A Chromaticity Analysis and PSF Subtraction Techniques for SCExAO/CHARIS Data

Benjamin L. Gerard[1,2] , Christian Marois[1,2] , Thayne Currie[3,4,5] , Timothy Brandt[6] , Jeffrey K. Chilcote[7] ,
Zachary H. Draper[1,2] , Tyler Groff[8] , Olivier Guyon[9,10] , Masahiko Hayashi[11], Nemanja Jovanovic[12], Gillian R. Knapp[13] ,
Tomoyuki Kudo[4] , Jungmi Kwon[14] , Julien Lozi[4] , Frantz Martinache[15] , Michael McElwain[8] ,
Motohide Tamura[10,16] , and Taichi Uyama[16]

[1] University of Victoria, Department of Physics and Astronomy, 3800 Finnerty Road, Victoria, V8P 5C2, Canada; bgerard@uvic.ca
[2] National Research Council of Canada, Astronomy & Astrophysics Program, 5071 West Saanich Road, Victoria, V9E 2E7, Canada
[3] NASA-Ames Research Center, Moffett Field, CA, USA
[4] Subaru Telescope, National Astronomical Observatory of Japan, 650 North A'ohōkū Place, Hilo, HI 96720, USA
[5] Eureka Scientific, 2452 Delmer Street Suite 100, Oakland, CA, USA
[6] Department of Physics, University of California-Santa Barbara, Santa Barbara, CA, USA
[7] Department of Physics, University of Notre Dame, 225 Nieuwland Science Hall, Notre Dame, IN, 46556, USA
[8] NASA-Goddard Space Flight Center, Greenbelt, MD, USA
[9] Steward Observatory, University of Arizona, Tucson, AZ 85721, USA
[10] Astrobiology Center, National Institutes of Natural Sciences, 2-21-1 Osawa, Mitaka, Tokyo, Japan
[11] NAOJ, 2-21-1 Osawa, Mitaka, Tokyo 181-8588, Japan
[12] Department of Astronomy, California Institute of Technology, 1200 E. California Boulevard, Pasadena, CA 91125, USA
[13] Department of Astrophysical Science, Princeton University, Peyton Hall, Ivy Lane, Princeton, NJ 08544, USA
[14] ISAS/JAXA, 3-1-1 Yoshinodai, Chuo-ku, Sagamihara, Kanagawa 252-5210, Japan
[15] Université Côte d'Azur, Observatoire de la Côte d'Azur, CNRS, Laboratoire Lagrange, France
[16] Department of Astronomy, Graduate School of Science, The University of Tokyo, 7-3-1, Hongo, Bunkyo-ku, Tokyo, 113-0033, Japan

Abstract

We present an analysis of instrument performance using new observations taken with the Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS) instrument and the Subaru Coronagraphic Extreme Adaptive Optics (SCExAO) system. In a correlation analysis of our data sets (which use the broadband mode covering the J band through the K band in a single spectrum), we find that chromaticity in the SCExAO/CHARIS system is generally worse than temporal stability. We also develop a point-spread function (PSF) subtraction pipeline optimized for the CHARIS broadband mode, including a forward modeling-based exoplanet algorithmic throughput correction scheme. We then present contrast curves using this newly developed pipeline. An analogous subtraction of the same data sets using only the H-band slices yields the same final contrasts as the full JHK sequences; this result is consistent with our chromaticity analysis, illustrating that PSF subtraction using spectral differential imaging (SDI) in this broadband mode is generally not more effective than SDI in the individual J, H, or K bands. In the future, the data processing framework and analysis developed in this paper will be important to consider for additional SCExAO/CHARIS broadband observations and other ExAO instruments which plan to implement a similar integral field spectrograph broadband mode.

Key words: planets and satellites: general – techniques: high angular resolution – techniques: image processing

1. Introduction

Ground-based exoplanet imaging is currently uniquely poised to understand how giant planets form and evolve over time. Distinct from other exoplanet detection methods, exoplanet imaging is currently sensitive to wide separations (greater than about 10 au), young ages (less than about 1 Gyr), and high masses (greater than about one Jupiter mass; e.g., Nielsen et al. 2019). However, the push to closer separations, older ages, and lower masses is constantly improving with new instrumentation and processing algorithms.

Post-processing of high-contrast imaging data is currently a crucial component of reaching the necessary contrasts in order to detect new exoplanets by subtracting the correlated speckle noise around their host star. There are three main point-spread function (PSF) subtraction observing strategies used to improve the exoplanet signal-to-noise ratio (S/N): spectral differential imaging (SDI; Rosenthal et al. 1996; Racine et al. 1999; Marois et al. 2000; Sparks & Ford 2002), angular differential imaging (ADI; Marois et al. 2006a), and the commonly used reference differential imaging (RDI) algorithm. Although these strategies are often combined with more advanced least-squares

subtraction algorithms (e.g., Lafrenière et al. 2007; Soummer et al. 2012), all three methods remain fundamentally limited by PSF chromaticity (SDI and RDI) and/or temporal stability (ADI and RDI). Thus, understanding and mitigating the level of chromatic and temporal decorrelation in high-contrast imaging data sets is a crucial step toward improving contrast to detect fainter and/or lower mass exoplanets.

Although significant efforts have been made to understand the relative impact of temporal stability versus chromaticity from a numerical perspective (e.g., Marois et al. 2008; Vigan et al. 2010), a similar comprehensive on-sky analysis with second generation high-contrast imaging instruments that are equipped with an integral field spectrograph (IFS) has not yet been completed. Hinkley et al. (2007) and Milli et al. (2016) have carried out temporal analyses on the quasi-static speckle lifetime, but with no corresponding chromaticity analysis. Following from the work of Marois et al. (2010), temporal correlation analyses have informed some new ADI+SDI PSF subtraction pipelines to select reference images ranked by correlation (e.g., Currie et al. 2012; Wang et al. 2015). Marois et al. (2014) presented an analyses of time-dependent versus

Table 1
Table of Our Target Observing Parameters

UT Date (yy mm dd)	Name (HD)	m_V	t_{exp} (s)	n_{frames}	Δ_{PA} (deg)
2017 Apr 9	70313	5.51	15	107	25
2017 Apr 9	98673	6.43	31	91	22
2017 Apr 9	109085	4.3	13	93	18
2017 Apr 9	125162	4.18	20	93	21
2017 Apr 9	141378	5.53	31	70	21
2018 Jan 7	23267	6.88	60	50	32
2018 Jan 8	" "	" "	" "	40	24
2018 Jan 8	38206	5.73	31	118	27

wavelength-dependent correlation in individual target images. In this paper we will extend this analysis to multiple target images; later we will show that such correlation can vary significantly depending on observing conditions and on location in the coronagraphic image.

