



Constrained Reference Star Differential Imaging: Enabling High-fidelity Imagery of Highly Structured Circumstellar Disks*

Kellen Lawson¹ , Thayne Currie^{2,3,4} , John P. Wisniewski¹ , Tyler D. Groff⁵ , Michael W. McElwain⁵ , and Joshua E. Schlieder⁵

¹ Department of Physics and Astronomy, University of Oklahoma, Norman, OK, USA; kellenlawson@gmail.com

² Subaru Telescope, National Astronomical Observatory of Japan, 650 North A'ohōkū Place, Hilo, HI 96720, USA

³ Department of Physics and Astronomy, University of Texas-San Antonio, 1 UTSA Circle, San Antonio, TX, USA

⁴ Eureka Scientific, 2452 Delmer Street, Suite 100, Oakland, CA, USA

⁵ NASA-Goddard Space Flight Center, Greenbelt, MD, USA

Received 2022 March 30; revised 2022 July 22; accepted 2022 July 29; published 2022 August 17

Abstract

High-contrast imaging presents us with the opportunity to study circumstellar disks and the planets still embedded within them, providing key insights into the formation and evolution of planetary systems. However, the postprocessing techniques that are often needed to suppress stellar halo light typically result in significant and variable loss of circumstellar light, even when using relatively conservative approaches like reference star differential imaging (RDI). We introduce “constrained reference star differential imaging” (constrained RDI), a new class of RDI point-spread-function (PSF) subtraction techniques for systems with circumstellar disks. Constrained RDI utilizes either high-resolution polarized-intensity (PI) images or disk models to severely limit or even eliminate the signal loss due to oversubtraction that is common to RDI. We demonstrate the ability of constrained RDI utilizing polarimetric data to yield an oversubtraction-free detection of the AB Aurigae protoplanetary disk in total intensity. PI-constrained RDI allows us to decisively recover the spectral signature of the confirmed, recently discovered protoplanet, AB Aurigae b. We further demonstrate that constrained RDI can be a powerful analysis tool for soon-to-be-acquired James Webb Space Telescope coronagraphic imaging of disks. In both cases, constrained RDI provides analysis-ready products that enable more detailed studies of disks and more robust verification of embedded exoplanets.

Unified Astronomy Thesaurus concepts: [Circumstellar disks \(235\)](#); [Protoplanetary disks \(1300\)](#); [Debris disks \(363\)](#); [Direct imaging \(387\)](#); [Exoplanet detection methods \(489\)](#); [Polarimetry \(1278\)](#); [Coronagraphic imaging \(313\)](#)

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 substellar objects that may cause them (e.g., Keppler et al. 2018). Combined with multiwavelength or integral field spectrograph (IFS) data, we have the technology to conduct incredibly detailed spatial and spectral analyses of these systems.

A key challenge for these studies is the isolation of the 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 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., LOCI (Lafrenière et al. 2007) or KLIP (Soummer et al. 2012)—better model the starlight but also cause some CSS to be lost (or “attenuated”; see Figure 1, left column). Because this CSS loss is neither spatially nor spectrally (for multiwavelength 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, because 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 nonuniform loss of signal in total intensity precludes this measurement.

In some scenarios, the challenges presented by disk signal loss and the confusion of disk and planet signals can be circumvented using forward-modeling techniques (e.g., Currie et al. 2015; Pueyo 2016; Lawson et al. 2020; Mazoyer 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

* Based on data collected at Subaru Telescope, which is operated by the National Astronomical Observatory of Japan.



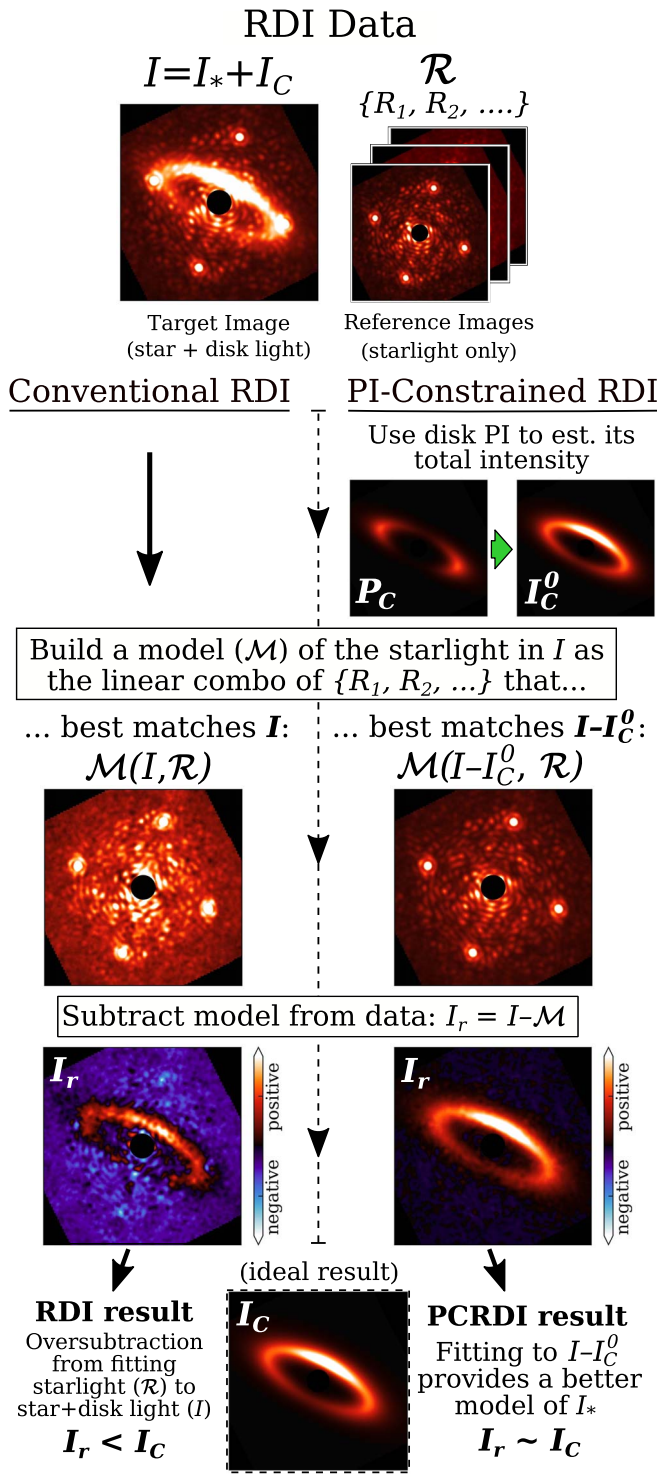


Figure 1. Top center: RDI data for a (synthetic) disk system. Left column: conventional RDI PSF subtraction, yielding significantly attenuated results due to oversubtraction (“RDI result,” lower left). Right column: PI-constrained RDI (PCRDI; Section 3.1). The PCRDI result (lower right) avoids attenuation by using PI disk imagery to estimate and suppress disk flux while modeling the diffraction pattern of the starlight in total intensity.

