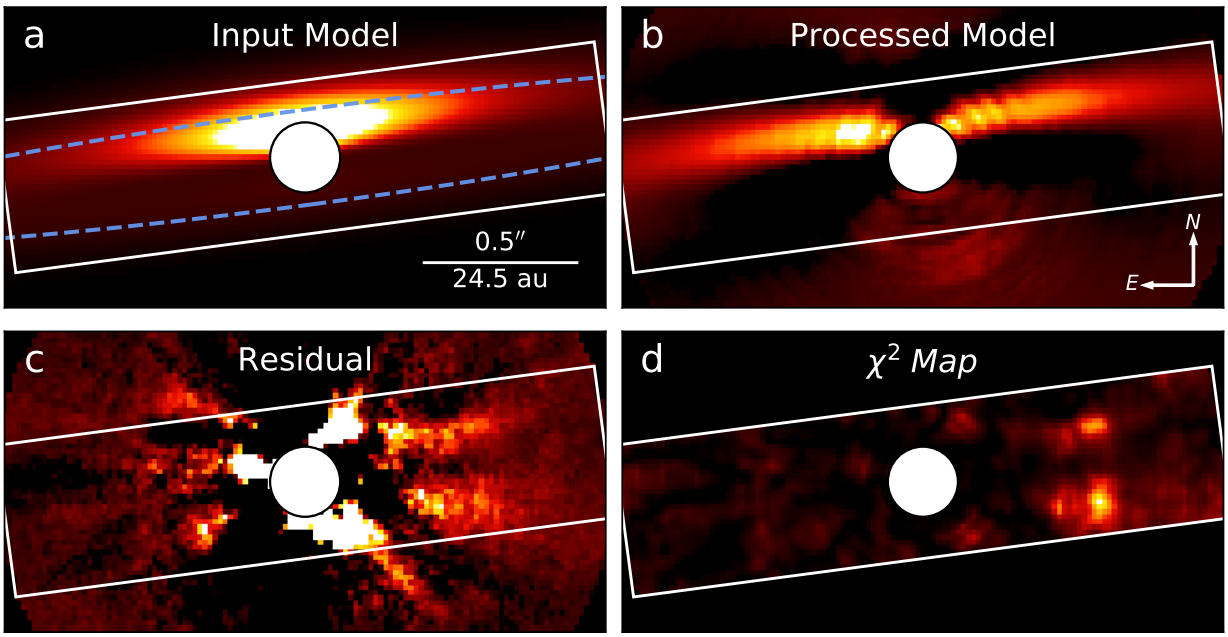

Figure 2.9: The lowest $\chi^2_{\nu,\mathrm{tot}}$ single ring model ($\chi^2_{\nu,\mathrm{tot}} = 1.166$) identified by the procedure outlined in Section 2.4.1, depicted for our KLIP reduction of August 30 data. Model parameters utilized can be found in Table 2.2. In all panels, the white rectangle indicates the "region of interest" for the purpose of χ^2 calculation, with the white circle indicating the inner exclusion radius. **a**) the input disk model convolved with the instrumental PSF. A schematic of the disk model is overlaid as a dashed blue ellipse having the same radius, inclination, and PA as the model disk. **b**) the disk model after application of the forward modeling procedure for the August 30 2017 KLIP reduction, displayed exactly as the images of Figure 2.2 (linear stretch over $-0.61 - 6.10\ mJy\ arcsec^{-2}$). **c**) Residuals for the August 30 2017 KLIP data product (upper left image in Figure 2.2) and the processed disk model, also displayed as Figure 2.2 (linear stretch over $-0.61 - 6.10\ mJy\ arcsec^{-2}$). **d**) unbinned χ^2 map for this model displayed with linear stretch over 0 to 19.46 in χ^2 (roughly equivalent to $0 - 0.036$ in χ^2_ν).

Table 2.3. Two Ring Model Optimization Results

Parameter	Bounds	Best	Acceptable
$R_{0,2}$ (au)	20.0 – 70.0	40.9	36.3 – 62.1
$\alpha_{in,2}$	2.0 – 10.0	3.0	2.0 – 8.2
$\alpha_{out,2}$	-10.0 – -2.0	-5.3	-9.4 – -4.6
$H_{0,2}/R_{0,2}$	0.01 – 0.10	0.03	0.02 – 0.08
PA_2 (deg)	277.2 – 285.0	281.6	280.2 – 283.8
i_2 (deg)	70.0 – 87.3	80.1	79.0 – 85.0
$F_{max,2}/F_{max,1}$	0.5 – 10.0	2.0	1.2 – 2.7
$\chi^2_{\nu,\mathrm{tot}}$	—	1.151	≤ 1.181

Note. — As Table 2.2, but for two ring model optimization.

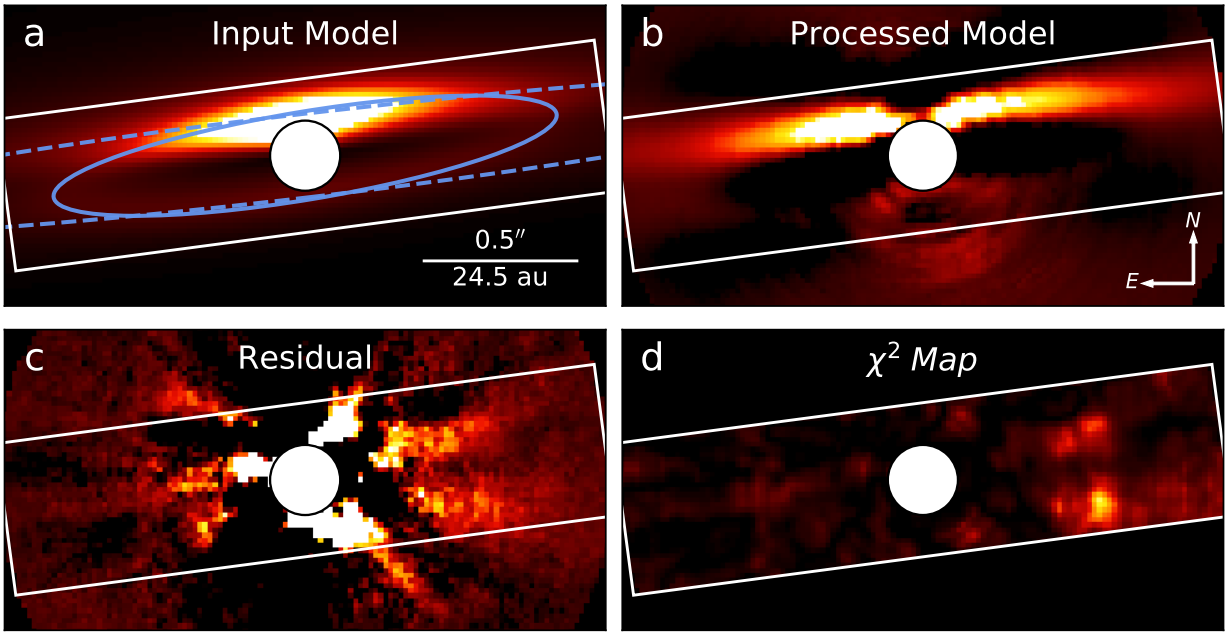

Figure 2.10: The lowest $\chi^2_{\nu,\mathrm{tot}}$ two ring model ($\chi^2_{\nu,\mathrm{tot}} = 1.151$) identified by the procedure outlined in Section 2.4.1. The schematic overlaid in panel a) shows the inner ring as a solid blue ellipse, and the outer as a dashed blue ellipse. The χ^2 map shown in panel d) is displayed with identical scaling to that of Figure 2.9. Otherwise, as Figure 2.9.

2.4.5. Modeling Discussion

With the optimal one and two ring models identified producing comparable χ^2_ν metrics of 1.166 and 1.151 respectively, both geometries appear statistically consistent with our data – making it difficult to rule out either scenario. However, our exploration of the model parameter spaces using differential evolution allows us to place some disambiguating constraints.

A few noteworthy observations can be made regarding the one ring model results. Firstly, the corner plot of our differential evolution procedure (Figure 2.7) reveals a preference for a very slow density change interior to R_0 (small α_{in}), with a very rapid change exterior to R_0 (large α_{out}). For both of these parameters, the optimization converges at, or very near to, the boundaries (2 and -10 respectively). While this could suggest true values for these parameters beyond the boundaries we've set, this seems unlikely given the results of previous studies of HD 15115's disk; for example, the observed extent of the disk with wider fields of view is inconsistent with a radial density profile having $\alpha_{out} \sim -10$. We note, however, that models more consistent with prior results produce acceptable fit metrics as well (see Table 2.2). Further, given our narrow field-of-view and the disk's high inclination, we should expect that our results are not as sensitive to changes in these particular parameters anyway. Additionally, the PA identified is seemingly distinct from values typically found in previous studies of the disk, with our procedure finding PA $= 279.^{\circ}8^{+0.7}_{-0.4}$ [10] compared to PA $= 278.^{\circ}9 \pm 0.^{\circ}1$ from Engler et al. (2019). However, the value identified here is consistent with the value of 280.02 ± 0.27 identified from ellipse fitting of the spine in Section 2.3.

From the schematic of our best two ring model (Figure 2.13, top panel), we see that our data is best explained by an inner ring with a projected semi-minor axis similar to that of the outer ring ($b_{proj} \sim 6.7\ au$ for $R_0 = 96\ au$ and i $= 86°$). Looking at the corner plot for the two ring optimization (Figure 2.8), perhaps the strongest apparent correlation occurs between the inner ring's radius (equivalently its projected semi-major axis, a) and its inclination. Noting that the projected semi-minor axis, b, of the inner ring is related to its inclination and projected semi-major axis by $b_{proj} = a_{proj} \cdot cos(\mathrm{i})$, the correlation in $R_{0,2}$ versus i_2 subplot falls very nearly along the line corresponding to $b_{2,proj} = 6.7\ au$ (see Figure 2.11). In fact, for the full set of acceptable two ring models following differential evolution, all have projected semi-minor axes between 4.8 and 8.0 au – while values from 0.9 to 23.9 au are permitted by our parameter bounds. The revised set of acceptable models following modeling of all four CHARIS reductions reduces this range even further, to 5.4 – 8.0 au. Combined with our analysis of the disk's spine and surface brightness profiles, our results appear to suggest a lack of any statistically significant distinct inner ring spine.

Overall, the results of our modeling procedure can be interpreted in a number of ways:

[10] Uncertainties here are roughly approximated as simply the upper and lower limits for acceptable models as presented in Table 2.2.

1. From the strong preference for $b_{1,proj} \sim b_{2,proj}$: a distinct inner ring exists but its brightest features happen to roughly line up along our line of sight with the canonical outer ring. This would result in a two ring geometry that is statistically indistinguishable from one with a single ring. A distinct inner ring that is coplanar (or nearly coplanar) with the outer ring is consistent only as the inner ring's radius approaches that of the outer, where $b_{2,proj}$ can near $b_{1,proj}$ while maintaining a matching inclination.

2. In the CHARIS field of view ($\sim 0\farcs2 - 1\farcs0$) either the outer or inner ring is substantially brighter, such that the other is not recovered in our data. The differing PA identified for the one ring model compared to literature could suggest that we're seeing the latter. However, we remark that our optimization procedure for two ring models allowed values of $F_{max,2}/F_{max,1}$ as large as 10; if the underlying system is well-described within CHARIS's field of view as an extremely faint outer ring with a misaligned bright inner ring, solutions with large ring flux ratios should have manifested.

3. Perhaps the system is truly better described as a single ring. The tendency of the single ring DE procedure toward parameter values which are seemingly at odds with prior results could be caused by inaccuracies in assumptions made by our models. e.g. if the true scattering phase function departs slightly from the Hong phase function assumed, if the disk is non-negligibly eccentric, or if the disk features a non-linear flaring profile, models matching the true underlying parameters may not coincide with the minimum χ^2_ν in our analysis. In this case, even if no second ring exists, a second ring model might serve to mitigate these inaccuracies sufficiently to result in competitive fitness metrics.

4. The disk has a geometry distinct from any probed here and that is difficult to diagnose as a result of the nearly edge-on orientation (e.g. a debris disk with spiral arms or significant warping).

We also point out the presence of the very small separation ($\rho \sim 0\farcs25$) residual signal that is not fit by our models (bright residuals appearing just beyond the software mask in panel c of Figures 2.9 and 2.10). While this feature could be evidence of the ~ 4 au "warm-dust" disk suggested by Moór et al. (2011) through SED fitting, perhaps the more likely explanation is that it is simply residual speckle noise. While "aggressive processing" is often cited as a cause of spurious features in debris disks (Duchêne et al., 2020), poor speckle suppression due to *insufficiently* aggressive PSF subtraction can also cause spurious features at small angular separations. Indeed, separate tests with more aggressive A-LOCI and KLIP settings for our data appear to confirm that this signal is due to residual speckles, but these approaches compromise our detections of the disk at $0\farcs25$–$1''$ and thus are not used.

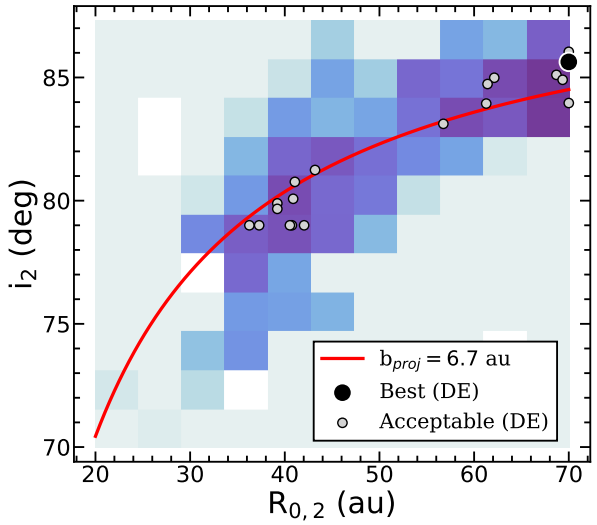

Figure 2.11: Reproduced from Figure 2.8: best χ^2 for two ring models as a function of inner disk radius and inclination. Values corresponding to the best model and acceptable models following the DE procedure of Section 2.4.1 (i.e. χ^2_ν accounting for only the A-LOCI reductions) are shown in black and grey respectively. The red line indicates values of radius and inclination that produce a projected semi-minor axis equal to that of the adopted outer ring. Acceptable two ring models are identified only for inner disk parameters resulting in nearly the same projected semi-minor axis (see Section 2.4.5).

2.4.6. Comparison with Recent Studies

Based on SPHERE polarimetry, Engler et al. (2019) suggest the possibility of a misaligned/non-coplanar inner ring with a radius of $\sim 1\farcs3$ – ultimately finding a slightly better fit to their non-polarized data for this geometry than for a one-ring or coplanar two-ring geometry. They further investigate the merit of the additional parameters of the two-ring model by comparing the Bayesian Information Criterion (BIC) for the two models, concluding that the BIC metric supports their best-fit non-coplanar two-ring model. While our modeling shows that CHARIS data is consistent with a misaligned inner ring for some combinations of inclination and PA (see Sections 2.4.5 and 2.4.4), we find no clear evidence indicating the presence of an inner ring oriented as hypothesized from SPHERE imagery (PA $= 276°$, i $= 80°$; see bottom panel of Figure 2.13). Carrying out forward modeling on the best-fitting misaligned two ring model identified in Engler et al. (2019) appears to reinforce this, with the model producing a $\chi^2_{\nu,\mathrm{tot}}$ of 2.13 (for the same assumption of M=10 free model parameters that they indicate) after forward modeling for our four reductions, compared with 1.151 for the overall best two ring model we identify (see fourth row of Figure 2.A.1). Even allowing freedom for the other parameters, our two ring optimization identified no strong solutions having the inner ring oriented similarly (see Figure 2.8). The best coplanar two ring model they identify fits our data somewhat better, resulting in a $\chi^2_{\nu,\mathrm{tot}}$ of 1.7 (see bottom row of

Figure 2.A.1). Models with a similar inner ring radius and roughly coplanar orientation manifest in our final set of acceptable models (Table 2.3) when using the Hong scattering phase function instead.

From ALMA observations, MacGregor et al. (2019) favor a disk model composed of two coplanar rings or a single ring with a Gaussian gap. Given the differences in parametizations between our models, it is difficult to unambiguously translate their results for direct application to our data. However, the gap suggested by their models in either case is small enough (~ 14 au) that the profile should manifest consistently with the profile we observe (e.g. with the appearance of a single spine in our imagery; see Section 2.4.5 for relevant discussion regarding this constraint). We also note that our spine trace (Figure 2.6) shows a $\sim 1\sigma$ shift around $0\farcs8$ on either side which is roughly coincident with the inner edge of the inner ring they propose. By applying a simple 3-pixel rolling weighted average to the spine trace and disk projected FWHM measurements, this feature becomes more clear (see Figure 2.12). While we find no significant evidence to support the presence of the ~ 14 au gap that they favor (given that it falls outside of our field of view, with $\rho \sim 1\farcs2$), their interpretation appears generally consistent with CHARIS imagery.

The smoothed disk spine in Figure 2.12 also appears remarkably similar to the single profile spine trace of NZ Lup's disk (another highly inclined debris disk) reported in Boccaletti et al. (2019) (their Figure 3). Boccaletti et al. (2019) ultimately favor a mutually inclined ($\Delta i \sim 5°$) two ring model for NZ Lup with an apparent gap that is roughly coincident with the dip in spine position seen in their spine trace. If a comparable explanation is assumed for the $\sim 0\farcs75 - 0\farcs80$ feature in our smoothed spine trace, the result would be a two ring disk with a gap similar to the one suggested by MacGregor et al. (2019), but at a somewhat smaller separation than their best fit. Though, given that the ALMA observations trace a significantly different dust population than ours, these results may be fully consistent with one another. Notably, this interpretation manifests similarly to the overall best two ring model identified in Section 2.4.4 (see Table 2.3), with a fiducial inner ring radius of 40.9 au ($0\farcs83$) and a mutual inclination of $5\fdg9$. The slight difference in the location of peak FWHM measured for the disk between the east and west sides in Figure 2.12 ($\sim 0\farcs75$ and $\sim 0\farcs82$ respectively) might be explained by the difference in PA between the rings suggested by our best two ring model ($\Delta\mathrm{PA} = 2\fdg7$).

2.5. Disk Surface Brightness and Color

2.5.1. Disk Photometry

To analyze the brightness and color of HD 15115's disk, we follow the general procedure of Goebel et al. (2018) to produce surface brightness profiles in CHARIS broadband and J, H, and K bands. However, instead of fitting a fourth order polynomial to the identified disk spine (see Section 2.3), we adopt the positions of the best-fit ellipse as the location of the disk spine for all imagery

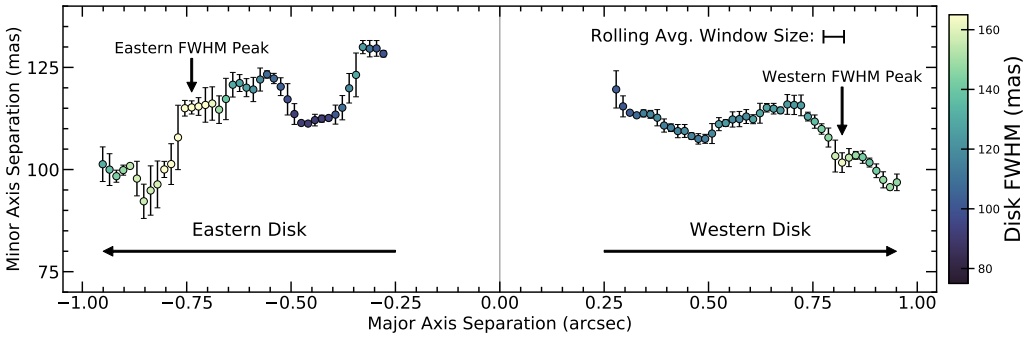

Figure 2.12: Disk spine position and projected FWHM measurements made in Section 2.3 with a 3 pixel wide rolling weighted average applied. At $\sim 0\farcs75 - 0\farcs80$ on either side, a $\sim 1 - 2\sigma$ dip in minor axis separation is coincident with an increase in measured disk FWHM. Further, a peak in FWHM can be seen on both sides, but occurs slightly asymmetrically.

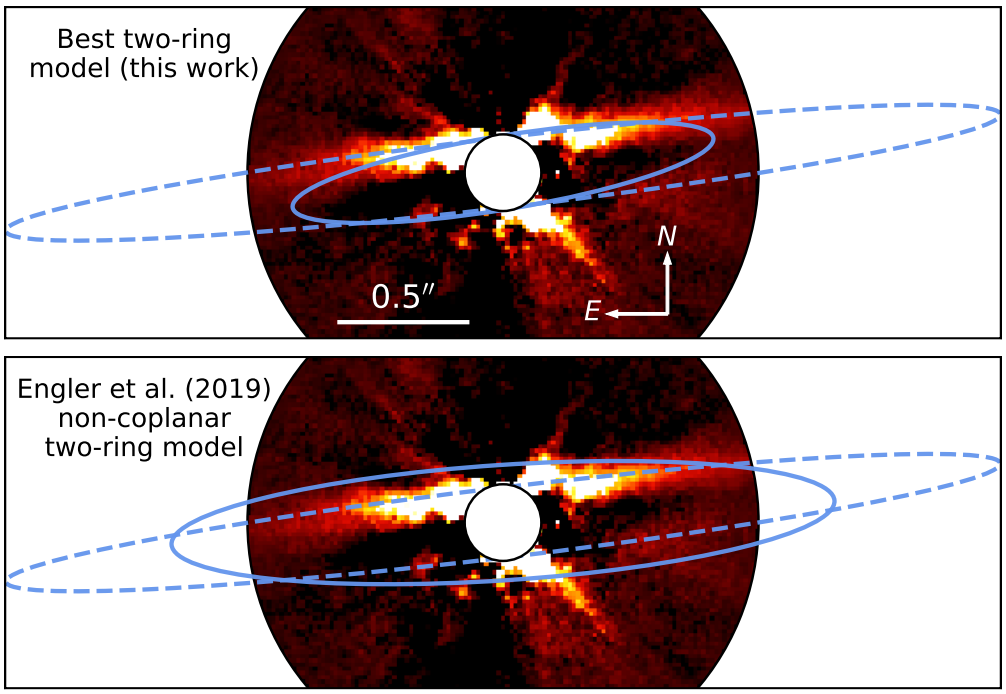

Figure 2.13: **Top:** Image of A-LOCI processed Aug 30 CHARIS data with overlaid ellipses corresponding to the outer (dashed) and inner (solid) rings of our best-fit two-ring solution (see Section 2.4.2). While our solution's inner ring has an inclination offset of $5\fdg9$ and a PA offset of $2\fdg7$ with respect to the outer ring, the best-fit inner ring's radius results in an inner ring whose features predominantly coincide with those of the outer ring along our line of sight. **Bottom:** As above, but with schematic depicting the inner and outer rings of the best-fit misaligned two-ring geometry from the results of Engler et al. (2019) (see also: the second row from the bottom of Figure 2.A.1). An inner ring oriented as posited by Engler et al. (2019) is not evident in the CHARIS data; to exist at such an orientation and still be consistent with our data, it would need to be substantially dimmer than the outer ring at similar projected separations.

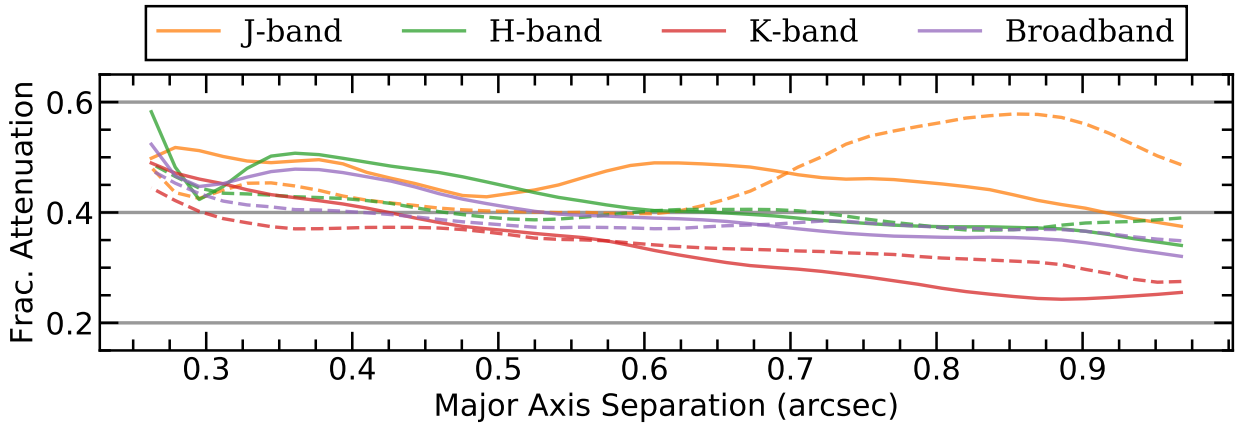

Figure 2.14: Estimates of fractional disk flux attenuation as a function of separation along the major axis. Solid and dashed lines show the values along the western and eastern extents of the disk respectively. Values manifest similarly in all four reductions, but are shown here for our A-LOCI reduction of August 30 data. These values correspond to $1 - \frac{1}{C}$, where C is the effective attenuation correction applied during measurement of surface brightness (see Section 2.5)

.

utilized [11]. For measuring surface brightness, imagery is not rotated as it was in Section 2.3. Rather, the spine locations measured for the rotated images are transformed to the native (north-up) image orientations, eliminating the possibility of image rotation interpolation affecting surface brightness measures. Flux attenuation cubes are created by dividing the PSF subtracted best disk model by the pre-PSF subtracted disk model (see: Section 2.4.1). Bhowmik et al. (2019) show that a comparable procedure results in erratic attenuation estimates for their KLIP reduction of SPHERE data for the highly inclined debris disk system HD 32297. However, this behavior does not manifest in our case (see Figure 2.14). These attenuation measurements are then used to correct the PSF subtracted cubes produced with A-LOCI and KLIP reductions of August 30 and September 07 data. As in Goebel et al. (2018), we see fractional attenuation that tends to increase at smaller separations and further from the disk spine. Along the spine, disk flux in CHARIS broadband is attenuated by $\sim 45 - 55\%$ at 0.″25 separation, and by $\sim 30 - 35\%$ at 0.″75. This attenuation varies by wavelength, with J-band typically being most attenuated ($\sim 65\%$ and 45% at 0.″25 and 0.″75 resp.), followed by H-band (with values comparable to those in broadband), and with K-band suffering the least attenuation ($\sim 45\%$ and 30% at 0.″25 and 0.″75 resp.).

Following this, the channels of the attenuation-corrected reduction products are merged to create images corresponding to J (channels $1 - 5$, $1.16 - 1.33$ μm), H (channels $8 - 14$, $1.47 - 1.80$ μm)

[11] Though different observing wavelengths may trace distinct dust populations, resulting in different spine positions, testing showed that the utilized positions fall very near to locations we fit for HST/STIS imagery where meaningful fitting was feasible. For the purpose of surface brightness measurements, the spine identified from CHARIS broadband imagery appears to be a reasonable approximation of the spine for all bands we analyze.

and K (channels $16 - 21$, $1.93 - 2.29$ μm) bands. For each (x,y) pixel position along the spine, we take the nominal surface brightness to be the average flux value within a circular aperture with diameter $0\farcs12$ (approximately the narrowest observed disk FWHM in CHARIS broadband imagery; this aperture size is used for all imagery). To accommodate the inclusion of partial pixels, we take the average in an aperture to be the weighted mean of the values with weights equal to the exact fraction of each pixel that is included in the aperture.

The uncertainty for each surface brightness measurement is determined as follows using non-attenuation-corrected images (attenuation maps become extremely noisy at the faint edges of the disk, where both attenuated and unattenuated models have values near zero). If the surface brightness, F_s, is measured at spine position (x_s, y_s) with corresponding stellocentric polar coordinates (r_s, θ_s), we make additional measurements the same way within apertures at an array of positions (r_s, θ_i), with θ_i placed every 10°. Any of these measurements whose aperture overlaps with any spine aperture are removed. We then compute the standard deviation of this array of measurements as $\sigma_{0,s}$. Since $\sigma_{0,s}$ is representative of the uncertainty in the average surface brightness for the aperture at position (x_s, y_s) **before** applying the attenuation correction, we additionally compute the nominal surface brightness at (x_s, y_s) in the uncorrected image, $F_{0,s}$. The effective attenuation correction applied is then $C_s = F_s \,/\, F_{0,s}$. From this, our final uncertainty for F_s is taken to be $\sigma_s = C_s \cdot \sigma_{0,s}$.

The procedure above is repeated to get arrays of surface brightness and corresponding uncertainties for each filter and reduction. For a given filter, the final nominal surface brightness at each position is taken to be the inverse variance weighted average of the corresponding surface brightness measurements for each of the four reductions utilized. The corresponding uncertainty is taken to be the standard error on the weighted average (see footnote, Section 2.3). Surface brightness is measured in the same manner for HST/STIS imagery of the system, except that no attenuation correction is necessary and only one reduction is used. The results of this procedure are depicted in Figure 2.15. We point out here that, although the disk appears to be recovered only marginally in J-band imagery (Figure 2.3), making SB measurements over a large ($0\farcs12$) aperture and averaging measurements for multiple reductions results in J-band surface brightness measurements with reasonable signal-to-noise.

The surface brightness measurements for each bandpass are then combined with measurements of the stellar flux (for CHARIS data: from analysis of satellite speckles during spectrophotometric calibration of the cubes, and for HST/STIS data: as reported in (Schneider et al., 2014)) to compute the local surface brightness of the disk relative to the stellar flux (see Figure 2.16).

The disk color in isolation from the stellar color can then be analyzed by taking the difference of the relative magnitudes computed above (see Figure 2.18), while the east-west asymmetry can be assessed by comparing measurements of opposing sides in a particular bandpass (see Figure 2.17). These results are discussed in Section 2.5.2).

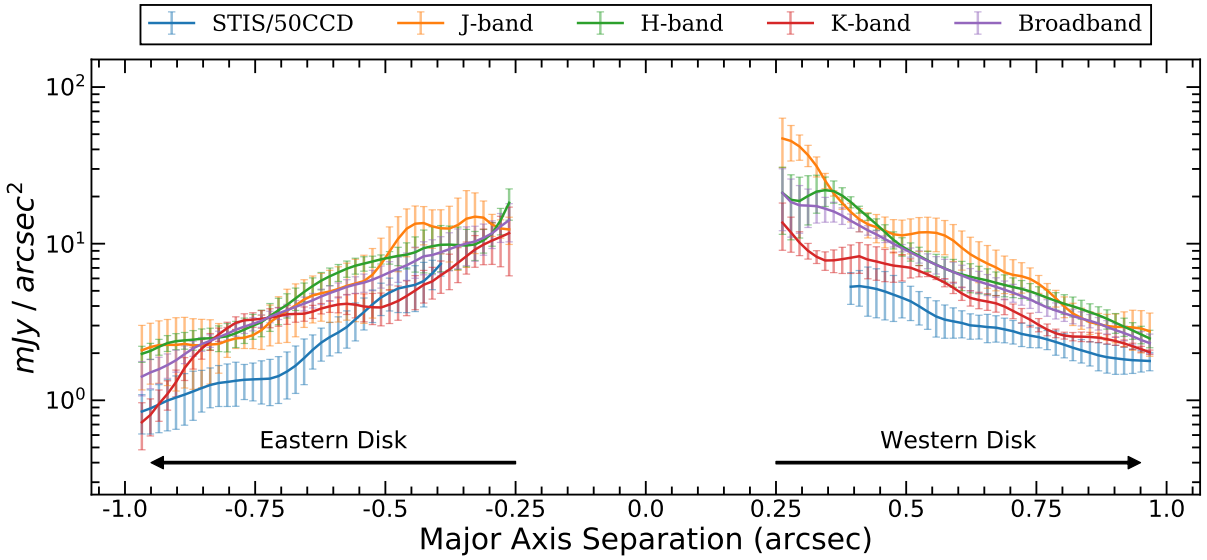

Figure 2.15: Measurements of surface brightness for HD 15115's disk as a function of stellocentric separation along the major axis in CHARIS and HST/STIS data. Here (and for other figures from Section 2.5), "broadband" refers to CHARIS's broadband. We note that the measurements here and in other Section 2.5) figures are not contiguously independent as a result of the choice to use a single aperture size for all photometric bands analyzed.

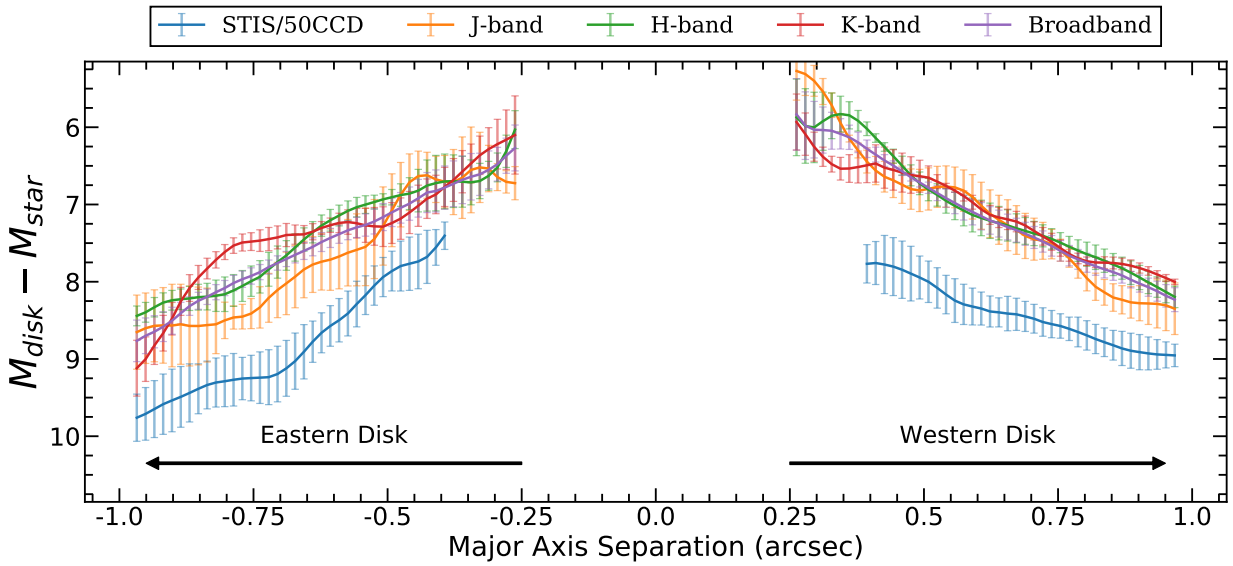

Figure 2.16: Measurements of surface reflectance for HD 15115's disk in CHARIS and HST/STIS data.

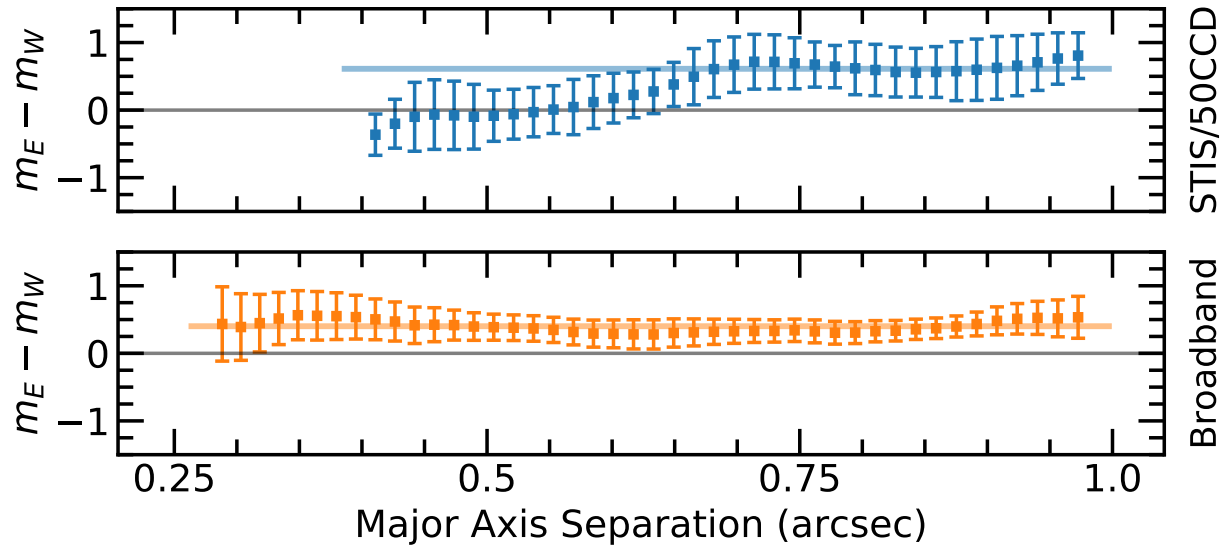

Figure 2.17: Relative surface brightness between the eastern and western extents of the disk for STIS/50CCD and CHARIS broadband. In each subplot, the horizontal colored line indicates the weighted average of the constituent flux measurements. The asymmetry in J, H, and K manifests similarly to the CHARIS broadband, albeit with larger uncertainties.