The main focus of this paper is thus to present a detailed analysis of chromaticity and temporal stability using our observations from the Subaru adaptive optics system (AO188; Minowa et al. 2010), Subaru Coronagraphic Extreme Adaptive Optics (SCExAO; Jovanovic et al. 2015b), and Coronagraphic High Angular Resolution Imaging Spectrograph (CHARIS; Groff et al. 2016, 2017). We also present new modifications to the ADI+SDI-based template locally optimized combination of images (TLOCI; Marois et al. 2014) and PSF subtraction pipeline for these data sets. Our SCExAO/CHARIS data sets use a Lyot coronagraph in the broadband *JHK* mode, covering the full 1.15–2.39 μm range at $R \sim 18$ (Groff et al. 2016); this new low-resolution broadband mode, currently unique to CHARIS among current high-contrast imaging instruments, is thought to provide a potential gain in achievable contrast through SDI PSF subtraction for an achromatic system (Marois et al. 2008, 2014), and so our performance analysis and PSF subtractions methods will focus on the impacts of using this broadband mode.

The structure of this paper is as follows. In Section 2 we present details of our target observations, in Section 3 we present the architecture of our PSF subtraction pipeline (Section 3.2) and then use the same reference image selection and correlation calculation procedures to carry out a detailed chromaticity analysis (Section 3.3), in Section 4 we present final contrast curves and discuss instrument performance implications for the future, and we conclude in Section 5. An additional correlation analysis is presented in the Appendix.

2. Observations

We observed a handful of young, nearby two temperature-component debris disks, the extrasolar analogs of the asteroid and Kuiper belts in our solar system (Wyatt 2008). We identified targets in Kennedy & Wyatt (2014) that had not yet been observed by the Gemini Planet Imager (GPI; Macintosh et al. 2014) or Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE; Beuzit et al. 2008) and/or are inaccessible to either for being too far north. We also observed two additional targets in 2018A through ancillary time (Hodapp 2018). Our observed targets are listed in Table 1, including dates of observations, exposure times, and field-of-view (FOV) rotation.

Observations were made in the standard pupil tracking mode for ADI. The coronagraphic setup in this broadband mode utilized the standard Lyot coronagraph, applying the 217 mas diameter focal plane mask (FPM), for which the inner working angle (IWA; separation at which an off-axis companion would have 50% optical throughput relative to the center of the star) is 113 mas; at the time, no other coronagraph options offered sufficient performance over the full broadband mode of CHARIS. The deformable mirror waffle spots (Marois et al. 2006b; Sivaramakrishnan & Oppenheimer 2006; Jovanovic et al. 2015a), or satellite spots, are applied with a nonstandard amplitude for the 2017 April data, and so instead an on-sky noncoronagraphic data sequence (using a neutral density filter but with the same Lyot stop) is taken to measure the star-to-spot ratio directly (see Section 3.1). The 2018 January data uses the standard 50 nm amplitude spots and corresponding precomputed star-to-spot ratio as in Currie et al. (2018b).