reproduced. This enables the robust assessment of disk geometry in the presence of signal loss, while also allowing an approximate correction of disk surface-brightness (SB) 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)—has been developed for the purpose of exoplanet and disk imaging. These include techniques that focus on more thoroughly removing starlight and/or better recovering point sources (e.g., TLOCI, A-LOCI, PACO; Currie et al. 2012; Marois et al. 2014; Flasseur et al. 2018; Currie et al. 2022a), 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 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 more extended disks or disks that dominate small separations (where temporal PSF changes are most significant). More recently, “source-separation” algorithms, such as MAYONNAISE (Pairet et al. 2021) and REXPACO (Flasseur et al. 2021), have provided sophisticated tools for isolating CSS and limiting the negative effects of PSF subtraction. Source-separation techniques attempt to simultaneously model each significant source of signal throughout a data sequence, e.g., separately considering the starlight and any rotating (in an ADI sequence) CSS, as well as the time-variable PSF, the coronagraph, and noise. As both tools significantly leverage the continuous spatial+temporal variations throughout an ADI (or ADI+SDI) sequence to disentangle sources of light, it is unclear if they will be effective for space-based observations, such as those from JWST, for which disk targets are typically observed at only two distinct roll angles. Further, while they are intended to precisely isolate CSS, they do not provide any clear methods for assessing, quantifying, or correcting any inaccuracy that might remain. As it stands, despite the exceptional capabilities of current and upcoming observatories, our ability to study circumstellar disks (and the young planets within them) is inhibited by the loss of CSS signal during postprocessing 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 2 and 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 4) and to on-sky CHARIS observations of the AB Aurigae protoplanetary disk

system, whose embedded protoplanet is verified by this approach (Section 5; Currie et al. 2022b). Additionally, using synthetic disk models as constraints, we apply constrained RDI to simulated JWST NIRC*am* observations (Section 3.2) to demonstrate the approach’s efficacy for upcoming JWST observations. 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 m class telescopes, and others; Section 6). Moreover, constrained RDI can be implemented in existing pipelines with only minor changes (Section 6) and is extremely unlikely to induce spurious (false-positive) circumstellar features (Section 7), and its results can be assessed using standard methods familiar to the disk-imaging community (e.g., forward modeling; Appendix C).

2. Disk Signal Attenuation in Reference Star Differential Imaging

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: $\mathcal{R} = \{R_1, R_2, \dots, R_N\}$. 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 as⁶

$$I_r = I - \mathcal{M}(I, \mathcal{R}). \quad (1)$$

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 Equation (1):

$$I'_C = I_C - \mathcal{M}(I_C, \mathcal{R}). \quad (2)$$

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).

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, because 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

$$I_r = I - \mathcal{M}(I - I_C^0, \mathcal{R}). \quad (3)$$

Notably, 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). The estimate of I_C , I_C^0 , is used exclusively in determining the optimal combination of reference images to use in the PSF model and is not directly used in the final product; the effect of the constraint must always be some combination of reference images. This provides substantial insulation against the introduction of false-positive circumstellar features. A particularly poor estimate of the circumstellar signal contained in the data could result in worse PSF-subtraction (e.g., by causing undersubtraction). However, to induce spurious features from the estimate in the final result, such features would need to be reproduced by some combination of the reference images and at a sufficient number of parallactic angles to remain after the sequence is derotated and averaged. For data with nonnegligible field rotation, using small/truncated reference sequences (e.g., retaining the most correlated images from a larger reference library) and using large/full-frame PSF-subtraction regions, this is very unlikely (see Section 7 for further discussion). Moreover, as with conventional RDI reductions, any features that do manifest can be verified using forward-modeling techniques (Appendix C). As such, this approach provides an extremely safe means by which to improve RDI reductions of circumstellar disk targets.

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 (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

⁶ For simplicity, the notation herein refers to a single target image. In application, there are generally multiple such images, which would each be PSF-subtracted in the same manner before being derotated to a north-up orientation and averaged together to form the final result.

scattered light is often described by Rayleigh polarization (e.g., Whitney et al. 2013; Stolker et al. 2016; Gomez Gonzalez et al. 2017):

$$\frac{\text{PI}}{I} = F_{\text{pol}} = \frac{1 - \cos^2 \phi}{1 + \cos^2 \phi}. \quad (4)$$

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. By inverting Equation (4), this can be used to transform P_C to an estimate of I_C :

$$I_C^0 = P_C \cdot \left(\frac{1 + \cos^2 \Phi}{1 - \cos^2 \Phi} \right) = \frac{P_C}{F_{\text{pol}}^0}, \quad (5)$$

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 Equation (3) to carry out constrained RDI PSF subtraction. This process is visualized in the right column of Figure 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 3.3). See Section 7 for a discussion of additional considerations.

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 (Section 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), 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 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 approximate attenuation corrections (e.g., Bhowmik et al. 2019), such corrections become extremely noisy in regions where the processed model approaches zero and also scale any noise or residual starlight when correcting CSS (meaning signal-to-noise is not improved this way; see Section 4.1). Using the same model to carry out MCRDI instead provides the same benefits but avoids inflating stellar residuals and noise—enabling accurate analysis over more of

the field of view (FOV) and potentially revealing fine morphological features and faint extended signals that would otherwise not be recovered.

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:

$$y = (I - I_C^0) - \mathcal{M}(I - I_C^0, \mathcal{R}). \quad (6)$$

This is simply the conventional RDI residual calculation (Equation (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:

$$\begin{aligned} \text{As } \mathcal{M}(I - I_C^0, \mathcal{R}) &\rightarrow I_* \text{ and } I_C^0 \rightarrow I_C, \\ y &\rightarrow (I - I_C) - I_* = I_* - I_* = \hat{0}. \end{aligned} \quad (7)$$

Because 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 straightforward objective function for optimization is then simply the sum of the squares of these residuals binned to the spatial resolution of the data.⁷ This objective function can then be combined with an optimization algorithm of choice to automatically determine the optimal estimate of CSS. This procedure might be thought of as an extension of the negative PSF imputation technique for point sources introduced by Marois et al. (2010), in that both use the CSS- and PSF-subtracted residuals of a data sequence to gauge the quality of a CSS model.

4. Throughput Assessment Using Simulated Data

4.1. PCRDI Throughput

To test the throughput of PCRDI in application to CHARIS IFS data (Groff et al. 2016), we generated multiwavelength synthetic disk image cubes in Stokes Q , U , and I using HO-CHUNK 3D (Whitney et al. 2013) for a spiral-armed disk of similar morphology to the disk of AB Aurigae (see Section 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: $\mathcal{I} = \{I_1, I_2, \dots\}$), and the other which will be left unchanged (the “reference” sequence, \mathcal{R}).⁸ 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 full sequence.⁹ After convolution, Q and

⁷ 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, because 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.

⁸ For more information regarding the size of the target and reference sequences, see Table 1.

⁹ 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 NIRCам.




Figure 2. Wavelength-averaged results for the simulated data described in Section 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 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 nonphysical over most of the field of view ($F_{\text{pol}} > 1$ or $F_{\text{pol}} < 0$). Though the “smooth surface” F_{pol}^0 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.

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 ($\sim 0.''16$ – $0.''45$) for all 22 wavelength channels. This optimization region is also used for the calculation of the goodness of fit for each trial result (Equation (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 2 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 Figure 3).¹⁰

When using attenuated disk images for analysis of disk SB and color, forward modeling is often used to derive approximate attenuation corrections (e.g., Goebel et al. 2018; Betti et al. 2022). To provide a more direct comparison of constrained RDI with such techniques, we derive and apply attenuation corrections for the conventional RDI reduction in a similar manner (see Appendix A for details) using two different disk models. For the first model, we use the “ground-truth” disk model contained in the simulated data—referred to hereafter as the “ideal” case. This should be regarded as an absolute upper limit on the accuracy of corrections derived this way; for similarly complicated disks, modeling results do not typically approach this accuracy and often do not attempt to reproduce nonaxisymmetric structures at all (e.g., Lomax et al. 2016; Betti et al. 2022). The second model instead emulates a more realistic modeling result, consisting of a simple two-ring disk model without any spiral structures but which generally reproduces the bulk of the disk’s shape and brightness—referred to hereafter as the “realistic” case. Both models are visualized in Figure 10 of Appendix C (as “Model 1” and “Model 2,” respectively), with the forward-modeled results for RDI presented in the top row.