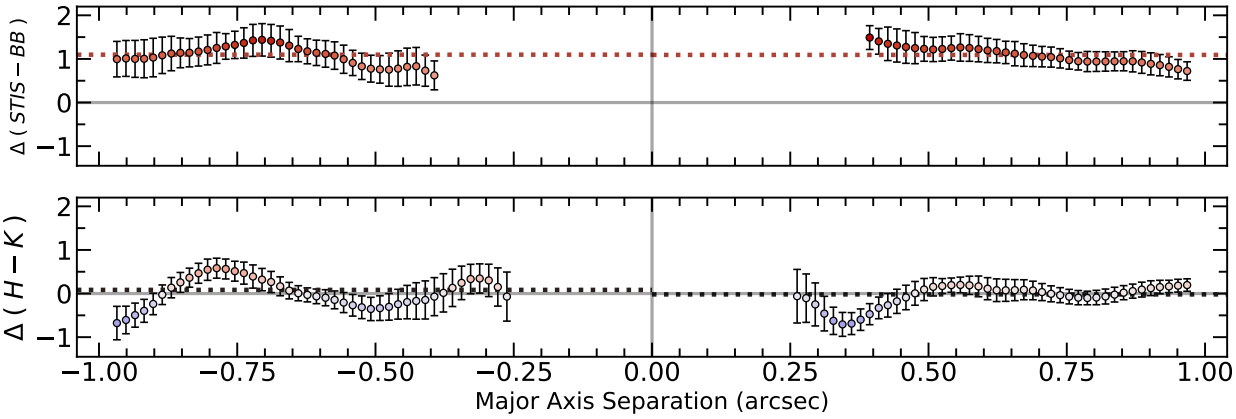

Figure 2.18: Disk color as a function of stellocentric separation along the disk major axis as identified from Section 2.3 (see Section 2.5), with $\Delta(H-K) = (H-K)_{disk} - (H-K)_{star}$, etc. Here, "STIS" refers to meaurements in STIS/50CCD, while "BB" refers to measurements in CHARIS's broadband (wavelength-collapsed) imagery. The horizontal dotted line in each subplot indicates the weighted average of the constituent flux measurements for that side of the disk. Major axis separation < 0 corresponds to the eastern extent of the disk. Other combinations of NIR filters (e.g. (J-K)) show a predominantly neutral color similar to our (H-K).

2.5.2. Surface Brightness Results

Disk Color – Though the nearly unprecedented field of view probed by our CHARIS observations precludes numerical comparisons of photometry with most prior studies, more quantitative comparisons can be made. Our red (STIS/50CCD - BB) color and neutral NIR color measured for the disk (Figure 2.18) appear generally consistent with prior literature that diagnosed the inner disk region in optical and NIR. e.g. combining the original discovery observations of Kalas et al. (2007) with new HST/NICMOS imagery, Debes et al. (2008) showed the disk's optical-NIR color becoming redder toward smaller separations.

Numerical simulations in Boccaletti et al. (2003) model disk colors for infrared bandpasses as a function of the dust size distribution's minimum grain size (a_{min}) and porosity (P). With P$= 0$, for 1.6 μm ($\sim$ H) versus 2.2 μm ($\sim$ K), they show:

1. a blue color for $a_{min} \lesssim 0.25\ \mu m$
2. a red color for $0.25\ \mu m \lesssim a_{min} \lesssim 2\ \mu m$
3. a neutral color for $a_{min} \gtrsim 2\ \mu m$ (and briefly for $a_{min} \sim 0.25\ \mu m$, as the color changes from blue to red)

Rodigas et al. (2012) found a predominantly gray ($K_s - L'$) color (2.1 and 3.8 μm respectively) across the disk from 1.″1 to 1.″45. Comparing this result with grain-color models, they suggest a distribution comprised of grains from $\sim 3 - 10\ \mu m$. The results of Boccaletti et al. (2003) show that a minimum grain size of $\sim 3 - 10\ \mu m$ should also produce a neutral color for (H-K), consistent with our measurements (Figure 2.18) and the suggestions of Rodigas et al. (2012).

However, by comparing measurements in the STIS/50CCD bandpass with our CHARIS broadband measurements, we find a definitively red color throughout the region of overlap (0.″4 $-$ 1.″0). While a wide range of minimum grain size values can produce a neutral color, a much smaller range result in a strong red color. Given that the redder filters analyzed in Boccaletti et al. (2003) predict no significantly red colors for minimum grain sizes larger than $\sim 1\ \mu m$, the measurement of a very red (STIS/50CCD $-$ CHARIS broadband) color suggests a smaller minimum grain size, $\sim 0.25 - 1.0$ μm, is needed to simultaneously produce the red (STIS/50CCD $-$ CHARIS broadband) and gray IR colors that we observe.

Rodigas et al. (2012) compute a blow-out size, a_{BO}, of $\sim 1-3\ \mu m$ for HD 15115, indicating that grains of the minimum grain size that we estimate above would likely be expelled from the system. While a larger porosity would increase the estimated minimum grain size, with $a_{min} \propto (1-P)^{-1}$ (Boccaletti et al., 2003), it should also increase the blowout size by a comparable factor, with $a_{BO} \propto (1-P)^{-1}$ (Arnold et al., 2019). A minimum grain size below the theoretical blow-out size can

be explained in a number of ways. Hughes et al. (2018) note that this phenomenon is commonly observed and suggest that it is likely the result of a change in grain collision physics near the limit of small grains. Alternatively, this could manifest if some mechanism for continually replenishing smaller grains is present, such as planetesimal collisions (e.g. Hahn 2010).

Disk Asymmetry – Numerous mechanisms have been proposed previously to explain the observed flux asymmetry of HD 15115's ring-like disk and extended halo. In the debris disk's discovery paper, Kalas et al. (2007) suggested the possibility of a past encounter with nearby star HD 12545 perturbing planetesimal orbits to cause the asymmetry. However, MacGregor et al. (2019) points out that the spatial motion of the two objects makes this encounter unlikely. Debes et al. (2009) explored the possibility of disk sculpting through interaction with the interstellar medium (ISM) to explain the bluer western color observed at large separations as well as the observed east-west brightness asymmetry. Since HD 15115's motion lies primarily in the direction of its apparently truncated eastern extent, pressure from clumps of ISM gas might redistribute dust from the eastern side to the western side (assuming motion of the ISM gas itself is favorable for this scenario); with smaller grains being more susceptible to this mechanism, this should result in both a bluer and brighter western disk. The substantial asymmetries in the outer disk halo uniquely revealed by Schneider et al. (2014) STIS imaging provide further evidence that the outer disk and halo are being perturbed. However, using the equations and parameters provided in Debes et al. (2009) gives an approximate 'deflection radius' (the stellocentric radius beyond which dust grains are likely to be significantly affected by the interaction) of $100-200$ au. Given CHARIS's $\sim 5-50$ au field of view, it seems unlikely that ISM interactions could be responsible for the asymmetry we observe. Moreover, this interaction should preferentially redistribute smaller grains to the western side– which is not supported by our disk color and surface brightness asymmetry measurements (See Figures 2.17 & 2.18); the fact that we measure a similar overall asymmetry in the STIS and CHARIS broadband data within the CHARIS field of view ($\Delta m \sim 0.6$ and 0.4 mags respectively) suggests that the phenomenon at work changes the overall dust density between the east and west, without significantly affecting the shape of the grain size distribution.

The results of Mazoyer et al. (2014) showed that while the eastern and western extents are significantly asymmetrical in brightness, the system features a symmetrical ring. This casts doubt on explanations of the brightness asymmetry which would necessitate an observable geometric asymmetry. More recently, Sai et al. (2015) reported an eccentricity for HD 15115's disk of $e = 0.06$, which could contribute to the asymmetry we observe. By itself, this eccentricity does not appear capable of producing an asymmetry of the observed size, with limited testing showing an induced east-west asymmetry of $\lesssim 10\%$. However, beyond asymmetry resulting directly from the eccentricity, Hahn (2010) notes that such a system may manifest with asymmetric dust distributions as a result of the difference in orbital velocity between apsides effectively enforcing differing

ejection criteria. It is unclear if this mechanism would be capable of producing asymmetry of the observed magnitude.

A number of studies have proposed the possibility of asymmetry resulting from dynamical interactions with an embedded planet – both for HD 15115 (e.g. Sai et al. 2015) and for similar nearly edge-on systems (e.g. HD 111520, Draper et al. 2016). Sai et al. (2015) suggests the possibility of planetsimals being trapped in the Lagrange points of an embedded planetary companion. Such an embedded planet might also induce other disk structures: as mentioned in Section 2.4.5, complicated disk structures such as spiral arms could be present here but self-obscured by the system's steep inclination. Such a geometry might result in asymmetries similar to those we observe.

Asymmetries might also be induced by major collisions within the disk (Hahn, 2010). The possibility of a minimum grain size below the blow-out size from our color analysis could be explained by this. In contrast to suggestions of planet "signposts" in similar systems, Thebault et al. (2012) simulates the interactions of debris disks and planets and concludes that for edge-on systems only weak asymmetries will typically result from planet interactions.

2.6. Limits on Planets

Following the procedure described in Currie et al. (2018) for planet forward modeling, we computed 5σ contrast limits in CHARIS broadband for the planet detection reductions outlined in Section 2.2.2. We then mapped these contrasts to planet detection limits using the hot-start, solar metallicity, hybrid cloud, synthetic planet spectra provided by Spiegel & Burrows (2012).

Model planet spectra corresponding to an array of distinct determinations for the system's age are utilized. These age determinations include: possible membership in TW Hydrae association from Banyan Σ (98 % likelihood; Gagné et al. 2018) with age 10 ± 3 Myr (Bell et al., 2015), possible membership in the β Pictoris moving group (Moór et al., 2006) with age 24 ± 3 Myr (Bell et al., 2015), and various other methods summarized in Rhee et al. (2007) which yield an age of ~ 100 Myr. Each planet spectrum was convolved with the filter profile for CHARIS's broadband filter and integrated to determine the photometric bandpass flux. The flux measured for HD 15115 was then converted to an absolute flux (to match the planet spectra) to determine the contrast at which each planet model would manifest. These values are indicated along the right axis of Figure 2.19. While planet contrast is intrinsically more favorable in the K-band, contrasts achieved are superior for CHARIS broadband imagery to the extent that the broadband offers the strongest constraints on the presence of planets.

The results of this procedure show that our August 30 data reduction is sensitive to $10\ M_j$ companions at the lower and upper suggested ages to separations of ~ 7.5 au and ~ 16 au respectively. We note that, given the small mass of the possible companion proposed by MacGregor et al. (2019), $0.2\ M_J$, we can place no constraints regarding its appearance anywhere within our

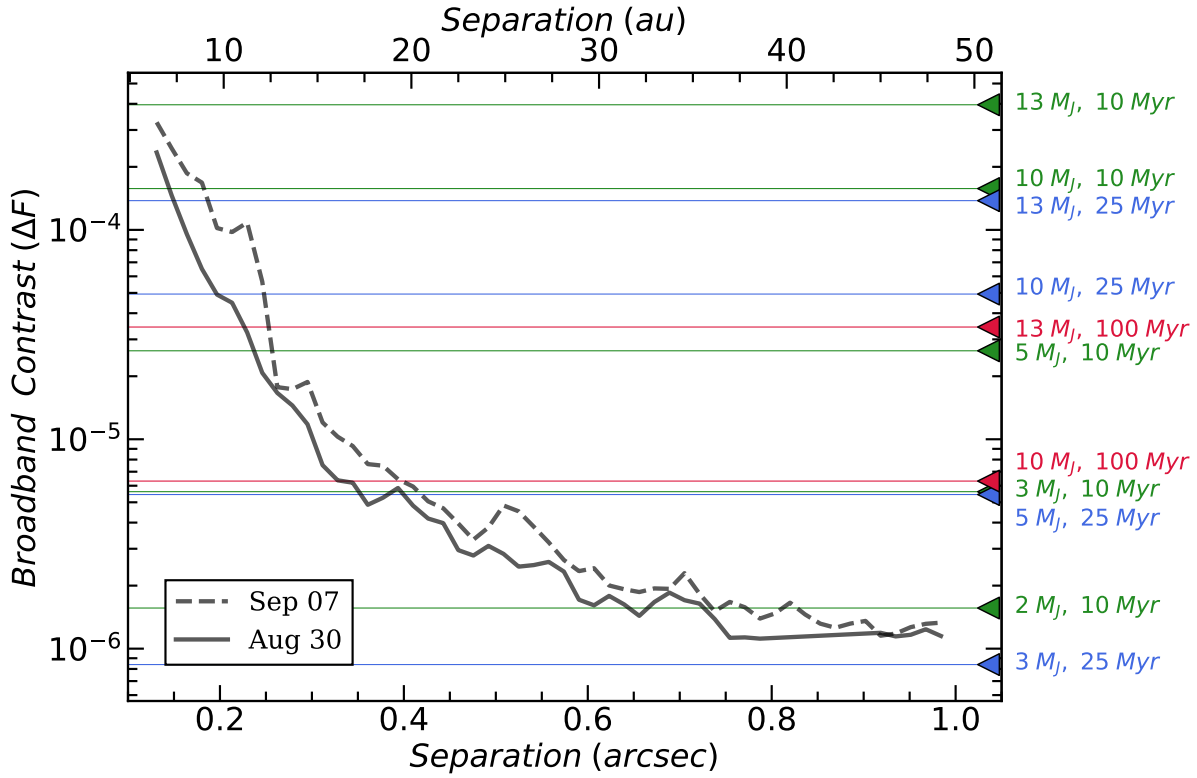

Figure 2.19: Broadband ($1.13 - 2.39$ μm) contrast curves for planet detection reductions of CHARIS HD 15115 data of August 30 and September 07 outlined in Section 2.2.2. 5σ contrast is given as a function of stellocentric angular separation (arcsec, lower x-axis) and projected separation (au, upper x-axis). Carets along the right edge of the figure and corresponding colored lines give the planet detection limits for 10 Myr (green), 25 Myr (blue), and 100 Myr (red) planets, based on hot-start, solar-metallicity, hybrid cloud planet evolution models of Spiegel & Burrows (2012). The displayed ages are chosen to correspond to various determinations for the HD 15115 system (see Section 2.6).

field of view. On the other hand, for the scenario of a 12 M_J companion at a separation of 45 au discussed in Sai et al. (2015), we can rule out the planet over the majority of its orbit (e.g. $\sim 93\%$ of its orbit for an age of 25 Myr).

2.7. Conclusions and Future Work

CHARIS imagery of the HD 15115 system has revealed the inner regions of the disk in remarkable detail and probed substantially further than any previous scattered light data (to $\rho \sim 0\farcs2$). This imagery revealed no direct evidence of planetary mass compansions and has allowed for new constraints to be placed on the possibility of a yet-unseen substellar companion in the disk. Combined with the differential evolution algorithm, CHARIS imagery has enabled us to conduct a thorough exploration of the recently proposed inner ring through forward modeling. In doing so, we

find a poor fit for a significantly non-coplanar inner ring, but reasonable fits for both a single ring and two rings aligned along our line of sight (either coplanar or manifesting with similar projected semi-minor axes). These data, combined with HST STIS imagery, have allowed for measurement of the disk's color and asymmetry at separations from $0\farcs25$ to $1\farcs0$ and spanning wavelengths from 0.6 μm to 2.3 μm. These measurements suggest a minimum grain size in the CHARIS field of view of $\lesssim 1.0$ μm, and thus smaller than previous estimates at larger separations.

The CHARIS observations presented here provide the first clear view of the system within $\rho \sim 0\farcs4$. In general, follow-up observations probing this region of the system will better substantiate the results of our analysis. Follow-up observations with CHARIS would enable further constraints to be placed on the presence of inner disk features or companions, as well as gauging the significance of the $\rho \lesssim 0\farcs25$ feature we note in Section 2.4.5. The use of CHARIS's new polarimetric integral field spectroscopy mode would allow for measurement of the disk's fractional polarization in CHARIS's field of view, a key diagnostic of the disk's dust properties (Perrin et al., 2015), while also allowing more rigorous assessment of any planet candidates that might be identified. High SNR mid-IR spectra of HD 15115 (e.g. with JWST/MIRI) could better constrain the dust composition within the disk by identifying the signatures of both silicates and non-silicate species using spectral decomposition software (Hughes et al., 2018).

Appendix

2.A. Scattering Phase Function Comparison

Figure 2.A.1 shows the results of forward modeling for models of various phase functions for our August 30 data. Overall, a simple HG phase function seems to very poorly describe the brightness profile that we observe.

2.B. Python implementation of differential evolution

Here, we provide a simple Python implementation of the differential evolution algorithm (Storn & Price, 1997) as described in Section 2.4.3. This code favors readability and simplicity over perfect computational efficiency and has no dependencies besides the NumPy[12] module. In comparison with grid searches, this implementation of DE will result in both a superior fit and orders of magnitude fewer model evaluations for the overwhelming majority of cases.

```
import numpy as np

def differential_evolution(objective_fn, converged, bounds,
    mutation=(0.5,1.), P=0.7, popsize=10):
    '''
    A simple implementation of the differential evolution algorithm using the
        'best1bin' strategy and allowing a 'dithered' mutation constant.

    Parameters
    ----------
    objective_fn : callable
        A function that takes the model parameters (1d array of length
            (popsize*K)) as its argument and returns the value to be
            minimized (typically chi-squared). This function should: generate
            the appropriate model from the list of parameters, propagate the
            model through your forward modeling routine, and then compare the
            model to your data to determine its fitness. You will probably
            want to have this function save the input and output models to
            disk as well.

    converged : callable
        A function that takes the current (normalized) population (2d array
            of shape (popsize*K, K) for K parameters) and their fit metrics
```

[12] https://numpy.org/

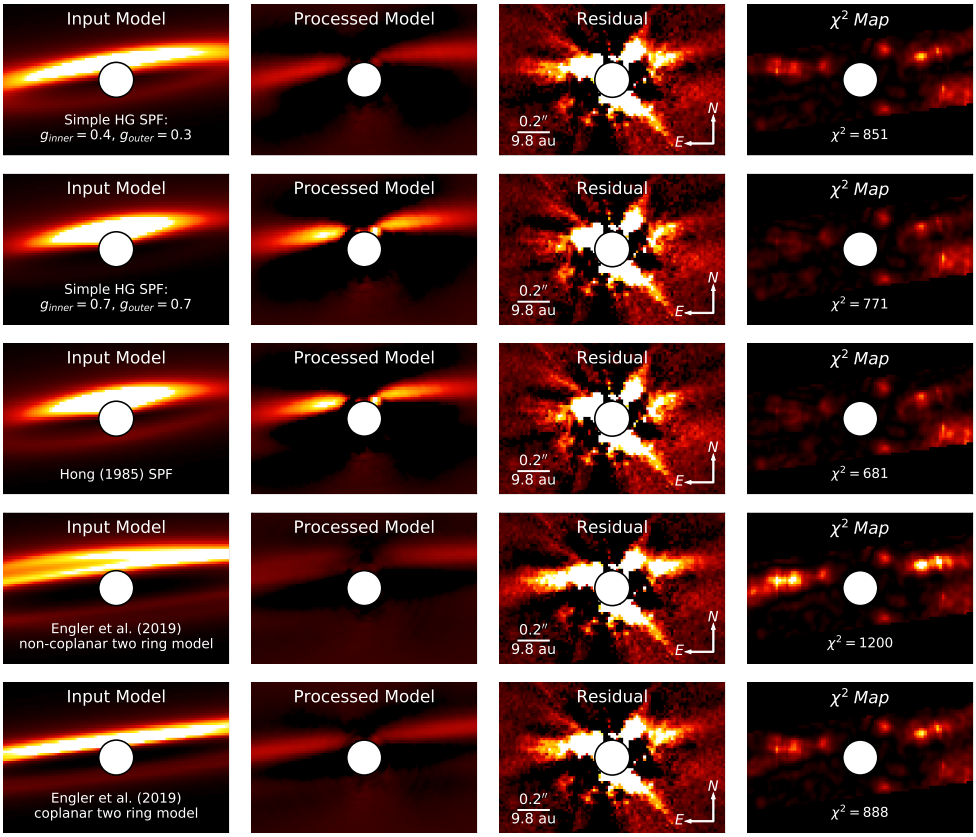

Figure 2.A.1: Each row's images depict (left to right): an initial two ring disk model, the model following attenuation by the forward-modeling procedure for our Aug 30 A-LOCI reduction and rescaling to minimize χ^2_ν (see Section 2.4.1), the residual after subtracting the model from the data, and the corresponding χ^2 map. The value of χ^2 given in the last panel is for the reduction shown only. In each case, images are shown at the same linear display stretch as the corresponding images in Figure 2.9. The first three rows show models in which only the phase function changes – the parameters are otherwise identical and correspond to the best overall two ring disk model identified in Section 2.4.1. The model in the first row utilizes the same phase function implemented by Engler et al. (2019) for their two ring models: a simple HG phase function with asymmetry parameter $g = 0.4$ for the inner ring and $g = 0.3$ for the outer. The second row's model changes the asymmetry parameter of both rings to $g = 0.7$, matching the highest weighted term in the Hong phase function. The third row's model utilizes the phase function of Hong (1985), as adopted for our modeling procedure (see Section 2.4.2). The models in the fourth and fifth row adopt the non-coplanar and coplanar (respectively) best fitting two ring models reported in Section 5.2 of Engler et al. (2019), which feature the same phase function as the model of the first row, but with differing parameters elsewhere. Note: many of the models utilizing a simple HG phase function appear especially dim in the "processed model" panels as a result of the rescaling applied at the end of the forward modeling procedure; this is simply the scaling of the model that best minimizes the weighted residuals.

```
        (1d array of length (popsize*K)), returning True if some
        convergence criteria has been met and fitting should cease and
        False otherwise.

bounds : numpy array of shape (K,2) where K is the number of model
    parameters
    Each entry of 'bounds', bounds[i,:], should provide the lower and
        upper bound for a parameter.

mutation : float or tuple(float, float), optional
    The mutation constant to utilize. Storn & Price (1997) suggest that
        values in the range [0.4, 1.0] are typically more favorable. If
        given as a tuple, the mutation constant is randomly selected each
        generation from the uniform distribution spanning the two values
        given.

P : float, optional
    The crossover probability to utilize. The value of P should be in the
        range (0,1].

popsize : int, optional
    The number of population members per free parameter to utilize.

Returns
-------
: array, float
    The set of best fitting parameters and the associated fitness metric.
'''
N,K = bounds.shape[0]*popsize, bounds.shape[0] # number pop. members and
    parameters
bmin, brange = bounds[:,0], np.diff(bounds.T, axis=0) # lower lims and
    range for each param
x = np.random.rand(N, K) # Generate initial (normed) population array
fx = np.array([objective_fn(xi) for xi in x*brange+bmin]) # The initial
    pop's fitness
indices = np.arange(N) # Define indices corresponding to population
    members
while not converged(x,fx): # Loop until converged(x,fx) returns True
    if type(mutation) == tuple: m = np.random.uniform(*mutation) # For
        dithered m
    else: m = mutation
    xtrial = np.zeros_like(x)
```

```
        j = np.argmin(fx) # For best1bin method, j is the index of the best
            member
        for i in indices:
            k,l = np.random.choice(indices[~np.isin(indices, [i,j])], 2,
                replace = False)
            xmi = np.clip(x[j] + m*(x[k]-x[l]), 0, 1) # ith mutant vector,
                clipped to bounds
            xtrial[i] = np.where(np.random.rand(K) < P, xmi, x[i]) # Get
                trial pop. vector
        fxtrial = np.array([objective_fn(xi) for xi in xtrial*brange+bmin]) #
            Fitness of trial pop.
        improved = fxtrial < fx # Boolean array indicating which trial
            members were improvements
        x[improved], fx[improved] = xtrial[improved], fxtrial[improved] #
            Replace improved members
    return x[np.argmin(fx)]*brange+bmin, np.min(fx) # Return the best params
        and fitness
```

The code as presented can be easily adapted for parallel processing with minor changes to the two lines that evaluate the fitness for a set of model parameters; e.g. using the Joblib module[13], the 4th line of code in the function could be replaced with (likewise for the 15th line):

```
from joblib import Parallel, delayed
fx = np.array(Parallel()(delayed(objective_fn)(xi) for xi in
    x*brange+bmin))
```

[13] https://joblib.readthedocs.io

CHAPTER 3

Integral Field Spectroscopy of the HD 36546 Debris Disk System[1]

3.1. Introduction

Gas depleted, dust-dominated *debris disks* are critical targets for understanding the composition, structure, and (typically late-stage) formation of planetary systems (Wyatt, 2008; Kenyon & Bromley, 2008; Hughes et al., 2018). Scattered-light imaging of debris disks can reveal and characterize signatures of planets responsible for shaping the disks' morphologies and help constrain the properties of the planets themselves (e.g. Kalas et al., 2005; Lagrange et al., 2010). "Extreme adaptive-optics" (exAO) facilities, such as SPHERE (Beuzit et al., 2019), GPI (Macintosh et al., 2015), and SCExAO (Jovanovic et al., 2015; Lozi et al., 2018; Currie et al., 2020a), now deliver high-fidelity images of numerous debris disks at sub-arcsecond separations (e.g. Milli et al., 2017; Duchêne et al., 2020; Esposito et al., 2020; Lawson et al., 2020).

Modeling of multi-wavelength imaging can enable analysis of the properties and composition of the constituent dust (e.g. Boccaletti et al., 2003; Fitzgerald et al., 2007; Goebel et al., 2018; Lawson et al., 2020), providing a context for the chemistry and evolution of our own Kuiper belt (Currie et al., 2015b; Milli et al., 2017; Chen et al., 2020). However, the assumed scattering phase function (SPF) critically affects resulting model disk images and thus our ability to infer disk properties. The traditionally adopted Henyey-Greenstein (H-G) SPF (Henyey & Greenstein, 1941) can reproduce scattered light images at large scattering angles but is not physically motivated and is often inconsistent with data at smaller scattering angles (e.g. Hughes et al., 2018). Combined with evidence suggesting that the best-fit H-G SPF from disk modeling is correlated with the scattering angles probed by the data (rather than any more meaningful physical properties of the system), Hughes et al. (2018) suggest that some other SPF – capable of explaining debris disk imagery irrespective of the scattering angles probed – is needed. Lawson et al. (2020) showed that the Hong (1985) SPF, which is derived from observations of solar system zodiacal dust, reasonably reproduced near-infrared (NIR) imagery of the debris disk of HD 15115. This finding is consistent with the suggestion in Hughes et al. (2018) that a near-universal SPF may exist to explain observed scattered light of a particular wavelength originating from nearly any circumstellar dust. Moreover, Lawson et al. (2020) found that the adopted SPF can fundamentally change the interpretation of scattered light images, e.g. the number of separate dust populations required to reproduce the data. If scattered light imagery of additional systems can be well-modeled using this SPF, it could provide

[1]This chapter is reproduced from Lawson et al. (2021a) with permission.

further evidence of such a universal SPF.

In this study, we present the first multi-wavelength scattered light imaging of the debris disk around HD 36546 (Currie et al., 2017), using the Subaru Coronagraphic Extreme Adaptive Optics (SCExAO) system and the Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS) integral field spectrograph in broadband (spanning NIR J, H, and K bands, $1.13 - 2.39$ μm) mode (Groff et al., 2016). HD 36546's disk is highly inclined, bright in scattered light, and has an extremely high fractional luminosity ($L_{\rm IR}/L_{\star} \sim 4{\times}10^{-3}$) comparable to long-studied debris disks around β Pictoris and HR 4796A and new benchmarks such as HD 115600 (Smith & Terrile, 1984; Schneider et al., 2014; Currie et al., 2015b; Millar-Blanchaer et al., 2015; Gibbs et al., 2019).

The HD 36546 system lies at a distance of 101.35 ± 0.73 pc (Gaia Collaboration et al., 2018), slightly foreground to the Taurus-Auriga star-forming region and is a member of the 3–10 Myr old 118 Tau association (99.1% probability, Gagné et al. 2018)[2]. Due to the system's extreme youth, the HD 36546's debris disk represents a rare opportunity to study debris disk structure and composition at very early stages where protoplanetary disk accretion ceases, colder gas dissipates, and collisions between icy planetesimals tracing Kuiper belt-analogue formation are predicted to yield detectable debris at peak luminosity (Currie et al., 2008, 2009; Cloutier et al., 2014; Ribas et al., 2015; Kenyon & Bromley, 2008).

Previous work on HD 36546 found evidence for a disk with an extremely large scale height and strong forward-scattering: an H-G g value of ~0.7–0.85, among the highest of any known debris disk (Currie et al., 2017; Hughes et al., 2018). Revisiting HD 36546's disk with our new, comparable-quality data obtained over a wider wavelength range will provide improved constraints on its structure and the first look at its composition. Much like for the case with HD 15115 (Lawson et al., 2020), modeling these data with a different SPF may lead to revisions in our understanding of the HD 36546 disk's properties.

Our new data resolve the disk down to separations as small as $\sim 0\farcs25$ (~ 25 au) and enable the first spectrophotometric analysis of the disk. From these data, we provide surface brightness, surface contrast, and asymmetry profiles in CHARIS broadband and J, H, and K bands – as well as disk color profiles in $\rm J - H$ and $\rm H - K$. Through detailed forward modeling, we provide a refined schematic for HD 36546's disk. Finally, we provide constraints for the presence of companions within the CHARIS field of view (FOV; angular separations of $0\farcs11 \lesssim \rho \lesssim 1\farcs0$).

3.2. Data

3.2.1. Observations

We observed HD 36546 (A0–A2V, V=6.95, H=6.92; Currie et al., 2017; Lisse et al., 2017) on 2019 January 12 using the Subaru Telescope's SCExAO running at 2 kHz paired with the CHARIS

[2]This association is labeled as "Mamajek 17" in Currie et al. (2017).

integral field spectrograph (Groff et al., 2016) operating in low-resolution (R $\sim$ 20), broadband (1.16–2.37 μm) mode, and utilizing SCExAO's Lyot coronagraph with 217 mas diameter occulting spot. Conditions were photometric with good 0.″4–0.″5 seeing and 16 km/hour winds[3]. The data were collected in *angular differential imaging* mode (ADI; Marois et al. 2006) and achieved total parallactic angle rotation of 99° with total integration time of t_{int} = 50 minutes. The data set is made up of 66 individual exposures, with 65 having exposure times of 45.73 seconds and one having an exposure time of 30.98 seconds. Sky frames were also obtained following the science data to enable sky subtraction.

A second target, HR 2466 (A2V; V = 5.21, H = 5.07), was observed on the same night with the same instrument configuration to enable *reference star differential imaging* (RDI). These data include 143 individual exposures of 20.65 seconds each, for a total integration time of 49 minutes. For all reduction and analysis, we adopt the revised CHARIS pixel scale of 0.″0162 and the North-up PA offset of 2.°2 as reported in Currie et al. (2018).

3.2.2. CHARIS Data Reduction

CHARIS data cubes were extracted from raw CHARIS exposures using the CHARIS Data Reduction Pipeline (Brandt et al., 2017), using modifications to read noise suppression listed in Currie et al. (2020a). Extracted data take the form of image cubes with dimensions $(N_\lambda, N_x, N_y) = (22, 201, 201)$ (i.e. 201 $\times$ 201 pixel images for each of 22 wavelength channels). Subsequent basic image processing – e.g. sky subtraction, image registration, spectrophotometric calibration – was carried out as in Currie et al. (2011, 2018). Inspection of registered data cubes showed that the quality of the AO correction varied between exposures taken before (good), during (worst[4]), and after (best) transit. Particularly low quality frames were identified by first measuring the peak-to-halo flux ratio of the satellite spots in each exposure[5]. These values were then averaged over the 4 satellite spots and 22 wavelength channels, and the median and median absolute deviation of the values across the full data sequence were computed. Any exposures falling more than 3 median absolute deviations below the median peak-to-halo ratio were removed. We removed three frames near transit which met this criteria.

[3]However, AO performance was compromised compared to typical performance due to a significant (4–5x) loss in the IR secondary mirror's reflectivity at blue optical wavelengths where AO188 (which provides a needed low-order correction for SCExAO) does wavefront sensing. This loss resulted from oxidization of the mirror's silver coating from sulfur dioxide stemming from the Kilauea eruption in May 2018. For wavefront sensing, HD 36546 was effectively a magnitude 10.5 star. The mirrors were realuminized with a fresh coating in late 2019.

[4]This is a result of HD 36546 transiting very near zenith from Subaru, making alt/az tracking through transit challenging.

[5]The peak flux is defined as the flux within an aperture with a radius of one instrumental full-width at half-maximum (FWHM), centered on the satellite spot. The halo flux is the flux within an annulus centered on the satellite spot, with inner and outer radii of 2·FWHM$-dr/2$ and 2·FWHM$+dr/2$ respectively, where dr is the radial extent that results in an annulus enclosing the same area as the peak aperture.

PSF subtraction to recover HD 36546's disk was performed using RDI with the *Karhunen-Loève Image Projection* (KLIP; Soummer et al. 2012) algorithm (RDI-KLIP), as well as ADI with both KLIP (ADI-KLIP) and the *Adaptive, Locally Optimized Combination of Images* (A-LOCI; Currie et al. 2012a, 2015a) algorithm (ADI-ALOCI) which is based on the LOCI algorithm (Lafrenière et al., 2007). To search for planets, we exploit both ADI and *spectral differential imaging* (SDI; Sparks & Ford, 2002) (see e.g., Beuzit et al., 2019), using a more aggressive implementation of A-LOCI (as in e.g., Currie et al., 2018, 2020a,b; Lawson et al., 2020). Unlike prior implementations, we performed an ADI-based reduction on the post-SDI residuals instead of SDI on the post-ADI residuals. We found that for data sets with somewhat variable AO corrections leading to temporally decorrelated images, the former approach is slightly superior. The utilized algorithm parameters are provided in Table 3.2.1.

3.2.3. Results

In broadband (wavelength-collapsed) imagery, all three PSF-subtraction techniques result in clear detection of the bright side of HD 36546's disk to separations as small as $\sim 0''.25$ (Figures 3.2.1, 3.2.2). Each result also shows a partial, weak detection of the fainter side of the disk in the west, which is strongest in the ADI-ALOCI result, with signal-to-noise ratio per resolution element (SNRE) ~ 2.5. As in the discovery imagery of the disk from SCExAO/HiCIAO (Currie et al., 2017), SCExAO/CHARIS imagery of HD 36546 reveals a highly inclined, significantly asymmetrically scattering disk. ADI-based reductions reveal no obvious evidence of ansae, indicative of gaps or a ring-like structure, within the field of view. The RDI-KLIP reduction, however, appears to show a change in the disk profile at $\rho \sim 0''.75$, consistent with the fiducial radius of the best-fitting disk model identified in Currie et al. (2017).

The SNRE of the disk is calculated following the standard procedure (e.g. Goebel et al. 2018); the signal per resolution element is determined by replacing each pixel with the sum of pixels within an aperture of one PSF full-width at half maximum (FWHM), while the noise per resolution element is the radial noise profile of the aforementioned signal with a finite-element-correction applied (Mawet et al., 2014). The quotient of these is taken to be the SNRE. The detection in the broadband images (Figure 3.2.3) is strongest for the ADI-ALOCI reduction, with a SNRE along the trace of the disk of $\sim 4-9$ for regions exterior to $0''.25$[6]. The ADI-KLIP reduction yields slightly weaker SNRE (typically ~ 1 less than the same locations in the ADI-ALOCI result). SNRE maps for the RDI-KLIP reduction suggest a somewhat weaker detection having SNRE of $\sim 3-4$ at small separations, dropping to ~ 2 beyond $0''.4$. We note, however, that the noise calculation for the RDI-KLIP reduction is more heavily affected by the disk's signal since RDI subtraction tends to be more conservative of disk flux. Therefore, the resulting SNRE depends on the size of the noise mask

[6] In the SNR calculation, we use a software mask to reduce the amount of disk signal included in the noise estimation.