3. Data Analysis

3.1. Data Cube Extraction and Setup

We used the CHARIS data reduction pipeline (DRP; Brandt et al. 2017) to generate data cubes from recorded IFS detector images, which uses a least-squares method to extract the flux in each slice of the data cube (Draper et al. 2014). A background subtraction, flat-field correction, and bad pixel mask are also applied in the extraction process as described in Brandt et al. (2017). All images are then prepared for subsequent PSF subtraction (Section 3.2) and correlation analysis (Section 3.3). The satellite spots are used for image registration, magnification with wavelength to align speckles, flux normalization to remove the wavelength dependence of the broadband transmission filter and stellar spectrum, and contrast calibration. We chose to magnify the images with a wavelength centered around slice 16 of 22 ($\lambda_0 = 2.0$ μm). Slice 16 was chosen instead of slice 11 ($\lambda = 1.7$ μm; in the middle of the bandpass) because of wavelength-dependent sampling issues. The CHARIS PSF in the *J* and *H* bands are slightly under-sampled, and we found that cubic spline interpolation, used to rotate and magnify the images for our ADI+SDI PSF subtraction algorithm (Section 3.2), provided a more accurate centroiding precision from numerical interpolation in the *K* band. The four satellite spots all lie at 15.91 λ_0/D separation, and so measured satellite spot positions in the registered, magnified images are used to calculate a plate scale for contrast curve measurements (Section 4). We were not able to use any on-sky astrometric data to validate these derived plate scale values. Images are then high-pass filtered using a 4.1 $\times$ 4.1 λ_0/D median boxcar filter to attenuate the smooth adaptive optics (AO) halo while retaining a close-to-unity algorithmic throughput of λ/D-scale static speckles and any possible planet(s). The λ^2 dependence of the satellite spot intensity is also removed during the flux normalization procedure. The satellite spots are omitted from reference image selection in Section 3.2 and correlation calculations for analysis in Section 3.3. Finally, all registered, magnified, flux-normalized, high-pass-filtered images are converted into raw contrast units (i.e., the pixel value-to-star ratio throughout the whole image) using measured and precalibrated star-to-satellite spot ratios, whose data acquisition is described at the end of Section 2. The same registration, magnification, flux normalization, and high-pass filtering
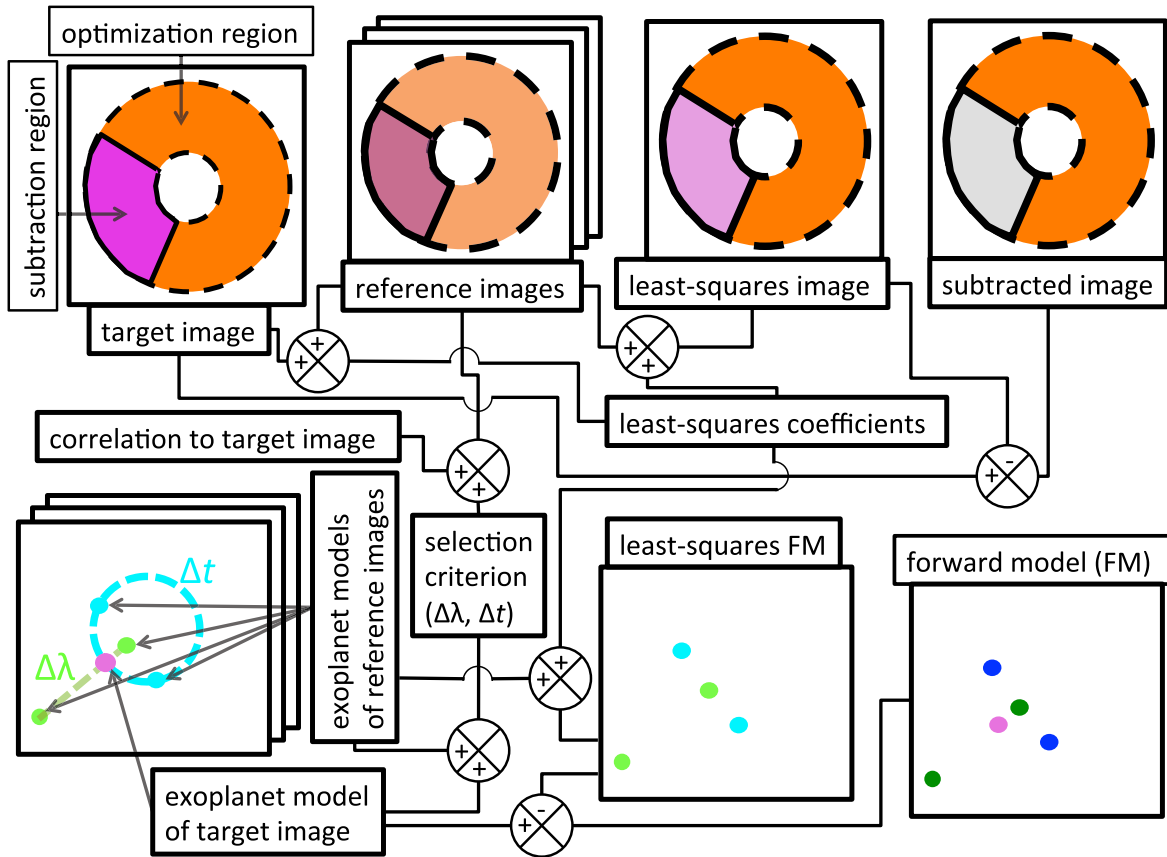
Figure 1. Block diagram illustrating the process of least-squares-based PSF subtraction and forward modeling. The junctions combine the two boxes (connected to a "+" or "-") and output to a third box. The junctions with a "-" sign indicate that the output box is the difference between the two images represented by input boxes. All other boxes, where the two junctions are "+" signs, represent a more abstract process of combining two input concepts (i.e., an image, a vector of images, a vector of coefficients, and/or a selection criterion) to produce an output concept. The inputs, on the lower left of the diagram, include a model of both the exoplanet target image and all other reference images in the sequence, both in time and wavelength, including the assumed spectrum of the planet. These models are then used to define a selection criterion, which is combined with the correlation of each reference image in the sequence to the target image in order to select the best reference images to use in subtracting the target image. After a few intermediate steps involving the calculation of least-squares coefficients, the outputs are illustrated on the right side of the diagram, including a subtracted image along with a noiseless forward model of the planet signal in the subtracted image.

procedure is applied to the custom 2017 April star-to-spot calibration data before computing the ratios.

3.2. PSF Subtraction

Our ADI+SDI PSF subtraction algorithms build on the least-squares-based locally optimized combination of images (LOCI; Lafrenière et al. 2007) algorithm, further developed by Marois et al. (2010) and Marois et al. (2014) as template LOCI (TLOCI). A conceptual illustration of the procedure used in our TLOCI pipeline to subtract a single region of a single target image is shown in Figure 1.

In addition to the process of generating least-squares coefficients to subtract a region of the target image, Figure 1 also illustrates the principle of forward modeling (FM), where the same coefficients are also used to model the algorithmic exoplanet throughput during the subtraction. As shown, a selection criterion, or aggressiveness, specifies that reference images be selected beyond a radial and/or azimuthal gap in wavelength and/or time, respectively, from the target image in order to minimize algorithmic exoplanet throughput losses. The highest aggressiveness of one (using dimensionless units calculated from relative aperture photometry between the target and reference image models) would allow the use of any reference image in the full target sequence, including the target

image. In contrast, the lowest aggressiveness of zero would not allow any reference images in the target sequence to be selected. Aggressiveness is also illustrated in Figure 2. In this figure, a range of times and wavelengths around the target image are not selectable because, if used, a planet present in the target image would be partially subtracted by the same planet in the reference images. This unselectable region of Figure 2, illustrated by a white color, would be either larger or smaller for a lower or higher value of user-defined aggressiveness, respectively. The size of this unselectable region also depends on the separation in the image, assumed exoplanet spectrum, wavelength, and parallactic angle relative to the target image. The target image is highlighted in the center of this region with a correlation of one. Then, reference images are chosen from the remaining sample based on correlation, also illustrated in Figure 2. Generally, the most correlated reference images are near the target image time or wavelength, although this depends on the temporal stability and chromaticity of the sequence. A more correlated set of reference images (i.e., those with a more similar speckle pattern to the target image) provides a better basis for the least-squares algorithm and will yield a deeper final contrast without changing any of the other free parameters (Marois et al. 2014). We will examine the chromatic and temporal dependence of these correlations
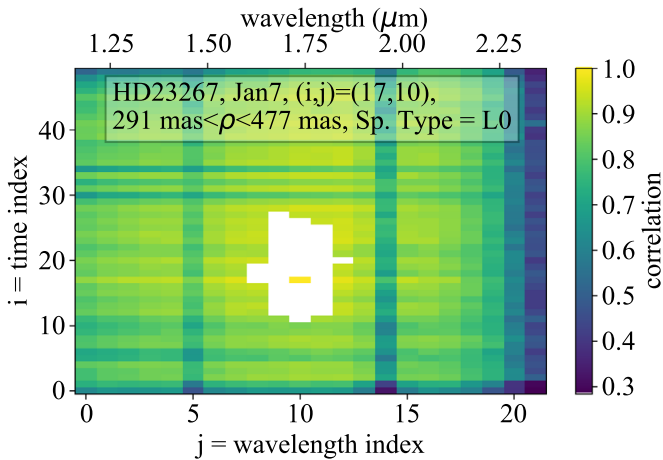
Figure 2. Correlation of selectable reference images to a single target image. The target image is illustrated with a correlation of one in the center of the white region. Various correlation trends illustrate the stability of the sequence in time and wavelength. Darker vertical lines illustrate wavelength-dependent effects, including the atmospheric water bands at wavelength slices 5 and 14, a degradation of performance toward the red edge of the K band, while darker horizontal lines illustrate time-dependent effects from variable AO and/or quasi-static stability.