Figure 4 compares the ground-truth SB and color profiles with those measured from the corrected RDI results, as well as the uncorrected RDI and the PCRDI results. For this purpose, we bin each of the final image cubes along the wavelength axis to approximate near-infrared (NIR) J , H , and K bands.

¹⁰ Note: to provide attenuation estimates that are not significantly dependent on the quality/stability of the AO correction throughout the observations, this assessment and the distributions plotted in Figure 3 exclude any residual starlight such that oversubtraction is the sole source of any change from the ground truth.




Figure 3. For the simulated data of Section 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.

Additionally, we take the average over the entire CHARIS wavelength range to produce a single image which we refer to simply as “broadband.” Overall, PCRDI significantly outperforms model-corrected RDI—even in the unlikely scenario in which modeling exactly recovers the underlying disk. The inaccuracy that remains even in this “ideal” correction scenario is the result of a key assumption made by the forward-modeling correction strategy: that only attenuated disk signal remains in the image being corrected. In fact, there is also residual starlight and noise—both of which are inflated when the disk signal is corrected, leading to the inaccuracy measured here. For the $J - H$ color profiles, comparing measurements for PCRDI with those of RDI with “ideal” corrections: PCRDI is, on average, closer to the ground truth by a factor of ~ 2 for the horizontal profile and by a factor of ~ 4 for the vertical profile (with uncertainties that are smaller by factors of ~ 4 and ~ 5 ,

respectively). Forward-modeling the PCRDI CSS estimate, I_C^0 , to make corrections for RDI yields comparable accuracy to the “ideal” corrections in the inner disk, which become somewhat worse beyond the region in which the estimate was optimized ($\sim 0''.45$). These results demonstrate that, in any regions where starlight and/or noise are nonnegligible compared to the residual CSS, the application of corrections to a conventional RDI result will yield less accurate measurements than constrained RDI. In other words, while correcting RDI this way cannot improve the ratio of CSS to (a) noise or (b) residual starlight at a given location, the use of constrained RDI can.

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

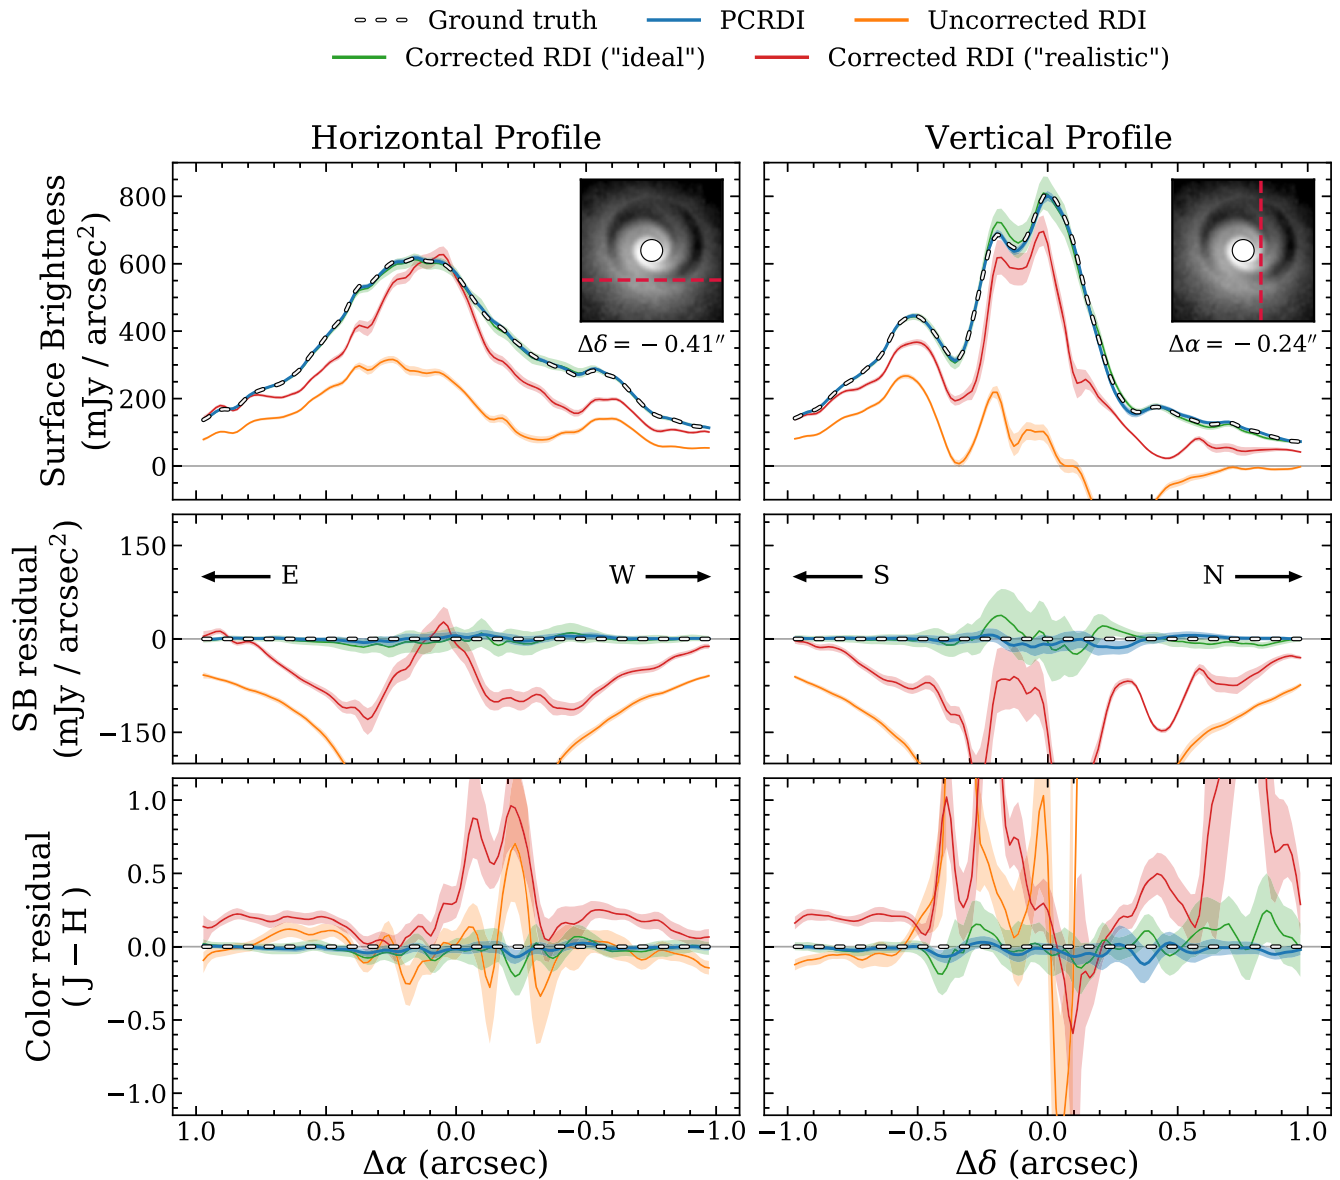


Figure 4. SB and color measurements for the results of Section 4.1 using an aperture with a 2 pixel radius (diameter $\sim 0''.065$, slightly larger than the FWHM of the longest CHARIS wavelength channel: $0''.062$). Shaded regions indicate approximate 3σ confidence intervals following the procedure of Appendix B. Top row: broadband SB along a horizontal (left) and vertical (right) profile (the dashed red line in the inset images). Center row: The difference between each reduction and the ground-truth profile for the measurements in the top row. Bottom row: offset of the $J - H$ disk color from that of the ground-truth model for the same products and positions. Gaps in the data (particularly for the vertical profile of the uncorrected RDI reduction) correspond to undefined color measurements resulting when one of the two filters measured a negative SB.