Table 3.2.1. PSF Subtraction Algorithm Settings

Method	Param. Tuning	Rad. Prof. Sub.	r_{min}	r_{max}	Δr_{sub}	Algorithm Parameters
RDI-KLIP	disk	False	15	70	55	$\mathrm{N_{PCA}} = 10$, $\mathrm{N_{zones}} = 1$
ADI-KLIP	disk	True	15	65	50	$\mathrm{N_{PCA}} = 15$, $\mathrm{N_{zones}} = 1$, $\delta_{\mathrm{FWHM}} = 2.5$
ADI-ALOCI	disk	True	15	70	5	$\mathrm{N_A} = 800$, $g = 0.1$, $\delta_{\mathrm{FWHM}} = 2.5$, SVD $= 0.02$
SDI+ADI-ALOCI	planet	True	5	70	2.5	$\mathrm{N_A} = 50$, $g = 0.25$, $\delta_{\mathrm{FWHM}} = 0.55$, SVD $= 1.05 \times 10^{-6}$

Note. — Algorithm settings for PSF subtraction with each of the three techniques utilized. 'Rad. Prof. Sub.' indicates whether each reduction was carried out on radial profile subtracted data. r_{min} and r_{max} refer to the inner and outer radii of the subtraction region for each reduction (in pixels). Δr_{sub} indicates the radial width of individual subtraction annuli. 'g' refers to the aspect ratio of the optimization regions (which are annular subsections, defined in polar-coordinates by radial and azimuthal extents) with $g < 1$ producing azimuthally elongated regions and $g > 1$ producing radially elongated regions. 'N_A' refers to the area of optimization regions in units of PSF cores. 'δ_{FWHM}' indicates the minimum rotation gap in units of PSF FWHM (for both A-LOCI and KLIP). 'Δr_{sub}' gives the radial size of subtraction regions in units of pixels (for both A-LOCI and KLIP). 'N_{PCA}' indicates the number of principal components utilized in construction of the model PSF. 'N_{zones}' is the number of subsections into which each KLIP optimization annulus was divided (with a value of 1 corresponding to full annuli). 'SVD' refers to the threshold for normalized singular values below which A-LOCI coefficients are truncated. For the ADI-ALOCI reduction, the algorithm is not restricted in the number of images that may be used in constructing the PSF model. For the SDI+ADI-ALOCI parameter tuning, we also employed a moving pixel mask over the subtraction zone (Marois et al., 2010; Currie et al., 2012a) and restricted the reference PSF construction for each target image to the 50 best-correlated frames (Currie et al., 2012b).

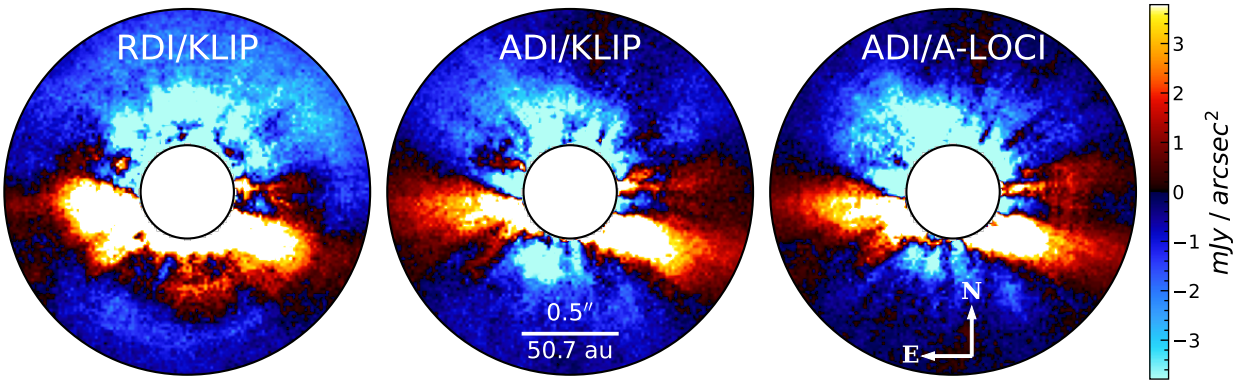

Figure 3.2.1: Wavelength-collapsed final images for all three reduction methods noted in Section 3.2.2. The inner and outer masks in each panel have radii of 0.″25 and 1.″0 respectively.

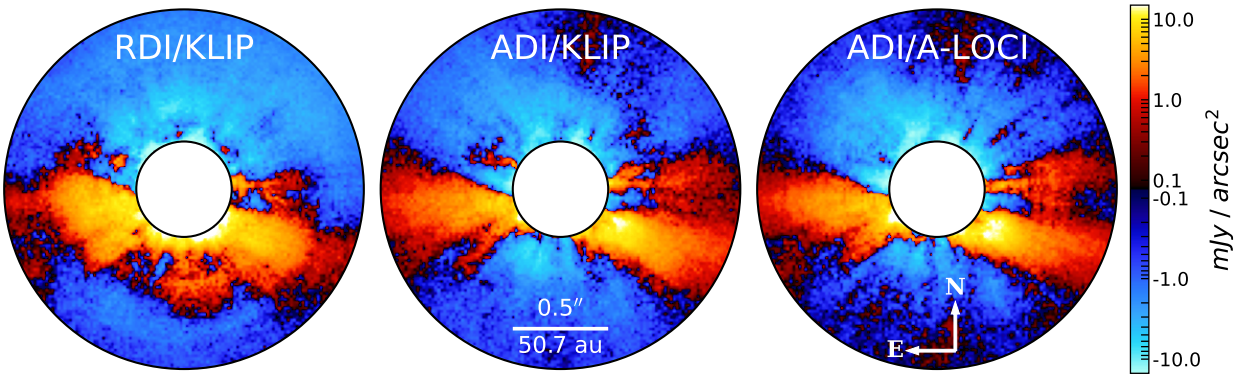

Figure 3.2.2: As Figure 3.2.1, but with a symmetric logarithmic color scale. This scaling is linear in the range [-0.1, 0.1], and logarithmic otherwise.

utilized. Using a slightly larger mask results in SNRE of $\sim 5-6$ out to 0.″6, and ~ 2.5 beyond (Figure 3.2.3, upper right). The detection is clearest in H-band, where the disk appears much like the broadband result. Detections in J-band and K-band are weaker, with J-band suffering especially from prominent residual speckles at small separations, and K-band from weaker scattering and higher background flux (Figure 3.2.4).

3.3. Modeling the Debris Disk of HD 36546

3.3.1. Disk Forward Modeling

To probe the morphology of HD 36546's disk, we adopt a strategy of forward-modeling synthetic disk images as described in Currie et al. (2019); Lawson et al. (2020), focused on KLIP instead of A-LOCI due to the former's much faster computational speed[7]. We create synthetic scattered-light images using a version of the GRaTeR software (Augereau et al., 1999). This model image is

[7]In this case, KLIP is faster because reductions are performed in full annuli, requiring far fewer matrix operations than A-LOCI which is performed in much smaller annular segments.

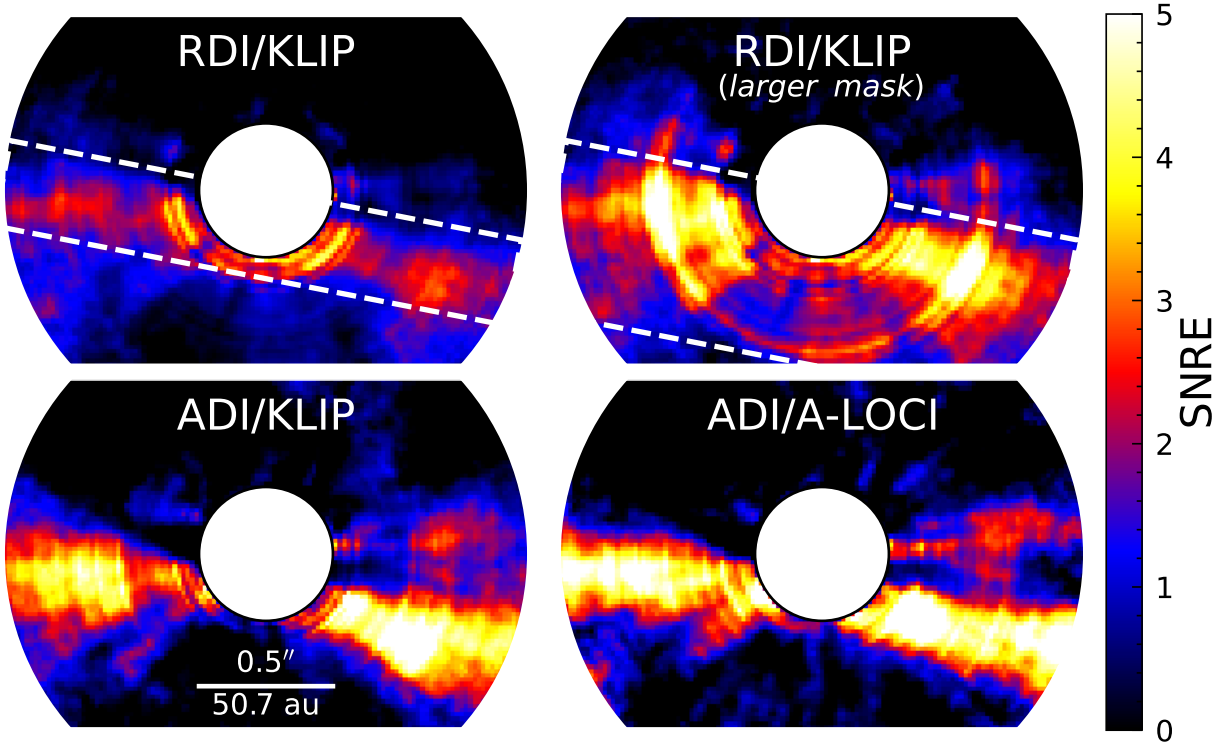

Figure 3.2.3: Signal-to-noise ratio per resolution element for the reductions shown in Figure 3.2.1. For all panels, a mask is applied over the disk when the noise is calculated. The RDI reduction is shown using two different mask sizes (overlaid dashed white boxes), as the chosen size was found to have a strong impact on the resulting SNRE. Both ADI reductions use the smaller mask.

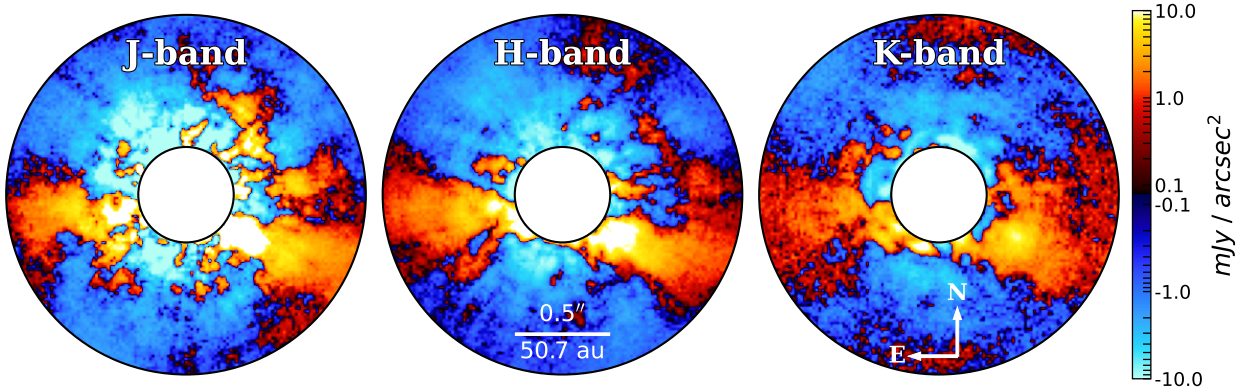

Figure 3.2.4: ADI-KLIP reduction of SCExAO/CHARIS HD 36546 data, with wavelength channels combined to produce images comparable to J (channels $1-5$, $1.16-1.33$ μm), H (channels $8-14$, $1.47-1.80$ μm) and K (channels $16-21$, $1.93-2.29$ μm) bands. Images are displayed with a symmetric logarithmic color scale which is linear in the range [-0.1, 0.1], and logarithmic otherwise.

then rotated to an array of parallactic angles matching those of the data, and convolved with the instrumental PSF for each of the 22 CHARIS wavelength channels. Here, the PSF is estimated empirically by median combining normalized cutouts of the satellite spots for each wavelength over the full data sequence. Using forward-modeling (Pueyo, 2016), we simulate annealing of the disk model for each CHARIS data cube in our sequence. In this procedure, over-subtraction (where the presence of disk signal results in an over-bright PSF model), direct self-subtraction (inclusion of the disk signal in basis vectors), and indirect self-subtraction (perturbation of basis vectors by the disk signal) terms are computed from the Karhunen-Loève modes used for the science data reduction and their perturbations by the disk.

Once this process is completed, we compare the disk model to the wavelength-collapsed science image to assess the goodness of fit. For this purpose, we utilize χ^2_ν, which is calculated here within a region of interest as described in Lawson et al. (2020). The region of interest considered is a rectangular box centered on the star with un-binned dimensions of 120 pixels $\times$ 50 pixels ($\sim 2\farcs0 \times 0\farcs8$) and rotated 11° clockwise to approximately align the region with the disk's major axis. Models which are acceptably consistent with the overall best model's fitness, $\chi^2_{\nu,min}$, are defined as those having $\chi^2_\nu \leq \chi^2_{\nu,min} + \sqrt{2/\nu}$ (Thalmann et al., 2013). Here, ν is the degree of freedom, defined as the difference between the number of bins in the region of interest when binned to resolution and the number of free model parameters. The disk models are restricted to those with simple one-ring geometries with linear flaring ($\beta = 1$), and a Gaussian vertical density distribution ($\gamma = 2$).

It is common to describe the angular distribution of scattered light in debris disks using the physically-unmotivated Henyey-Greenstein (H-G) scattering phase function (Henyey & Greenstein, 1941). The H-G SPF, parameterized by a single variable – the asymmetry parameter g[8], can reproduce scattered light images at large scattering angles but often fails at smaller angles (e.g. Hughes et al., 2018). Moreover, best-fit single H-G scattering asymmetry parameters are well correlated with the scattering angles probed by observations, while lacking correlation with the properties of the systems (Hughes et al., 2018). This appears to suggest that some other SPF, capable of reproducing observations irrespective of the scattering angles probed, is necessary. Linear combinations of two or more H-G phase functions of differing asymmetry parameter could provide a suitable solution specific to the details of a particular disk (e.g., its composition). However, optimization of a disk model with an SPF formed from the linear combination of N H-G phase functions will introduce either $2N$ additional free parameters, with an asymmetry parameter (g_i) and weight (w_i) for each, or $2N - 1$ free parameters for a constrained optimization problem (with the constraint that the weights sum to unity, allowing the weight for one asymmetry parameter to

[8]The asymmetry parameter is defined as the average of the cosine of the scattering angle, weighted by the (normalised) phase function, over all directions.

be assumed). Without providing additional constraints (e.g. that $w_i \geq w_{i+1}$), these weights and asymmetry parameters will also have significant correlations or degeneracies with one another. Since the utilized SPF is expected to alter the best-fit value of other disk parameters, the SPF cannot be optimized independently. Thus, exploration of such phase functions would introduce significant, possibly intractable, model complexity for attempts to optimize disk models.

Another possibility, evidenced by measurements of both Solar-system and extrasolar dust, is that there exists a nearly universal SPF to describe circumstellar dust in a given wavelength regime (Hughes et al., 2018). In such a scenario, this SPF could be measured independently, such as by observation of Solar-system dust, and then simply adopted for use by groups modeling scattered light from disks. Lawson et al. (2020) showed that imagery of HD 15115 could be reasonably reproduced by adopting the fixed Hong (1985) SPF – modeled as the linear combination of three separate H-G phase functions based on observations of solar system zodiacal dust. Here, we again adopt the Hong (1985) scattering phase function, with asymmetry parameters and weights fixed to the values identified in Hong (1985): $g_1 = 0.7$, $g_2 = -0.2$, and $g_3 = -0.81$, with weights $w_1 = 0.665$, $w_2 = 0.330$, and $w_3 = 0.005$.

Altogether, our prescription results in a model having 6 free parameters:

1. R_0, the radius of peak grain density in au
2. $\frac{H_0}{R_0}$, the ratio of disk scale height at R_0 to R_0
3. α_{in}, the power law index describing the change in radial density interior to R_0
4. α_{out}, the power law index describing the change in radial density exterior to R_0
5. *i*, the inclination of the disk in degrees
6. *PA*, the position angle of the disk in degrees

To explore these parameters, we utilize the differential evolution (DE) optimization algorithm as described in Lawson et al. (2020). In DE, a population of trial solutions is iteratively improved by adding the weighted difference of two random solutions to a third and keeping only those "mutated" population members that result in an improvement in the fitness metric. Algorithm meta-parameters are set as follows: $N_{pop} = 15$, resulting in a population size of 90 (6 free parameters $\times 15$), a dithered mutation constant $m = [0.5, 1.0]$, and a crossover probability $P = 0.7$. We test for convergence in the current population at the end of each generation. We consider the population to be converged if either of the following criteria is met:

1. For each varying parameter, the standard deviation among the set of normalized parameter values for the current population falls below a threshold value, δ_{σ_X}

2. The standard deviation of the fitness of the current population, divided by the median fitness of the current population, falls below a threshold value, $\delta_{\sigma_{\rm FX}}$, and more than 10 generations have completed.

Effectively, the first criterion is met by a population whose parameters have converged to a single minimum position, while the second accommodates a population with diverse parameter values but comparable goodness of fit, as might occur with weak or degenerate parameters. The additional requirement that 10 generations have finished for the second criteria prevents the unlikely case in which convergence might otherwise be achieved by a population in the first few generations with approximately uniform but weak fitness. Both thresholds, $\delta_{\sigma_{\rm X}}$ and $\delta_{\sigma_{\rm FX}}$, are set to 0.005[9].

We carry out several distinct procedures for optimizing the disk model. Our primary results focus on a procedure in which the model is optimized for both the RDI-KLIP and ADI-KLIP reductions simultaneously (i.e. seeking the model that minimizes the combined χ^2 for both reductions). However, since the adopted noise profile for the RDI reduction suggests substantially larger noise levels than that of the ADI-KLIP reduction, its contribution to a combined χ^2_ν metric will be minor; effectively, if the global best-fit solution differs considerably between the RDI and ADI reductions, this would not likely be apparent for such a procedure. As such, we also carry out procedures optimizing the model for each individual reduction. A secondary motivation for this is to validate the consistency of the solutions provided by DE. Toward this latter point, we also run two separate optimization procedures with identical settings for the RDI-KLIP data[10].

3.3.2. Model Results

The key results of the four modeling procedures (two for RDI-KLIP, one for ADI-KLIP, and one combined) are provided in Table 3.3.1. A corner plot for the combined procedure is provided in Figure 3.3.1. For analysis hereafter (e.g. for attenuation correction in Section 3.4), we adopt the best-fitting model resulting for the combined procedure, having parameters: $\mathrm{R}_0 = 82.65$ au, $H_0/R_0 = 0.005$, $\alpha_{\rm in} = 1.00$, $\alpha_{\rm out} = 1.51$, $\mathrm{i} = 79\overset{\circ}{.}06$ and $\mathrm{PA} = 80\overset{\circ}{.}09$ (visualized in Figures 3.3.2, 3.3.3). The results of the other modeling procedures are visualized and compared in Appendix 3.A.

3.3.3. Modeling Discussion

The results of our forward-modeling procedure echo many of the themes of the results in Currie et al. (2017): favoring a highly inclined disk model with shallow radial density profiles (small α_{in},

[9]Typically, larger values ~ 0.01 are chosen. However, we opt for a stricter convergence requirement here purely for the sake of producing full corner plot visualizations – the solutions show no statistically significant changes between a threshold of 0.01 and 0.005.

[10]Since DE involves a randomly initialized set of trial solutions which are then randomly mutated over subsequent generations, it is conceivable that two DE procedures could converge to distinct solutions.

Table 3.3.1. Disk Model Optimization Results

	R_0 (au)	H_0/R_0	α_{in}	α_{out}	i (deg)	*PA* (deg)	χ^2_ν
RDI-KLIP (1)	75.4	0.007	1.0	-1.4	78.9	79.0	0.45
N=1980	[60.0, 103.2]	[0.005, 0.143]	[1.0, 5.0]	[−2.1, −1.0]	[75.0, 81.4]	[75.1, 82.4]	<0.50
RDI-KLIP (2)	76.4	0.012	1.0	-1.4	78.9	78.9	0.45
N=2340	[60.0, 102.7]	[0.005, 0.138]	[1.0, 5.4]	[−2.0, −1.0]	[75.0, 81.9]	[75.5, 82.4]	<0.50
ADI-KLIP	87.4	0.005	1.0	-1.6	79.3	80.6	0.76
N=1800	[70.3, 114.6]	[0.005, 0.050]	[1.0, 1.9]	[−2.0, −1.3]	[76.9, 81.3]	[78.7, 82.6]	<0.81
ADI & RDI	82.7	0.005	1.0	-1.5	79.1	80.1	0.61
N=2160	[64.1, 109.0]	[0.005, 0.052]	[1.0, 1.9]	[−2.0, −1.2]	[77.3, 80.7]	[78.6, 82.0]	<0.65
Currie et al. (2017) [a]	75.6	0.118	3	-3	75	75	1.03
	[66.7, 84.5]	[0.070, 0.160]	[3, 10]	[−3, −3]	[70, 75]		<1.07
Bounds	[60, 120]	[0.005, 0.2]	[1, 10]	[−10, −1]	[75, 85]	[75, 85]	–

Note. — Results for each of the four distinct DE optimization runs outlined in Section 3.3.1, as well as the forward modeling of HD 36546 in Currie et al. (2017). Each parameter column provides the value for the best model for the run indicated in the leftmost column (with N indicating the total number of models evaluated for each procedure), with the following line providing the range of values resulting in an acceptably fitting model. The final column gives the corresponding run's minimum χ^2_ν and the threshold for acceptable models. The permitted boundary for each parameter in modeling for CHARIS data (the modeling procedure of Currie et al. (2017) used different bounds) is given on the final line. We adopt the best model for the combined run ("ADI & RDI") for analysis hereafter. a) In Currie et al. (2017), models are discussed in terms of H_0 (ξ_0 or ksi_0 in the text) rather than $\mathrm{H}_0/\mathrm{R}_0$; for simplicity we present these results in terms of $\mathrm{H}_0/\mathrm{R}_0$ instead. Additionally, we have updated relevant values for the difference in the assumed distance to the system — 114 pc in Currie et al. (2017), versus 101.35 pc here.

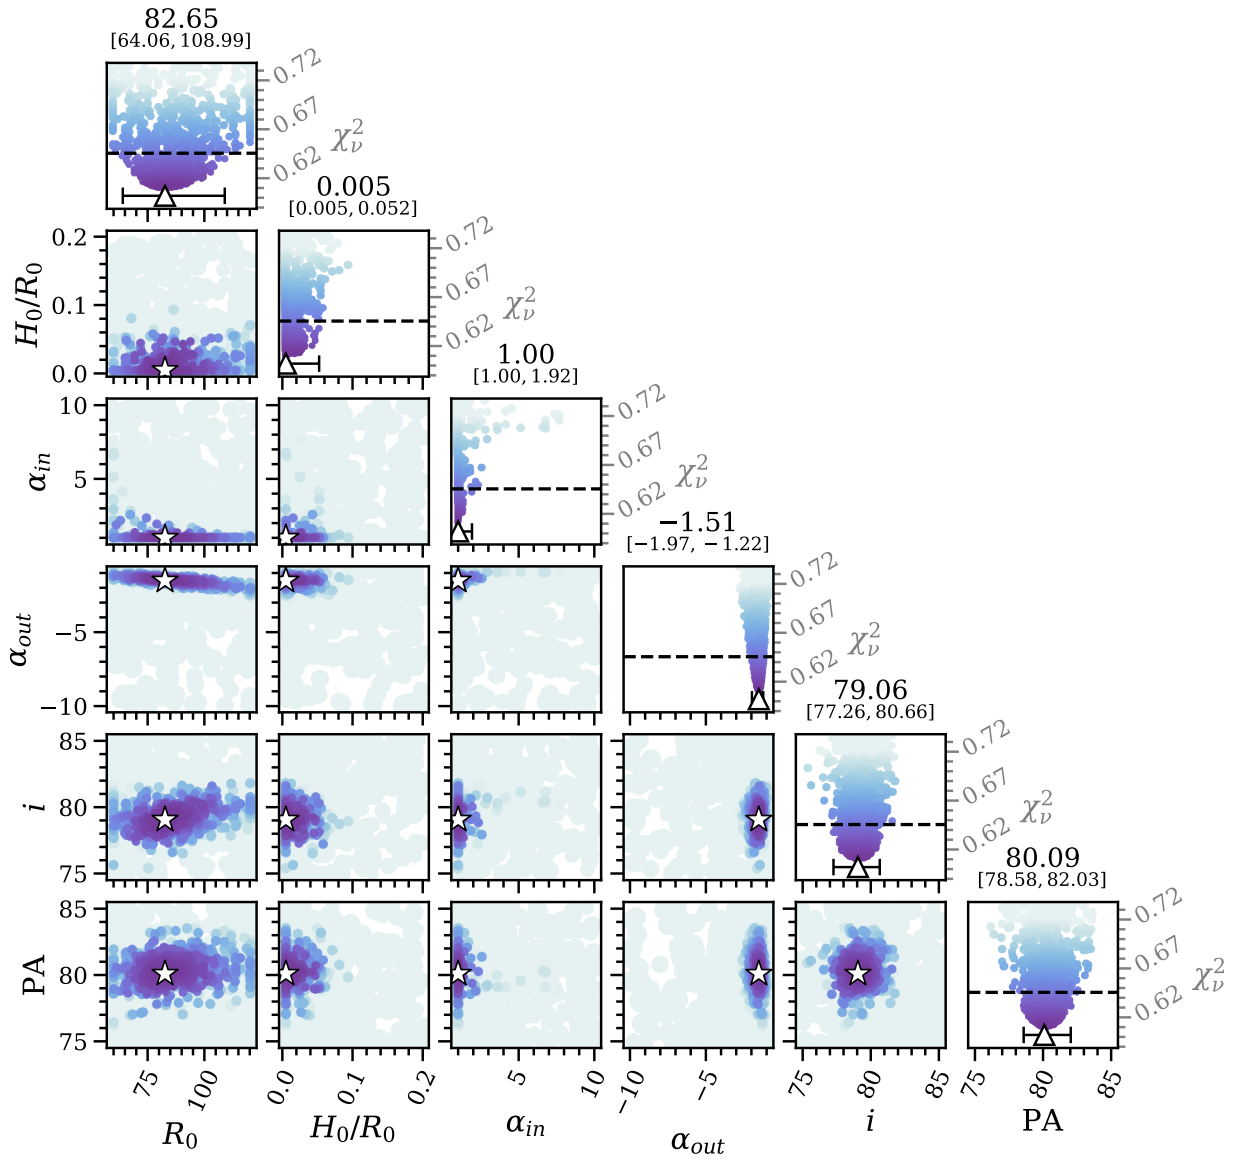

Figure 3.3.1: A variation of the corner plot for optimization of the disk model to both the RDI-KLIP and ADI-KLIP reductions simultaneously (corresponding to the row labeled 'ADI & RDI' in Table 3.3.1). Each off-diagonal panel visualizes solutions as a function of two of the parameters. Each point in the subplots corresponds to a single sample, colored according to its χ^2_ν score. The color mapping scales from $\chi^2_{\nu,min}$ to $\chi^2_{\nu,min} + 3 \cdot \sqrt{2/\nu}$ with darker colors indicating more optimal values of χ^2_ν. Diagonal elements provide a one-dimensional view of each of the parameters, while also serving as a color bar for the off-diagonals. The threshold fitness for acceptable solutions ($\chi^2_\nu \leq \chi^2_{\nu,min} + \sqrt{2/\nu}$) is indicated by a horizontal black dashed line in each diagonal plot. The black-outlined star (triangle) marker in each off-diagonal (diagonal) panel indicates the position of the overall best-fitting model. The text at the top of each diagonal subplot provides the best-fitting value for the corresponding parameter, with the range beneath indicating the smallest and largest value producing an acceptable χ^2_ν; these limits are also depicted using error-bars for the 'best-fit' marker in each diagonal panel.

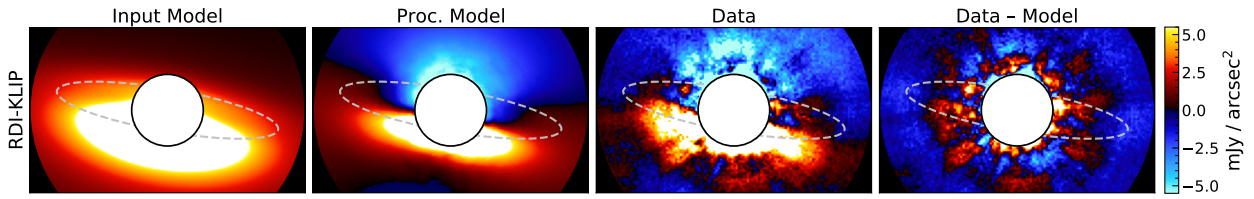

Figure 3.3.2: From left to right, for the RDI-KLIP reduction: the PSF-convolved input model, attenuated model, data, and residual for the overall best model resulting from the combined DE run. The grey dashed line overlaid in each panel is the ellipse corresponding to the radius, PA, and inclination of the model.

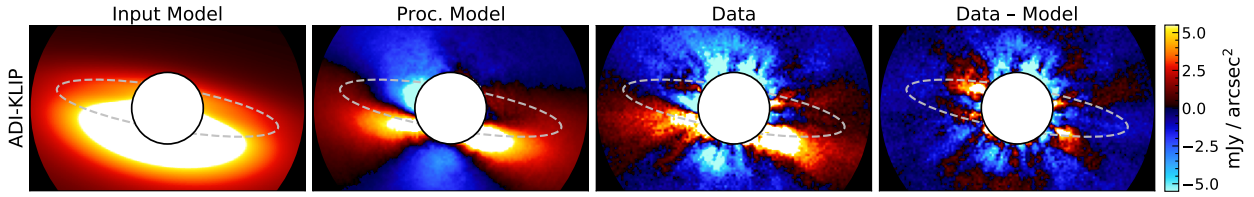

Figure 3.3.3: As Figure 3.3.2, but for the ADI-KLIP reduction.

α_{out}), suggestive of an extended disk as opposed to the ring-like geometry often found for debris disks.

We emphasize here that, while the nominal parameter values are expected to very closely correspond to the global minimum of the parameter space, the acceptable parameter ranges from DE provide only a rough approximation of parameter uncertainties. Unlike ensemble samplers (e.g. Markov chain Monte Carlo), which draw samples from a parameter space in proportion to the statistical likelihood of those samples, DE is designed to rapidly approach the global minimum with little regard for sub-optimal combinations of parameters; in cases where the algorithm identifies a strong value for a particular parameter early in the run, the values of that parameter which yield "acceptable" results may be poorly explored. If a stronger approximation of parameter uncertainties is desired, repeated DE runs (or follow-up runs near the best DE solution with "local" optimization algorithms, such as damped least-squares) can provide thorough exploration of acceptable values while still requiring orders of magnitude fewer samples than would be required by an ensemble sampler.

While highly inclined disks are expected to be relatively insensitive to changes in the α parameters (see, e.g., Lawson et al. 2020), these parameters are perhaps the best constrained of any in our modeling (see Figure 3.3.1). This might be explained by the particularly small scale height favored by our modeling: $\mathrm{H_0/R_0} = 0.005$, an order of magnitude smaller than the median value of 0.06 provided by Hughes et al. (2018) and the lower boundary for values permitted by our procedure. However, we note that scale heights up to $H_0/R_0 = 0.052$ are also acceptable solutions. In fact,

the diagonal subplot for H_0/R_0 in Figure 3.3.1 shows that the minimum χ^2_ν changes very little as a function of H_0/R_0 for $H_0/R_0 \lesssim 0.05$. Given that the corner plot shows no obvious strong correlations between H_0/R_0 and other parameters, this appears to suggest that our data simply cannot discriminate between H_0/R_0 values in this range. This is not particularly surprising, given that $H_0/R_0 = 0.05$ for a fiducial radius of $R_0 \sim 83$ au corresponds to $H_0 \sim 4.2$ au, which projects to $\sim 0''.04$ or 2.8 pixels, compared to our resolution of 2.67 pixels. Considering this apparent degeneracy for H_0/R_0, our results for this parameter should be interpreted as $H_0/R_0 \lesssim 0.05$. Additionally, we note that the value of α_{in} converges to the lower bound for permitted values ($\alpha_{in} = 1$). This result, indicating an exceptionally slow decline in density interior to r_0, could suggest that the disk is better described by a geometry with a nearly flat density slope over some intermediate region, followed by typical power-law slopes interior and exterior to that region. An alternative explanation for the strong α_{out} constraint that manifests here is addressed briefly in Section 3.4.1.

The best-fit parameter values and acceptable ranges resulting from forward modeling in Currie et al. (2017) are included for comparison in Table 3.3.1. Notably, this modeling was conducted on a coarse grid; the acceptable ranges provided are much less likely to correspond to complete ranges of acceptably-fitting parameters. Additionally, we note that, while the value for PA identified in Currie et al. (2017) is at the lower bound of allowed values for our presented DE runs, in earlier modeling runs not presented here the range of allowed PA values extended as low as 60° but showed very poor fits for PA $\lesssim 75°$. As such, the final procedure restricted PA to the provided range to speed convergence. As stated previously, our results are thematically consistent with the picture of the disk suggested by the analysis in Currie et al. (2017). Many of the specific parameter values identified, however, appear inconsistent; e.g., despite evaluating models with $\alpha_{\rm out} = -1$, the best-fit in Currie et al. (2017) is achieved for $\alpha_{\rm out} = -3$. Forward modeling the best fit solution from Currie et al. (2017) for our reductions results in a combined χ^2_ν value of 0.85 (well beyond the limit for acceptable models of $\chi^2_\nu = 0.65$). This apparent difference in identified optimal parameters might be explained by some combination of a few key differences in the modeling procedures:

1. the data itself, e.g. with the CHARIS reduction making a partial detection of the fainter side of the disk where HiCIAO data did not,

2. the scattering phase functions, with our choice to adopt the Hong phase function, as opposed to fitting an asymmetry parameter as in Currie et al. (2017); in prior work modeling the disk of HD 15115 (Lawson et al., 2020), this was found to have a significant impact on the optimal parameters,

3. the choice in Currie et al. (2017) to not vary PA, instead adopting the PA resulting from spine tracing,

4. the parameter bounds, e.g. with values of α_{in} in Currie et al. (2017) not including the best-fit value that we identify,

5. the parameter optimization, with Currie et al. (2017) identifying best-fit parameters through a much less thorough grid search.

The best-fit model for the combined ("ADI & RDI") DE run results in a reduced χ^2_ν score for each reduction that is very similar to the best scores achieved in the DE runs for individual reductions. The DE run for only the ADI-KLIP reduction results in a best-fit model having $\chi^2_\nu = 0.758$ (with "acceptable" models having $\chi^2_\nu \leq 0.808$), while the best-fit model of the combined run results in $\chi^2_\nu = 0.760$ when considering only the ADI-KLIP data — well within the limit to be considered an acceptable solution. Similarly, the DE runs for only the RDI-KLIP reduction result in a best-fit model having $\chi^2_\nu = 0.435$ (with $\chi^2_\nu \leq 0.504$ being acceptable) while the best-fit model of the combined run results in $\chi^2_\nu = 0.459$ for the RDI-KLIP reduction. In other words, the adopted disk model from our combined modeling run is also a reasonable explanation when considering the imagery of each data reduction individually.

3.4. Disk Spine Trace and Surface Brightness

To measure the disk's spine and surface brightness (SB), we proceed largely as detailed in Lawson et al. (2020) and summarized hereafter. In place of the RDI-KLIP reduction utilized for modeling, we perform spine fitting and surface brightness measurements on an RDI-KLIP reduction in which radial profile subtraction is utilized (with all other parameters remaining unchanged). This results in a generally weaker detection of the disk at very small separations and in wavelength collapsed imagery, but provides a more consistent detection across the field of view and in individual wavelength channels. For the sake of attenuation correction, we forward-model the adopted best-fitting model of Section 3.3.1 for this additional reduction. The wavelength-collapsed image for this reduction and the corresponding forward modeling results are provided in Appendix 3.B.

3.4.1. Disk Spine

The positions of peak radial brightness along the disk are referred to as the disk's spine, and are typically used as locations for surface brightness measurements and to approximate the geometric paramaters of the disk (e.g., fiducial radius, inclination, and position angle). We fit the spine position using wavelength-collapsed products and applying an initial image rotation of -11° to approximately horizontally align the disk's major-axis. We then measure the spine at each pixel column position by fitting Lorentzian profiles to the image values weighted by the inverse of the radial noise profile. The fit Lorentzian peak and associated fit uncertainty are taken as the spine position and uncertainty. As a departure from the method of Lawson et al. (2020), the fitting is

carried out on both the ADI-KLIP and RDI-KLIP images simultaneously to identify the spine position most consistent with both images (rather than taking the weighted average of the results).