further in Section 3.3. Additional details are described below about the setup and new modifications of our TLOCI pipeline.

To help define, identify, and remove uncorrelated images from each data set we will introduce a new free parameter into our PSF subtraction pipeline called a correlation cut. For a single target image, the value of a correlation cut will define the minimum level of correlation that is rendered acceptable in a reference image used to subtract the target image speckle pattern. A given target image/region will be discarded from the subtracted sequence if there are not at least N_{ref} reference images above the correlation cut value, all of which must pass a selection criterion, where N_{ref} and aggressiveness are also free parameters. For example, if correlation cut $= 0.5$ and $N_{ref} = 10$ but there are only five reference images that are greater than 50% correlated to the target image using a moderate aggressiveness, the target image/region will be discarded from the sequence. The target image would be discarded in this case even if, e.g., one of those five reference images is 99% correlated to the target image. Thus, it is clearly important to use a physically motivated correlation cut so that a maximal number of target images in an observing sequence can be subtracted while still rejecting a maximal number of uncorrelated reference images. Marois et al. (2010) and subsequent papers have shown that using correlated reference images is crucial to prevent unnecessary noise propagation in the covariance matrix inversion step of a least-squares/ principal component analysis subtraction algorithm. The correlation cut values are dimensionless by definition and have a minimum and maximum of zero and one, respectively. Analysis in the Appendix inform the decision to use the following correlation cut values for the four different radial sections of an image, progressing from the innermost to the outermost separation: 0.8, 0.74, 0.67, and 0.59. Although the specific angular separations of these regions varies for each target based on our method of measuring the pixel scale using the satellite spots (Section 3.1), the average separation and standard deviation across all targets are: 104 ± 2 mas $< \rho < 286 \pm 4$ mas,

286 ± 4 mas $< \rho < 468 \pm 7$ mas, 468 ± 7 mas $< \rho < 649 \pm 10$ mas, and $\rho > 649 \pm 10$ mas. A graphical illustration of these zones is illustrated later in Section 3.2. Note that we chose to fix the radial separation of subtraction regions based on pixel separation, which is why the derived angular separation varies slightly from target to target. In pixels, the radial separations for each annulus for all targets are $7 < \rho < 19$, $19 < \rho < 32$, $32 < \rho < 44$, and $\rho > 44$.

Based on the definitions and work in Marois et al. (2010) and Johnson-Groh (2016), for all subsequent analyses in this paper we use an L0 or T6 input spectrum (J02281101+2537380 and J02281101+2537380 from Burgasser et al. 2004, 2008, respectively, subsequently convolved and binned to the appropriate CHARIS spectral resolution), an aggressiveness of 0.5, and 15 reference images per least-squares subtraction. Gerard & Marois (2016) found, on GPI data sets, that optimizing aggressiveness and number of reference images yielded a negligible gain in contrast, and so here we use the same aggressiveness, number of reference images per subtraction, and input spectral types (to define a selection criterion and for subsequent FM). We use a singular value decomposition (SVD) cutoff value for the covariance matrix inversion of 1×10^{-4}, as in Gerard & Marois (2016), where a maximum value of one would set all the LOCI coefficients to zero. Although this is a low value compared to other algorithms (e.g., Currie et al. 2018a), note that there are many degeneracies between the number of reference images used, reference image correlation, and SVD cutoff (Marois et al. 2010); our approach here uses a lower SVD cutoff but compensates by selecting more correlated reference images.

In a new addition to our TLOCI pipeline, the central region of each subtraction region is defined by both the parallactic angle and wavelength at any given timestamp, and so the location of a single region will both rotate with time and magnify with wavelength over a single observing sequence, similar to the approach in Wang et al. (2015). These definitions are used so that the subtraction regions line up after derotation and de-magnification, which allows the generation of an FM map for each target. The LOCI optimization regions are defined for each subtraction region as the other two azimuthal regions at the same radial separation. This subtraction/optimization region geometry ensures that the least-squares coefficients are never generated from the subtraction zone where a planet would be. S/N maps are calculated using back-rotated images, as in Gerard & Marois (2016). A few example data products of a single target sequence from our PSF subtraction pipeline are shown in Figure 3, including a subtracted collapsed data cube, corresponding S/N map, and corresponding FM map.