4.2. MCRDI Throughput

To demonstrate the MCRDI technique, we carried out simplified simulations of JWST NIRCcam F335M observations of the debris disk system HD 10647.¹¹ The NIRCcam PSF was simulated using WebbPSF¹² (Perrin et al. 2014) while the debris disk was simulated using DiskDyn (Gaspar & Rieke 2020). The NIRCcam 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 ($\Delta\text{PA} = 10^\circ$) and a five-point dither of the reference star, but ignore other factors such as noise and thermal effects. For simplicity, the

PSF reference star’s spectrum is assumed to be identical to that of the target—corresponding to a perfect color and flux match. 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/NIRCcam 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 brightness) and using a Hong-like scattering phase function (Hong 1985; in contrast to the DiskDyn model, in which scattering properties are simulated for a disk of astronomical silicate dust for particles evenly spaced in mass/size log space between $0.1 \mu\text{m}$ and 1cm). Using a different tool than the one used to generate the “true” input disk model is intended to

¹¹ Based on upcoming JWST Cycle 1 Guaranteed Time Observations (GTO) observations of this system (Program 1183).

¹² <https://github.com/spacetelescope/webbpsf>




Figure 5. Simulated results for JWST NIRCcam 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 Figure 6).

better emulate real observations, where a model 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 parameters 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). 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 RDI,¹³ KLIP RDI,¹⁴ and model-constrained RDI are shown in Figures 5 and 6. The optimization region used for both KLIP and MCRDI excludes only the region within 11 pixels of the star ($0''.69$, the approximate inner working angle). 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.

In this case, we do not compare MCRDI with model-corrected RDI results as we did in the previous section. Because our simulated data do not include realistic noise, this

would not provide a reasonable estimate of the quality that can be achieved with such corrections.

5. Application of PCARDI 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 astronomical units; Grady et al. 1999; Fukagawa et al. 2004) and small scales (tens of astronomical units; Hashimoto et al. 2011; Boccaletti et al. 2020) in optical and NIR imagery, and with a large, depleted inner cavity in submillimeter (van der Marel et al. 2021).

As a further demonstration of the PCARDI technique, we utilize Subaru Coronagraphic Extreme Adaptive Optics (SCEXAO)/CHARIS observations of AB Aur in total intensity ($t_{\text{int}} \approx 116$ minutes) and polarized intensity ($t_{\text{int}} \approx 74$ minutes), originally reported in Currie et al. (2022b) and Lawson et al. (2021b) respectively. These data were taken over two consecutive nights in October 2020, with both polarized and total-intensity sequences using CHARIS in low-res broadband mode ($R \sim 19$, $1.15\text{--}2.39 \mu\text{m}$, producing 22 wavelength channels per exposure) and utilizing a 113 mas Lyot coronagraph. Currie et al. (2022b) also observed a reference star (HD 31233) before and after the AB Aur total-intensity sequence ($t_{\text{int}} \approx 21$ minutes), which we also use here to enable RDI. See Currie et al. (2022b) 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 $r = 10\text{--}25$ pixels ($\sim 0''.16\text{--}0''.41$).¹⁵ For these 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 $r \sim 0''.5$ (see

¹³ For classical RDI, we use the median of the reference sequence as the PSF model. More so than KLIP or MCRDI, the efficacy of classical RDI depends on factors that are difficult to generalize here. For example, the amount of residual starlight is dictated in part by the randomly determined positions of the target and reference stars behind the coronagraph (the pointing offset). For the small number of images used, a given sequence could plausibly perform noticeably better or worse than the result here. However, the qualitative relationship will remain—i.e., that classical RDI provides a high CSS throughput, but worse starlight suppression at smaller separations.

¹⁴ 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 ($N = 1\text{--}15$). Using the unconstrained RDI procedure of Section 4.1 instead yields comparable results.

¹⁵ The conventional RDI reduction is different from (but comparable to) the one in Currie et al. (2022b), which used KLIP with larger optimization regions instead.

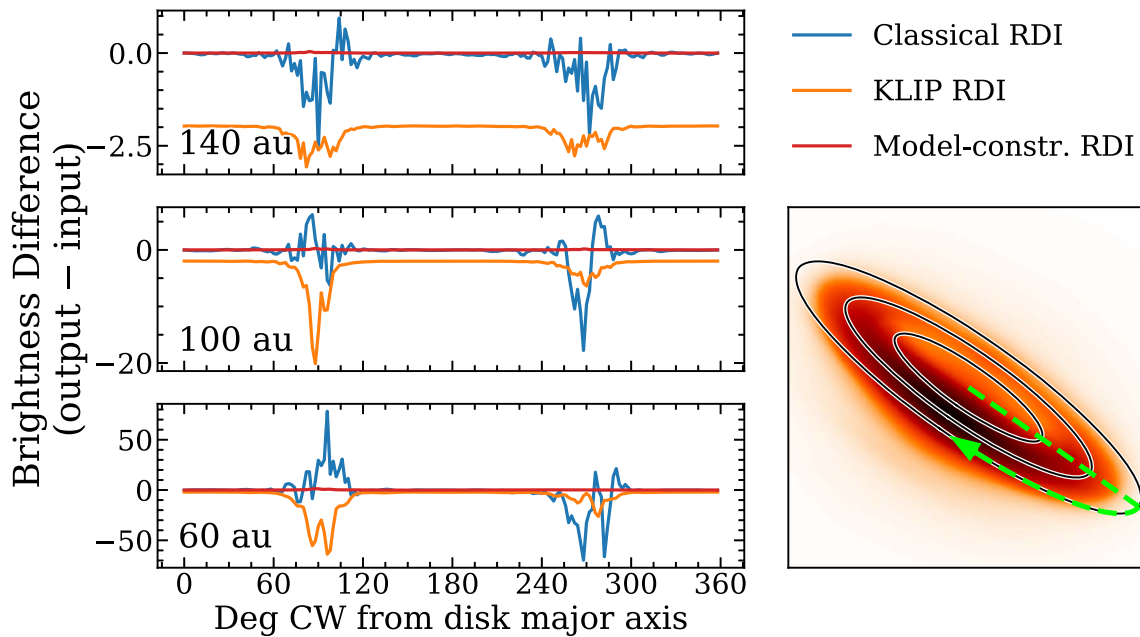


Figure 6. For each reduction of Figure 5: the difference between the output and input brightness of the disk measured in apertures with diameter equal to the PSF FWHM. The lower-right panel provides a visualization of the three radial separations measured (black ellipses) and also shows the origin and 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.




Figure 7. The components of the PCRDI procedure (Section 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_C^0); right: and the final sequence-combined residual image from PCRDI.