Notably, the fit spine (Figure 3.4.1) appears to grow further from the major axis at larger separations, suggesting a possible "swept-back wing" geometry (see e.g. Hughes et al. 2018), in which a disk's outer dust ring appears to bend (rather than remaining planar) at increasing separations. As a result of this, the measured spine is extremely poorly described by an ellipse and thus cannot easily provide a meaningful assessment of the disk's inclination or radius. While it may be possible to estimate the position angle of the disk's major axis from this spine profile, it is unclear how the value would be affected by the noted phenomenon. As such, we provide no estimate of the disk geometry based on spine fitting. For the locations of surface brightness measurements, we use the fit spine profile in place of the best-fit ellipse profile used for our analysis in Lawson et al. (2020) — correcting the initial rotation of $-11°$ using a rotational transformation of the spine positions.

Interestingly, applying this spine fitting procedure to our best-fitting models also results in a divergent spine trace (both before and after attenuation is induced by forward modeling). After comparison to the complete sample of models, this behavior manifests only if $|\alpha_{out}| < 2$. This could suggest that the observed spine shape is simply the result of a disk with a dust distribution meeting this criteria and viewed nearly edge-on. Alternatively, models having such values of α_{out} may have been the only accessible means by which to reproduce this feature of the observations – with the true cause being more similar to the explanations summarized by Hughes et al. (2018) (e.g. ISM interaction).

3.4.2. Disk Surface Brightness

Peak surface brightness is measured for both the RDI-KLIP and ADI-KLIP reductions at the aforementioned spine locations using an aperture of diameter $0\farcs12$ (7.4 pixels). As in Lawson et al. (2020), a) we apply wavelength-dependent attenuation corrections based the adopted disk model (determined by comparison of the PSF-convolved input model cube and the attenuated model cube after forward modeling), and b) we approximate the uncertainty for each SB measurement as the standard deviation of an array of like-measurements taken at the same stellocentric radius but with azimuthal angles placed every $10°$. These measurements are made on the "broadband" wavelength-collapsed images, as well as images corresponding to J, H, and K bands. The final surface brightness profile for a given filter is taken to be the weighted average of the results for the two individual reductions. The results of this procedure are visualized in Figure 3.4.3.

Disk surface contrast (the surface brightness of the disk relative to the stellar flux) is then computed by combining disk SB with measurements of the stellar flux for each filter (Figure 3.4.4).

Any east-west asymmetry is measured by comparing the surface brightness of opposing sides in a given filter (Figure 3.4.5).

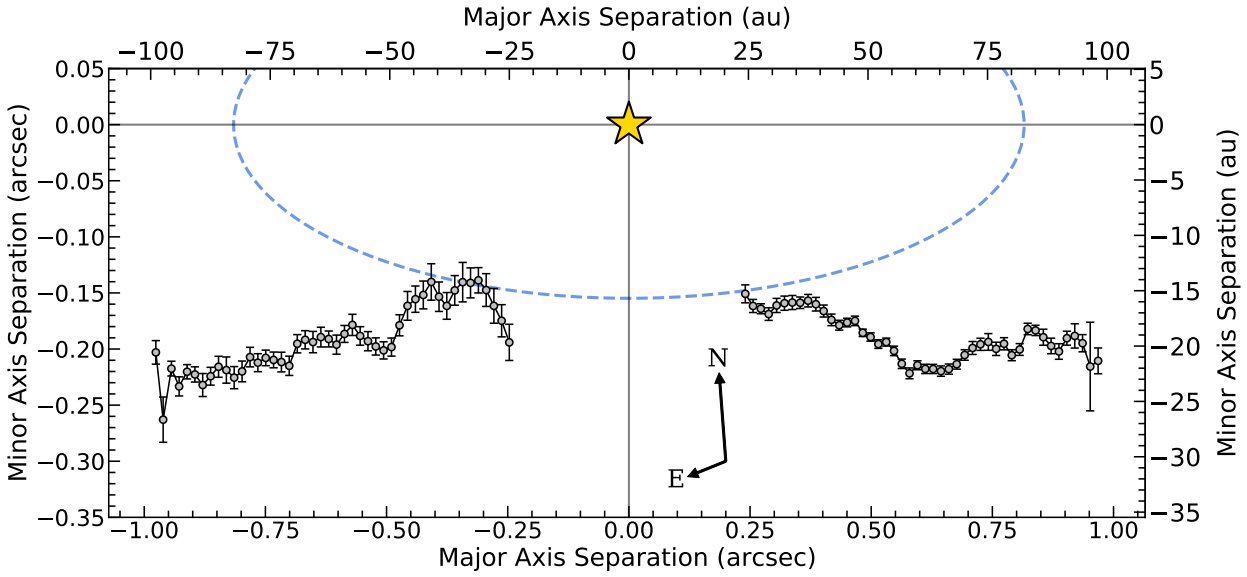

Figure 3.4.1: Stellocentric separation of disk spine fits along the major and minor axes (for a major axis position angle of 80.°09). The compass provided indicates both the orientation of the figure and the difference in scale between the axes (which results in $\boldsymbol{N} \not\perp \boldsymbol{E}$, as depicted). The spine positions are identified by fitting Lorentzian profiles to both wavelength-collapsed KLIP images (ADI and RDI) simultaneously, and may indicate a swept-back wing profile (see Section 3.4.1). The dashed blue ellipse corresponds to the adopted best-fitting model of Section 3.3.1 ($r_0 = 82.65$ au, $i = 79.°06$, $PA = 80.°09$), and is shown to demonstrate that the ellipsoidal profile very poorly describes the spine profile measured.

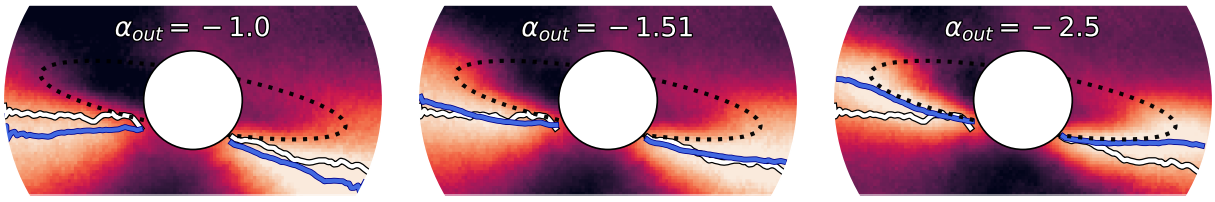

Figure 3.4.2: Disk models forward-modeled for our ADI-KLIP reduction with the parameters of the overall best model identified in Section 3.3.1 but with differing values of α_{out}. To emphasize the apparent disk spine, a $1/r^2$ correction and arbitrary color scale have been applied to each image. The dotted black ellipse in each panel corresponds to the radius, PA, and inclination of the model, while the blue line shows the spine trace found for the model following the procedure used for the real data (Section 3.4.1). The spine trace for the real data is overlaid in white. For $\alpha_{out} \gtrsim -2$, the fit spine deviates considerably from the profile of the ellipse. The model having $\alpha_{out} = -1.51$ corresponds to the overall best model identified in Section 3.3.1.

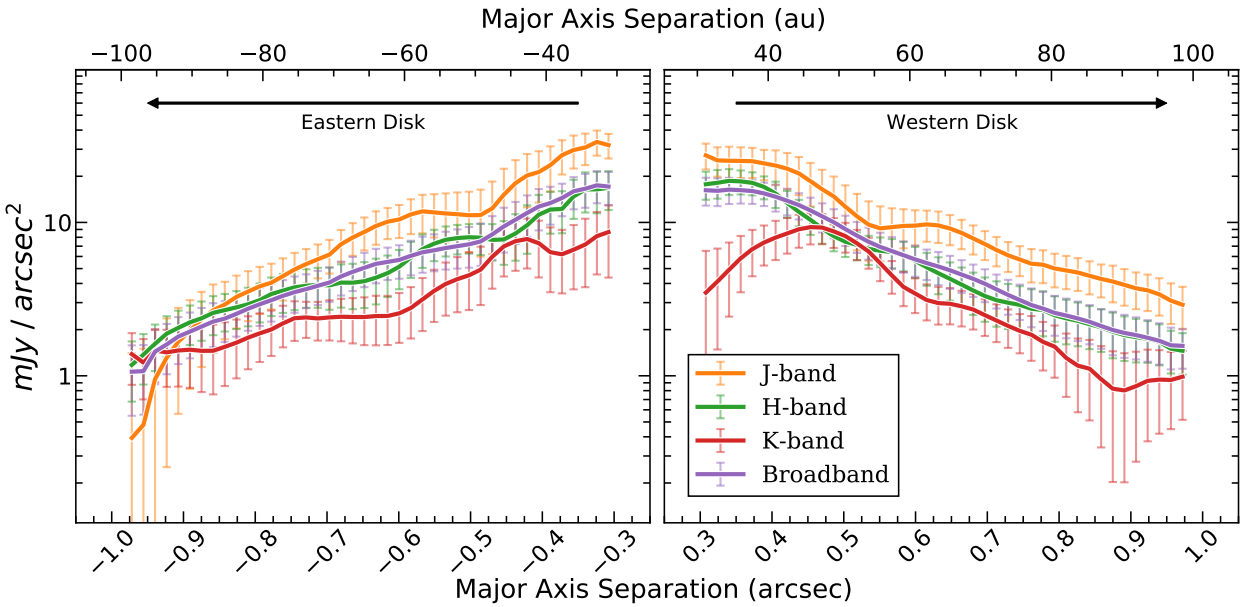

Figure 3.4.3: Measurements of peak surface brightness for HD 36546's disk as a function of projected stellocentric separation in arcseconds and au along the major axis in CHARIS data. Here, "broadband" refers to CHARIS's broadband. Since measurements are made using an aperture of radius 0.″06, comparable SB measurements can only be provided to within 0.″06 of the smallest separations at which the disk is recovered (0.″25) — hence, SB is provided only to $\rho \sim 0.''31$. We note that the measurements here and in other Section 3.4 figures are not statistically independent from neighboring measurements as a result of the 0.″12 diameter aperture. The apparent dip in western K-band surface brightness at small separations is not likely physical — given that no comparable behavior manifests in other bands, we conclude that this is probably the result of the somewhat compromised K-band data quality that was noted in Section 3.2.3.

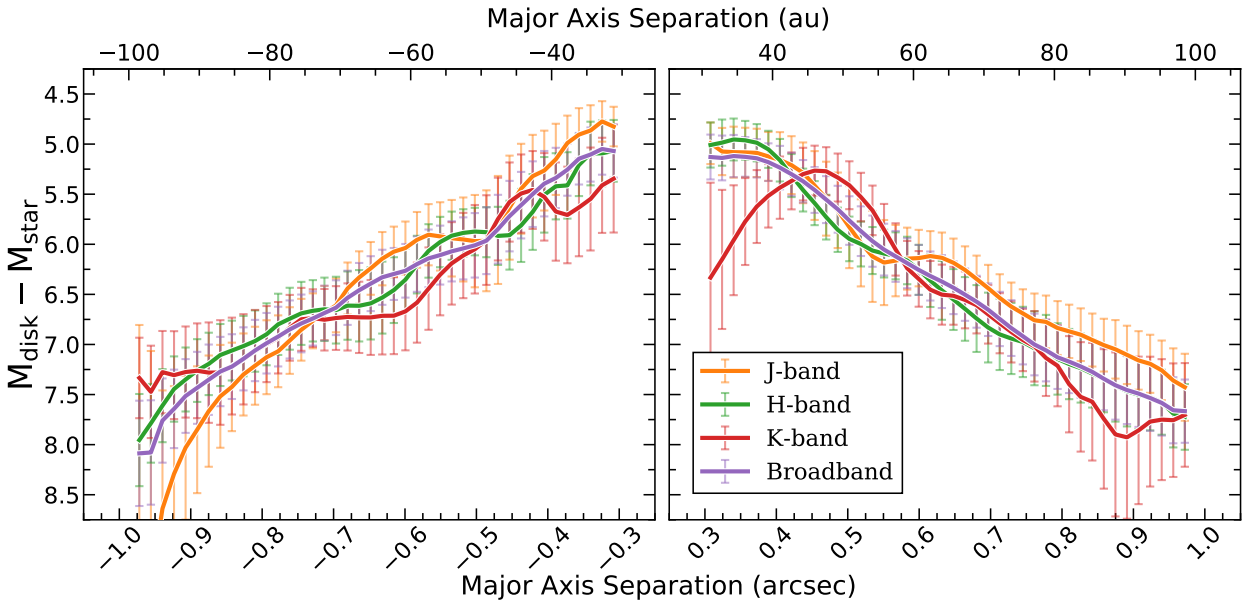

Figure 3.4.4: Measurements of surface contrast for HD 36546's disk in CHARIS data.

Finally, the disk color in isolation from the stellar color is measured by converting surface contrasts to magnitudes and taking the difference of measurements between filters (see Figure 3.4.6), e.g. $(J_{\rm disk} - J_{\rm star}) - (K_{\rm disk} - K_{\rm star}) = (J-K)_{\rm disk} - (J-K)_{\rm star} = \Delta(J-K)$, where $\Delta(J-K)$ is the $J-K$ color of the disk without contribution from the star.

3.4.3. Surface Brightness Power-Law

For the broadband (wavelength-collapsed) images, we make an additional set of measurements to facilitate fitting a power-law to the disk's radial surface brightness profile. For this purpose, we follow the SB measurement procedure of Section 3.4.2, except:

1. measurements are made at positions with an array of radial separations but corresponding to a single scattering angle (25°[11], assuming $PA = 80\overset{\circ}{.}09$ and $i = 79\overset{\circ}{.}06$ based on the results of Section 3.3.2),

2. images are binned to resolution ($\mathrm{FWHM} = 2.67$ pixels $= 0\overset{''}{.}043$),

3. the aperture diameter is set to one resolution element

The first change enables characterization of the radial dependence of the disk's SB for comparison with results in the literature, while the latter two changes result in approximately statistically

[11]The particular scattering angle is chosen arbitrarily in an attempt to probe a scattering angle providing radial coverage with both generally high disk signal and well defined attenuation estimates in the east and west. Values within $\sim 5°$ provide comparable results.

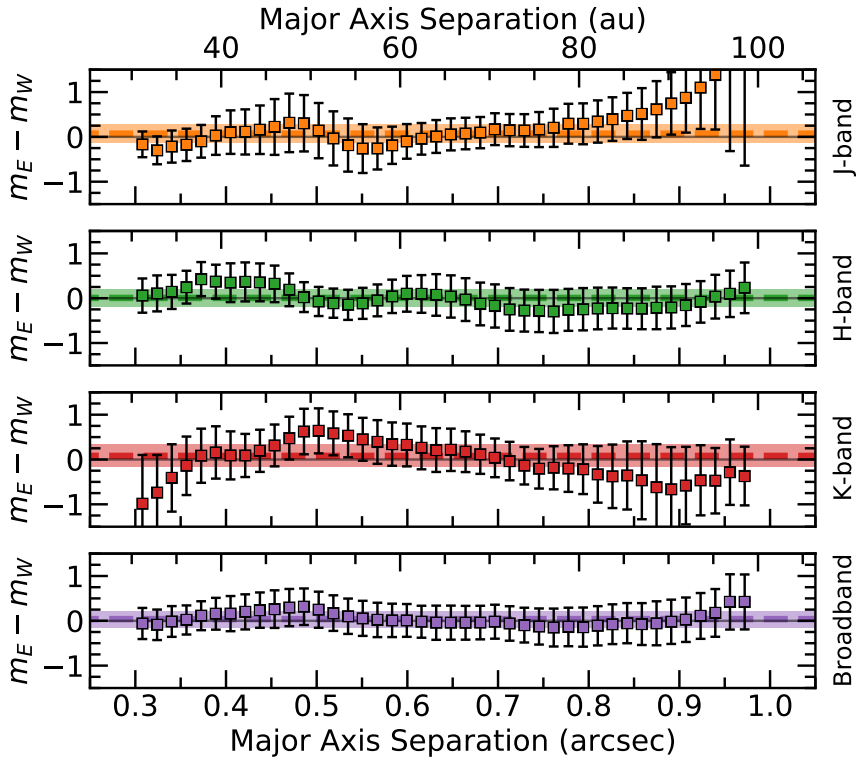

Figure 3.4.5: Relative surface brightness between the eastern and western extents of the disk for CHARIS J, H, K, and broadband. In each subplot, the horizontal colored dashed line indicates the weighted average of the full set of constituent flux measurements, while the colored region corresponds to the 3-σ confidence interval for the average. The averages for all four filters are within 1-σ of zero – i.e. consistent with an overall symmetric disk flux.

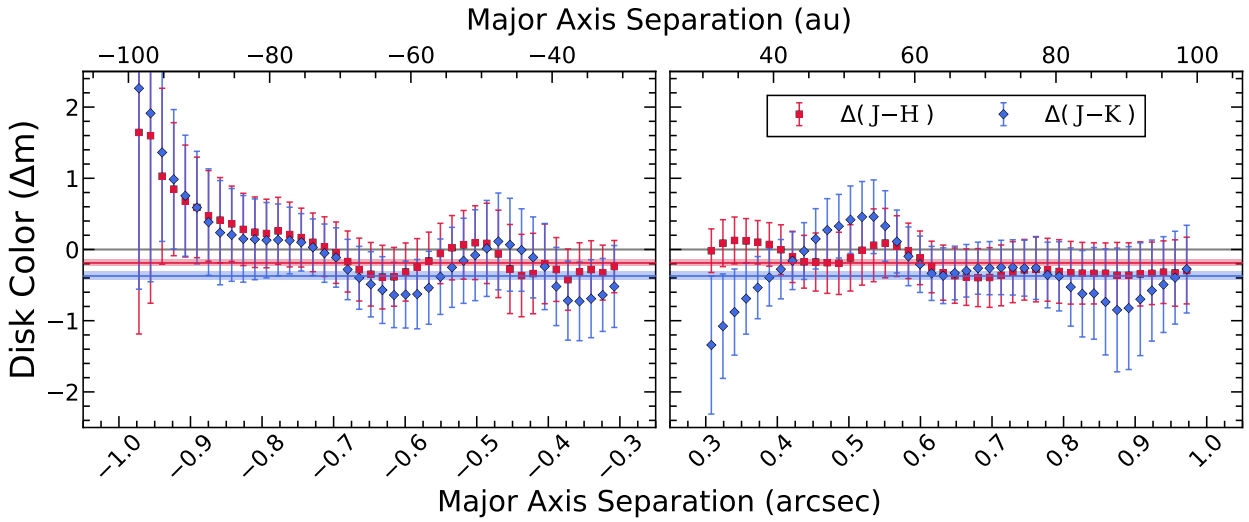

Figure 3.4.6: Disk color as a function of stellocentric separation along the disk major axis (see Section 3.4), with $\Delta(J-K) = (J-K)_{\rm disk} - (J-K)_{\rm star}$, etc. Major axis separation < 0 corresponds to the eastern side of the disk. Horizontal lines and shaded regions of the corresponding color represent the overall average and 1σ confidence respectively. The averages suggest an overall slightly blue color with $\geq 3\sigma$ confidence. $\Delta(H-K)$ manifests similarly to $\Delta(J-H)$.

independent SB measurements for fitting. The locations of these measurements are overlaid on the binned ADI-KLIP image in the top panel of Figure 3.4.7. We then fit power-laws to these SB measurements as a function of deprojected stellocentric separation[12] for the eastern and western extents of the disk separately (Figure 3.4.7). We find best fitting power law slopes of $k = -2.45 \pm 0.14$ and $k = -2.50 \pm 0.21$ for the west and east respectively.

3.4.4. Surface Brightness Discussion

Disk Color – Our individual measurements of disk color along the apparent spine of the disk (Figure 3.4.6) appear generally consistent with a disk scattering profile that is grey in NIR. However, taking the inverse-variance weighted average of the individual measurements for each color suggests a slightly blue color for the disk in both cases: $J-H = -0.19 \pm 0.05$ and $J-K = -0.37 \pm 0.06$. [13] Notably, the locations of the apparent disk spine at which this blue color was measured fall beyond the projected location of peak density for the disk. To assess the properties of the bulk population of dust in the disk, rather than just those of the dust pushed to larger separations by pressure from

[12] Derived deprojected stellocentric separations consider only the determined PA and inclination of the system, and assume that our line of sight probes light scattered from the disk midplane. A more detailed assessment would take scale height, flaring, and similar effects into account, but is beyond the scope of this work.

[13] This general result appears to be robust to the specific disk model adopted for the sake of attenuation correction. Calculating the average disk colors using each of the 691 acceptable disk models for attenuation correction instead resulted in color ranges of $J-H = [-0.23, -0.16]$ and $J-K = [-0.43, -0.32]$ with uncertainties comparable to the values found using the best-fit disk for attenuation correction.

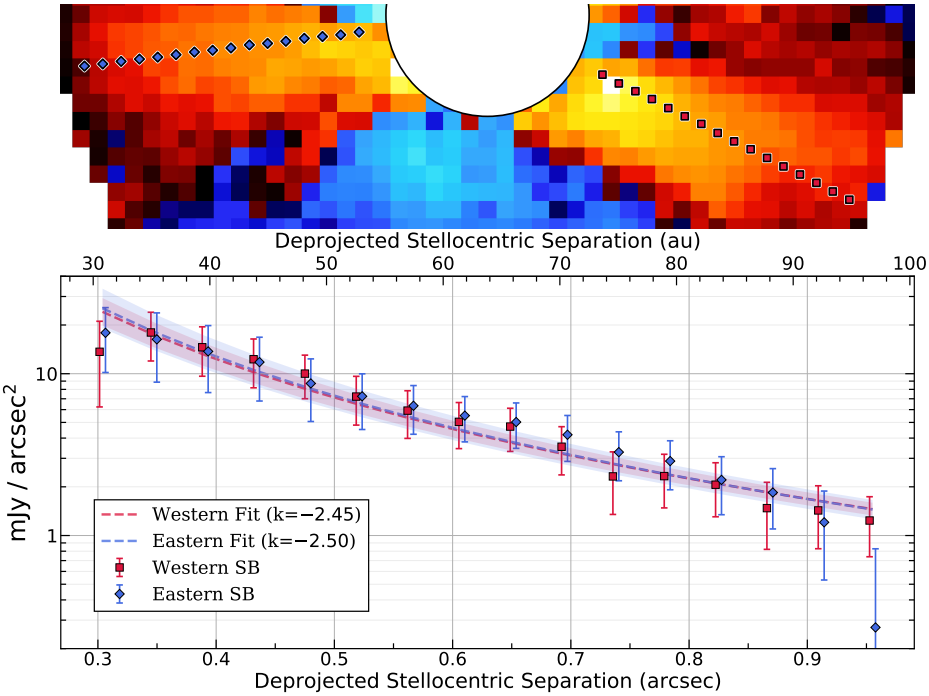

Figure 3.4.7: Broadband radial surface brightness (SB) profiles created following the procedure outlined in Section 3.4.3. The top panel shows the uncorrected ADI-KLIP image binned to resolution, with markers indicating the locations of SB measurements. The lower panel shows the radial surface brightness measurements, with east/west markers offset very slightly from one-another in the x-axis direction to avoid overlap. Power-law fits to the data are shown as dashed lines of the corresponding color (with shaded regions indicating the $1-\sigma$ confidence intervals). Separate sides of the disk are generally in good agreement, with confidence intervals overlapping significantly.

the parent star, we make additional measurements of the disk's color in the vicinity of the disk's peak radial density. For this purpose, we measure surface brightness in apertures of the same size (0.″06) along the brighter side of the disk at the projected location of the adopted disk model's density peak (i.e., along the southern edge of the ellipse overlaid in Figure 3.4.1 – corresponding to a single deprojected stellocentric separation). From these measurements, we proceed to average color measurements as before, finding: $J - H = -0.13 \pm 0.08$, $H - K = -0.30 \pm 0.09$, and $J - K = -0.40 \pm 0.11$. These results again suggest a slightly blue NIR color for the disk.

Comparison with the numerical simulations in Boccaletti et al. (2003), which model NIR disk colors as a function of the dust size distribution's minimum grain size (a_{min}) and porosity (P), can enable some constraints regarding the size distribution of the dust in HD 36546's disk. With $P = 0$, for 1.1 μm versus 2.2 μm ($\sim J - K$), a blue disk color is achieved only for $a_{min} \lesssim 0.25$ μm. Similarly, for 1.6 μm versus 2.2 μm ($\sim H - K$) a blue color is achieved for $a_{min} \lesssim 0.3$ μm. Finally, for 1.1 μm versus 1.6 μm ($\sim J - H$), a blue color is achieved for $a_{min} \lesssim 0.2$ μm[14] Based on these simulations, we can estimate a minimum grain size (assuming P=0) of $a_{min} \lesssim 0.2$ μm to produce the blue color that we observe. The minimum grain size producing a blue color tends to increase for higher porosity, e.g. with $a_{min} \lesssim 3$ μm producing a blue J−K color for P=0.95.

Disk Asymmetry – Both individual measurements and the overall averages for east-west flux asymmetry show no statistically significant evidence for a disk that is asymmetric in brightness in any of the four filters analyzed (Figure 3.4.5). While the processed imagery of the disk may initially suggest some asymmetry is present – with the western extent of the disk appearing brighter in all three presented reductions in Figure 3.2.2 – no asymmetry is evident in SB measurements once attenuation corrections are applied. Thus, we conclude that the apparent asymmetry in CHARIS imagery is an artifact of PSF-subtraction and is not truly astrophysical in nature.

Disk Power-Law – The results of power-law fitting (Figure 3.4.7) are in good agreement with the results of our asymmetry measurements (Figure 3.4.5) — again suggesting that the disk's flux is largely symmetrical at NIR wavelengths in the $\rho = 0.″3 - 1.″0$ region of CHARIS's detection.

Comparing our best-fitting power law slope of $k \sim -2.5$ for HD 36546 ($\rho \sim 25 - 100$ au) to values in the literature for other debris disks provides further evidence that HD 36546 has a particularly shallow density slope. For the debris disk of HD 32297 ($\rho \sim 30 - 280$ au), Currie et al. (2012b) find $k \sim -6$ and for the debris disk of HR 4796 A ($\rho \sim 80 - 120$ au), Thalmann et al. (2011) report $k \sim -10$.

[14]The models of Boccaletti et al. (2003) are presented in terms of $\sim J - K$ and $H - K$. For our purposes, we combine the provided measurements to present these results in terms of $J - H$ as well.

3.5. Limits on Planets

For the planet detection reduction outlined in Section 3.2.2, we computed 5σ contrast limits in CHARIS broadband, following the procedure of Currie et al. (2018) for computing planet throughput corrections.

Contrasts are about a factor of 2–3 worse at $\rho < 0\farcs75$ than good performance for bright diskless stars (Currie et al., 2020a) and slightly poorer than our previous limits for the debris disk-bearing HD 15115 (Lawson et al., 2020). However, despite a loss of wavefront sensing sensitivity (see Section 3.2), we reach contrasts at mid spatial frequencies slightly better than those achieved in Currie et al. (2017) with SCExAO/HiCIAO (e.g. 4.4×10^{-6} at $0\farcs45$ with CHARIS, versus $\sim 8 \times 10^{-6}$ with HiCIAO). This gain is likely due to the use of SDI for additional speckle suppression.

To convert from contrast to planet mass limits, we integrated the hot-start, solar metallicity, hybrid cloud, synthetic planet spectra provided by Spiegel & Burrows (2012) over the CHARIS bandpass and divided by the integrated stellar spectrum. We consider two mappings between contrast and mass, corresponding to the lower and upper age limits estimated for the system (3 Myr and 10 Myr).

The right axis of Figure 3.5.1 denotes these mass limits. For $r \gtrsim 35$ au ($\rho \sim 0\farcs35$), our reduction is sensitive to ~ 2 M_{Jup} planets at ~ 3 Myr, and ~ 5 M_{Jup} planets at ~ 10 Myr. Despite achieving generally inferior contrasts compared to the best reduction of CHARIS data for the debris disk-bearing star HD 15115 (Lawson et al., 2020), we are able to probe slightly less massive planets at a given projected separation due to HD 36546's relative youth.

While our data do not result in the detection of any candidate planets, HD 36546 may show tentative, indirect evidence for a massive and (in-principle) imageable companion. The *Hipparcos-Gaia* Catalog of Accelerations (HGCA; Brandt 2018) provides calibrated proper motion data for $\sim 115,000$ nearby stars spanning the ~ 24 year baseline between *Hipparcos* and *Gaia* astrometry. Accelerations based on this data have been used to measure dynamical masses of directly imaged exoplanets (e.g. Brandt et al., 2019), as well as to provide evidence of gravitational interaction with yet-unseen companions for targeted direct imaging searches (Currie et al., 2020b; Steiger et al., 2021). Brandt et al. (2021, submitted) updates HGCA, replacing *Gaia* DR2 measurements with those from *Gaia* eDR3, which yield a factor of ~ 3 improvement in astrometric precision. Using these updated *Gaia* measurements, HGCA for HD 36546 (Table 3.5.1) indicates a marginally significant acceleration relative to a model of constant proper motion — $\chi^2 = 7.65$, or 2.3σ for two degrees of freedom — that *could* be indicative of a companion.

Following equations in Brandt et al. (2019) to estimate companion mass as a function of projected separation using HGCA accelerations, we can assess the completeness of our search for the HD 36546 companion that would induce this acceleration. Since radial velocity (RV)

Table 3.5.1. HGCA Measurements of HD 36546

Source	μ_α	σ_α	μ_δ	σ_δ	Corr	t_α	t_δ
Hip	8.513	0.974	-41.876	0.472	0.257	1991.15	1990.76
HG	7.368	0.028	-41.226	0.014	0.087		
Gaia	7.508	0.056	-41.305	0.040	-0.137	2016.00	2016.50

Note. — HGCA calibrated proper motion measurements for HD 36546 (Brandt, 2018) as updated with *Gaia* eDR3 measurements in Brandt et al. (2021, submitted).

measurements are not available for HD 36546, we adopt an RV acceleration of zero to compute a minimum mass instead. Mapping of the resulting minimum masses to contrasts for 3 Myr and 10 Myr are overlaid on Figure 3.5.1. These results suggest that we would likely recover the HGCA companion for a projected separation beyond $\sim 0\farcs15$ assuming an age of 3 Myr, or $\sim 0\farcs18$ for an age of 10 Myr.

3.6. Conclusions and Future Work

Our study provides the first multi-wavelength imaging of the HD 36546 disk and recovers disk signal to separations as small as $\sim 0\farcs25$. Through detailed modeling implementing a Hong (1985) scattering phase function and using the differential evolution optimization algorithm, we provide an updated schematic of HD 36546's disk. This schematic suggests a disk with a particularly shallow radial dust density profile, a fiducial radius of ~ 83 au, an inclination of $\sim 79°$, and a position angle of $\sim 80°$. Through spine tracing, we find a disk spine that is consistent with our modeling results, but also consistent with a "swept-back wing" geometry. Surface brightness measurements show no significant flux asymmetry between the eastern and western extents of the disk, and slightly blue NIR colors on average. While we report no evidence of directly imaged companions, we provide contrast limits to constrain the most plausible remaining companions.

Notably, prior work modeling the debris disks of HD 36546 and HD 15115 in NIR total intensity data identified very distinct optimal HG SPF asymmetry parameters of $g = 0.70 - 0.85$ (Currie et al., 2017) and $g = 0.3 - 0.4$ (Engler et al., 2019) respectively. Alongside the results of Lawson et al. (2020), we have shown that imagery of both disks can be reasonably reproduced by models using the same Hong (1985) SPF instead. As is suggested in Hughes et al. (2018), this may support the existence of a nearly universal SPF for circumstellar dust. Additional data to study smaller scattering angles and additional disk systems is necessary to clarify this possibility. If such an SPF is verified,

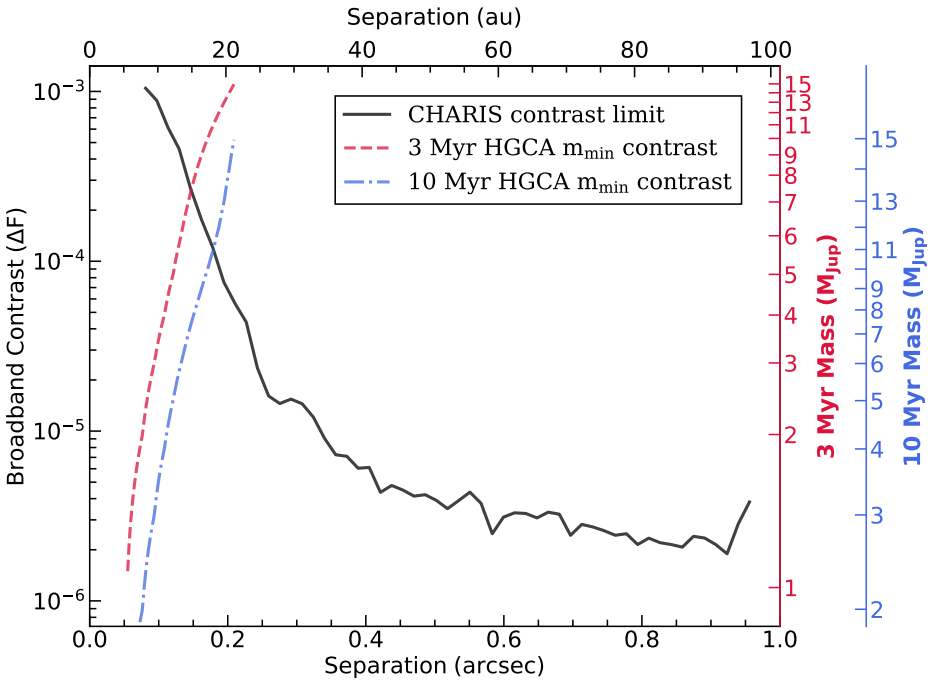

Figure 3.5.1: Broadband ($1.13 - 2.39$ μm) contrast curve (solid black line) for the planet detection reduction of CHARIS HD 36546 data, as outlined in Section 3.2.2. 5σ contrast is given as a function of stellocentric angular separation (arcsec, lower x-axis) and projected separation (au, upper x-axis). The right edge of the figure gives contrast mapped to companion masses for 3 Myr (red) and 10 Myr (blue) planets, based on hot-start, solar-metallicity, hybrid cloud planet evolution models of Spiegel & Burrows (2012). Overplotted lines of the corresponding color indicate the minimum mass and projected separations corresponding to the companion implied by HGCA measurements, mapped to contrasts based on the same models. The displayed ages represent the lower and upper limits of age estimates for HD 36546.

it will enable both better exploration of model parameters (by reducing the number of varying parameters) and tighter constraints on parameter values (by mitigating parameter degeneracies).

Deeper follow-up observations with SCExAO/CHARIS would provide a) further constraints on companions within the disk and b) higher SNR measurements of the disk color, enabling a more detailed assessment of the dust properties of the disk. The upcoming replacement of Subaru's AO188 system and additional commissioned improvements to wavefront sensing are expected to improve SCExAO/CHARIS contrasts at small separations by roughly a factor of 10 by late 2021 (Currie et al., 2020a). Assuming a system age of 3 Myr, these contrasts would enable detection of $\sim$1 M_{Jup}-mass planets at $\sim$ 25 au. A future non-detection would likely restrict the location of any super-Jovian mass planet to $\rho \lesssim 0\farcs1$.

Observations measuring the polarized intensity of the disk at comparable scales – or the total intensity of the disk in a wider field of view – could clarify the underlying cause of the noted diverging spine trace. CHARIS's integral field spectropolarimetry mode or space-based imaging of the system could provide these data. Further, polarized intensity imaging would enable even tighter constraints on disk model parameters by better probing regions of the disk where the polarized intensity phase function results in more favorable sensitivity.

Appendix

3.A. Additional Model Results

The results of our additional differential evolution disk model optimization procedures are included here, with the information of Table 3.3.1 graphically summarized in Figure 3.A.1. Individual reductions are in close agreement with one another; while the optimal values for some parameters (e.g. R_0) differ noticeably compared to the permitted parameter bounds, each reduction's best fit values fall within the acceptable ranges for the other reductions.

Comparison of the samples for α_{in} between the two RDI-KLIP DE runs (the first two rows of Figure 3.A.1) illustrates our prior warning regarding approximation of parameter uncertainties from DE (see Section 3.3.3). While the runs ultimately yield very similar acceptable ranges, the comparable upper limit achieved for "RDI-KLIP (2)" is the result of a single sample; except for the lone acceptable model with $\alpha_{in} \sim 5.5$, the acceptable range of α_{in} values for this run would manifest much more comparably to that of the ADI-KLIP run.

As a result of the larger noise levels in the RDI-KLIP reduction, the RDI-KLIP modeling procedure (Fig. 3.A.2) shows much weaker constraints than the ADI-KLIP (Fig. 3.A.3) modeling procedure. Despite some differences in the resulting parameter values, the best-fit solutions for these procedures manifest very similarly to those of the combined run (e.g. Figures 3.3.2, 3.3.3).