For each target we also run a bootstrapping procedure, illustrated in Figures 4 and 5. As in Marois et al. (2010, 2014), we implement an FM algorithmic throughput correction procedure in each subtraction region of the image, including a model for the radial variation of the algorithmic throughput within each region. In addition to the center, four additional FM throughput corrections using the same LOCI coefficients are computed at smaller and larger separations within each subtraction region. A fifth-order polynomial is then fit to the radial variation of throughput correction values to generate a continuous function across the region. As a result of this throughput correction procedure, a bootstrapped planet at any location in the image (i.e., not necessarily at the radial or azimuthal center of a subtraction region), is extracted at the
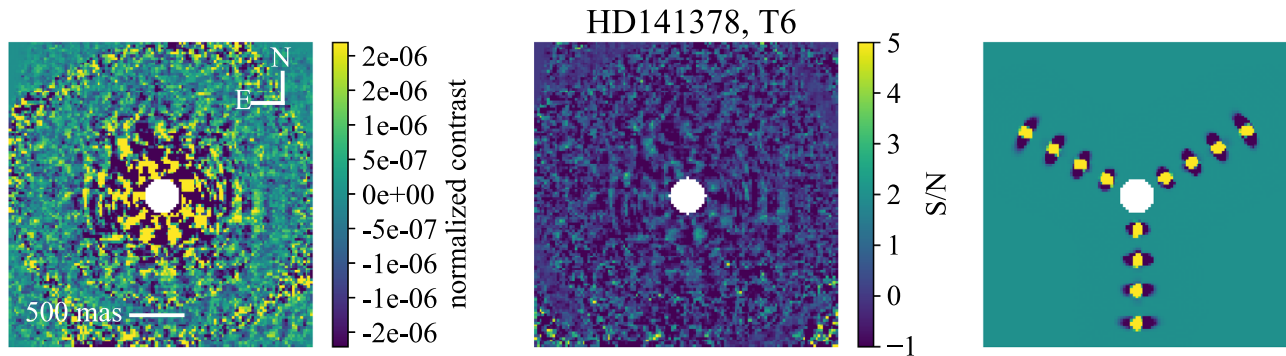
Figure 3. PSF subtracted data products for a single target, HD 141378. Left: an image of a collapsed PSF subtracted sequence. Middle: an S/N map of the left image. Right: a collapsed forward model map. Note the relative absence of dark shadows in the radial relative to azimuthal direction. This suggests that images at the same wavelength and different time are more correlated than images at a different time and same wavelength, causing the former to be selected as reference images more often than the latter.

input value and does not need any additional throughput correction. This is illustrated in Figure 5. The simulated planets were injected throughout the image, one per subtraction region, such that none lie at the radial or azimuthal center of the region, as in the left panel of Figure 4. Almost all recovered fluxes match the input value within their expected scatter, materialized by the 5σ contrast value at the location of the planet.

3.3. Chromaticity Analysis

As discussed in Section 1, in a perfectly achromatic system, the larger lever arm of wavelength coverage over the *JHK* filter of CHARIS should provide a better algorithmic exoplanet throughput than a narrower bandpass, and so if both cases provide the same level of speckle subtraction, the final throughput-corrected contrast should be better in the former case. However, in reality broadband systems are not achromatic, and the level of wavefront decorrelation, or chromaticity (i.e., separate from wavelength magnification), will ultimately determine the limits of speckle subtraction via SDI processing (e.g., Marois et al. 2008). In this section we will measure the chromaticity in our CHARIS data sets, which will then inform the SDI speckle subtraction limits in Section 4.

A generalization of Figure 2 extends to Figure 6, which shows the median correlation of all of our targets as a function of time and wavelength. Figure 6 is generated as follows.

1. For a single target image at a single time, wavelength, user-defined separation, spectral type, and aggressiveness, a selection criterion is applied to all of the images in the observing sequence; images that do not meet this criterion are discarded.
2. Correlation of the target image with every remaining image from step 1 is computed and then normalized to unity using the standard deviation in both the target image and reference images. Steps 1–2 are illustrated in Figure 2.
3. In order to create a common zero-point for every subsequent target image (i.e., so the central index of Figure 6 is zero for every target image), each correlation value from step 2 is saved in a two-dimensional array of $2i_t \times 2j_t$ in size, where i_t and j_t are the total number of time stamps and wavelength slices, respectively (i.e., so both the negative and positive indices can be shown in Figure 6). The x and y indices that fill a quarter of the total area of the two-dimensional array for each

correlation value from step 2 are $(2i_t - i_{\text{target}})$ and $(2j_t - j_{\text{target}})$, respectively, where i_{target} and j_{target} are the indices of time and wavelength for the target image. This index definition ensures that the target image is in the center, or zero-point, of each $2i_t \times 2j_t$ array. Both the region removed by the selection criterion and the other three quarters of the array remain empty and are filled with "not a number" (NaN) placeholders.
4. Steps 1–3 are repeated over time and wavelength for every target image within a single observing sequence, generating a vector of two-dimensional arrays of $i_t j_t \times 2i_t \times 2j_t$ in size. The autocorrelation of each target image is in the center of each $2i_t \times 2j_t$ array.
5. A median is taken across the zeroth axis of the vector in step 4 (i.e., the axis of dimension $i_t j_t$), ignoring NaN values, to generate a two-dimensional array of $2i_t \times 2j_t$ in size, representing the median correlation between the target image and every other image in the sequence. A slice along the x direction of this array at $y = j_t$ (the zero-point of the wavelength direction) reveals the median correlation of the target image (at $x = i_t$) versus time, whereas a slice along the y direction of this array at $x = i_t$ reveals the median correlation versus wavelength.
6. Steps 1–5 are repeated over each user-defined spectral type.
7. Steps 1–6 are repeated over each user-defined separation.
8. Steps 1–7 are repeated for each target, computing the median and standard deviation of correlation over time and wavelength for each separation.

The outcome of the above procedure is illustrated in Figure 6, whereas steps 1–7 are illustrated in Figure 7, for which we used a separate data set from our sample, the κ And data from Currie et al. (2018a). Rather than a single target in our sample, the Currie et al. (2018a) data set was chosen for Figure 7 to understand how the conclusions from Figure 6 change with higher quality AO correction.