- Lawson et al. 2021b). Limiting the optimization region to separations with full azimuthal coverage substantially simplifies and accelerates PSF model optimization.
- Hashimoto et al. (2011) report a misalignment between AB Aur’s inner disk ($\sim 0''.21$ – $0''.43$) and outer disk ($\sim 0''.63$ – $0''.84$), which would inhibit our smooth scattering surface estimate for an optimization region extending to the outer disk.
 - AB Aur shows a substantial total-intensity enhancement at $r \sim 0''.6$ (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. (2022b), where it is determined to be an 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_{pol}^0), the corresponding PI-based CSS estimate (I_C^0), and the final PCRDI result (I_r) is provided in Figure 7. Figure 8 compares the PCRDI results with the conventional RDI results in an array of wavelength ranges—where, as for the simulated data (Section 4), PCRDI provides an improvement in CSS throughput that is immediately apparent. Notably, in the PCRDI imagery, it is evident that the morphological feature $\sim 0''.6$ south of the image center (AB Aur b; Currie et al. 2022b) has a distinct (bluer) spectral energy distribution relative to the nearby disk appearing at a much higher contrast to the disk in the J band than in the 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. This effect is visualized in Figure 9

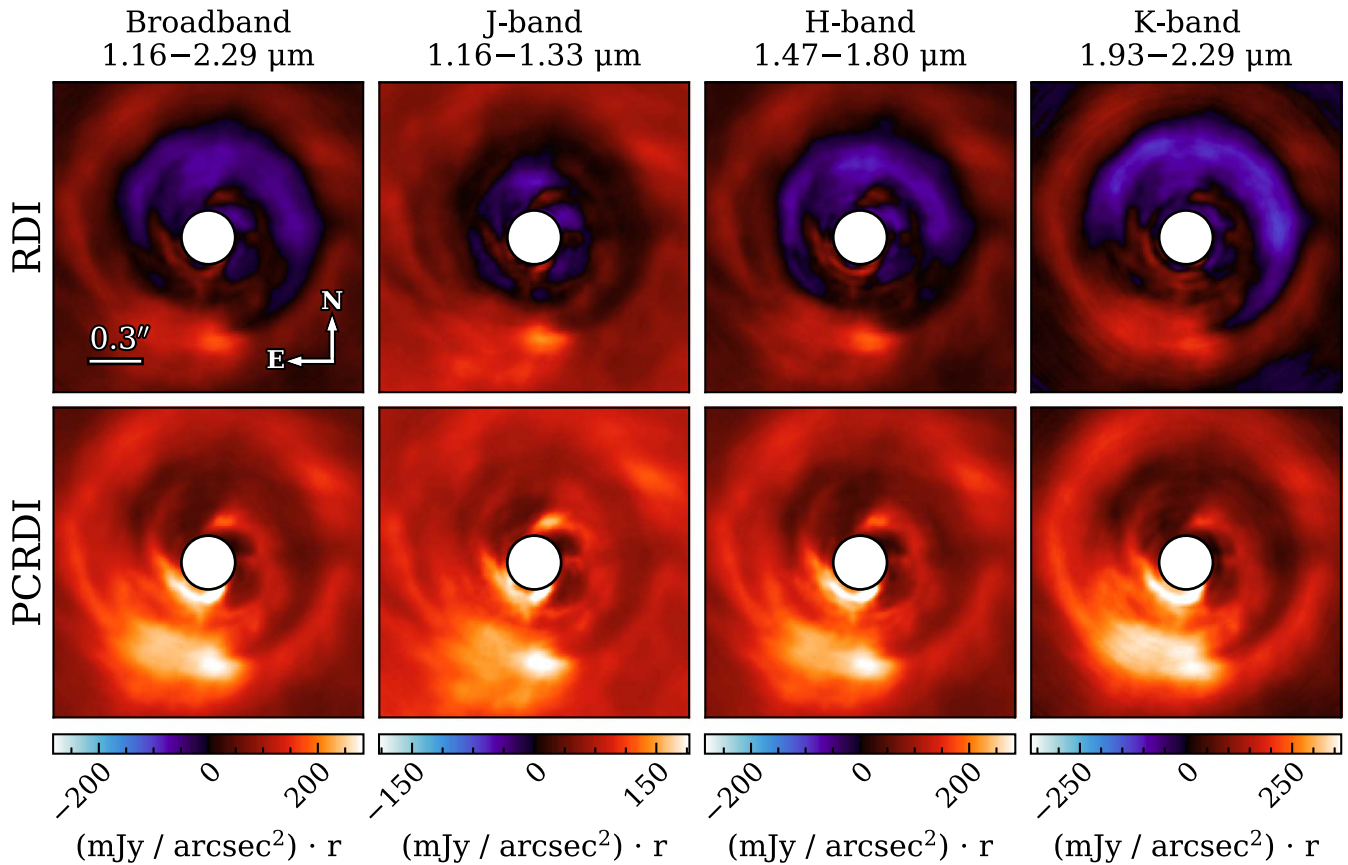


Figure 8. A comparison of PSF-subtracted results for SCEXAO/CHARIS observations of AB Aur (see Section 5) using a conventional RDI procedure (top row) and PI-constrained RDI (PCRDI; bottom row). The PCRDI results (see Sections 3.1 and 5) use the same PSF-subtraction algorithm and algorithm settings as 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).

using measurements of SB and color across the position of AB Aur b.

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 D.

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, constrained RDI 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 runtimes for the reductions carried out in this work are provided in Appendix E.

Though we have demonstrated the use of PCRDI with near-contemporaneous PI data from the same instrument, the same principles can be applied using noncontemporaneous 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; Rich 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 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). In such a scenario, the procedure for constrained RDI forward modeling described in Appendix C could be adopted. Here, faster scattered-light models could be used to optimize the MCRDI result, with forward modeling of more physically motivated models (perhaps using the MCRDI model parameters to initialize the procedure) then being used as the basis for interpretation of the imagery.

Additionally, for RDI data for which a full forward-modeling procedure has been performed, MCRDI can be applied retroactively using the model identified from forward modeling




Figure 9. A comparison of SB and color profiles across the position (east-west) of the confirmed protoplanet AB Aur b for the conventional RDI (orange) and PCRDI (blue) reductions; see Section 5. The shaded regions are the approximate 3σ confidence intervals (see Appendix B). Top: broadband SB measured using an aperture with a radius of 4 pixels ($0''.065$; for consistency with Currie et al. 2022b). The two results show similar overall shapes, but with the standard RDI result being significantly fainter due to oversubtraction. Bottom: as above, but for the $J-K$ color of the disk in magnitudes (with the color of the parent star subtracted). The PCRDI result shows a bluer color at the position of AB Aur b relative to the disk on either side. While the RDI result also shows a blue dip, it is (a) of much lower significance and (b) manifests as a smoother color transition that is not clearly distinct around the location of AB Aur b.

to improve the RDI results. In fact, the residuals when evaluating an MCRDI CSS estimate (Equation (6)) are identical to the forward-modeled residuals (processed data – model) when adopting the MCRDI constraint as the model—so long as $\mathcal{M}(I, \mathcal{R})$ is linear in I (e.g., KLIP and LOCI). That is, if $\mathcal{M}(I_1 - I_2, \mathcal{R}) = \mathcal{M}(I_1, \mathcal{R}) - \mathcal{M}(I_2, \mathcal{R})$, then (starting from the right-hand side of Equation (6))

$$\begin{aligned} & (I - I_C^0) - \mathcal{M}(I - I_C^0, \mathcal{R}) \\ &= (I - I_C^0) - (\mathcal{M}(I, \mathcal{R}) - \mathcal{M}(I_C^0, \mathcal{R})) \\ &= (I - \mathcal{M}(I, \mathcal{R})) - (I_C^0 - \mathcal{M}(I_C^0, \mathcal{R})), \end{aligned} \quad (8)$$

which is the standard equation for forward-modeled residuals for model I_C^0 with a conventional RDI procedure. In other words, model parameters from forward-modeling optimization should be identical to those from MCRDI optimization. Following from the analysis and discussion comparing

constrained RDI and model-corrected RDI in Section 4.1, when a strong CSS model has been identified, it appears that an MCRDI reduction using that model as the constraint should improve measurement accuracy over the application of model-based corrections to the conventional RDI product.