3.B. RDI-KLIP Reduction using Radial Profile Subtraction

The results of forward modeling the adopted best-fitting model for the radial profile subtracted RDI-KLIP reduction (see Section 3.4) are provided in Figure 3.B.1.

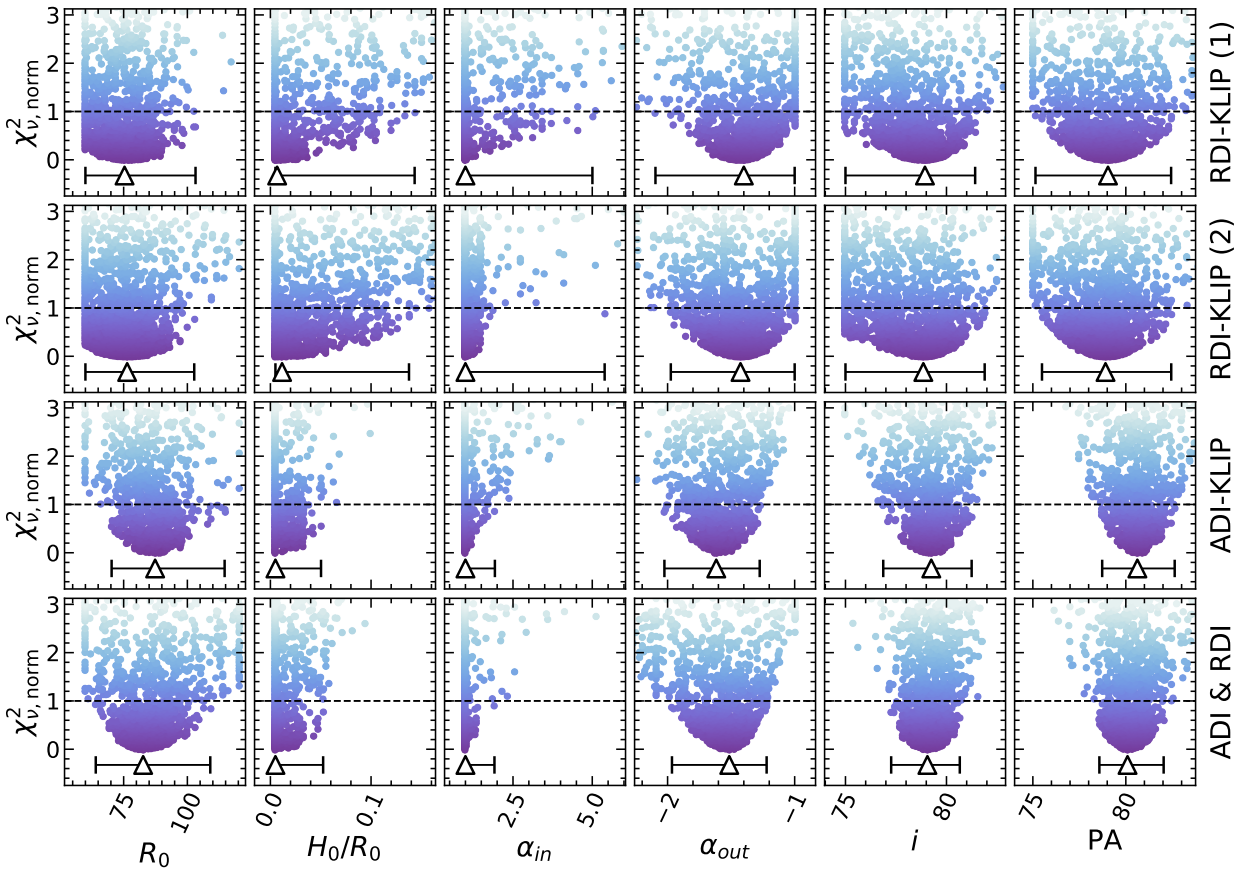

Figure 3.A.1: The diagonal (one dimensional) elements of corner plots for each of the four DE procedures (as indicated to the right of each row of subplots), as summarized in Table 3.3.1. χ^2_ν is presented normalized as $\chi^2_{\nu,norm} = (\chi^2_\nu - \chi^2_{\nu,min})/\sqrt{2/\nu}$, such that the best model for each run falls at $\chi^2_{\nu,norm} = 0$ and acceptable models have $\chi^2_{\nu,norm} \leq 1$. Parameter bounds are truncated to better show the distributions in the vicinity of the minima.

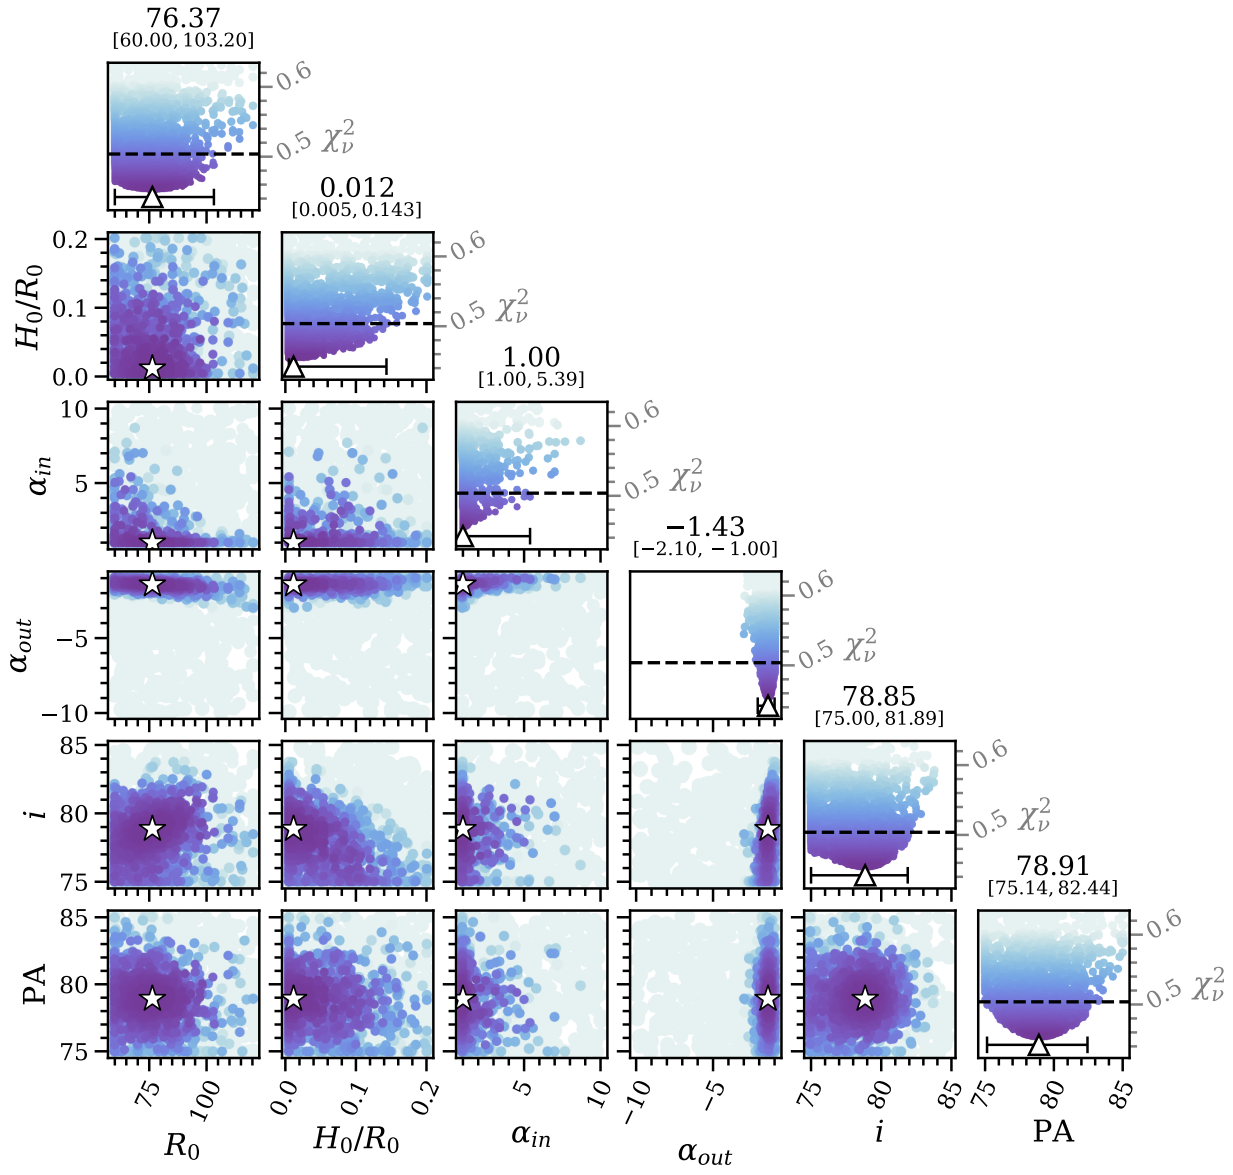

Figure 3.A.2: As Figure 3.3.1, but for optimization of the disk model to only the RDI-KLIP reduction of HD 36546. This figure includes both runs for the RDI-KLIP data (corresponding to the entries for "RDI-KLIP (1)" and "RDI-KLIP (2)" in Table 3.3.1).

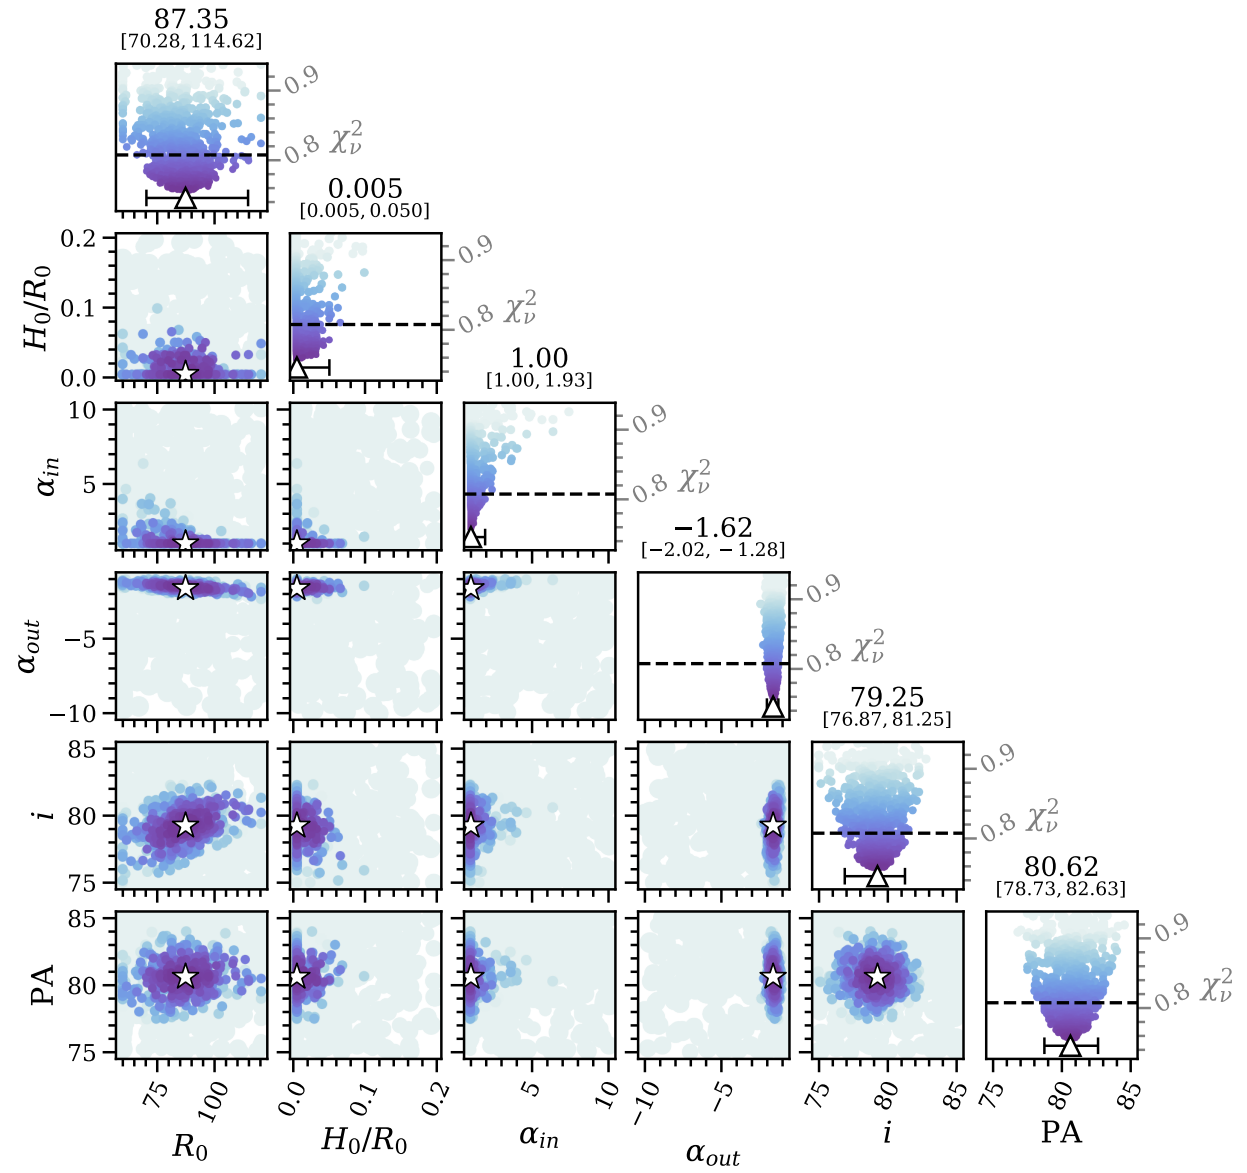

Figure 3.A.3: As Figure 3.3.1, but for optimization of the disk model to only the ADI-KLIP reduction of HD 36546 (corresponding to the entry for "ADI-KLIP" in Table 3.3.1). Compared to that of the RDI-KLIP data, the ADI-KLIP optimization results in stronger constraints for the model parameters by virtue of a higher signal-to-noise ratio.

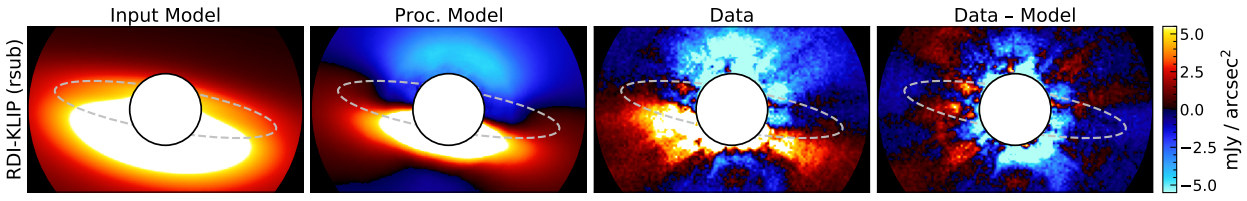

Figure 3.B.1: As Figure 3.3.2, but for the RDI-KLIP reduction with radial profile subtraction that was utilized for spine tracing and surface brightness measurements (Section 3.4).

CHAPTER 4

Integral Field Spectropolarimetry of Planet-forming Disks[1]

4.1. Introduction

Protoplanetary disks serve as benchmark systems to study how and where exoplanets form. With the advent of ground-based extreme adaptive optics (AO) facilities, such as Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE; Beuzit et al., 2019), Gemini Planet Imager (GPI; Macintosh et al., 2015), and Subaru Coronagraphic Extreme Adaptive Optics (SCExAO; Jovanovic et al., 2015; Lozi et al., 2018; Currie et al., 2020a), we have entered a revolutionary era in the study of young protoplanetary disks. It is now feasible to both a) spatially resolve the morphological signatures within disks that could be caused by newly formed/forming sub-stellar objects and planets (such as gaps or spirals), and b) identify young candidate planets and sub-stellar objects that may cause them (e.g. Lagrange et al. 2010; Keppler et al. 2018).

Differential imaging techniques for high-contrast imaging in total intensity, such as angular differential imaging (ADI; Marois et al., 2006), spectral differential imaging (SDI; Sparks & Ford, 2002), and reference star differential imaging (RDI), attempt to disentangle the light of the bright stellar point spread function (PSF) from that of any circumstellar sources by modeling and subtracting the starlight in a data sequence. While these techniques are generally reasonably well suited for the recovery and characterization of isolated point-like companion candidates, the presence of signal from circumstellar disks can significantly inhibit their efficacy – resulting in products in which disk signal is substantially attenuated. For ADI and SDI, this results from self-subtraction (where disk signal is erroneously included in the model PSF) and oversubtraction (where the presence of disk signal in the data results in a brighter PSF model than if it were not present)(Pueyo, 2016; Currie et al., 2019). While RDI is expected to perform better for disk imaging – by eliminating the possibility of self-subtraction – oversubtraction still occurs, and can be significant for bright disks or disks which dominate the field of view (FOV). As a result, any analysis of disks recovered from total intensity differential imaging will generally require the attenuation to be quantified. Further, rigorous assessment of planet candidates embedded in disk material requires a means by which to differentiate a planet's signal from that of disk substructure (e.g. Rich et al. 2019; Currie et al. 2019). Forward-modeling of synthetic disk models is a common approach for approximating disk attenuation (e.g. Currie et al. 2019; Lawson et al. 2020). Combined with integral field spectroscopy or multi-wavelength imaging, this can also provide a means by

[1]This chapter is reproduced from Lawson et al. (2021b) with permission.

which to validate or invalidate candidate embedded planets. However, given the complexity of total intensity PSF-subtraction and of the geometry of many protoplanetary disks, forward-modeling can be prohibitively time-consuming.

High-contrast near-infrared (NIR) polarimetric imagers – such as Gemini South's Gemini Planet Imager (GPI; Macintosh et al., 2015; Perrin et al., 2015), VLT's SPHERE-InfraRed Dual-band Imager and Spectrograph (IRDIS; de Boer et al., 2020; van Holstein et al., 2020), and Subaru's High-Contrast Coronographic Imager for Adaptive Optics (HiCIAO; Suzuki et al., 2010) – have been prolific in detecting and characterizing circumstellar disks through the use of another differential imaging technique: polarimetric differential imaging (PDI). While PDI is limited to the recovery of polarized flux, it provides an important complement for total intensity differential imaging by producing a nearly unattenuated view of the disk. Additionally, since self-luminous sub-stellar objects and exoplanets are expected to be unpolarized while light from spatially resolved circumstellar disks is $25 - 50\%$ polarized (Perrin et al., 2009; Uyama et al., 2017), comparison of total and polarized intensity imaging can serve as another powerful tool for disentangling planet signals from disk sub-structure. Since the PDI procedure is computationally simple, forward modeling disks to simulate the effects of PDI is substantially less time consuming – thus allowing for more detailed investigation of disk model parameters.

Since the decommissioning of Subaru's HiCIAO, the Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS) NIR integral field spectrograph (IFS) (Groff et al., 2016), paired with SCExAO, has served as Subaru's NIR total intensity high-contrast imager. More recently, CHARIS was upgraded to include spectropolarimetric capabilities as well (Lozi et al., 2020). Unprecedented among extreme AO imagers, this new spectropolarimetric imaging mode allows polarimetric observations to be conducted at the same array of wavelengths as its classical (total intensity) observing mode. This capability makes accessible many novel investigations relevant for protoplanetary disk studies, such as characterization of the wavelength dependence of the scattering phase function in disks.

We present preliminary science products from the novel integral field spectropolarimetry mode, or PDI-mode, for Subaru's SCExAO/CHARIS. Combining the benefits of multi-wavelength IFS imaging with those of polarimetric imaging, this new observing mode provides an exciting tool for groups studying circumstellar disks. Herein, we: 1) provide a summary of the CHARIS PDI-mode, outline its available configurations, and make recommendations for those observing with this mode (Section 4.2), 2) provide a description of the data processing and calibration procedures for CHARIS PDI data, implemented in a new module for the CHARIS Data Processing Pipeline (DPP) (Section 4.3), 3) present preliminary science products for two protoplanetary disk targets observed with CHARIS in PDI-mode (Section 4.4), 4) demonstrate disk forward-modeling for CHARIS PDI data (Section 4.5), and 5) conclude by summarizing compelling future use-cases for CHARIS's

PDI-mode in application to the study of circumstellar disks (Section 4.6).

4.2. CHARIS Spectropolarimetry Mode

The spectropolarimetric observing mode for CHARIS is enabled by the addition of a field stop and Wollaston prism upstream of the imager in combination with the existing half-wave plate (HWP) – originally added to enable polarimetry with HiCIAO. The Wollaston prism serves to split incoming light to two orthogonally polarized states, while the HWP enables shifting of the polarization direction being measured, and the field stop eliminates cross-talk between the two polarizations[2]. Together, the Wollaston prism and HWP enable measurement of linear polarization parameters (Stokes-Q and U; in addition to the total intensity, Stokes-I, measured by the classical-mode), at CHARIS's typical array of wavelengths, at the cost of reducing CHARIS's FOV from $\sim 2'' \times 2''$ to $\sim 1'' \times 2''$. In the typical PDI procedure, circumstellar polarized flux is isolated by making observations at HWP angles of 0° and 45° or 22.°5 and 67.°5 and using "double-differencing" (see a full description of PDI in Section 4.3.4).

A complete model for the correction of instrumental polarization for CHARIS PDI-mode is near completion (see van Holstein et al. 2020) and will be implemented in the data processing pipeline for CHARIS soon after. In lieu of this correction, results herein should be considered preliminary.

4.2.1. Half-wave Plate Cycles

Switching of the HWP requires ~ 6 seconds. While the most straight-forward strategy might be to change the HWP angle between each exposure, this could introduce significant overheads for targets requiring short exposure times. Especially in these cases, taking multiple exposures at a given HWP position before each switch may improve the quality of results by enabling more time on target overall. This introduces two possible costs, however. The first is that the suppression of unpolarized stellar flux may be compromised by evolution of the diffraction pattern over the relatively protracted HWP cycles. Additionally, since SCExAO operates exclusively in pupil-tracking mode, exposures at paired HWP positions for double differencing will be at slightly different parallactic angles. This could result in noticeable smearing of circumstellar features for targets observed near zenith (though the impact of this is unlikely to be significant, see: Section 4.5). As such, target properties, observing conditions, and the circumstellar features of interest should all be considered when planning observations.

4.2.2. Instrument Configurations

With some exceptions, the available configurations for CHARIS's PDI-mode are the same as for its classical-mode. These configurations are summarized hereafter.

[2] see van Holstein et al. 2020; Lozi et al. 2020 for a more complete description of the instrumental design of the CHARIS spectropolarimetry mode

Astrogrid – As with CHARIS's classical imaging mode, CHARIS PDI-mode allows for use of the astrogrid during observations. The astrogrid introduces a set of satellite spots which are attenuated copies of the unocculted stellar PSF. The astrogrid enables high quality image registration and flux calibration even when the star is obscured by a coronagraph. While the same array of astrogrid spacings are available as for CHARIS's classical mode, the standard classical-mode spacing of 15.9 λ/D places spots outside of the narrower $2'' \times 1''$ FOV of PDI-mode by the latter channels of H-band. Thus a spacing of ~ 11.2 λ/D or smaller should be utilized for the PDI-mode (Lozi et al., 2020). While a limited sequence of exposures utilizing the astrogrid might be sufficient for reasonable registration and flux calibration, use of the astrogrid throughout the entirety of an observing sequence is strongly recommended unless a compelling science case exists to preclude this – such as simultaneous use of the Visible Aperture Masking Polarimetric Imager for Resolved Exoplanetary Structures (VAMPIRES) instrument, which cannot be used with the astrogrid on.

Filters and Spectral Resolution – J-band ($1.18 - 1.33$ μm), H-band ($1.49 - 1.78$ μm), and K-band ($2.02 - 2.38$ μm) are available with high spectral resolution ($R \sim 70 - 90$, yielding $15 - 20$ wavelength channels in extracted image cubes). CHARIS broadband ($1.15 - 2.39$ μm), and Cold ND3 ($1.17 - 2.37$ μm) are available with low spectral resolution ($R \sim 19$, yielding 22 wavelength channels in extracted image cubes).

Coronagraphs – the Lyot coronagraphs[3] can be used with CHARIS PDI-mode for any available filter. Additionally, the phase-induced amplitude apodization complex mask coronagraph (PIAACMC) can be used when observing in H-band. Other coronagraphs that are generally available for SCExAO, e.g. vector Apodizing Phase Plate (vAPP) and vortex coronagraphs, induce intractable changes in the polarization state of incident light and are thus unsuitable for observations in PDI-mode.

4.2.3. Calibration Data

With the complete Mueller matrix model for instrumental polarization implemented, observations of polarized/unpolarized standard stars will not be directly necessary for calibration of science products. However, these observations are still recommended, as they require very little time, can be used to verify the integrity of the instrumental polarization corrections, and can be useful in PSF-subtraction of total-intensity products by enabling RDI.

4.2.4. Observations

To demonstrate the use of the methods described herein and to showcase this new observing mode, we include preliminary results of observations of two protoplanetary disk-bearing stars. In both cases, we reserve detailed analysis, discussion, and target-specific conclusions for future

[3] See https://www.naoj.org/Projects/SCEXAO/ for information regarding available coronagraph sizes.

publications.

AB Aurigae– AB Aurigae (1–3 Myr , Kenyon & Bromley 2008; $d = 156$ pc, Gaia Collaboration et al. 2016, 2021) is a pre-main sequence star that hosts a complex protoplanetary disk bearing spiral arms at scales from tens of au(Hashimoto et al., 2011; Boccaletti et al., 2020) to hundreds(Grady et al., 1999; Fukagawa et al., 2004).

On 2020 October 4, we observed AB Aur with SCExAO/CHARIS in low-res ($R \sim 19$), broadband ($1.15 - 2.39\ \mu m$), PDI-mode, using a 113 mas radius Lyot coronagraph and a four-spot astrogrid of spacing $11.2\ \lambda/D$. AO performance was average / below-average, with a strong wind-driven halo. We collected 73 science exposures of 60.48 seconds each for a combined integration time of ~ 73.6 minutes, achieving a total field rotation of $79.^{\circ}95$. Sky frames were also collected at the end of the sequence.

Additional observations were made earlier in the same night using an astrogrid spacing of $15.9\ \lambda/D$. Since the satellite spots are not recoverable in every channel at this spacing, image calibration for this data is much less precise. As such, we present only the observations with $11.2\ \lambda/D$ astrogrid spacing herein.

TW Hydrae– TW Hydrae (3–10 Myr, Hoff et al. 1998; Barrado Y Navascués 2006; Vacca & Sandell 2011; van Boekel et al. 2017; $d = 60$ pcGaia Collaboration et al. 2016, 2021) is a nearby T-Tauri star with a nearly face-on disk which shows multiple radial gaps that could suggest the presence of planets or their ongoing formation (Akiyama et al., 2015; van Boekel et al., 2017).

On 2021 March 20, we observed TW Hya with CHARIS in the same configuration as for AB Aur, but using a smaller Lyot coronagraph (78 mas radius). The AO performance for TW Hya was worse than that of AB Aur, largely as a result of TW Hya's fainter apparent magnitude and lower altitude ($\sim 35^{\circ}$ versus $\gtrsim 74^{\circ}$ for AB Aur). We collected 52 science exposures of 60.48 seconds each for a combined integration time of ~ 52.4 minutes, achieving a total field rotation of $15.^{\circ}84$. No sky frames were collected with these data.

4.3. CHARIS PDI Data Reduction and Processing

4.3.1. Cube Extraction

CHARIS PDI data were extracted from raw IFS exposures using the CHARIS Data Reduction Pipeline (DRP) (Brandt et al., 2017) with modifications to read noise suppression provided in Currie et al. (2020a). The 22 wavelength channels of a single extracted AB Aur image cube are shown in Figure 4.3.1.

4.3.2. Preprocessing

Preliminary processing of data for PDI-mode – e.g. sky subtraction, image registration, and spectrophotometric calibration – is carried out largely as typical for CHARIS classical-mode data

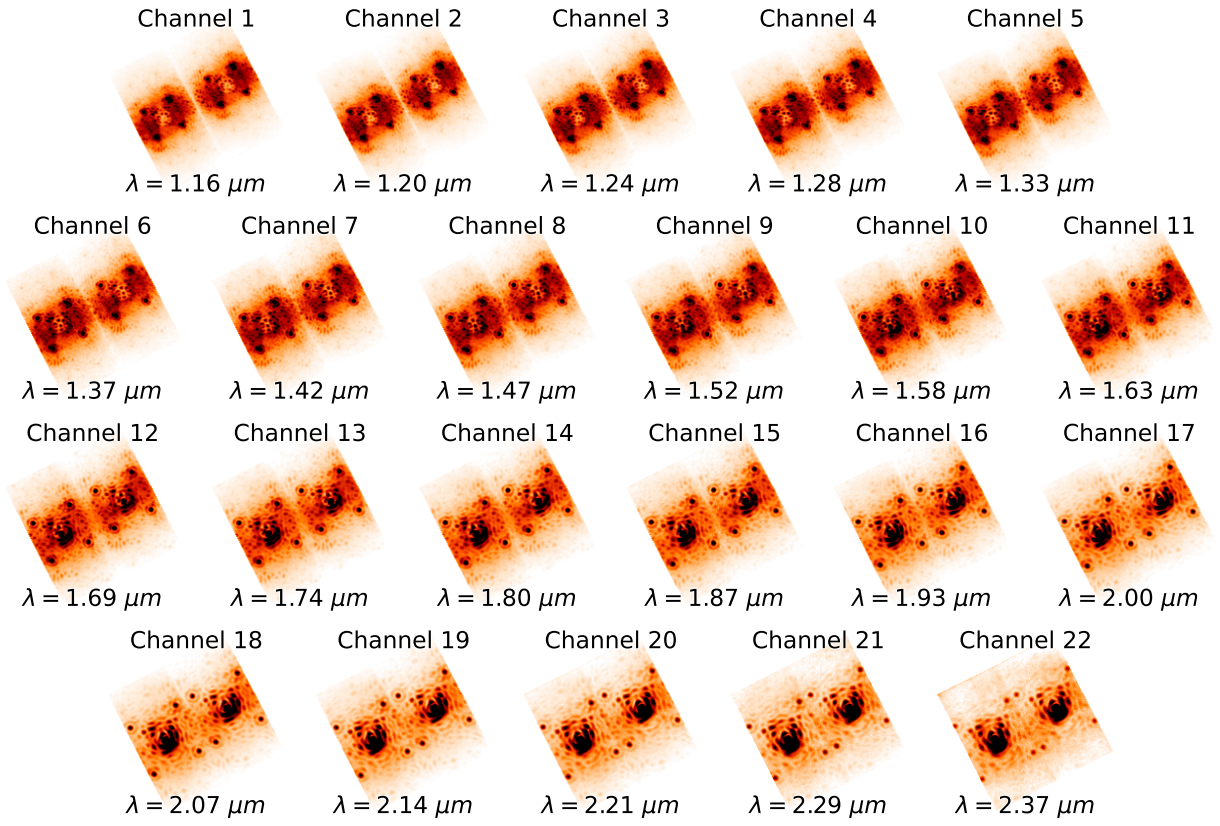

Figure 4.3.1: The 22 individual wavelength channels of an extracted low-res, broadband, PDI-mode CHARIS exposure of AB Aur – showing both left and right polarizations. For this exposure, a four-satellite-spot astrogrid with spacing of 11.2 λ/D, and a Lyot-coronagraph of radius 113 mas were utilized.

(Currie et al., 2018), with the following caveats.

1. **Registration** – the left and right polarized halves of each image cube are first split into separate files. The sets of left and right frames are then registered independently, but otherwise as standard for CHARIS data (e.g. using satellite spots for occulted data, or PSF centroiding for unocculted data).

2. **Spectrophotometric Calibration** – for this purpose, the sets of registered left and right polarization cubes are added to form a sequence of single sum total intensity images. These images are then flux calibrated in the same manner as classical-mode data – resulting in a set of scaling factors for each of the 22 wavelength channels. These scaling factors are then applied to the original constituent left and right polarized image cubes.

4.3.3. Half-Wave Plate Cycle Matching

To enable PDI using double-differencing, we group observations into sets containing exposures at each of the four HWP positions. While sets of two images (with HWP angles of 0° and 45° or 22.5 and 67.5) would be sufficient to enable double-differencing (see Section 4.3.4), forming groups of complete HWP cycles allows polarized intensity outputs for each individual cycle – which can be helpful in evaluating the validity of marginally detected features in products for the fully-combined observing sequence.

Here, each exposure is used only once, and is grouped to minimize parallactic angle differences between the exposures in a cycle. Ultimately, this procedure results in a number of full HWP cycles equal to the minimum number of observations for any one HWP position after low-quality or otherwise problematic exposures are eliminated. Any unmatched exposures are not utilized in the PDI procedure.

4.3.4. Polarimetric Differential Imaging

Reduction of the preprocessed data sequence to final polarimetry products is carried out using the typical double differencing procedure for each HWP cycle (see e.g. de Boer et al. 2020). First, single sums and differences are computed for each exposure (adopting the notation of van Holstein et al. 2020):

$$X^{\pm} = I_{det,L} - I_{det,R}\,, \tag{4.1}$$

$$I_{X^{\pm}} = I_{det,L} + I_{det,R}\,, \tag{4.2}$$

where $I_{det,L}$ and $I_{det,R}$ are the left and right images respectively, $X^{\pm}$ are the single differences – Q^{+}, U^{+}, Q^{-}, U^{-} – and $I_{X^{\pm}}$ are the single sums – $I_{Q^{+}}$, $I_{U^{+}}$, $I_{Q^{-}}$, $I_{U^{-}}$ – for exposures at HWP angles of 0°, 22.5, 45°, and 67.5 respectively. From these, double sums and differences are computed as:

$$X = \frac{1}{2}(X^{+} - X^{-})\,, \tag{4.3}$$

$$I_X = \frac{1}{2}(I_{X^{+}} + I_{X^{-}})\,, \tag{4.4}$$

where the double difference, X, corresponds to Stokes Q or U, and the double sum, I_X, corresponds to the components of Stokes I: I_Q or I_U. At this stage, in lieu of a final model for instrumental polarization, we make a first order approximation of instrumental polarization following the procedure of de Boer et al. (2020) and summarized hereafter. For each double difference image cube, X, resulting from the procedure above, the ratio with the corresponding double-sum image cube, I_X, is computed: $x = X/I_X$. Then, in each wavelength channel, i, we calculate the median value of x within a circular stellocentric aperture (excluding the inner occulted region if a coronagraph was utilized): c_X^i. The first-order instrumental polarization subtracted cube is then:

$$X_{\rm IPS} = X - c_X \cdot I_X\,, \tag{4.5}$$

where c_X is the array of coefficients for the N_λ wavelength channels, $c_X = [c_X^1, c_X^2, ..., c_X^{N_\lambda}]$. Since this procedure assumes that the polarization measured in the aperture should be zero in the absence of instrumental polarization, the coverage of the aperture should be chosen to minimize inclusion of circumstellar material. An example of a single wavelength channel of Stokes parameter images resulting from this procedure is shown in Figure 4.3.2.

Once $Q_{\rm IPS}$ and $U_{\rm IPS}$ (hereafter simply Q and U) are attained, the products can be derotated to a "north-up" orientation. Noting that a) X^+ and X^- are not perfectly contemporaneous, and b) SCExAO observes exclusively in pupil-tracking mode: the constituent images will have slightly different angular offsets from north-up. Here, we simply use the average of the parallactic angles of the two constituent frames to compute the rotation needed for a north-up orientation. The effect of this difference in parallactic angle is assessed in Section 4.5.