The main empirical conclusion from Figures 6 and 7 is that SCExAO/CHARIS data sets are generally more correlated as a function of time than as a function of wavelength, most notably at small and large separations. This is particularly important at small separations, where PSF subtraction algorithms generally underperform compared to the rest of the image. For a given difference in parallactic angle and/or wavelength between a target and reference image, the respective azimuthal and/or radial movement of an exoplanet in between the two images is
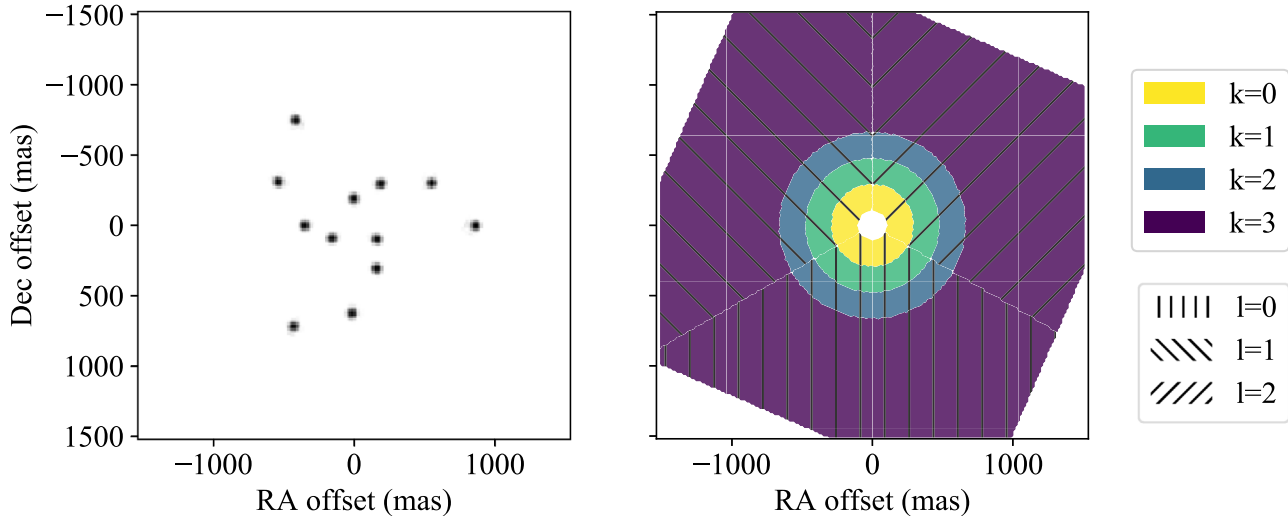
Figure 4. Right: definitions of the different subtraction regions defined in the final collapsed image. "k" and "l" represent the radial and azimuthal indices for each subtraction region, respectively. Left: the locations of simulated planets (one per subtraction).

larger at larger separations. In other words, for a given aggressiveness, more reference images near the target image, both in time and wavelength, have to be discarded at lower separations compared to higher separations. Although a longer wavelength lever arm would help overcome this problem, Figures 6 and 7 illustrate that this would generally add unnecessary noise in the covariance matrix inversion. A more correlated image at the same wavelength and different time will be selected instead. Later in Section 4 we will test the predictions from Figure 6 on our PSF subtraction pipeline.

With that said, at small separations in Figure 6 and at all separations in Figure 7, reference images are actually more correlated at the closest wavelengths than at the closest times. Although these images will be preferentially selected for SDI processing, this approach will not necessarily reach the deepest (throughput-corrected) final contrasts. Selecting a closer wavelength will reach a better nonthroughput-corrected contrast but will also cause more self-subtraction, therefore requiring a larger throughput correction. Thus, the optimal reference selection based on final throughput-corrected contrast will ultimately depend on optimizing the trade-off between correlation and self-subtraction, as in Gerard & Marois (2016). Additionally, although the chromaticity is still worse than the temporal stability in Figure 7, note that there is a significant improvement in both temporal stability and chromaticity between Figures 6 and 7. Although the improvement from temporal stability is most likely from improved AO performance (strehl ratios for the Currie et al. 2018a data set are around 0.92), the origin of chromaticity improvement is not clear and will require further investigation and analysis, beyond the discussion below.

Although the specific origin of the observed chromaticity is beyond the scope of this paper, we discuss a few possibilities here that could explain these results. In general, Fresnel propagation generates a chromatic evolution of the wavefront from any out-of-pupil plane optics, particularly for optics that are transmissive and/or near the focal plane (e.g., Marois et al. 2008), as well as from atmospheric scintillation (Guyon 2005), although both of these effects are expected to be very weak at the current level of raw contrasts (but see Madurowicz et al. 2018, who show that in some cases scintillation effects are

observed in GPI data). Additional sources of the observed behavior may be algorithmic in nature, arising from either numerical interpolation errors and/or DRP extraction errors. These numerical interpolation errors are wavelength-dependent when aligning data cubes, since, as discussed in Section 3.1, the full CHARIS bandpass goes from near the Nyquist sampling limit in the J band to more oversampled in the K band. Although numerical interpolation errors are more likely to disproportionally effect the J band (from under-sampling effects) while DRP extraction errors are more likely to affect the K band (from increased thermal background effects), further investigation is needed to understand the relative impact of either effect. At small separations there are two additional possible origins of the observed chromaticity: the atmospheric dispersion corrector (ADC) and the FPM IWA. The ADC does not sufficiently offset differential atmospheric tip/tilt to center the FPM on the star at the red edge of the K band. For the FPM IWA, because our chromaticity analysis uses magnified data cubes to align speckles as a function of wavelength, the IWA on magnified data cubes varies between about 2 and 4 λ/D at the red and blue ends of the bandpass, respectively, which could explain the observed chromaticity in the innermost radial zone.

4. Contrast Curves

Our 5σ contrast curves are shown in Figure 8. Final contrasts are calculated using the standard deviation at each separation in a three-pixel-wide annulus on the collapsed, back-rotated images. Raw contrasts are calculated on a median image of the stack of the registered, magnified, flux-normalized, high-pass-filtered images (i.e., the same image preparation used before the correlation analysis in addition to a median stack of the full sequence). Also note that the frames used to calculate the raw contrast are not derotated to align the parallactic angle, thus keeping quasi-static aberration aligned during the median collapse.