Unlike examples presented for REXPACO (Flasseur et al. 2021) and MAYONNAISE (Pairet et al. 2021), the constrained RDI examples presented here make no effort to directly disentangle distinct sources of CSS (e.g., separating the planet from the disk signal). In application to AB Aur, we have simply avoided the vicinity of AB Aur b while constructing the PSF model and then relied on spectrophotometric analysis (afforded by the IFS data) to distinguish AB Aur b from the disk. In principle, however, a point source could be added to the CSS estimate used by constrained RDI, optimizing its location and brightness alongside any parameters governing the estimate of the disk. Alternatively, constrained RDI could conceivably be combined with one of these source-separation tools, e.g., by using the constrained RDI PSF model in place of the PSF model from the “fixed-point algorithm” in MAYONNAISE.

7. Considerations and Limitations

As noted, our PI-based estimate of the total intensity of CSS assumes a simple “bell-shaped” fractional polarization curve that peaks at a scattering angle of 90° (Equation (4)). This prescription is generally appropriate for NIR observations of disks, irrespective of the relation between the dominant dust size and the observing wavelength (with larger grains tending to produce diminished fractional polarization, but still maintaining a bell shape and center about 90° ; Benisty et al. 2022). However, even if the true fractional polarization peaks at a different scattering angle or manifests with a skewed shape (e.g., the C01 model of Takami et al. 2013), the utilized model may still be sufficient for a strong PCRDI solution. By tuning the parameters governing the assumed scattering surface (away from those of the “true” surface), the utilized curve can emulate fractional polarization maps for other polarization curves over a region of interest. For example, an inclined disk with a particular surface flare (c) and a polarization curve with a peak skewed from 90° toward higher or lower scattering angles could be reproduced by increasing or decreasing (respectively) the assumed flare of the scattering surface without altering the assumed polarization curve. By allowing our optimization to explore a wide range of scattering surfaces, the polarization model of Equation (4) can provide strong solutions even for disks with polarization that is not truly well described in this way.

Nevertheless, care should be taken when applying this model to highly inclined and/or highly flared disks (where a wider range of scattering angles may be probed, e.g., IM Lup; Avenhaus et al. 2018), which also have unconstrained or otherwise non-bell-shaped polarization curves. In such a scenario, results may be improved by utilizing a more appropriate polarization curve in place of Equation (4). This may be particularly relevant for observations at optical wavelengths, where fractional polarization curves can take more complex “rippled” shapes for some compositions (e.g., for the case of amorphous carbons in Tazaki & Dominik 2022).

Given these considerations, we would also recommend against interpreting the best-fit scattering surface parameters from optimization of PCRDI (inclination, etc.) as the “true” values unless the disk’s fractional polarization curve is well

constrained (and an appropriate model utilized in fitting) and its surface is consistent with the assumed scattering surface model. Moreover, unless both the scattering surface and the fractional polarization curve for a given disk are well constrained, utilizing the optimization method outlined in Section 3.3, rather than adopting literature parameters, will typically be more effective for eliminating oversubtraction. In cases where there may be differences in the flux calibration between the polarized and total intensity, the peak fractional polarization can be optimized to offset any impact this might otherwise have on PCRDI throughput.

The strictest requirement for use of PCRDI is simply that the disk be detected in PI within the optimization region to be used for total-intensity PSF subtraction. Specific limits with respect to disk morphology and orientation will depend significantly on the details of the system and the data. For example, because PCRDI scales the PI image, any noise will be significantly inflated in regions where the assumed fractional polarization is low (e.g., scattering angles far from 90° for a bell-shaped polarization curve). Though this noise will not directly enter the final PCRDI result,¹⁶ it may complicate optimization for disks with weak PI signal-to-noise ratio and low fractional polarization in the optimization region. We relegate a more complete investigation of any geometric and morphological limits related to PCRDI to a later study.

In the case of MCRDI, the sole prerequisite is that the underlying circumstellar signal can be superficially reproduced with synthetic disk models; there are no additional considerations related to the morphology or orientation of the disk beyond those relevant for typical forward-modeling procedures. As for PCRDI, we hazard against directly adopting the disk parameters that result from MCRDI optimization unless a fully physically motivated model has been used.

Finally, we note that our examples and analysis utilize small, near-contemporaneous sequences of reference images and target data with significant field rotation. Additional considerations may be motivated for applications to dissimilar data. When using extremely large PSF libraries spanning many nights (e.g., such as those used for ALICE; Choquet et al. 2014; Hagan et al. 2018), it may be possible to manifest features of the CSS estimate that are not truly present in the data from noise or artifacts in the reference data. 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) and by adopting frame-selection practices similar to those developed for ALICE (e.g., to remove reference images containing CSS or artifacts that might mimic circumstellar structure; Choquet et al. 2014; Hagan et al. 2018). Likewise, use of NMF (Ren et al. 2018) with constrained RDI (in place of LOCI or KLIP) may help to further insulate against false-positive CSS features for such data; because an erroneous feature in the CSS estimate would manifest as a negative imprint in the CSS-subtracted data for which the starlight model is constructed, reconstruction of such a feature from noise in reference images should be much less likely if negative reference coefficients are not permitted. Ultimately, the best cases for application of constrained RDI are those showing significant but badly oversubtracted CSS when using conventional PSF-subtraction techniques. If the conventional RDI result is effectively nulled, then any

significant CSS must apparently be well matched by the PSF model with the permitted model freedom, and so any features appearing following the application of constrained RDI may be unreliable. In such scenarios, the forward-modeling procedure of Appendix C could also be used to test against false-positive features resulting this way, e.g., by forward modeling a model lacking the feature and checking if said feature manifests in the result anyway due to the constraint and the freedom of the PSF model. We emphasize, however, that this possibility is purely speculative; this has not occurred for any of our testing thus far.

8. 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 (with both also allowing more accurate measurement of disk brightness and color than with model-based corrections). 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 planets embedded in disks, like AB Aur b (Currie et al. 2022b) and to conduct detailed studies of disk composition and morphology.

Proliferation of these techniques, or others that yield comparable results, is 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 m LUVOIR/HabEx “hybrid” and 30 m ground-based telescopes), it is more important still to ensure that available postprocessing 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 E).

We thank our referee, whose comments helped us to improve both the content and clarity of this manuscript.

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.

The development of SCEXAO was supported by the Japan Society for the Promotion of Science (Grant-in-Aid for Research #23340051, #26220704, #23103002, #19H00703, and #19H00695), the Astrobiology Center of the National Institutes of Natural Sciences, Japan, the Mt Cuba Foundation and the

¹⁶ As noted in Section 3, the CSS estimate is never directly used in the final result.

directors contingency fund at Subaru Telescope. We acknowledge funding support from the NASA XRP program via grants 80NSSC20K0252 and NNX17AF88G. T.C. was supported by a NASA Senior Postdoctoral Fellowship.

This work was directly enabled by Sigma Xi Grants In Aid of Research (GIAR), which provided support for the GPU-equipped workstation used in formulating the techniques presented here.

Software: 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), DiskDyn (Gaspar & Rieke 2020), WebbPSF (Perrin et al. 2014), CHARIS Data Reduction Pipeline (Brandt et al. 2017), CHARIS Data Processing Pipeline (Currie et al. 2020).