From the north-up polarization products, we can compute a few distinct polarized intensity products. Generally:

$$\mathrm{PI} = \sqrt{\mathrm{Q}^2 + \mathrm{U}^2}\,, \tag{4.6}$$

where PI is the polarized intensity. For a single HWP cycle:

1. a corresponding PI cube can be created,

2. a selection of wavelength slices of the Q and U cubes can be combined to approximate common NIR filters (e.g., J, H or K bands), from which a corresponding PI image can be calculated,

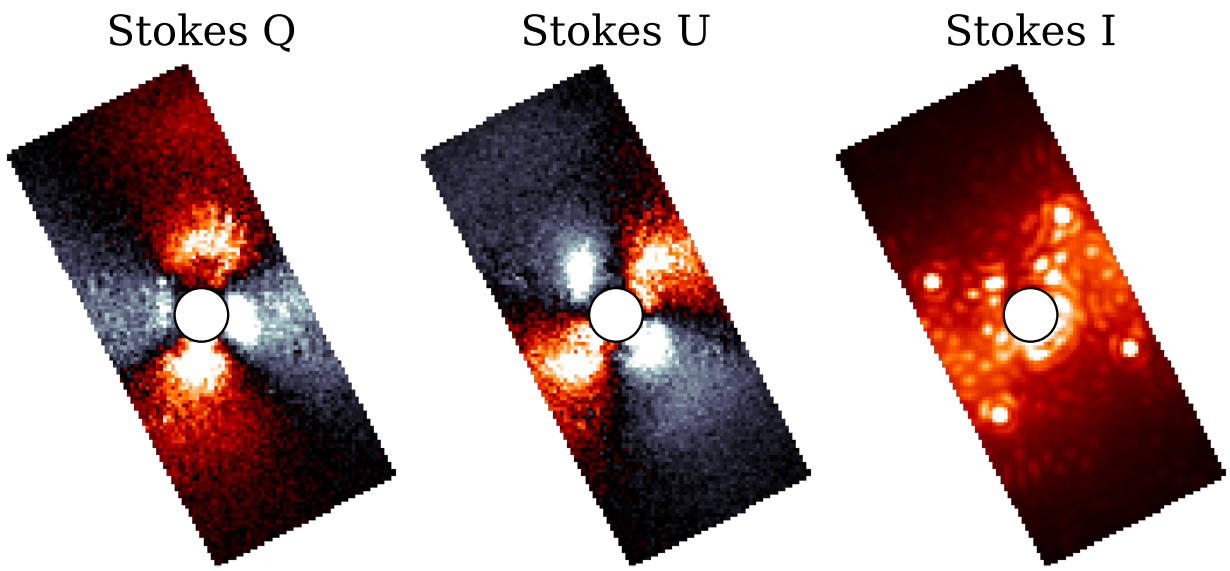

Figure 4.3.2: A single wavelength channel (channel 11, $\lambda = 1.63\ \mu m$) showing Stokes Q, U, and I images resulting for one HWP cycle of CHARIS AB Aur data (including the exposure of Figure 4.3.1). I is computed from the I_Q and I_U products for the cycle as $I = 0.5(I_Q + I_U)$. Q and U images have a first order instrumental polarization correction applied. Color stretches are symmetric about zero, with the range for I differing from that of Q and U for visualization. The frames have not yet been derotated.

3. all wavelength slices for Q and U cubes can be combined to enable calculation of a higher signal-to-noise ratio (SNR) PI image.

Additionally, we can average the full set of Q and U image cubes over all cycles to produce a sequence-combined Q and U image cube. From this, the same array of options for presenting PI are available. Since the nominal $1'' \times 2''$ FOV has been derotated by varying amounts, a sequence-combined product will have full coverage over $r \lesssim 0\farcs5$ and partial coverage of a larger area – approaching a $\sim 1''$ circle of partial coverage when a total field rotation of $\Delta\mathrm{PA} \sim 125°$ is achieved. For $\Delta\mathrm{PA} \lesssim 125°$, the region of zero coverage in a sequence-combined product manifests as two opposing wedges extending from $r \gtrsim 0\farcs5$ (see Figure 4.4.1). We note, however, that the exact FOV recovered will depend on the details of the dataset.

For each of these cases, we can also calculate the angle of linear polarization (AOLP):

$$\mathrm{AOLP} = \frac{1}{2}\arctan\left(\frac{\mathrm{U}}{\mathrm{Q}}\right), \tag{4.7}$$

as well as the azimuthal Stokes parameters:

$$Q_\phi = -Q\cos(2\phi) - U\sin(2\phi), \tag{4.8}$$

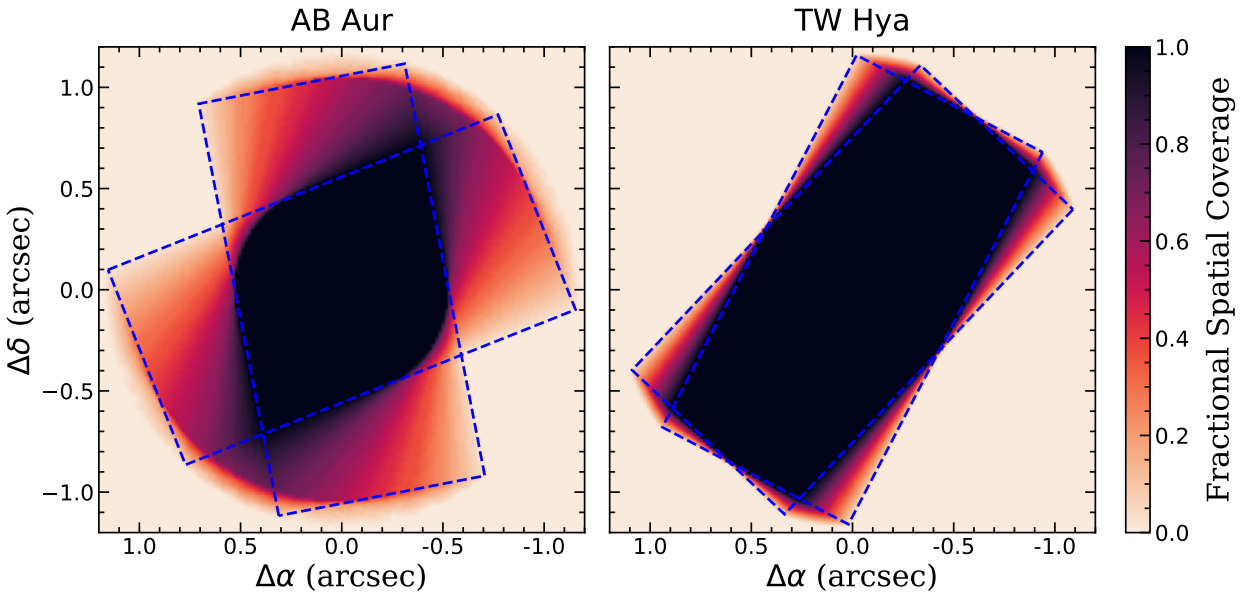

Figure 4.4.1: Approximate spatial coverage maps for our sequence-combined AB Aur and TW Hya products. The dashed blue rectangles indicate the extents of the derotated first and last image in each sequence. TW Hya has significantly more area with zero coverage but also more area with full coverage as a result of the TW Hya data's lesser field rotation.

$$U_\phi = Q \sin(2\phi) - U \cos(2\phi)\,, \tag{4.9}$$

where ϕ is the pixel-wise azimuthal angle relative to the image center.

4.4. Preliminary CHARIS Spectropolarimetry Results

We apply the methods of Section 4.3 to the CHARIS PDI data of AB Aur and TW Hya (see Section 4.2.4). Both targets were registered and flux calibrated using the induced astrogrid. Approximate coverage maps for the final sequence-combined products are visualized in Figure 4.4.1.

For both targets, a first order instrumental polarization correction was computed within an annulus extending from $r = 10$ to $r = 30$ pixels ($\sim 0\farcs16 - 0\farcs49$). While significant polarized disk signal is present here, a) this is the case throughout the CHARIS FOV for these targets, and b) this approximation is reasonable only with full azimuthal coverage over the annulus, which occurs for $r \lesssim 30$ pixels ($\sim 0\farcs5$) with CHARIS PDI data. However, we note that the inclusion of disk material in the annulus means that this correction is very likely over-estimating the contribution of instrumental polarization; this correction is used only as a placeholder until the formal Mueller Matrix correction is available.

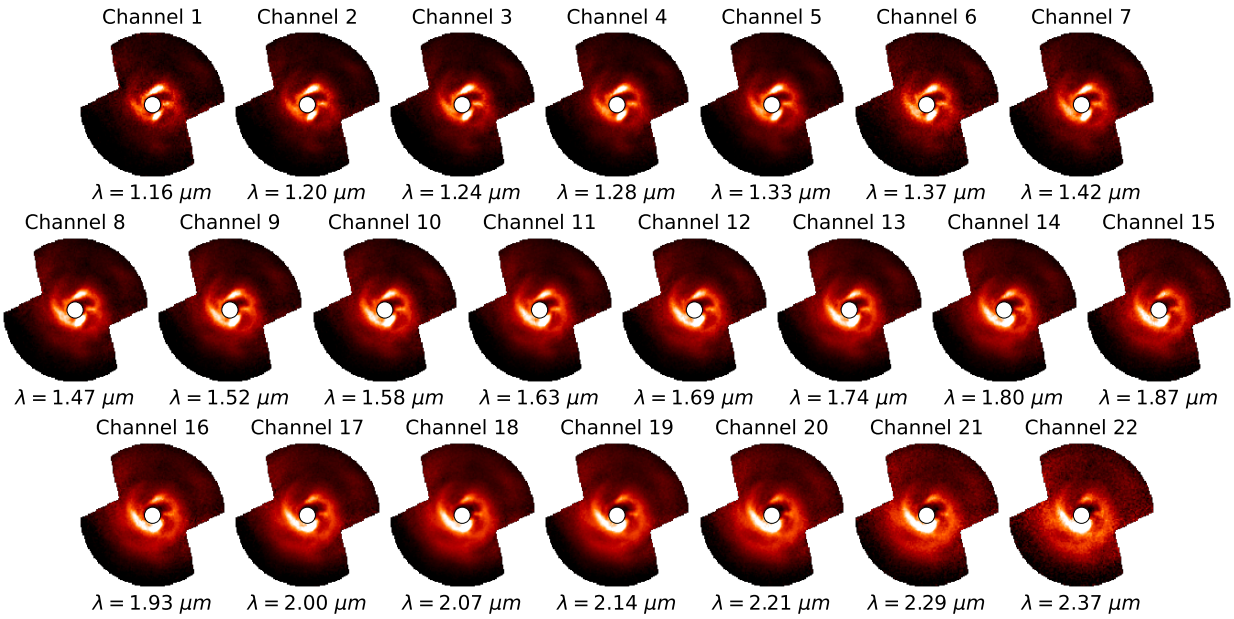

Figure 4.4.2: Polarized intensity in 22 wavelength channels for the final, sequence-combined CHARIS PDI-mode observations of AB Aur. The color stretch is linear throughout, with range chosen arbitrarily for each channel in order to maximize the visibility of disk structures. North is up and east is left, with the FOV extending to $r \sim 1\farcs15$.

4.4.1. AB Aurigae

Following preprocessing, AB Aurigae observations were matched based on HWP position to produce 16 complete cycles. We then applied the double-differencing PDI technique of Section 4.3.4 to produce: a) a sequence-combined polarized intensity (PI) image cube (Figure 4.4.2), b) J, H, and K-band PI images (Figure 4.4.3), and c) a wavelength collapsed broadband PI image (Figure 4.4.4). These results show unambiguous detections of the complex spiral-armed disk of AB Aur – not only in wavelength-collapsed products, but also in all 22 individual wavelength channels.

4.4.2. TW Hydrae

HWP matching for TW Hya results in 11 complete cycles. Following the double-differencing procedure, we mask any pixels with coverage in less than 25% of derotated images to mitigate noisy edges in the products. Figure 4.4.5 shows PI images for TW Hya, including a comparison of the H-band PI products with and without first-order instrumental polarization correction. K-band is not included; the lack of sky frames for this night combined with the weaker AO correction resulting from the target's faintness produce only a marginal detection of disk features in K-band.

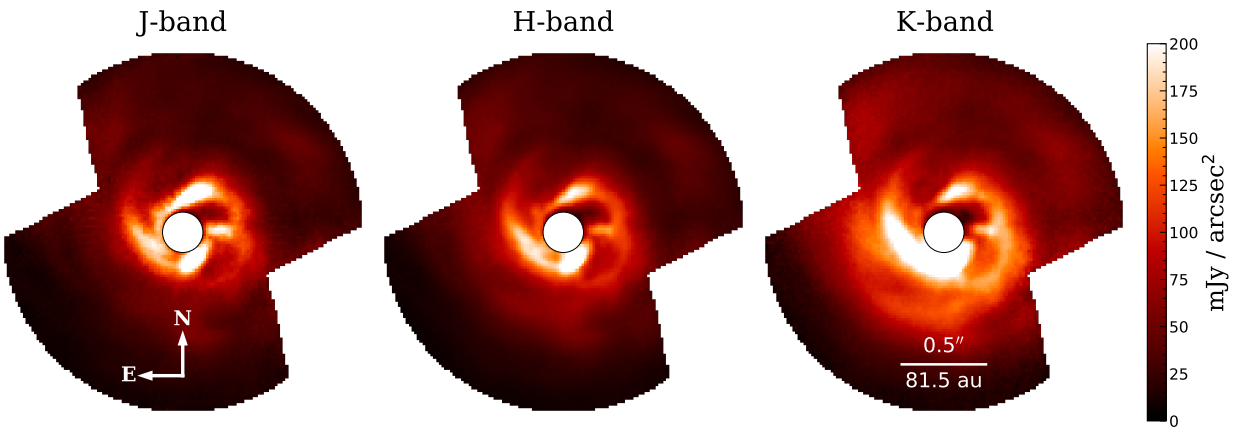

Figure 4.4.3: AB Aur CHARIS PI images created by binning sequence-combined Q and U image cubes along the wavelength axis to approximate NIR J-band (channels $1-5$, $1.16-1.33\ \mu m$), H-band (channels $8-14$, $1.47-1.80\ \mu m$) and K-band (channels $16-21$, $1.93-2.29\ \mu m$) images. These Q and U images are then combined to produce PI using equation 4.6.

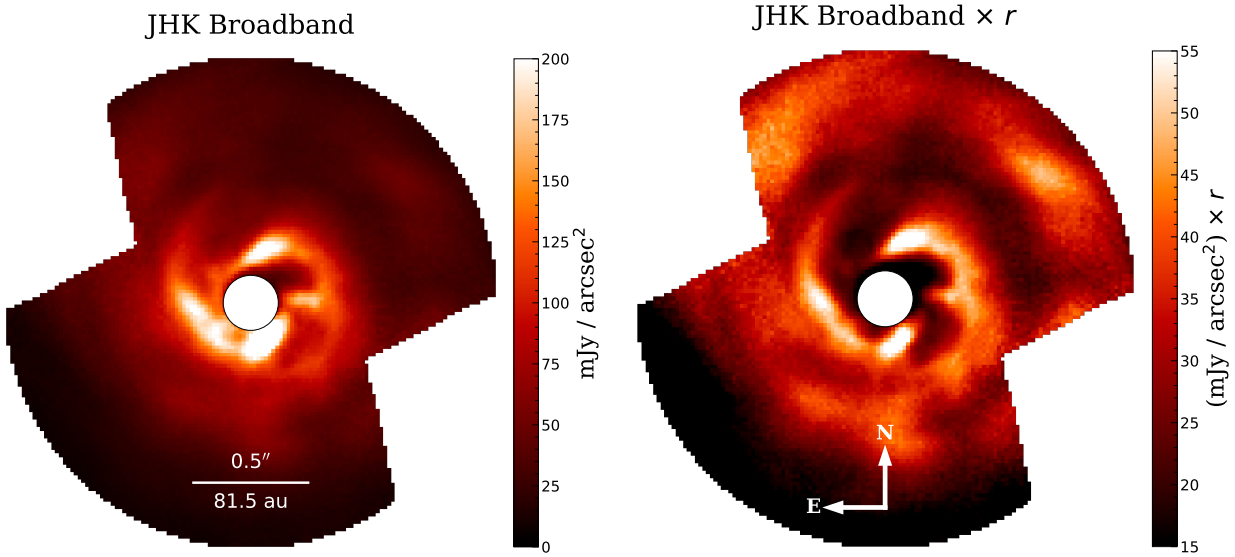

Figure 4.4.4: CHARIS wavelength-collapsed PI image of AB Aur created by combining all wavelength channels of the Q and U image cubes before application of equation 4.6. The image on the left is presented as the images of Figure 4.4.3, while the image on the right has been multiplied by the projected stellocentric separation in arcsec (assuming an inclination of $30°$ and a position angle of $60°$).

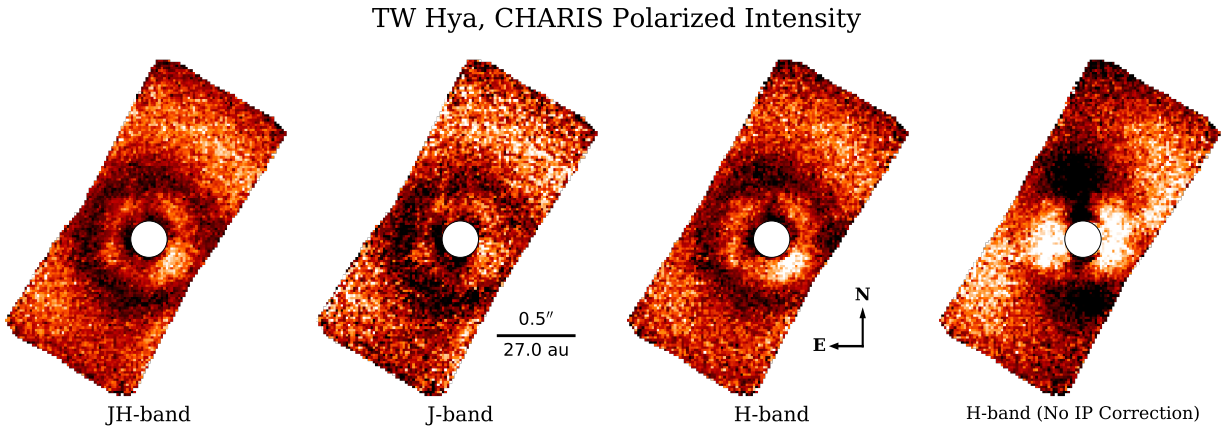

Figure 4.4.5: TW Hya CHARIS PI broadband images created by binning sequence-combined Q and U image cubes along the wavelength axis before application of equation 4.6. The leftmost image, labeled JH-band, is created by combining both J-band (channels $1-5$, $1.16-1.33\ \mu m$) and H-band channels ($8-14$, $1.47-1.80\ \mu m$). The third and fourth images provide a comparison of the result in H-band with and without first order instrumental polarization correction. The correction is necessary for a clear detection in this case.

4.5. CHARIS PDI Disk Forward Modeling

For other varieties of differential imaging (e.g. ADI, RDI, SDI), detailed and computationally intensive forward-modeling is needed to quantify the (typically significant) erroneous attenuation of disk signal that occurs during nulling of the stellar PSF. For PDI, the observed disk image is generally much closer to the true disk, and forward-modeling is relatively simple.

For this purpose, we consider two factors affecting the resulting disk signal:

1. the PSF during our observations
2. the slight difference in orientation between X^+ and X^- used in double differencing (see Equation 4.3)

To simulate these effects for a disk model, we proceed as follows. First, synthetic Q and U images are generated for the model. These images are then rotated to match the array of parallactic angles of the appropriate $X^\pm$ data (i.e. those of the utilized images at HWP angles of $0°$ and $45°$ for Q or $22.^\circ5$ and $67.^\circ5$ for U). This results in two three-dimensional sets of model image sequences. Each frame is then convolved with the empirical PSF model for each wavelength channel (created from either the satellite spots or the unocculted stellar PSF in the data), and frames corresponding to X^- are multiplied by -1. Finally, values falling beyond the CHARIS PDI-mode FOV are set to zero. This produces a sequence of image cubes for $Q^\pm$ and $U^\pm$ comparable to the polarized astrophysical signal contained in the data following single-differencing (Equation 4.1).

Once synthetic data sequences containing only disk signal have been created, the double-differencing procedure is applied (Equation 4.3), using the same HWP cycle matches as for the data. Since the orientations of the model images are offset by the same amount as the real data, any effects that result from this will be propagated to the final product from forward-modeling. Following this, the resulting sequences are derotated and combined in the same manner as the data.

Examples of input models alongside corresponding PSF-convolved models and fully processed models resulting from this procedure are shown in Figure 4.5.1. These results show little difference between the PSF-convolved and fully processed models. This suggests that, even with the relatively high rate of field rotation for the utilized data, combining exposures at slightly different parallactic angles in double differencing has little effect on the products. This is true even in the case of fine azimuthal features – such as the highly inclined disk model – where we should expect the impact to be most significant. Additionally, while we expect no change as a result of the parallactic angle difference in the case of the face-on disk, the change between the PSF-convolved and processed surface brightness (SB) profiles for this model appears comparable with the change for inclined models. Based on this, it seems that the changes we see are predominantly from other sources (e.g. interpolation during image rotation). In general, the impact of the offset in parallactic angle between double-difference pairs will be negligible, but may warrant consideration when observing targets very near zenith.

The relative simplicity of disk forward modeling for PDI data confers a significant strength for disk studies; since the procedure is much less computationally expensive – requiring only image rotation, PSF convolution, and simple algebra – disk models can be forward-modeled orders of magnitude more quickly than for other disk imaging techniques. Combined with swift optimization algorithms, such as differential evolution (DE; see e.g., Lawson et al. 2020), PDI data enables a much more detailed analysis of disk morphology and composition. This is especially so in the case of debris disks, which can be reasonably approximated with simple scattered-light models (rather than more expensive radiative transfer models). Along with the unique ability of CHARIS to conduct high contrast polarimetry simultaneously at an array of wavelengths, the inherent strength of PDI for disk studies is further magnified.

4.6. Summary and Conclusions

The recently-introduced spectropolarimetry mode for the Subaru Telescope's SCExAO/CHARIS provides an exciting new tool for groups studying planet-forming circumstellar disks. This new observing mode enables simultaneous polarized intensity imaging at an array of NIR wavelengths in a $1'' \times 2''$ FOV. We have summarized available instrument configurations and outlined the general procedure used to process CHARIS PDI data. In this procedure, we leverage proven existing software for classical-mode CHARIS data processing alongside purpose-built tools for carrying

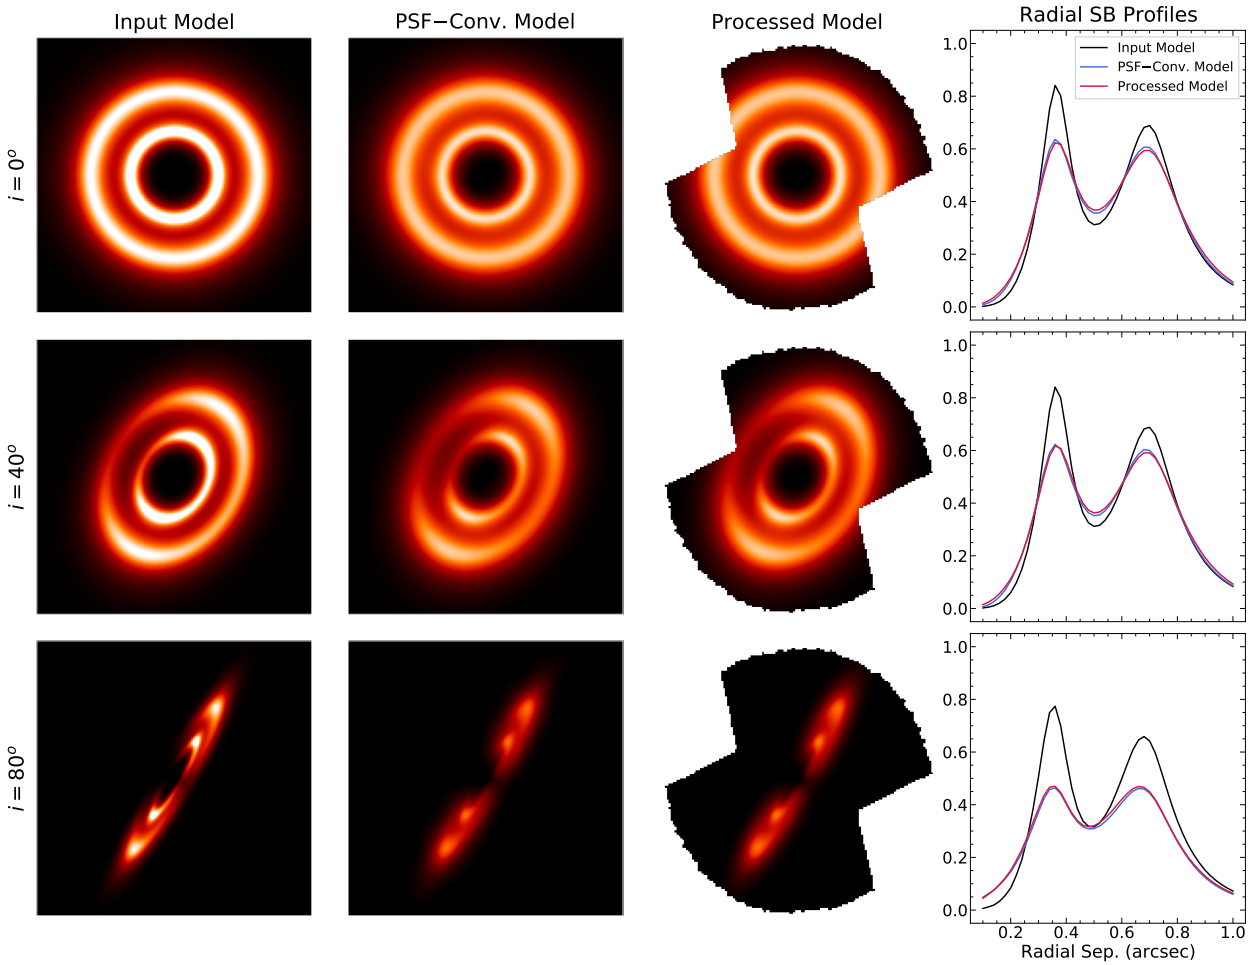

Figure 4.5.1: A set of simple synthetic scattered light disk models in polarized intensity are shown throughout the forward modeling procedure for CHARIS PDI-mode. For this purpose, we use the PSF, parallactic angles, and HWP matches from the observations of AB Aur. Each row corresponds to a distinct input model of differing inclination. The final column displays the radial SB of the images at each stage, measured in apertures of 1 full width at half maximum (FWHM) in diameter and along the major axis of the model.

out HWP matching and double-differencing procedures. With the upcoming implementation of a full instrumental polarization correction (van Holstein et al., 2020), this new observing mode for SCExAO/CHARIS will enable numerous exciting avenues of study relevant to planet-forming disks.

Forward-modeling for disks in PDI data, using the simple procedure outlined in Section 4.5, is extremely time efficient compared to the procedures for other differential imaging techniques. As such, CHARIS PDI-mode will be uniquely well-suited to studying parameters accessible through forward modeling that are expected to have wavelength dependence, such as scattering phase function.

CHARIS PDI-mode will also have utility in helping to understand the sources of spiral arms in disks which exhibit them. From continued observation of such disks, the movement of the features can be measured, which can then be used to discriminate between their two likely causes: yet-unseen companions and gravitational instability (Dong et al., 2018). While total intensity observations might also enable such investigations, it would be much more challenging given that the amount and location of disk attenuation will likely vary between epochs. Being free of significant attenuation, PDI data will significantly strengthen such investigations.

Finally, comparison of polarized intensity and total intensity PDI products (e.g. from application of RDI) enables measurement of fractional polarization (PI/I) across CHARIS's array of wavelengths. These measurements can then be used to identify significant sources of thermal emission relative to other parts of the disk, such as embedded protoplanets, as well as to inform us regarding the properties of the dust within the disk, such as porosity (Hughes et al., 2018).

CHAPTER 5

Constrained Reference Star Differential Imaging: Enabling High Fidelity Imagery of Even Highly Structured Circumstellar Disks[1]

5.1. Introduction

Circumstellar disk systems serve as benchmarks for the study of how and where exoplanets form. With the advent of ground-based extreme adaptive optics (AO) facilities, such as Subaru's SCExAO (Jovanovic et al., 2015), VLT's SPHERE (Beuzit et al., 2019), Gemini's GPI (Macintosh et al., 2015), and Magellan's MagAO-X (Males et al., 2018), and with the recent launch of the James Webb Space Telescope (JWST), high-contrast imaging studies of disks have reached an exciting new era. We can now spatially resolve the morphological signatures within disks (e.g., gaps or spirals) thought to be caused by newly formed/forming companions while also directly identifying & characterizing the young planets and sub-stellar objects that may cause them (e.g., Keppler et al., 2018). Combined with multi-wavelength or integral field spectrograph (IFS) data, we have the technology to conduct incredibly detailed spatial and spectral analysis of these systems.

A key challenge for these studies is the isolation of circumstellar signal (CSS) from the bright pattern of diffracted starlight from the host star (the stellar point spread function or PSF). For disk-focused studies, this is commonly achieved using reference star differential imaging (RDI). "Classical RDI", in which a reference star image is directly subtracted from the target image, can leave significant residual starlight where the PSF changes significantly between exposures (e.g., at narrow separations). More advanced techniques — e.g., Locally Optimized Combination of Images (LOCI; Lafrenière et al., 2007) or Karhunen-Loève Image Projection (KLIP; Soummer et al., 2012) — better model the starlight but also cause some CSS to be lost (or "attenuated"; see Fig. 5.2.1, top row). Since this CSS loss is neither spatially nor spectrally (for multi-wavelength data) uniform (e.g., Figure 2 in Betti et al., 2022), features identified within disks – including planets – can be challenging to validate. Polarimetric differential imaging (PDI) complements conventional total intensity (I) disk studies by producing unattenuated polarized intensity (PI) imagery of disks. However, since emission from young planets is not expected to be significantly polarized, planets cannot generally be detected in PI alone. Though the comparison of I and PI should highlight planets due to their diminished polarization relative to highly polarized disk material, the non-uniform loss of signal in total intensity precludes this measurement.

[1]This chapter is reproduced from Lawson et al. (2022, submitted) with permission.

In some scenarios, the challenges presented by disk signal loss and the confusion of disk and planet signal can be circumvented using forward modeling techniques (e.g., Currie et al., 2015a; Pueyo, 2016; Mazoyer et al., 2020; Lawson et al., 2020). In disk forward modeling, signal loss from PSF-subtraction is induced on synthetic disk images of varying parameters until the observed (attenuated) CSS is reasonably reproduced. This enables robust assessment of disk geometry in the presence of signal loss, while also allowing approximate correction of disk surface brightness measurements for some parts of the disk (e.g., Bhowmik et al., 2019; Betti et al., 2022). However, as this requires generating and processing models spanning the breadth of plausible disk parameters (Pueyo, 2016), it is generally infeasible for structurally complicated disks (e.g., having spiral arms) for which more model parameters must be explored and for which individual models can be more expensive to generate. In the case of simpler disks where exhaustive modeling is feasible, it nevertheless introduces intractable uncertainties for analysis and is often a significant computational bottleneck.

In recent years, an array of advancements for PSF-subtraction — leveraging RDI as well as angular differential imaging (ADI) and spectral differential imaging (SDI) — have been developed for the purpose of exoplanet and disk imaging. These include techniques which focus on more thoroughly removing starlight and/or better recovering point sources (e.g., *TLOCI*, *A-LOCI*, *PACO*; Marois et al., 2014; Currie et al., 2012a; Flasseur et al., 2018), but which still face steep challenges in addressing the attenuation of extended CSS through forward-modeling. More relevant for disk studies, Non-negative Matrix Factorization (NMF) with the Best Factor Finding (BFF) procedure (Ren et al., 2018) and data imputation using sequential nonnegative matrix factorization (DI-sNMF; Ren et al., 2020) are intended to mitigate CSS loss for disk systems. However, they depend on the availability of regions free of CSS to completely eliminate signal loss and are thus less effective for disks which dominate the field of view (FOV). Recent advancements, such as *MAYONNAISE* (Pairet et al., 2021) and *REXPACO* (Flasseur et al., 2021), provide particularly sophisticated tools for isolating circumstellar signal. However, incorporation of such complex techniques into existing data-processing pipelines would likely require significant reworks of those pipelines. Despite the exceptional capabilities of current and upcoming observatories, our ability to study circumstellar disks (and the young planets within them) is currently inhibited by the loss of CSS signal during post-processing and the difficulty of quantifying this loss.

In this work, we describe a new class of RDI PSF-subtraction techniques for circumstellar disk systems: constrained RDI. In constrained RDI, available information regarding a disk (e.g., through polarimetry) is used to prevent CSS signal loss during PSF-subtraction. By tuning constraints to best explain a target's observations, constrained RDI can be optimized to effectively eliminate CSS loss during PSF-subtraction. We generalize these techniques for use with any RDI-based PSF-subtraction technique in which reference images are combined in some manner to minimize

residuals with the target data (e.g., LOCI, KLIP, NMF, etc.; Sections 5.2, 5.3).

Using PI-based constraints, we demonstrate constrained RDI in application to simulated Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS) IFS observations of a spiral armed disk system (Section 5.4) and to on-sky CHARIS observations of the AB Aurigae protoplanetary disk system — whose embedded protoplanet is verified by this approach (Section 5.5; Currie et al. 2022, Nature Astronomy, in press). Additionally, using synthetic disk models as constraints, we apply constrained RDI to simulated JWST NIRCam observations (Section 5.3.2). These techniques are broadly applicable for nearly any disk imaging study, including those using data from: ground-based observatories, Hubble Space Telescope (HST), JWST, and future observatories (e.g., Roman Space Telescope, 30-meter class telescopes, and others; Section 5.6). Moreover, constrained RDI is extremely unlikely to induce spurious (false-positive) detections, and can be implemented in existing pipelines with only minor changes.

5.2. Disk Signal Attenuation in Reference Star Differential Imaging

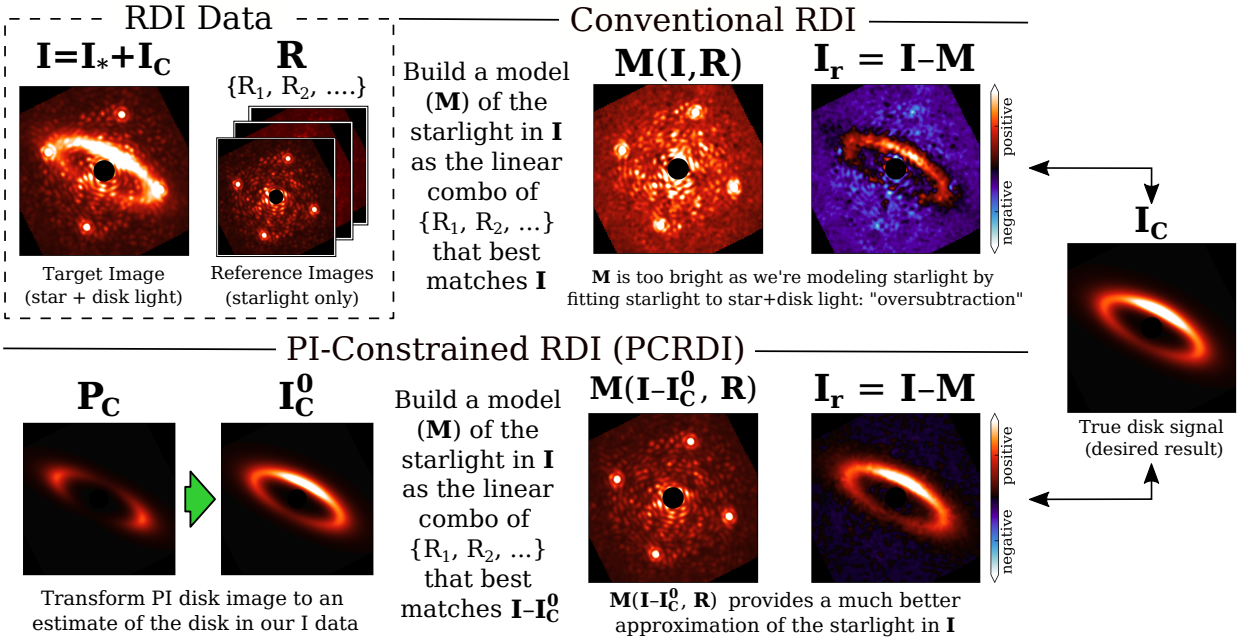

Figure 5.2.1: Top row: Conventional RDI PSF-subtraction for a (synthetic) disk system – yielding significantly attenuated results (right panel). **Bottom row**: PI-constrained RDI (PCRDI; Section 5.3.1). PCRDI eliminates attenuation *and* reduces residual noise by better modeling the diffraction pattern of the starlight. Since the disk estimate (I_C^0) is *only* used in determining the optimal combination of reference frames for the stellar PSF model (i.e., is not directly used in the result), PCRDI avoids the risk of new spurious CSS features when implemented this way.

Consider a (total intensity) target image, I, containing both stellar signal, I_*, and circumstellar signal, I_C, such that $I = I_* + I_C$. The objective of PSF-subtraction is to determine the stellar

signal in the image so that it can be subtracted from the target frame to isolate the circumstellar signal (CSS). In reference star differential imaging (RDI), this problem is addressed by using observations of an additional star for which no CSS is present. To accommodate temporal changes to the PSF between the target and reference frames, it is common to utilize a set of reference images from which an optimal match to each target frame can be determined (e.g., Lafrenière et al., 2007; Soummer et al., 2012; Choquet et al., 2014; Hagan et al., 2018).