Although other deeper contrasts have been reached with CHARIS on more stable data sets (e.g., Currie et al. 2018a; Rich et al. 2019), the main purpose of this paper is instead to provide a detailed correlation analysis (Section 3.3) that can be
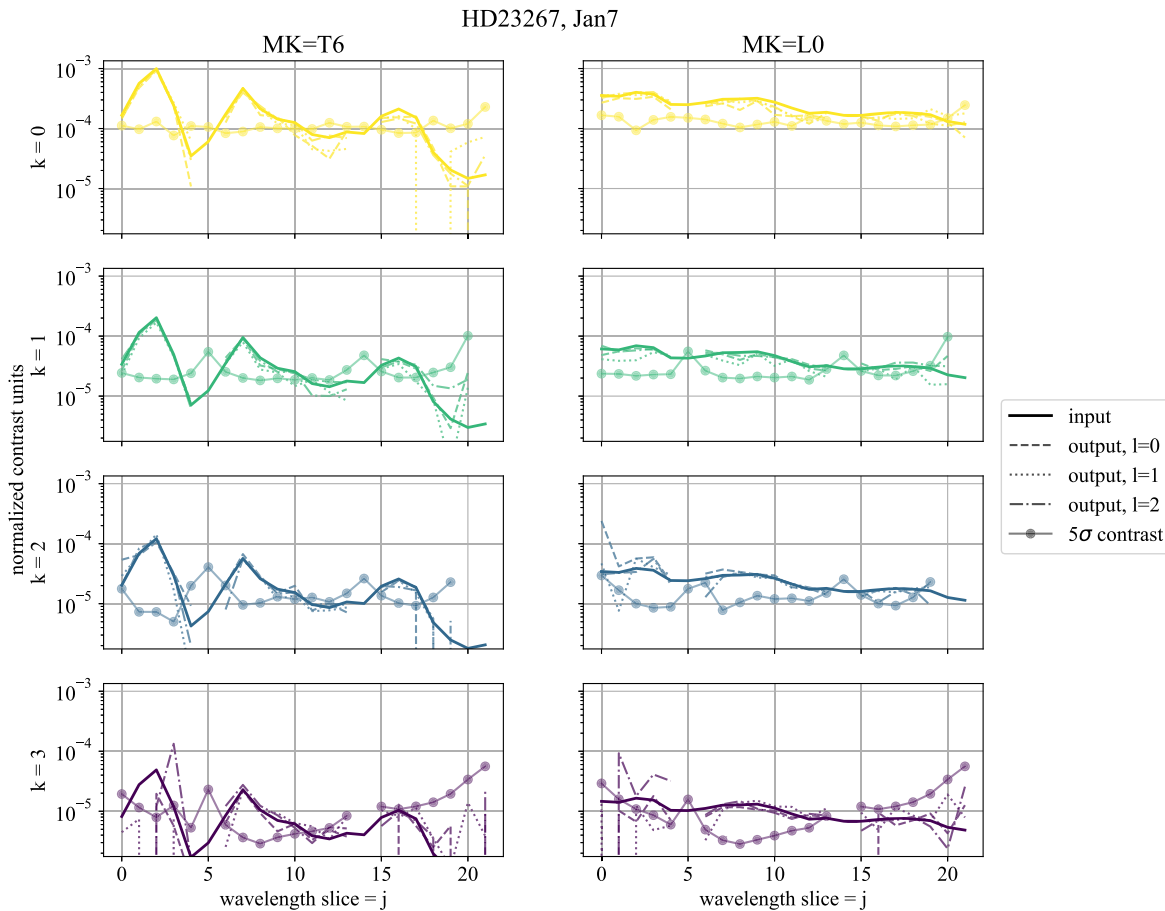
Figure 5. The injected and recovered fluxes of the simulated planets illustrated in Figure 4, as well as the 5σ contrasts corresponding to the separation of each simulated planet (calculated from the back-rotated images). The same conventions for "k" and "l" are used here as in Figure 4.

used as a reference for future CHARIS observations (e.g., considering the impact of SDI in broadband mode).

With that said, Figures 8(a) and (b) do show that T-type subtractions generally reach up to ~50% deeper contrasts than the corresponding L-type subtractions for the same target, illustrating that SDI *is* improving contrast. Contrasts in a single slice are shown in addition to the final broadband contrasts to control for the effects of noise averaging in comparing an L-type versus T-type reduction. In applying a weighted average using the assumed planet spectral type to collapse the final subtracted cube, noisier frames that lie further away from the bright part(s) of the spectrum are given less weight to the final collapsed cube (Marois et al. 2014), and so looking only at a single wavelength slice at the peak of the *H*-band T-type spectrum removes this effect. Although other wavelengths will not show as much of a performance gain further away from the T-type spectrum peak, a deeper contrast reached by a T-type versus L-type, at any wavelength, illustrates that SDI is being used to improve contrast. This improvement occurs because, within the 15 reference images that are chosen to subtract a target image at a given time and wavelength, some are at the same time but different wavelength as the target image. These results are consistent with our analysis in Section 3.3, which showed that chromaticity generally prevents the use of a reference image for SDI with a large wavelength lever arm, but reference images closer in wavelength were often more correlated than images closer in time. For a T-type subtraction, the relative lack of long dark shadows in the spectral (i.e.,

radial) direction of the right panel of Figure 3 is also consistent with these results; selecting the reference images closest in wavelength, especially for a T-type subtraction, are still often more correlated than the closest images in time.

Along these lines, because the main conclusions of Figures 6 and 7 indicate that reference images at a large wavelength difference from the target image are unlikely to be selected (i.e., separate from the hypothesis of whether or not SDI is improving contrast), we carried out the exact same PSF subtraction procedure as described in Section 3.2 on all seven targets using only the *H*-band wavelengths within the broadband data cubes (i.e., slices 6–13). The idea here is to see what level of contrast improvement is lost by removing the broadband feature of CHARIS while still using SDI in a standard ~20% bandpass. Note that using the same parameters as in Section 3.2 (e.g., $N_{ref} = 15$) means that there are less available reference images as a function of wavelength to select for the *H*-band-only sequence. However, because reference images are selected by correlation, our chromaticity analysis in Section 3.3 and Figure 6 have determined that images at these wider wavelength differences are generally not selected, and so we expect this discrepancy to have a negligible effect on the final contrast.