Appendix A Model-based RDI Corrections

RDI forward modeling for a single target exposure is carried out as follows. Let I_M be a synthetic disk model that has been rotated to the orientation of the target image and convolved with the PSF. Let the corresponding forward-modeled result be I'_M . For a conventional RDI procedure, I'_M can be found simply by substituting I_M in place of I_C in Equation (2):

$$I'_M = I_M - \mathcal{M}(I_M, \mathcal{R}). \quad (\text{A1})$$

Both I'_M and I_M are then derotated to align with north up. This process is repeated for each target image to create sequences of input (unattenuated) model images and processed (attenuated) model images matching the data sequence.

Typically, the input and processed model sequences are each averaged in the same manner as the real result to get final input and processed model images, whose ratio forms an attenuation correction map (e.g., Goebel et al. 2018; Lawson et al. 2021a; Betti et al. 2022)¹⁷ by which the RDI result is multiplied to produce a model-corrected result. For our purposes, we instead compute attenuation corrections for each exposure in the target sequence—dividing the full input model sequence by the output model sequence. In regions where the processed model image is very close to zero (e.g., when transitioning from positive to negative), large spurious values manifest in the corrections. To handle this, any correction factors with absolute values larger than 100 were replaced with a value of 1. This does not impinge on the nonspurious corrections (which are predominantly $\lesssim 5$) and significantly improves the accuracy of SB measurements for the corrected RDI results. Limited testing with other values between 10 and 500 showed no significant impact on the accuracy of the corrected results.

Once prepared, we multiply the PSF-subtracted and derotated data sequence with this correction sequence to form a sequence of model-corrected images. We then average this corrected sequence in the same manner as the uncorrected RDI result to reach the final model-corrected result. From testing with the simulated data of Section 4.1, we found that using exposure-by-exposure corrections—rather than a single averaged correction—improved the accuracy of the final results while also significantly reducing noise that otherwise remained in faint areas.

¹⁷ Some differences in the procedure are common to improve efficiency for data with a predominantly symmetric PSF.

Appendix B Surface-brightness Uncertainty

AB Aur and similar disks present a number of challenges when discussing signal-to-noise ratios and similar measurements of significance. Widely used methods for estimating the uncertainty of flux or SB measurements use the variance among the set of values in an image having the same stellocentric separation (e.g., Mawet et al. 2014; Lawson et al. 2021a). For disks that fill much of the FOV (or all of it, as is the case for the CHARIS AB Aur data), there is real and significant astrophysical variation in the values at a given separation that will inflate such noise estimates. Because our reductions have different amounts of astrophysical throughput, such noise estimates would be inflated by differing amounts, and so it is extremely challenging to make a fair comparison of significance between different reductions with these techniques.

To estimate the uncertainty of the SB and color profiles for our results, we adopt a strategy of statistical bootstrapping (Efron et al. 1994) as follows. The final image for a reduction is normally created by taking the average of the sequence of N_T target exposures (following PSF subtraction and derotation). Here, we instead form a “final” product by averaging a random selection, with replacement, of N_T images from the same sequence. We then carry out the relevant SB measurements on this result. This is repeated 10,000 times in each case, retaining each set of measurements. The uncertainty for a measurement at a given position is then measured from the distribution among the 10,000 bootstrapped results at that location. To ensure a fair comparison, each reduction being compared uses the same selection of frames for a given iteration. Because both PCRDI results (AB Aur and the simulated test case) use a number of exposures well in excess of the 50 samples recommended by Efron et al. 1994, bootstrapping should provide at least comparable (between reductions) estimates of uncertainty for these measurements. However, as there is a temporal correlation in any residual speckle noise, the variance between exposures for a specific location is not entirely uncorrelated, meaning that the statistics derived this way are not fully robust. Nevertheless, testing with an RDI reduction of data without a disk (otherwise using the same reference data used in the simulations of Section 4.1) showed that noise estimated this way is, on average, within a factor of ~ 1.2 of noise computed with the common radial method.

Appendix C Forward Modeling for Constrained RDI

For general use, formal forward modeling is not necessary for PCRDI products (and is mostly redundant for MCRDI); to assess geometry, the optimized PCRDI procedure typically provides sufficient throughput that products can simply be compared to PSF-convolved models. However, conventional forward-modeling and forward-model-based flux corrections can still be used. Let I_C^0 be the PI-based estimate used with the real data and let I_M be a PSF-convolved synthetic disk model to be forward modeled. The corresponding forward-modeled result can be computed as follows:

$$I'_M = I_M - \mathcal{M}(I_M - I_C^0, \mathcal{R}). \quad (\text{C1})$$

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




Figure 10. Top row, left to right: the standard RDI data result, model 1 (the ideal disk model) after RDI forward modeling, the residual map (data – processed model) for model 1, model 2 (the nonideal model) after forward modeling, and the residual map for model 2. Bottom row: as the top row, but for the nonoptimized PCRDI reduction (directly adopting the PI image as the constraint instead of attempting to transform it to total intensity). The small inset images show the corresponding model before forward modeling. The residual columns, which are shown at a different color scale than the other columns for visibility, show very nearly identical residuals between the two reductions.

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. Because attenuation corrections will also scale any residual starlight or noise (see Section 4.1), mitigating the majority of oversubtraction with PCRDI will enable higher signal to noise in corrected products than in corrected RDI products. Additionally, being less affected by biasing from CSS attenuation, the higher-throughput PCRDI products can be used to make much better estimates of disk parameters and so will likely require exploration of many fewer models overall.

As an example, we compare forward-modeled results for two reductions of the simulated data from Section 4.1: the standard RDI result (as in Figure 2: top row, center-left panel) and a nonoptimized PCRDI result (directly adopting the PI image as the CSS estimate). For these reductions, we apply forward modeling to the two disk models previously used for RDI attenuation corrections in Section 4.1. These results are visualized in Figure 10 and demonstrate that—even with the simplest constraint—PCRDI significantly improves throughput while still permitting the same tools for the interpretation of the results.

Appendix D

Validation Using Conventional RDI 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 the assessment of the geometry and properties of a real attenuated result; if a disk model can be identified that, 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 Equation (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 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). Figure 11 shows the result of this procedure. The negligible residual signal, in this case, provides additional evidence that our PCRDI optimization converged to a final result that is consistent with the CSS contained in the data.

It is important to note that forward modeling the unconstrained RDI result in the same manner will also result in similarly small residuals; a forward-modeled image being consistent with the RDI result only indicates that the difference between the input image and the RDI result is well approximated by some combination of reference images (or reference image eigenvectors, in the case of KLIP). In other words, a good agreement between a given forward-modeled result and the RDI result can only indicate that the input image is a suitable solution (rather than the suitable solution).¹⁸ In general, this test will not provide any better assessment of a

¹⁸ For any nonazimuthally symmetric PSF for data with nonnegligible field rotation, the forward-modeled constrained RDI result should yield smaller residuals; an oversubtracted result will have the imprint of the derotated and sequence-averaged PSF, which will not be properly aligned with the PSF when rotating the “model” to the parallactic angles of the data for forward modeling. In testing, the improvement was small but nonnegligible.




Figure 11. Left: the conventional RDI result for AB Aur (as in Figure 8, top left). Center: the PCRDI result of Section 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 8).

constrained RDI result than the objective function calculation used in constrained RDI optimization (Section 3.3), but may still be useful, e.g., for demonstrating that constrained RDI was implemented correctly.

Appendix E Supporting Software and Runtime

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

The public release of the Python software to support the 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.

E.1. Runtime

To carry out the 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 @ 3.70 GHz \times 20 GPU: NVIDIA GeForce RTX 3070.


The dimensions and runtimes for each data set utilized are provided in Table 1. Note: the number of function evaluations for optimization is generally detached from the disk complexity for PCRDI, but is strongly linked for MCRDI (e.g., with a multiring system requiring optimization of more parameters).