Let $\mathcal{R}$ be a sequence of N reference star images containing only stellar signal: $\mathcal{R} = \{R_1, R_2, \ldots, R_N\}$. Let $\mathcal{M}(I, \mathcal{R})$ be the PSF model constructed from the reference data $\mathcal{R}$ to minimize residuals with a target image I. $\mathcal{M}(I, \mathcal{R})$ could be a PSF model constructed from a linear combination of reference frames or with a principal component analysis (PCA) based technique such as KLIP.

For conventional RDI PSF-subtraction techniques (hereafter simply referred to as "RDI"), the residuals, I_r, are determined as:

$$I_r = I - \mathcal{M}(I, \mathcal{R}) \tag{5.1}$$

In an ideal scenario, the PSF model would perfectly reproduce the stellar signal in the target image, $\mathcal{M}(I, \mathcal{R}) = I_*$, and so the residuals would simply be the CSS in the target image: $I_r = I - I_* = I_C$. In practice, the presence of CSS in I directly influences the PSF model that is constructed; rather than constructing the stellar PSF model that minimizes the residuals with the target stellar PSF, we actually construct the model that minimizes the residuals with the stellar *and* circumstellar signal. This results in a PSF model that is brighter than the starlight in the target image, and thus the circumstellar signal in the residuals is artificially reduced in brightness: $I_r < I_C$. This is referred to as "oversubtraction" and is the sole source of systematic attenuation for RDI (Pueyo, 2016). In this notation, the attenuated CSS result, I'_C, for a given I_C is found simply by replacing I with I_C in Eq 5.1:

$$I'_C = I_C - \mathcal{M}(I_C, \mathcal{R}) \tag{5.2}$$

For point-like companions, this effect is generally minor. For extended CSS, such as from a circumstellar disk, the effect can be severe (e.g., Lawson et al., 2021a; Betti et al., 2022).

5.3. Mitigating RDI Oversubtraction: Constrained RDI

If the underlying circumstellar signal in I, I_C, was known *a priori*, oversubtraction could be eliminated by computing the residuals as: $I_r = I - \mathcal{M}(I - I_C, \mathcal{R})$, or in other words, $I_r = I - \mathcal{M}(I_*, \mathcal{R})$. Of course, since I_C is the desired product of RDI, this provides no immediate utility.

However, we can approximate I_C to mitigate oversubtraction in RDI products. Denoting an

estimate of I_C as I_C^0, the residuals are then:

$$I_r = I - \mathcal{M}(I - I_C^0, \mathcal{R}) \tag{5.3}$$

Notably, the estimate of I_C, I_C^0, is used exclusively in determining the optimal reference-based PSF model and is not directly used in the final product. As in standard RDI reductions, the PSF model itself is still constructed entirely from reference images (or the eigenvectors of the reference image covariance matrix in the case of KLIP). In other words, this technique only changes how we determine the coefficients without altering any of the imagery that goes into the final result. As such, this approach provides an extremely safe means by which to improve RDI reductions of circumstellar disk targets; a particularly poor estimate of the circumstellar signal contained in the data could result in a worse PSF-subtraction (e.g., by causing under-subtraction), but is extremely unlikely to produce spurious extended features in the final result. However, some care should be taken when using extremely large PSF libraries (e.g., such as those used for ALICE; Choquet et al., 2014; Hagan et al., 2018). In such a scenario, it may be possible to manifest features of the CSS estimate which are not truly present in the data from noise in a large but featureless reference library. In practice, this can be avoided by limiting the size of the PSF library in some way (e.g., by using KLIP and retaining only part of the KLIP decomposition of the reference library).

5.3.1. Polarized Intensity Constrained RDI

For targets also observed with polarimetric differential imaging (PDI), the unattenuated polarized intensity (PI) of circumstellar signal, P_C, can be attained using standard techniques (e.g., double-differencing; Kuhn et al., 2001). Given that, by definition, $I_C \geq P_C$, RDI results can be improved in such a case by adopting $I_C^0 = P_C$ and proceeding as in Equation 5.3. Generally, using P_C as a conservative estimate of I_C will offer an unambiguous improvement for RDI products – though the improvement may be small for particularly compact disks (where oversubtraction was already minor) or disks with low fractional polarization (where $I_C \gg P_C$).

However, we can do much better by using reasonable assumptions regarding the scattering properties of circumstellar material. As a function of scattering angle, ϕ, the ratio of polarized to total intensity, or fractional polarization (F_{pol}), for scattered light is often described by Rayleigh polarization (e.g., Whitney et al., 2013; Stolker et al., 2016; Gomez Gonzalez et al., 2017):

$$\frac{\mathrm{PI}}{\mathrm{I}} = F_{pol} = \frac{1 - \cos^2\phi}{1 + \cos^2\phi} \tag{5.4}$$

By assuming a particular scattering surface for the disk (e.g., as in *diskmap*; Stolker et al., 2016), a map of the corresponding scattering angles probed throughout the field, Φ, and thus the fractional

polarization, can be derived. Inverting Eq 5.4, this can be used to transform P_C to an estimate of I_C:

$$I_C^0 = P_C \cdot \left(\frac{1 + \cos^2 \Phi}{1 - \cos^2 \Phi} \right) = \frac{P_C}{F_{pol}^0}, \tag{5.5}$$

where F_{pol}^0 denotes the estimate of the fractional polarization for a particular assumed surface geometry. This total intensity CSS estimate, I_C^0, can then be used with Eq 5.3 to carry out constrained RDI PSF subtraction. This process is visualized in the bottom row of Figure 5.2.1. Hereafter, we refer to this as PI-constrained RDI or PCRDI.

Using *diskmap*, P_C is transformed to I_C^0 by assuming a smooth scattering surface with a radial profile defined by a, b, and c as $h(r) = a + br^c$, with a particular peak fractional polarization (s), and which is viewed at a particular orientation (inclination, i, and position angle, PA). For well-studied disks, disk modeling results from literature can be used to estimate these parameters and thus enable PCRDI. Alternatively, their values can be directly optimized (see Section 5.3.3).

5.3.2. Model Constrained RDI

For disks with simple geometries but which lack suitable PI imagery to enable PCRDI, synthetic disk models can be adopted as I_C^0 in place of the PI-based estimates used in PCRDI. We refer to this approach as Model Constrained RDI or MCRDI. If a suitable literature-based model is not available, the model's parameters can be optimized directly, much as for PCRDI (Sec. 5.3.3). We note that this may be significantly more time consuming than for PCRDI; not only does generating the CSS estimate for a particular set of parameters take longer (i.e., constructing an entire disk model for MCRDI versus doing some geometry for PCRDI), but there are also more parameters to consider. Moreover, these parameters can have significant degeneracies and local minima that may motivate a more thorough exploration of the parameter space if decent estimates of the parameters for the disk are not known *a priori*.

In general, the timescales required to reach an optimal MCRDI solution will often be comparable to those required to optimize a disk model using conventional forward-modeling techniques (e.g., Lawson et al., 2021a). Nevertheless, this technique provides significant utility over forward-modeling. While an attenuated model in good agreement with the data can provide attenuation corrections (e.g., Bhowmik et al., 2019), such corrections are typically only valid where the attenuated data remains positive, and are generally suitable only for analysis of surface brightness. Using the same model to carry out MCRDI instead provides the same benefits but enables brightness measurements over the full FOV and can also reveal fine morphological features and faint extended signal that would otherwise not be recovered.

5.3.3. Optimization of Constrained RDI

For a constrained RDI technique, the values of the parameters governing the CSS estimate (e.g., $s, a, b, c, i,$ and PA for PCRDI) can be directly optimized to identify the best match to the true underlying CSS. For this purpose, we compute the following:

$$y = (I - I_C^0) - \mathcal{M}(I - I_C^0, \mathcal{R}) \tag{5.6}$$

This is simply the conventional RDI residual calculation (Eq. 5.1) where I is replaced by $I - I_C^0$ throughout. As our models of the stellar and circumstellar signal approach the true stellar and circumstellar signal, the residual signal contained in the image y will generally decrease:

$$\begin{aligned} &\text{As } \mathcal{M}(I - I_C^0, \mathcal{R}) \to I_* \text{ and } I_C^0 \to I_C, \\ &\qquad y \to (I - I_C) - I_* = I_* - I_* = \hat{0} \end{aligned} \tag{5.7}$$

Since our stellar model will also generally improve as our estimate of the circumstellar signal improves, y can be used to assess the quality of our CSS estimate, I_C^0. A straight-forward objective function for optimization is then simply the sum of the squares of these residuals binned to the spatial resolution of the data[2]. This objective function can then be combined with an optimization algorithm of choice to automatically determine the optimal estimate of CSS.

5.4. Throughput Assessment using Simulated Data

5.4.1. PCRDI Throughput

To test the throughput of PCRDI in application to CHARIS IFS data (Groff et al., 2016), we generated multi-wavelength synthetic disk image cubes in Stokes Q, U, and I using HO-CHUNK 3-D (Whitney et al., 2013) for a spiral-armed disk of similar morphology to the disk of AB Aurigae (see Section 5.5). To simulate an RDI sequence for this model disk, we divided a large CHARIS reference star sequence into two sets – one of which will have the CSS added (the sequence of "target" images: $\mathcal{I} = \{I_1, I_2, \ldots\}$), and the other which will be left unchanged (the "reference" sequence, $\mathcal{R}$)[3]. We then scaled the wavelength slices of the target sequence to match the brightness of AB Aur. The disk model cubes – Q, U, and I – were convolved with the PSF for the target sequence, which was determined using the average shape of the calibration satellite spots over the

[2] In some contexts, it may be preferable to compute a χ^2 metric using pixel-wise uncertainties – as is common for disk modeling purposes (e.g., Thalmann et al., 2013; Currie et al., 2019). However, since the most common PSF-subtraction algorithms are solving some version of an unweighted least-squares problem, the unweighted objective function tends to converge more smoothly.

[3] For more information regarding the size of the target and reference sequences, see Table 5.C.1.

full sequence[4]. After convolution, Q and U were combined to form the disk PI, and the slices of the PI and I cubes (P_C and I_C, respectively) were scaled to approximately match the brightness of AB Aurigae's disk. Finally, the I_C cube was rotated to the parallactic angle of each exposure in the target sequence, $\mathcal{I}$, and added to the target cubes. We note here that the details of the disk-to-star contrast and how the initial reference star sequence is divided to form the simulated target and reference sequences will affect the relative residual noise level and thus the signal-to-noise per resolution element (SNRE) in the products, but will not affect the typical percentage of CSS attenuated by oversubtraction. Similarly, while the quality of the spectral match between the target and the reference star in this simulated sequence is unrealistic (given that they are the same star), a) this should again affect only the residual noise rather than the attenuation, and b) the effect would generally be minor in application to IFS data anyway (given that the PSF model is optimized separately for each narrow wavelength channel).

For PSF-subtraction, the PSF model for each target frame is constructed using a linear combination of the reference frames. To accelerate the PSF subtraction procedure, images are compared within a single annular optimization region spanning $r = 10 - 28$ pixels ($\sim 0\farcs16 - 0\farcs45$) for all 22 wavelength channels. This optimization region is also used for calculation of the goodness of fit for each trial result (Eq. 5.6). Allowing all six scattering surface parameters to vary in wide ranges, we conducted optimization using the Levenberg–Marquardt (or "damped least-squares") algorithm (as implemented in the Python package *LMFIT*; Newville et al., 2014).

The results of PCRDI optimization are shown in Figure 5.4.1 alongside those of the standard (unconstrained) RDI procedure. While the best-fit parameter set produces a fractional polarization model that is much less detailed than that of the (PSF-convolved) input model, it nevertheless enables an extremely high fidelity fractional polarization map to be computed using the disk PI and the PCRDI total intensity result. We note that small differences between the input disk images and the PCRDI total intensity images remain. However, binning the images to resolution and computing the percent difference per resolution element shows differences generally within $\sim 1\%$ of zero (see Fig. 5.4.2)[5].

A procedure for PCRDI forward modeling is provided in Appendix 5.A. However, for any typical application, analysis can be performed on PCRDI products without the need for extensive forward modeling.

[4]PSF-convolution of the model cubes before rotation to the parallactic angles needed for the synthetic data sequence assumes a predominantly azimuthally symmetric PSF. This is a reasonable assumption for the CHARIS PSF, but is less reasonable for systems such as JWST NIRCam.

[5]Note: to provide attenuation estimates which are not significantly dependent on the quality/stability of the AO correction throughout the observations, this assessment and the distributions plotted in Fig. 5.4.2 exclude any residual starlight – such that oversubtraction is the sole source of any change.

5.4.2. MCRDI Throughput

To demonstrate the MCRDI technique, we carried out simplified simulations of JWST NIRCam F335M observations of the debris disk system HD 10647[6]. The NIRCam PSF was simulated using WebbPSF[7] (Perrin et al., 2014) while the debris disk was simulated using DiskDyn (Gaspar & Rieke, 2020). The NIRCam PSF simulations include Gaussian telescope jitter ($\sigma = 7$ mas) and pointing error ($\sigma = 4$ mas per axis), and include two roll angles for the target ($\Delta\mathrm{PA} = 10°$) and a 5-point dither of the reference star, but neglect other factors such as noise and thermal effects. These simulations should provide a reasonable comparison of the relative performance of the PSF subtraction techniques, but are not representative of the absolute quality of JWST/NIRCam products.

To generate the disk models used as constraints in MCRDI, we utilized the GRaTeR debris disk modeling code (Augereau et al., 1999), parameterized as described in Lawson et al. (2021a) (with the addition of a parameter for the disk's brightenss). Using a different tool than the one used to generate the "true" input disk model is intended to better emulate real observations, where a model

[6]Based on upcoming JWST Cycle 1 Guaranteed Time Observations (GTO) observations of this system (Program 1183).

[7]https://github.com/spacetelescope/webbpsf

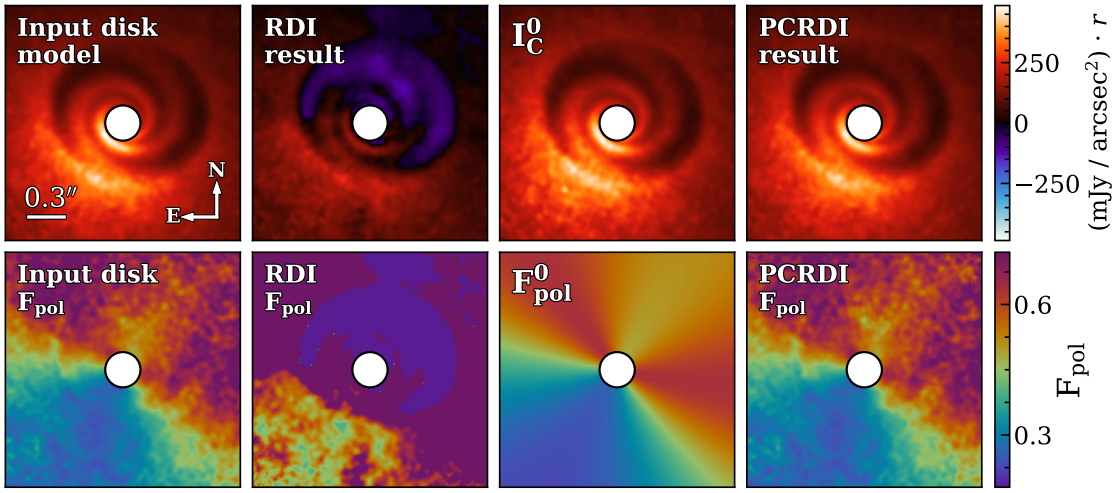

Figure 5.4.1: Wavelength-averaged results for the simulated data described in Section 5.4. **Top row**, from left to right: the PSF-convolved input disk model image (the desired result), the heavily-attenuated final image using RDI PSF-subtraction, the optimized PI-based estimate of the disk signal used by PCRDI, and the final image using PCRDI PSF-subtraction (Section 5.3.1). While the disk signal estimate used by PCRDI is imperfect, PCRDI is still able to recover an extremely high-fidelity image of the disk. **Bottom row**: the corresponding fractional polarization map for each result in the top row. The RDI result is non-physical over most of the field of view ($F_{pol} > 1$ or $F_{pol} < 0$). Though the "smooth surface" F_{pol} estimate (center-right) used to derive I_C^0 is much less detailed than the F_{pol} for the input disk model (left), the PCRDI result (right) still recovers the very fine polarization gradients seen in the input model.

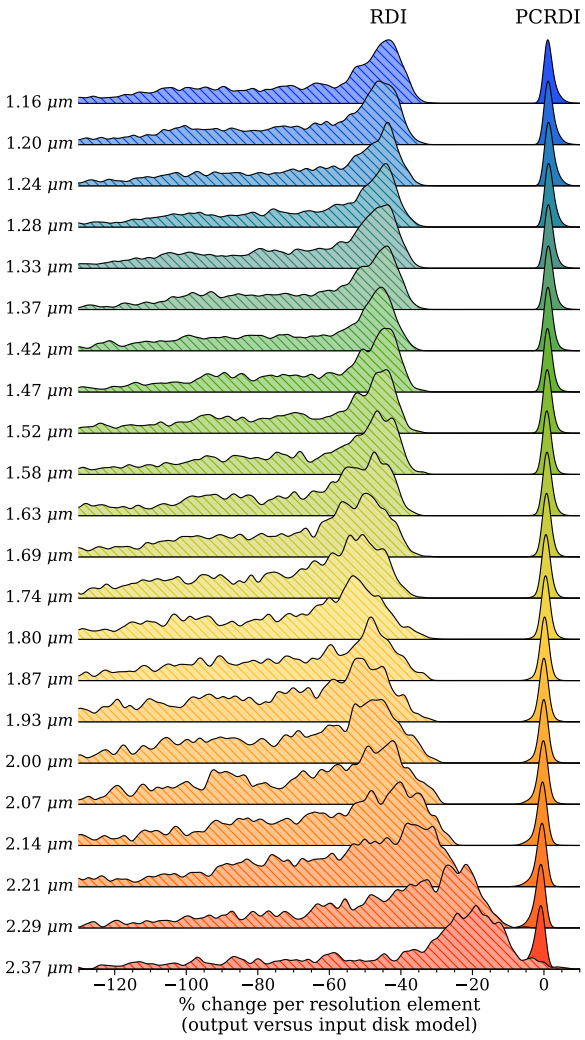

Figure 5.4.2: For the simulated data of Section 5.4: distributions (as kernel density estimates) of percent change in disk signal (relative to the 'ground truth' input disk model) due to attenuation for each CHARIS wavelength channel. For this purpose, all results are binned to the resolution of that wavelength channel prior to the calculation. Hatched and solid regions correspond to results for the standard and PI-constrained RDI reductions respectively. RDI results show significant and highly variable attenuation (within a single wavelength channel and between wavelength channels). Meanwhile, PCRDI results are uniformly within roughly a percentage point of the true values.

is unlikely to perfectly describe the data. For optimization of the disk signal estimate, we again used the Levenberg–Marquardt algorithm. Given the differences in parameterization between the two tools (and the fact that DiskDyn is a dynamical code, such that the values of some paramaters for the final evolved model differ from their initializations), the optimal GRaTeR parameter values are not known *a priori* — with the exception of inclination and position angle (PA). For the MCRDI optimization, we initialized inclination and PA with values offset from the true DiskDyn values by the 1σ uncertainties for these parameters reported in Lovell et al. (2021) (1° in both cases).

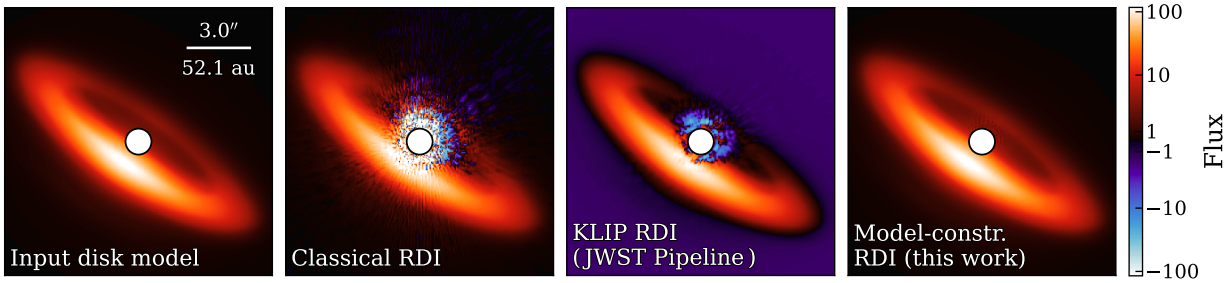

Figure 5.4.3: Simulated results for JWST NIRCam observations of the debris disk system HD 10647. The leftmost panel shows the input disk model (the ideal result), with the subsequent panels showing the result using the indicated subtraction technique. Model constrained RDI (MCRDI) meets or exceeds the fidelity of classical RDI at large separations, while performing substantially better than both classical and KLIP subtraction at small separations (see also: Fig. 5.4.4).

The remaining parameters are initialized using rough 'by-eye' estimates of their values from an unconstrained (i.e., conventional RDI) reduction of the data. The proximity of the initial values to the optimal values may affect the number of function evaluations required to reach a solution but will not affect the quality of the MCRDI result (unless the parameter space has multiple optima, in which case a global optimization algorithm would be warranted instead). For general use, values from disk modeling performed in prior studies could also be adopted.

A comparison of the results using classical, KLIP[8], and model constrained RDI are shown in Figures 5.4.3 and 5.4.4. These results show that, at moderate to large separations, classical and constrained reference subtraction both reasonably reproduce the "ground truth" brightness profile – with MCRDI performing somewhat better and KLIP performing significantly worse. At small separations, MCRDI performs substantially better than the other techniques.

5.5. Application of PCRDI to IFS Observations of AB Aurigae

AB Aurigae is a pre-main sequence star (1–3 Myr, Kenyon et al. 2008; $d = 156$ pc, Gaia Collaboration et al. 2016, 2021) hosting a highly structured protoplanetary disk that presents with spiral arms at both large scales (hundreds of au; Grady et al., 1999; Fukagawa et al., 2004) and small scales scales (tens of au; Hashimoto et al., 2011; Boccaletti et al., 2020) in optical and near-infrared (NIR) imagery, and with a large, depleted inner cavity in sub-millimeter (van der Marel et al., 2021).

As a further demonstration of the PCRDI technique, we utilize SCExAO/CHARIS observations of AB Aur in total intensity ($t_{int} \approx 116$ minutes) and polarized intensity ($t_{int} \approx 74$ minutes) – originally reported in Currie et al. (2022, in press) and Lawson et al. (2021b) respectively. These

[8]As implemented in the official JWST pipeline and retaining $N = 10$ KLIP modes; the quality of the result does not improve noticeably for other choices ($N = 1 - 15$).

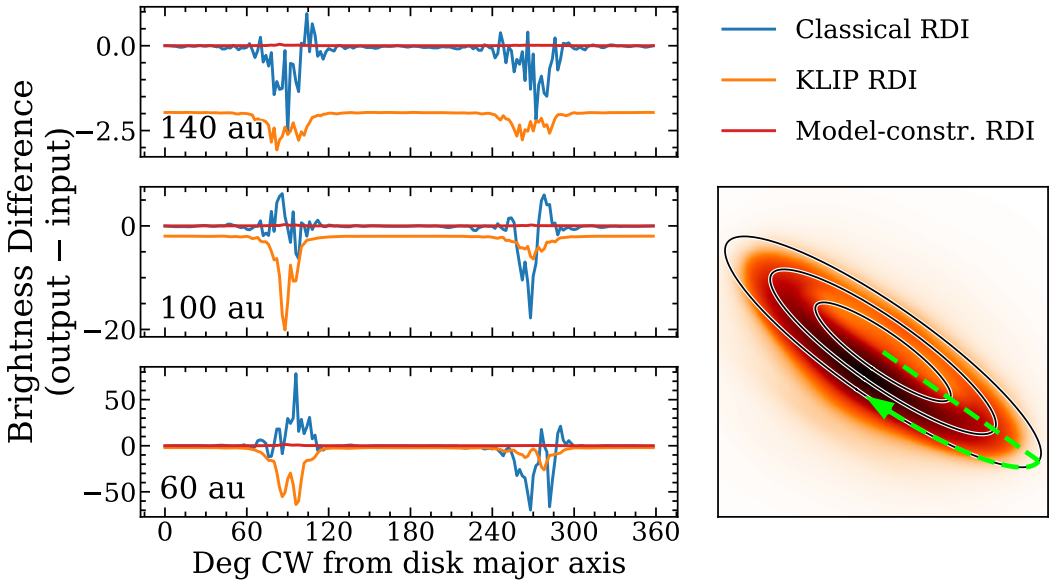

Figure 5.4.4: For each reduction of Fig. 5.4.3: the difference between the output and input brightness of the disk measured in apertures with diameter equal to the PSF full-width at half maximum. The lower-right panel provides a visualization of the three radial separations measured (black ellipses), and also shows the origin and direction of the measurements (green arrow). At $r = 140$ au and $r = 100$ au, classical RDI and model constrained RDI (MCRDI) both reasonably reproduce the majority of the "ground truth" brightness profile. MCRDI provides substantially improved measurements throughout the inner region ($r = 60$ au), corresponding to the inner asteroid belt for this system, and the most likely location of any yet-unseen companions.

data were taken over two consecutive nights in October 2020, with both polarized and total intensity sequences using CHARIS in low-res broadband mode ($R \sim 19$, $1.15 - 2.39\ \mu m$ – producing 22 wavelength channels per exposure) and utilizing a 113 mas Lyot coronagraph. Currie et al. (2022, in press) also observed a reference star (HD 31233) before and after the AB Aur total intensity sequence ($t_{int} \approx 21$ minutes), which we also use here to enable RDI. See Currie et al. (2022, in press) for a description of preprocessing (sky subtraction, image registration, etc.) used for the total intensity data, and Lawson et al. (2021b) for a description of the PDI data reduction for the PI data.

As with the synthetic data, RDI PSF-subtraction is performed by constructing the PSF model for each target frame using a linear combination of the reference frames and comparing the target and reference data within an annular optimization region spanning $r = 10 - 25$ pixels ($\sim 0\farcs16 - 0\farcs41$). For this data, the relatively narrow optimization region is necessary to accommodate the following:

1. CHARIS PDI data has a rectangular $\sim 1'' \times 2''$ FOV (compared to the square $2'' \times 2''$ FOV for the classical mode) and is collected in pupil-tracking mode. The final sequence-combined PI image lacks coverage of two wedge-shaped regions extending from $r \sim 0\farcs5$ (see Lawson et al., 2021b). Limiting the optimization region to separations with full azimuthal coverage substantially simplifies and accelerates PSF model optimization.

2. Hashimoto et al. (2011) report a misalignment between AB Aur's inner disk ($\sim 0\farcs21 - 0\farcs43$) and outer disk ($\sim 0\farcs63 - 0\farcs84$), which would inhibit our smooth-scattering surface estimate for an optimization region extending to the outer disk.

3. AB Aur shows a substantial total intensity enhancement at $r \sim 0\farcs6$ (nearly directly south of the star) which is not reproduced by scaling the PI image with reasonable fractional polarization models. This feature is analyzed further in Currie et al. (2022, in press), where it is determined to be a likely embedded protoplanet. Avoiding this region by optimizing at smaller separations enables a higher fidelity estimate of the total intensity.

A visualization of the polarized intensity, the final optimized fractional polarization estimate (F^0_{pol}), the corresponding PI-based CSS estimate (I^0_C), and the final PCRDI result (I_r) is provided in Figure 5.5.1. Figure 5.5.2 compares the PCRDI results with the conventional RDI results in an array of wavelength ranges – where, as for the simulated data (Section 5.4), PCRDI provides an improvement in CSS throughput that is immediately apparent. Notably, in the PCRDI imagery it is evident that the morphological feature $\sim 0\farcs6$ south of the image center (AB Aur b; Currie et al. 2022, in press) has a distinct (bluer) spectral energy distribution relative to the nearby disk – appearing at much higher contrast to the disk in J-band than in K-band. Similar conclusions are more difficult to draw from the conventional RDI imagery as a result of the substantial spatially- and spectrally- variable oversubtraction that occurred during PSF subtraction.

To further validate this result (or PCRDI results in general), an additional test using forward-modeling for the conventional RDI reduction is summarized in Appendix 5.B.

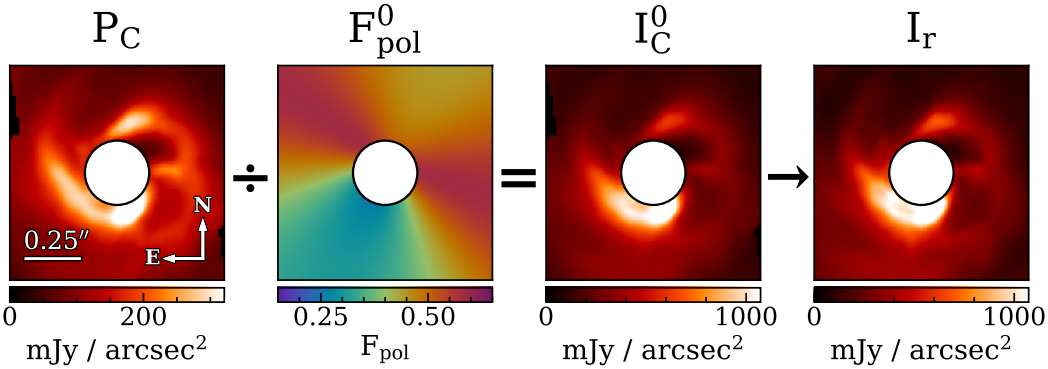

Figure 5.5.1: The components of the PCRDI procedure (Section 5.3.1) as utilized for AB Aur. Left: the polarized intensity disk image (P_C); center-left: the best-fit fractional polarization model by which P_C is divided to estimate the total intensity of the disk; center-right: the total intensity estimate of the disk (I^0_C); right: and the final sequence-combined residual image from PCRDI.

5.6. Broader Applications

Both variations of constrained RDI presented here – PCRDI and MCRDI – can be applied to substantially improve the products of many disk studies. Moreover, the constrained RDI strategy can be combined with many existing algorithms for construction of the stellar PSF model. As such, this strategy can be implemented in existing pipelines with very little alteration necessary; the only change needed in most cases is the ability to model the stellar PSF using a different (CSS-estimate-subtracted) target sequence than the one from which the PSF model is ultimately subtracted. For reference, the run-times for the reductions carried out in this work are provided in Appendix 5.C.

Though we have demonstrated the use of PCRDI with near contemporaneous PI data from the same instrument, the same principles can be applied using non-contemporaneous PI data from an entirely different instrument. For example, when using the H-band VLT/SPHERE PI imagery of AB Aur reported in Boccaletti et al. (2020) to carry out PCRDI on the CHARIS RDI data, the result is nearly identical. Considering the prolific extent of prior extreme adaptive optics PI disk surveys (e.g., with GPI and SPHERE; Esposito et al., 2020; Garufi et al., 2022), this feature provides PCRDI with broad utility for nearly any group conducting high-contrast imaging studies of disks. This provides a compelling application for data from instruments without polarimetric imagers as well – including upcoming JWST data – for which ground-based PI imagery can be used to improve PSF-subtraction.

Similarly, MCRDI can be applied for any system whose disk can be *superficially* reproduced

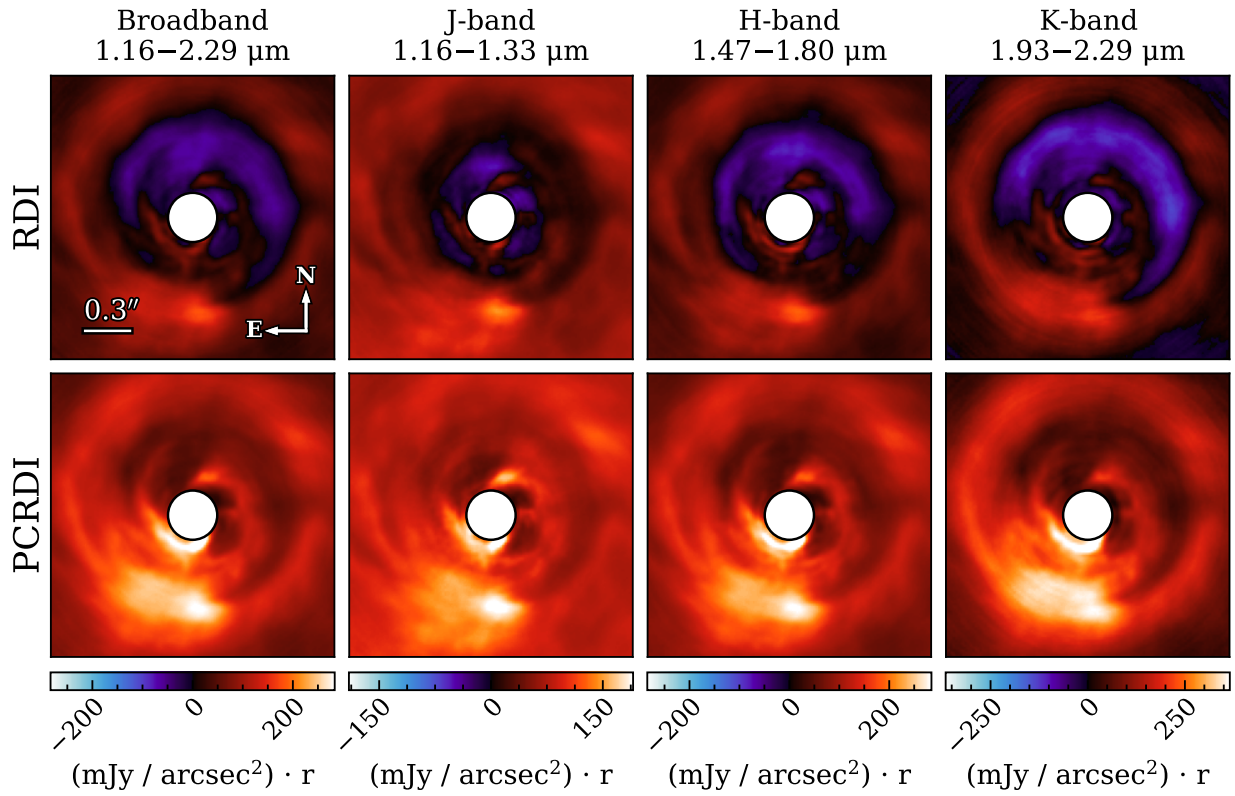

Figure 5.5.2: A comparison of PSF-subtracted results for SCExAO/CHARIS observations of AB Aur (see Section 5.5) using a conventional RDI procedure (top row) and PI-constrained RDI (PCRDI; bottom row). The PCRDI results (see Sections 5.3.1, 5.5) use the same PSF-subtraction algorithm and algorithm settings as the conventional RDI result, but adopt an optimized PI-based constraint for the circumstellar signal to mitigate oversubtraction. All images are multiplied by the stellocentric separation in units of arcseconds and assuming an orientation and geometry appropriate for the inner disk. The columns correspond to different binnings over the wavelength axis of the final CHARIS IFS image cube; from left to right: JHK broadband ($1.16 - 2.29$ μm), J-band (channels $1 - 5$, $1.16 - 1.33$ μm), H-band (channels $8 - 14$, $1.47 - 1.80$ μm) and K-band (channels $16 - 21$, $1.93 - 2.29$ μm).

with a synthetic model. In other words, the utilized model need not be fully physically motivated. For example, a simple scattered light disk model (physically appropriate for a debris disk) could be used to suppress oversubtraction for observations of a transitional disk (whose appropriate physical model would normally require a more involved radiative transfer framework). Additionally, for RDI data for which a full forward modeling procedure has been performed, MCRDI can be applied retroactively using the model solution identified from forward modeling to improve the RDI results.

5.7. Conclusions

Herein, we have presented "constrained RDI", a new class of RDI PSF subtraction techniques well-suited for circumstellar disk targets. Using variations with both PI-based (PCRDI) and model-

based (MCRDI) constraints, constrained RDI can effectively eliminate oversubtraction for RDI products. For the simple disk systems that can be feasibly modeled: PCRDI provides analysis-ready imagery orders of magnitude more quickly than the approximate corrections provided by forward-modeling, while MCRDI provides final products of higher quality and which are more conducive to detailed analysis of disk features and properties. For significantly extended and highly structured disks, PCRDI uniquely provides total intensity products which are unaffected by the significant and variable CSS loss that normally inhibits studies of these systems. In turn, these products enable the detailed spectral and spatial analysis needed to robustly identify embedded planets and to conduct detailed studies of disk composition and morphology.

Proliferation of these techniques, or others that yield comparable results, are paramount for realizing the capability of current high-contrast imaging systems. With the recent launch of JWST, and with many high-impact direct imaging missions on the horizon – including the Roman Space Telescope, and the observatories recommended by the 2020 Astronomy and Astrophysics Decadal Survey (the 6-meter LUVOIR/HabEx "hybrid", and 30-meter ground-based telescopes), it is more important still to ensure that available post-processing tools are not the limiting factor in the yield of exoplanet and disk studies.

Software supporting the application of these techniques is planned for release in late 2022 (see Appendix 5.C).

Appendix

5.A. Forward Modeling for PCRDI Products

For general use, formal forward modeling is not necessary for PCRDI products. For example, to assess geometry, PCRDI typically provides sufficient throughput that PCRDI products can simply be compared to PSF-convolved models.

However, if a measurement requires particularly high precision, conventional forward modeling and forward-model-based flux corrections can still be used. Let I_C^0 be the PI-based estimate used for the real data. Let the PSF-convolved synthetic disk model to be forward modeled be I_M. The corresponding forward modeled result, I'_M, can be computed as follows:

$$I'_M = I_M - \mathcal{M}(I_M - I_C^0, \mathcal{R}) \tag{5.8}$$

In other words, if I_M is truly the underlying CSS in the data, then $I_M - I_C^0$ is the amount of CSS remaining in the data when we construct the stellar PSF model – and thus determines the over (or under) subtraction that will occur.

Compared with conventional RDI reductions using forward modeling strategies, PCRDI forward modeling maintains a number of distinct advantages. Since attenuation corrections for brightness can only be valid where the final image remains non-negative, mitigating the majority of oversubtraction with PCRDI will enable valid corrections over much more of the FOV. Additionally, being less affected by biasing from CSS attenuation, the higher fidelity PCRDI products can be used to make much better estimates of disk parameters, and so will likely require exploration of many fewer models overall.

5.B. Validation of Constrained RDI Using Forward Modeling

In forward modeling, the effects of signal attenuation are induced on an input CSS image. For disk studies, it is common to forward model PSF-convolved synthetic disk models to enable assessment of the geometry and properties of a real attenuated result; if a disk model can be identified which, when attenuation is induced, closely matches the real attenuated data, then that model provides a reasonable explanation of the true unattenuated disk signal. For a standard RDI reduction, the attenuated result for a given input image can be found by evaluating Eq 5.2 with the input image in place of I_C.

To further validate a final PCRDI result in a manner more familiar to the disk-imaging community, we can simply carry out this forward modeling procedure with the PCRDI result as the input

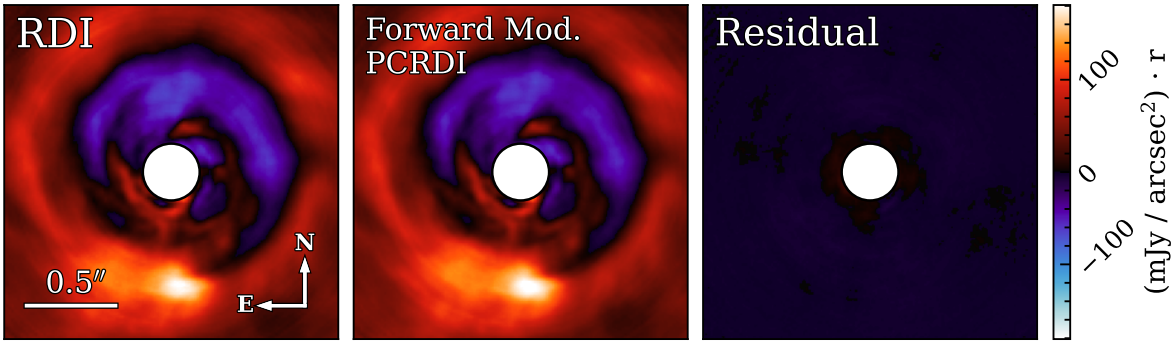

Figure 5.B.1: Left: the conventional RDI result for AB Aur (as in Figure 5.5.2, top left). **Center:** the PCRDI result of Section 5.5 after inducing standard RDI attenuation via forward modeling. **Right:** The difference between the RDI result and the forward-modeled PCRDI result, demonstrating that the PCRDI product is an accurate representation of the unattenuated CSS. All images have been wavelength-averaged and multiplied by the stellocentric separation in units of arcseconds (as in Figure 5.5.2).

image and then compare the result with the standard RDI product. If the attenuated RDI result is consistent with the PCRDI result when RDI attenuation is induced, then the PCRDI image provides a reasonable "model" for the unattenuated CSS.

To demonstrate, we carried this out using the PCRDI result for AB Aurigae (Section 5.5). Figure 5.B.1 shows the result of this procedure. The negligible residual signal in this case provides additional evidence that our PCRDI optimization converged to a CSS estimate sufficient for a very high throughput final image.

5.C. Supporting Software and Run-time

The software written to carry out optimization of PCRDI and MCRDI makes extensive use of vectorization, multi-threading, and GPU acceleration to carry out these procedures relatively quickly.

Public release of Python software to support generalized application of the techniques described here is planned for late 2022. Until then, groups interested in using this technique (or who are interested in other uses of accelerated data reduction for high-contrast imaging) are encouraged to contact the authors.

5.C.1. Run-time

To carry out optimization of PCRDI and MCRDI with the aforementioned software, a personal computer with the following specifications was used:

Memory: 31 GiB

Processor: Intel Core i9-10900X CPU @ 3.70GHz × 20

Table 5.C.1. PCRDI and MCRDI Dataset Dimensionality and Optimization Run-times

Dataset	$N_{\rm tar}$	$N_{\rm ref}$	Exposure dimensions	Opt. $N_{\rm fev}$	Opt. time (s)
PCRDI (sim. data)	107	36	[22 λ $\times$ 201 pix $\times$ 201 pix]	92	125
PCRDI (AB Aur)	168	51	[22 λ $\times$ 201 pix $\times$ 201 pix]	48	88
MCRDI (sim. data)	30	15	[320 pix $\times$ 320 pix]	90	29

Note. — For each dataset utilized in this work, the dimensions and total time required to reach the final constrained RDI product. $N_{\rm tar}$ and $N_{\rm ref}$ give the number of target and reference exposures, respectively – with the size of each image or image-cube given by "exposure dimensions". "Opt. $N_{\rm fev}$" and "Opt. time" indicate the number of function evaluations and total run-time (in seconds) required to reach the presented solution, respectively.

GPU: NVIDIA GeForce RTX 3070

The dimensions and run-times for each dataset utilized are provided in Table 5.C.1. Note: the number of function evaluations for optimization is generally detached from disk complexity for PCRDI, but is strongly linked for MCRDI (e.g., with a multi-ring system requiring optimization of more parameters).

CHAPTER 6

Conclusions

6.1. Chapter Summaries

Chapter 2

In Chapter 2, we presented observations of the edge-on debris disk system HD 15115 from the Subaru Telescope's extreme adaptive optics (AO) system, SCExAO, paired with its integral field spectrograph (IFS), CHARIS, which provides high-contrast imagery at 22 near-infrared (NIR) wavelengths simultaneously. With multi-wavelength imagery of such systems, we can constrain the properties of their dust and more confidently assess possible signposts of yet-unseen companions. However, this requires detailed measurement of disk signal loss during data processing. For simple disks, disk forward-modeling provides a way to approximate signal loss. However, forward-modeling for IFS data is historically very time consuming: $\sim$ one model per 10 minutes per processor core. Common techniques, like Markov chain Monte Carlo, require evaluation of $\sim 10^6$ models, or $\sim$ 19 years on a single core. As a result, the majority of prior ground-based disk studies searched only a coarse grid of disk model parameter values – likely leading to sub-optimal solutions. In application to these observations of HD 15115, we introduced the use of the differential evolution (DE) optimization algorithm for disk modeling; this technique provides robust, global solutions in only $\sim 10^3$ models ($\sim$ 1 week on a single core). Using DE, alongside a physically-motivated "Hong" scattering phase function (SPF), we were able to cast doubt on recent suggestions of a planet-driven misaligned inner ring for HD 15115 while also providing much improved constraints for the size distribution of dust in the disk, and showing that the disk manifests with a definitively red color throughout the CHARIS field of view (as opposed to the blue color at larger separations).

Chapter 3

In Chapter 3, we presented the first multi-band observations of the debris disk system HD 36546 – again using Subaru's SCExAO/CHARIS IFS. Though highly inclined like HD 15115, HD 36546 is distinct in its particularly young age. After application of a similar forward-modeling procedure as for HD 15115, we conducted detailed analysis of the morphology, surface brightness, and color of the disk. These results show that HD 36546 has a peculiar blue NIR color and a possible "back-swept wing profile", both of which may be indicative of planetary dynamics. Additionally, we provide constraints on the presence of a possible companion responsible for inducing a marginal proper motion anomaly on the parent star. Combined with our HD 15115 analysis, these results

demonstrate the value of swift optimization procedures such as DE rather than the grid searches employed by most similar studies, and lend credence to the possibility of a near-universal Hong-like SPF for circumstellar dust (as suggested by Hughes et al. 2018).

Chapter 4

Polarized intensity (PI) observations of disks provide an important complement for total intensity observations by providing an unattenuated view of disk PI and enabling additional tests on companion candidates detected in total intensity. In Chapter 4, we describe the unique integral field spectropolarimetry mode recently added to Subaru's SCExAO/CHARIS, which provides high-contrast polarimetry at the same array of wavelengths as its total intensity mode. We summarized the data processing software for this observing mode, detailed observing options and best practices, outlined unique science applications for disk imaging, and presented the first preliminary science results for two protoplanetary disk systems: AB Aur and TW Hya.

Chapter 5

In Chapter 5, we introduced constrained reference star differential imaging (RDI) — a new class of point spread function (PSF)-subtraction techniques for RDI data which uses available information (e.g., polarized intensity data) to mitigate the loss of circumstellar signal that plagues disk imaging studies. We demonstrated constrained RDI in application to simulated Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS) IFS observations of a spiral armed disk system, simulated JWST imagery of a debris disk, and to on-sky CHARIS observations of the AB Aur protoplanetary disk system. Using the applications to simulated data, we showed that constrained RDI reduces circumstellar signal (CSS) attenuation to negligible levels — enabling detailed spatial and spectral analysis of even particularly complex disk systems. These techniques are broadly applicable for nearly any disk imaging study, including those using data from: ground-based observatories, Hubble Space Telescope (HST), the James Webb Space Telescope (JWST), and future observatories (e.g., Roman Space Telescope, 30 meter class telescopes, and others). In application to CHARIS observations of AB Aur, this technique helped to confirm AB Aur b — one of a very small number of known embedded exoplanets (and the first concrete evidence of planet formation by gravitational instability). Finally, we summarized the broader applications of constrained RDI for the high-contrast imaging community.

6.2. Future Work

6.2.1. Applying Constrained RDI to Study Disks and Search for Embedded Planets

The constrained RDI techniques (model constrained RDI and polarized intensity constrained RDI) provide data products of sufficient photometric precision to not only confidently identify embedded planets, but also to characterize those planets and to study disk properties in unprecedented detail. We will apply constrained RDI techniques to observations of circumstellar disks

from both Subaru's Subaru Coronagraphic Extreme Adaptive Optics (SCExAO)/CHARIS IFS and to upcoming JWST NIRCam observations. SCExAO/CHARIS observations analyzed will include at least 18 young circumstellar disk systems (already in-hand or scheduled) – of which seven are known to have potentially companion-driven structure. JWST NIRCam observations will include dual-wavelength imaging of four known disk-bearing stars: AU Mic and Fomalhaut C as part of Guaranteed Time Observations (GTO) Program 1184, along with HR 4796A and HD 141569A as part of Director's Discretionary Early Release Science (DD-ERS) Program 1386.

Any candidate embedded companions identified using constrained RDI will be verified by leveraging the high-fidelity products to a) (PI-constrained RDI) check for diminished $\mathrm{F_{pol}}$ at the candidate's location, and b) (PI- & model-constrained RDI) compare spectrophotometry at the candidate position to that of other areas along the disk. The composition of each disk will also be assessed by leveraging available spectrophotometric information.

Any embedded exoplanets uncovered in these data will contribute to the extremely small number of objects at this critical early evolutionary stage, helping us to reach a better understanding of how planets form and how they interact with their circumstellar environments. Additionally, by publicly releasing fast, effective, and generalized software to carry out constrained RDI, we will enable more effective studies of disk systems by the broader disk imaging community.

6.2.2. Meta-parameter Optimization for PSF Subtraction

Beyond the actual detection of exoplanets, an important product of direct imaging studies is the constraints they place on the presence of planets. These constraints fill an otherwise vacant niche in our information regarding, e.g., the occurrence rate of planets as a function of mass and orbital separation. The strength of these constraints is dictated by the sensitivity of the final data products, and is thus governed not only by the precision of the instrumentation, but also by the efficacy of the techniques used in data processing. This efficacy is, in turn, determined by the choice of the settings for PSF-subtraction algorithms such as Locally Optimized Combination of Images (LOCI). However, the optimal settings, or "meta-parameters", cannot be generalized because they depend on countless uncontrollable and difficult to quantify variables of the data. A method for automated optimization of PSF subtraction settings is needed to maximize the value of direct imaging studies.

For meta-parameter optimization, the PSF-subtraction procedure is performed repeatedly while varying the values of these settings – testing the achieved contrasts each time – until the settings that produce the optimal contrasts are found. Though simple in principle, this task is complicated by the discrete nature of many parameters. For example, for the LOCI procedure, N_{ref} governs the maximum number of reference frames to use in modeling the PSF of each target frame – and thus only integer changes will affect the results. Other parameters, such as δ_{FWHM}, which gives the minimum rotational displacement required for each angular differential imaging (ADI) frame to be

considered a valid reference for the target frame, also have discrete effects whose thresholds depend on both the data and the other parameter settings. This results in a non-differentiable optimization problem and thus precludes the use of many common optimization algorithms (e.g. least-squares and related techniques). Differential evolution (DE), however, is particularly well-suited to this scenario: it makes no assumptions about the continuity or differentiability of the problem, performs well even in complicated parameter spaces with multiple local-optima, and typically requires only $\sim 10^3$ samples for convergence. Combining DE with GPU-based code for carrying out PSF subtraction will enable automated optimization of PSF subtraction settings on timescales comparable to a single reduction with currently available software.

We will apply this technique to upcoming JWST NIRCam imagery of a sample of young, nearby M dwarf stars. With optimized PSF subtraction, these data will be effective in probing sub-Saturn planetary masses – unlocking an entirely new discovery space for exoplanet searches by constraining the presence of wide-orbit, sub-Saturn mass companions. Finally, by releasing this tool for public use, we will strengthen the results of direct imaging studies with both current and next-generation observatories.

References

Akiyama, E., Muto, T., Kusakabe, N., et al. 2015, ApJL, 802, L17

Arnold, J. A., Weinberger, A. J., Videen, G., & Zubko, E. S. 2019, AJ, 157, 157, doi: 10.3847/1538-3881/ab095e

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: 10.1051/0004-6361/201322068

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f

Augereau, J. C., Lagrange, A. M., Mouillet, D., Papaloizou, J. C. B., & Grorod, P. A. 1999, A&A, 348, 557. https://arxiv.org/abs/astro-ph/9906429

Barrado Y Navascués, D. 2006, A&A, 459, 511

Beckers, J. M. 1993, ARA&A, 31, 13, doi: 10.1146/annurev.aa.31.090193.000305

Bell, C. P. M., Mamajek, E. E., & Naylor, T. 2015, MNRAS, 454, 593, doi: 10.1093/mnras/stv1981

Betti, S. K., Follette, K., Jorquera, S., et al. 2022, arXiv e-prints, arXiv:2201.08868. https://arxiv.org/abs/2201.08868

Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, A&A, 631, A155, doi: 10.1051/0004-6361/201935251

Bhowmik, T., Boccaletti, A., Thébault, P., et al. 2019, A&A, 630, A85, doi: 10.1051/0004-6361/201936076

Biller, B. A., Liu, M. C., Wahhaj, Z., et al. 2014, in Adaptive Optics Systems IV, ed. E. Marchetti, L. M. Close, & J.-P. Véran, Vol. 9148, International Society for Optics and Photonics (SPIE), 594 – 608, doi: 10.1117/12.2057339

Boccaletti, A., Augereau, J. C., Marchis, F., & Hahn, J. 2003, ApJ, 585, 494, doi: 10.1086/346019

Boccaletti, A., Thébault, P., Pawellek, N., et al. 2019, A&A, 625, A21, doi: 10.1051/0004-6361/201935135

Boccaletti, A., Di Folco, E., Pantin, E., et al. 2020, A&A, 637, L5

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Science, 327, 977, doi: 10.1126/science.1185402

Brandt, T. D. 2018, ApJS, 239, 31, doi: 10.3847/1538-4365/aaec06

Brandt, T. D., Dupuy, T. J., & Bowler, B. P. 2019, AJ, 158, 140, doi: 10.3847/1538-3881/ab04a8

Brandt, T. D., Rizzo, M., Groff, T., et al. 2017, Journal of Astronomical Telescopes, Instruments, and Systems, 3, 048002, doi: 10.1117/1.JATIS.3.4.048002

Carson, J., Thalmann, C., Janson, M., et al. 2013, ApJL, 763, L32, doi: 10.1088/2041-8205/763/2/L32

Caswell, T. A., Droettboom, M., Lee, A., et al. 2021, matplotlib/matplotlib: REL: v3.5.1, v3.5.1, Zenodo, doi: 10.5281/zenodo.5773480

Chabrier, G., Johansen, A., Janson, M., & Rafikov, R. 2014, in Protostars and Planets VI, ed. H. Beuther, R. S. Klessen, C. P. Dullemond, & T. Henning, 619, doi: 10.2458/azu_uapress_9780816531240-ch027

Chen, C., Mazoyer, J., Poteet, C. A., et al. 2020, ApJ, 898, 55, doi: 10.3847/1538-4357/ab9aba

Choquet, É., Pueyo, L., Hagan, J. B., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9143, Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave, ed. J. Oschmann, Jacobus M., M. Clampin, G. G. Fazio, & H. A. MacEwen, 914357, doi: 10.1117/12.2056672

Cloutier, R., Currie, T., Rieke, G. H., et al. 2014, ApJ, 796, 127, doi: 10.1088/0004-637X/796/2/127

Currie, T., Cloutier, R., Brittain, S., et al. 2015a, ApJ, 814, L27, doi: 10.1088/2041-8205/814/2/L27

Currie, T., Kenyon, S. J., Balog, Z., et al. 2008, ApJ, 672, 558, doi: 10.1086/523698

Currie, T., Lada, C. J., Plavchan, P., et al. 2009, ApJ, 698, 1, doi: 10.1088/0004-637X/698/1/1

Currie, T., Lisse, C. M., Kuchner, M., et al. 2015b, ApJL, 807, L7, doi: 10.1088/2041-8205/807/1/L7

Currie, T., Burrows, A., Itoh, Y., et al. 2011, ApJ, 729, 128, doi: 10.1088/0004-637X/729/2/128

Currie, T., Debes, J., Rodigas, T. J., et al. 2012a, ApJL, 760, L32, doi: 10.1088/2041-8205/760/2/L32

Currie, T., Rodigas, T. J., Debes, J., et al. 2012b, ApJ, 757, 28, doi: 10.1088/0004-637X/757/1/28

Currie, T., Guyon, O., Tamura, M., et al. 2017, ApJ, 836, L15, doi: 10.3847/2041-8213/836/1/L15

Currie, T., Brandt, T. D., Uyama, T., et al. 2018, AJ, 156, 291, doi: 10.3847/1538-3881/aae9ea

Currie, T., Marois, C., Cieza, L., et al. 2019, ApJL, 877, L3, doi: 10.3847/2041-8213/ab1b42

Currie, T., Guyon, O., Lozi, J., et al. 2019, in Techniques and Instrumentation for Detection of Exoplanets IX, ed. S. B. Shaklan, Vol. 11117, International Society for Optics and Photonics (SPIE), 285 – 300, doi: 10.1117/12.2529689

Currie, T., Guyon, O., Lozi, J., et al. 2020a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11448, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 114487H, doi: 10.1117/12.2576349

Currie, T., Brandt, T. D., Kuzuhara, M., et al. 2020b, ApJL, 904, L25, doi: 10.3847/2041-8213/abc631

de Boer, J., Langlois, M., van Holstein, R. G., et al. 2020, A&A, 633, A63, doi: 10.1051/0004-6361/201834989

Debes, J. H., Weinberger, A. J., & Kuchner, M. J. 2009, ApJ, 702, 318, doi: 10.1088/0004-637X/702/1/318

Debes, J. H., Weinberger, A. J., & Song, I. 2008, ApJ, 684, doi: 10.1086/592018

Dong, R., Najita, J. R., & Brittain, S. 2018, ApJ, 862, 103, doi: 10.3847/1538-4357/aaccfc

Dong, R., Hashimoto, J., Rafikov, R., et al. 2012, ApJ, 760, 111, doi: 10.1088/0004-637X/760/2/111

Draper, Z. H., Duchêne, G., Millar-Blanchaer, M. A., et al. 2016, ApJ, 826, 147, doi: 10.3847/0004-637X/826/2/147

Duchêne, G., Rice, M., Hom, J., et al. 2020, AJ, 159, 251, doi: 10.3847/1538-3881/ab8881

Engler, N., Boccaletti, A., Schmid, H. M., et al. 2019, A&A, 622, A192, doi: 10.1051/0004-6361/201833542

Esposito, T. M., Kalas, P., Fitzgerald, M. P., et al. 2020, AJ, 160, 24, doi: 10.3847/1538-3881/ab9199

Fitzgerald, M. P., Kalas, P. G., Duchêne, G., Pinte, C., & Graham, J. R. 2007, ApJ, 670, 536, doi: 10.1086/521344

Flasseur, O., Denis, L., Thiébaut, É., & Langlois, M. 2018, A&A, 618, A138, doi: 10.1051/0004-6361/201832745

Flasseur, O., Thé, S., Denis, L., Thiébaut, É., & Langlois, M. 2021, A&A, 651, A62, doi: 10.1051/0004-6361/202038957

Francis, L., & van der Marel, N. 2020, ApJ, 892, 111, doi: 10.3847/1538-4357/ab7b63

Fukagawa, M., Hayashi, M., Tamura, M., et al. 2004, ApJL, 605, L53

Gagné, J., Mamajek, E. E., Malo, L., et al. 2018, ApJ, 856, 23, doi: 10.3847/1538-4357/aaae09

Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, A&A, 616, A1, doi: 10.1051/0004-6361/201833051

—. 2021, A&A, 649, A1

Garufi, A., Dominik, C., Ginski, C., et al. 2022, A&A, 658, A137, doi: 10.1051/0004-6361/202141692

Gaspar, A., & Rieke, G. 2020, Proceedings of the National Academy of Science, 117, 9712

Gibbs, A., Wagner, K., Apai, D., et al. 2019, AJ, 157, 39, doi: 10.3847/1538-3881/aaf1bd

Goebel, S., Currie, T., Guyon, O., et al. 2018, AJ, 156, 279, doi: 10.3847/1538-3881/aaeb24

Gomez Gonzalez, C. A., Wertz, O., Absil, O., et al. 2017, AJ, 154, 7, doi: 10.3847/1538-3881/aa73d7

Grady, C. A., Woodgate, B., Bruhweiler, F. C., et al. 1999, ApJL, 523, L151, doi: 10.1086/312270

Groff, T. D., Chilcote, J., Kasdin, N. J., et al. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI, 99080O, doi: 10.1117/12.2233447

Guyon, O. 2009, in American Institute of Physics Conference Series, Vol. 1158, Exoplanets and Disks: Their Formation and Diversity, ed. T. Usuda, M. Tamura, & M. Ishii, 285–292, doi: 10.1063/1.3215871

Hagan, J. B., Choquet, É., Soummer, R., & Vigan, A. 2018, AJ, 155, 179, doi: 10.3847/1538-3881/aab14b

Hahn, J. M. 2010, ApJ, 719, 1699, doi: 10.1088/0004-637X/719/2/1699

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: 10.1038/s41586-020-2649-2

Hashimoto, J., Tamura, M., Muto, T., et al. 2011, ApJL, 729, L17

Hashimoto, J., Dong, R., Kudo, T., et al. 2012, ApJL, 758, L19, doi: 10.1088/2041-8205/758/1/L19

Hayano, Y., Takami, H., Guyon, O., et al. 2008, in Adaptive Optics Systems, ed. N. Hubin, C. E. Max, & P. L. Wizinowich, Vol. 7015, International Society for Optics and Photonics (SPIE), 255 – 262, doi: 10.1117/12.789992

Henyey, L. G., & Greenstein, J. L. 1941, ApJ, 93, 70, doi: 10.1086/144246

Hoff, W., Henning, T., & Pfau, W. 1998, A&A, 336, 242

Hong, S. S. 1985, A&A, 146, 67

Hughes, A. M., Duchêne, G., & Matthews, B. C. 2018, ARA&A, 56, 541, doi: 10.1146/annurev-astro-081817-052035

Hunter, J. D. 2007, Computing in Science & Engineering, 9, 90, doi: 10.1109/MCSE.2007.55

Janson, M., Brandt, T. D., Moro-Martín, A., et al. 2013, ApJ, 773, 73, doi: 10.1088/0004-637X/773/1/73

Joergens, V., & Neuhäuser, R. 2003, in ESA Special Publication, Vol. 539, Earths: DARWIN/TPF and the Search for Extrasolar Terrestrial Planets, ed. M. Fridlund, T. Henning, & H. Lacoste, 455–458. https://arxiv.org/abs/astro-ph/0306458

Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, PASP, 127, 890, doi: 10.1086/682989

Kalas, P., Fitzgerald, M. P., & Graham, J. R. 2007, ApJ, 661, L85, doi: 10.1086/518652

Kalas, P., Graham, J. R., & Clampin, M. 2005, Nature, 435, 1067, doi: 10.1038/nature03601

Kenyon, S. J., & Bromley, B. C. 2002, ApJL, 577, L35, doi: 10.1086/344084

—. 2008, ApJS, 179, 451, doi: 10.1086/591794

Kenyon, S. J., Gómez, M., & Whitney, B. A. 2008, in Handbook of Star Forming Regions, Volume I, ed. B. Reipurth, Vol. 4 (ASP Monograph Publications), 405

Keppler, M., Benisty, M., Müller, A., et al. 2018, A&A, 617, A44, doi: 10.1051/0004-6361/201832957

Kuhn, J. R., Potter, D., & Parise, B. 2001, ApJL, 553, L189, doi: 10.1086/320686

Kuzuhara, M., Tamura, M., Kudo, T., et al. 2013, ApJ, 774, 11, doi: 10.1088/0004-637X/774/1/11

Labroquère, J., Héritier, A., Riccardi, A., & Izzo, D. 2014, in Parallel Problem Solving from Nature – PPSN XIII, ed. T. Bartz-Beielstein, J. Branke, B. Filipič, & J. Smith (Cham: Springer International Publishing), 262–271

Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, ApJ, 660, 770, doi: 10.1086/513180

Lagrange, A. M., Gratadour, D., Chauvin, G., et al. 2009, A&A, 493, L21, doi: 10.1051/0004-6361:200811325

Lagrange, A.-M., Bonnefoy, M., Chauvin, G., et al. 2010, Science, 329, 57, doi: 10.1126/science.1187187

Lawson, K., Currie, T., Wisniewski, J. P., et al. 2020, AJ, 160, 163, doi: 10.3847/1538-3881/ababa6

—. 2021a, AJ, 162, 293, doi: 10.3847/1538-3881/ac2823

Lawson, K., Currie, T., Wisniewski, J. P., et al. 2021b, in Proc. SPIE, Vol. 118230, 118230D

Lawson, K. D., Wisniewski, J. P., Bellm, E. C., Kowalski, A. F., & Shupe, D. L. 2019, The Astronomical Journal, 158, 119, doi: 10.3847/1538-3881/ab3461

Lisse, C. M., Sitko, M. L., Russell, R. W., et al. 2017, ApJL, 840, L20, doi: 10.3847/2041-8213/aa6ea3

Liu, M. C., Wahhaj, Z., Biller, B. A., et al. 2010, in Adaptive Optics Systems II, ed. B. L. Ellerbroek, M. Hart, N. Hubin, & P. L. Wizinowich, Vol. 7736, International Society for Optics and Photonics (SPIE), 607 – 617, doi: 10.1117/12.858358

Lovell, J. B., Marino, S., Wyatt, M. C., et al. 2021, MNRAS, 506, 1978, doi: 10.1093/mnras/stab1678

Lozi, J., Guyon, O., Jovanovic, N., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10703, Proc. SPIE, 1070359, doi: 10.1117/12.2314282

Lozi, J., Guyon, O., Jovanovic, N., et al. 2020, in American Astronomical Society Meeting Abstracts, Vol. 235, American Astronomical Society Meeting Abstracts #235, 161.07

Lyot, B. 1939, Monthly Notices of the Royal Astronomical Society, 99, 580, doi: 10.1093/mnras/99.8.580

MacGregor, M. A., Weinberger, A. J., Nesvold, E. R., et al. 2019, ApJL, 877, L32, doi: 10.3847/2041-8213/ab21c2

Macintosh, B., Graham, J. R., Barman, T., et al. 2015, Science, 350, 64, doi: 10.1126/science.aac5891

Males, J. R., Close, L. M., Miller, K., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10703, Adaptive Optics Systems VI, ed. L. M. Close, L. Schreiber, & D. Schmidt, 1070309, doi: 10.1117/12.2312992

Marois, C., Correia, C., Galicher, R., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9148, Adaptive Optics Systems IV, ed. E. Marchetti, L. M. Close, & J.-P. Vran, 91480U, doi: 10.1117/12.2055245

Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, ApJ, 641, 556, doi: 10.1086/500401

Marois, C., Macintosh, B., & Véran, J.-P. 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7736, Proc. SPIE, 77361J, doi: 10.1117/12.857225

Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, ApJ, 792, 97, doi: 10.1088/0004-637X/792/2/97

Mazoyer, J., Boccaletti, A., Augereau, J. C., et al. 2014, A&A, 569, A29, doi: 10.1051/0004-6361/201424479

Mazoyer, J., Arriaga, P., Hom, J., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11447, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1144759, doi: 10.1117/12.2560091

Millar-Blanchaer, M. A., Graham, J. R., Pueyo, L., et al. 2015, ApJ, 811, 18, doi: 10.1088/0004-637X/811/1/18

Milli, J., Mouillet, D., Lagrange, A. M., et al. 2012, A&A, 545, A111, doi: 10.1051/0004-6361/201219687

Milli, J., Vigan, A., Mouillet, D., et al. 2017, A&A, 599, A108, doi: 10.1051/0004-6361/201527838

Moór, A., Ábrahám, P., Derekas, A., et al. 2006, ApJ, 644, 525, doi: 10.1086/503381

Moór, A., Pascucci, I., Kóspál, Á., et al. 2011, ApJS, 193, 4, doi: 10.1088/0067-0049/193/1/4

Müller, A., Keppler, M., Henning, T., et al. 2018, A&A, 617, L2, doi: 10.1051/0004-6361/201833584

Newville, M., Stensitzki, T., Allen, D. B., & Ingargiola, A. 2014, LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python, Zenodo, doi: 10.5281/ZENODO.11813

Okuta, R., Unno, Y., Nishino, D., Hido, S., & Crissman. 2017, in 31st Conference on Neural Information Processing Systems

Pairet, B., Cantalloube, F., & Jacques, L. 2021, MNRAS, 503, 3724, doi: 10.1093/mnras/stab607

Paulson, D. B., Cochran, W. D., & Hatzes, A. P. 2004, AJ, 127, 3579, doi: 10.1086/420710

Perrin, M. D., Schneider, G., Duchene, G., et al. 2009, ApJL, 707, L132

Perrin, M. D., Sivaramakrishnan, A., Lajoie, C.-P., et al. 2014, in Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave, ed. J. M. O. Jr., M. Clampin, G. G. Fazio, & H. A. MacEwen, Vol. 9143, International Society for Optics and Photonics (SPIE), 1174 – 1184, doi: 10.1117/12.2056689

Perrin, M. D., Duchene, G., Millar-Blanchaer, M., et al. 2015, ApJ, 799, 182, doi: 10.1088/0004-637X/799/2/182

Pueyo, L. 2016, ApJ, 824, 117, doi: 10.3847/0004-637X/824/2/117

Pueyo, L., Soummer, R., Hoffmann, J., et al. 2015, ApJ, 803, 31, doi: 10.1088/0004-637X/803/1/31

Ren, B., Pueyo, L., Chen, C., et al. 2020, ApJ, 892, 74, doi: 10.3847/1538-4357/ab7024

Ren, B., Pueyo, L., Zhu, G. B., Debes, J., & Duchêne, G. 2018, ApJ, 852, 104, doi: 10.3847/1538-4357/aaa1f2

Rhee, J. H., Song, I., Zuckerman, B., & McElwain, M. 2007, ApJ, 660, 1556, doi: 10.1086/509912

Ribas, Á., Bouy, H., & Merín, B. 2015, A&A, 576, A52, doi: 10.1051/0004-6361/201424846

Rich, E. A., Wisniewski, J. P., Currie, T., et al. 2019, ApJ, 875, 38

Rodigas, T. J., Hinz, P. M., Leisenring, J., et al. 2012, ApJ, 752, 57, doi: 10.1088/0004-637X/752/1/57

Sai, S., Itoh, Y., Fukagawa, M., Shibai, H., & Sumi, T. 2015, PASJ, 67, 20, doi: 10.1093/pasj/psu152

Schneider, G., Grady, C. A., Hines, D. C., et al. 2014, AJ, 148, 59, doi: 10.1088/0004-6256/148/4/59

Shinn, J.-H. 2018, ApJS, 239, 21, doi: 10.3847/1538-4365/aae3e5

Smith, B. A., & Terrile, R. J. 1984, Science, 226, 1421, doi: 10.1126/science.226.4681.1421

Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJL, 755, L28, doi: 10.1088/2041-8205/755/2/L28

Sparks, W. B., & Ford, H. C. 2002, ApJ, 578, 543, doi: 10.1086/342401

Spiegel, D. S., & Burrows, A. 2012, ApJ, 745, 174, doi: 10.1088/0004-637X/745/2/174

Steiger, S., Currie, T., Brandt, T. D., et al. 2021, AJ, 162, 44, doi: 10.3847/1538-3881/ac02cc

Stolker, T., Dominik, C., Min, M., et al. 2016, A&A, 596, A70, doi: 10.1051/0004-6361/201629098

Storn, R., & Price, K. 1997, Journal of Global Optimization, 11, 341, doi: 10.1023/a:1008202821328

Suzuki, R., Kudo, T., Hashimoto, J., et al. 2010, in Ground-based and Airborne Instrumentation for Astronomy III, ed. I. S. McLean, S. K. Ramsay, & H. Takami, Vol. 7735, International Society for Optics and Photonics (Proc. SPIE), 1184 – 1196, doi: 10.1117/12.857361

Tamura, M. 2009, in American Institute of Physics Conference Series, Vol. 1158, American Institute of Physics Conference Series, ed. T. Usuda, M. Tamura, & M. Ishii, 11–16, doi: 10.1063/1.3215811

Tamura, M. 2016, Proceedings of the Japan Academy, Series B, 92, 45, doi: 10.2183/pjab.92.45

Thalmann, C., Janson, M., Buenzli, E., et al. 2011, ApJL, 743, L6, doi: 10.1088/2041-8205/743/1/L6

—. 2013, ApJL, 763, L29, doi: 10.1088/2041-8205/763/2/L29

Thebault, P., Kral, Q., & Ertel, S. 2012, A&A, 547, A92, doi: 10.1051/0004-6361/201219962

Uyama, T., Hashimoto, J., Kuzuhara, M., et al. 2017, AJ, 153, 106

Vacca, W. D., & Sandell, G. 2011, ApJ, 732, 8

van Boekel, R., Henning, T., Menu, J., et al. 2017, ApJ, 837, 132

van der Marel, N., Birnstiel, T., Garufi, A., et al. 2021, AJ, 161, 33, doi: 10.3847/1538-3881/abc3ba

van Holstein, R. G., Girard, J. H., de Boer, J., et al. 2020, A&A, 633, A64

van Holstein, R. G., Bos, S. P., Ruigrok, J., et al. 2020, in Ground-based and Airborne Instrumentation for Astronomy VIII, ed. C. J. Evans, J. J. Bryant, & K. Motohara, Vol. 11447, International Society for Optics and Photonics (Proc. SPIE), 1113 – 1126, doi: 10.1117/12.2576188

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: 10.1038/s41592-019-0686-2

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261

Wang, J. J., Ruffio, J.-B., De Rosa, R. J., et al. 2015, pyKLIP: PSF Subtraction for Exoplanets and Disks. http://ascl.net/1506.001

Whitney, B. A., Robitaille, T. P., Bjorkman, J. E., et al. 2013, ApJS, 207, 30, doi: 10.1088/0067-0049/207/2/30

Williams, J. P., & Cieza, L. A. 2011, ARA&A, 49, 67, doi: 10.1146/annurev-astro-081710-102548

Wyatt, M. C. 2008, ARA&A, 46, 339, doi: 10.1146/annurev.astro.45.051806.110525