Table 1
PCRDI and MCRDI Data-set Dimensionality and Optimization Runtimes

| Data Set | N_{tar} | N_{ref} | Exposure Dimensions | Opt. N_{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 data set utilized in this work, the dimensions and total time required to reach the final constrained RDI product. N_{tar} and N_{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_{fev} ” and “Opt. time” indicate the number of function evaluations and total runtime (in seconds) required to reach the presented solution, respectively.

ORCID iDs

Kellen Lawson  <https://orcid.org/0000-0002-6964-8732>
 Thayne Currie  <https://orcid.org/0000-0002-7405-3119>
 John P. Wisniewski  <https://orcid.org/0000-0001-9209-1808>
 Tyler D. Groff  <https://orcid.org/0000-0001-5978-3247>
 Michael W. McElwain  <https://orcid.org/0000-0003-0241-8956>
 Joshua E. Schlieder  <https://orcid.org/0000-0001-5347-7062>

References

- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Augereau, J. C., Lagrange, A. M., Mouillet, D., Papaloizou, J. C. B., & Grorod, P. A. 1999, *A&A*, **348**, 557
- Avenhaus, H., Quanz, S. P., Garufi, A., et al. 2018, *ApJ*, **863**, 44
- Benisty, M., Dominik, C., Follette, K., et al. 2022, arXiv:2203.09991
- Betti, S. K., Follette, K., Jorquera, S., et al. 2022, *AJ*, **163**, 145
- Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, **631**, A155
- Bhowmik, T., Boccaletti, A., Thébault, P., et al. 2019, *A&A*, **630**, A85
- Boccaletti, A., Di Folco, E., Pantin, E., et al. 2020, *A&A*, **637**, L5
- Brandt, T. D., Rizzo, M., Groff, T., et al. 2017, *JATIS*, **3**, 048002
- 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
- Choquet, É., Pueyo, L., Hagan, J. B., et al. 2014, *Proc. SPIE*, **9143**, 914357
- Currie, T., Biller, B., Lagrange, A.-M., et al. 2022a, arXiv:2205.05696
- Currie, T., Cloutier, R., Brittain, S., et al. 2015, *ApJ*, **814**, L27
- Currie, T., Debes, J., Rodigas, T. J., et al. 2012, *ApJL*, **760**, L32
- Currie, T., Marois, C., Cieza, L., et al. 2019, *ApJL*, **877**, L3
- Currie, T., Guyon, O., Lozi, J., et al. 2020, *Proc. SPIE*, **11448**, 114487H
- Currie, T., Lawson, K., Schneider, G., et al. 2022b, *NatAs*, **6**, 751
- Efron, B., Tibshirani, R., & Tibshirani, R. J. 1994, *An Introduction to the Bootstrap* (Philadelphia, PA: Chapman and Hall), doi:10.1201/9780429246593
- Esposito, T. M., Kalas, P., Fitzgerald, M. P., et al. 2020, *AJ*, **160**, 24
- Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2018, *A&A*, **618**, A138
- Flasseur, O., Thé, S., Denis, L., Thiébaud, É., & Langlois, M. 2021, *A&A*, **651**, A62
- Fukagawa, M., Hayashi, M., Tamura, M., et al. 2004, *ApJL*, **605**, L53
- 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. 2021, *A&A*, **649**, A1
- Garufi, A., Dominik, C., Ginski, C., et al. 2022, *A&A*, **658**, A137
- Gaspar, A., & Rieke, G. 2020, *PNAS*, **117**, 9712
- Goebel, S., Currie, T., Guyon, O., et al. 2018, *AJ*, **156**, 279
- Gomez Gonzalez, C. A., Wertz, O., Absil, O., et al. 2017, *AJ*, **154**, 7
- Grady, C. A., Woodgate, B., Bruhweiler, F. C., et al. 1999, *ApJL*, **523**, L151
- Groff, T. D., Chilcote, J., Kasdin, N. J., et al. 2016, *Proc. SPIE*, **9908**, 990800
- Hagan, J. B., Choquet, É., Soummer, R., & Vigan, A. 2018, *AJ*, **155**, 179
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Hashimoto, J., Tamura, M., Muto, T., et al. 2011, *ApJL*, **729**, L17
- Hong, S. S. 1985, *A&A*, **146**, 67
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, **127**, 890
- Kenyon, S. J., Gómez, M., & Whitney, B. A. 2008, in *Handbook of Star Forming Regions, Volume I: The Northern Sky*, ed. B. Reipurth, Vol. 4 (San Francisco, CA: ASP), 405
- Kepler, M., Benisty, M., Müller, A., et al. 2018, *A&A*, **617**, A44
- Kuhn, J. R., Potter, D., & Parise, B. 2001, *ApJL*, **553**, L189
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, **660**, 770
- Lawson, K., Currie, T., Wisniewski, J. P., et al. 2020, *AJ*, **160**, 163
- Lawson, K., Currie, T., Wisniewski, J. P., et al. 2021a, *AJ*, **162**, 293
- Lawson, K., Currie, T., Wisniewski, J. P., et al. 2021b, *Proc. SPIE*, **11823**, 118230D
- Lomax, J. R., Wisniewski, J. P., Grady, C. A., et al. 2016, *ApJ*, **828**, 2
- Lovell, J. B., Marino, S., Wyatt, M. C., et al. 2021, *MNRAS*, **506**, 1978
- Macintosh, B., Graham, J. R., Barman, T., et al. 2015, *Sci*, **350**, 64
- Males, J. R., Close, L. M., Miller, K., et al. 2018, *Proc. SPIE*, **10703**, 1070309
- Marois, C., Correia, C., Galicher, R., et al. 2014, *Proc. SPIE*, **9148**, 91480U
- Marois, C., Macintosh, B., & Véran, J.-P. 2010, *Proc. SPIE*, **7736**, 77361J
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, **792**, 97
- Mazoyer, J., Arriaga, P., Hom, J., et al. 2020, *Proc. SPIE*, **11447**, 1144759
- 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, A. 2017, in *A NumPy-compatible Library for NVIDIA GPU Calculations. Proceedings of Workshop on Machine Learning Systems (LearningSys) in the Thirty-first Annual Conference on Neural Information Processing Systems*, http://learningsys.org/nips17/assets/papers/paper_16.pdf
- Pairet, B., Cantalloube, F., & Jacques, L. 2021, *MNRAS*, **503**, 3724
- Perrin, M. D., Sivaramakrishnan, A., Lajoie, C.-P., et al. 2014, *Proc. SPIE*, **9143**, 91433X
- Pueyo, L. 2016, *ApJ*, **824**, 117
- Ren, B., Pueyo, L., Chen, C., et al. 2020, *ApJ*, **892**, 74
- Ren, B., Pueyo, L., Zhu, G. B., Debes, J., & Duchêne, G. 2018, *ApJ*, **852**, 104
- Rich, E. A., Monnier, J. D., Aarnio, A., et al. 2022, arXiv:2206.05815
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJL*, **755**, L28
- Stolker, T., Dominik, C., Min, M., et al. 2016, *A&A*, **596**, A70
- Takami, M., Karr, J. L., Hashimoto, J., et al. 2013, *ApJ*, **772**, 145
- Tazaki, R., & Dominik, C. 2022, *A&A*, **663**, A57
- Thalmann, C., Janson, M., Buenzli, E., et al. 2013, *ApJL*, **763**, L29
- van der Marel, N., Birnstiel, T., Garufi, A., et al. 2021, *AJ*, **161**, 33
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Methods*, **17**, 261
- Whitney, B. A., Robitaille, T. P., Bjorkman, J. E., et al. 2013, *ApJS*, **207**, 30