Figure 9 illustrates the contrast difference on a single slice in the center of the *H* band for all of our targets between the full broadband mode and the *H*-band mode. Note that comparing the collapsed cubes between the broadband and *H*-band-only subtractions would introduce a systematic bias; in the final
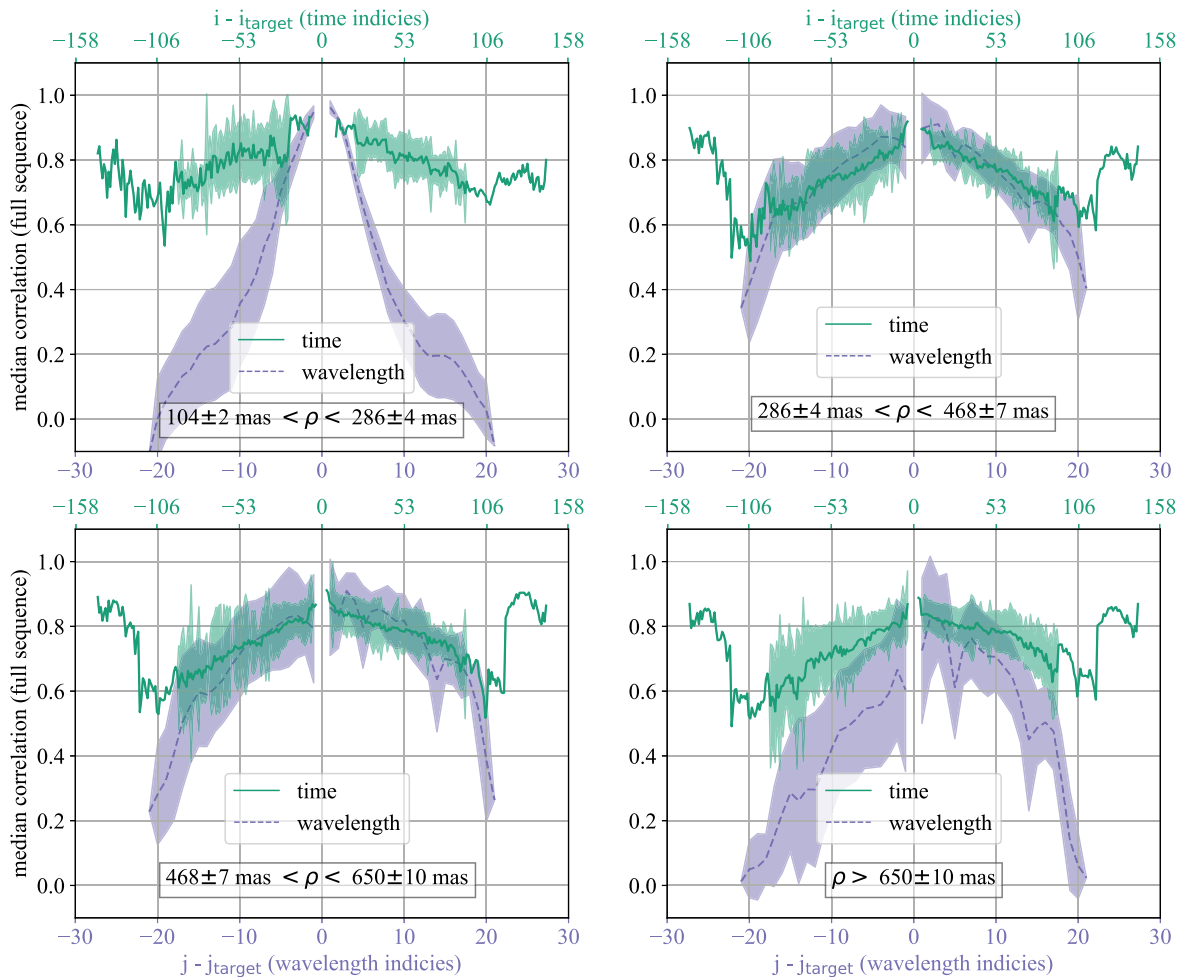
Figure 6. Median correlation of selectable references over all the observed targets, separated into time and wavelength. The shaded regions show the robust standard deviation over all targets. Each panel corresponds to a different separation in the image. For every target image, each selectable reference image is identified as either the target wavelength and a different time ($j = j_{target}$, $i \neq i_{target}$; top x-axis) or the target time and a different wavelength ($i = i_{target}$, $j \neq j_{target}$; bottom x-axis). The main conclusion from this figure is that there is significant chromaticity at small and large separations over the full JHK bandpass; at adjacent wavelengths SDI may still be advantageous, but this figure suggests that this is not the case for a large wavelength lever arm across the full bandpass (see Section 4 for a validation of this hypothesis).

wavelength collapse the broadband data set would average more slices than the H-band-only data set, producing a deeper final contrast for the former even if the individual subtracted cubes are the same for both sequences. In general, all of our 2017 April observations show little to no gain using the broadband data set, as expected from Figure 6. Interestingly, the January 2018 observations, which reach the deepest raw and final contrasts in Figure 8, instead consistently show a ~0.5–1 magnitude gain using the broadband mode. Also when using the broadband mode, in a few cases at separations less than about 200 mas, the L-type subtractions appear to decrease in contrast by up to a magnitude.

Ultimately, the origin of this observed chromaticity that may be improving and/or degrading contrast will need to be better understood before a detailed analysis leading to mitigation strategies can be done. As discussed in Section 3.3, many different possible sources of chromaticity could reproduce the results in Figures 6, 7, and 9. These sources may be algorithmic and/or optical in nature, and could be quasi-static and/or variable in time. However, independent of the origin, in this paper we illustrate that chromaticity in our broadband data sets is limiting the final achievable contrast, thus motivating the need for further work toward more achromatic ExAO systems.

5. Conclusions

In this paper, we have presented a detailed correlation analysis, data processing architecture, and final contrast curves of our SCExAO/CHARIS broadband data sets to the community. Our main findings are as follows.

1. A correlation analysis of the SCExAO/CHARIS system across all seven of our targets shows that for small ($\lesssim 300$ mas) and large ($\gtrsim 650$ mas) separations, chromaticity across the full JHK bandpass is worse than temporal stability across the full observing sequence (Figure 6). This result suggests that at these separations, the large wavelength lever arm of our CHARIS broadband data sets could not be used advantageously to improve SDI post-processing (i.e., slices in the J band can generally not be used to subtract slices in the K band, and vice versa).

2. With that said, we find that our T-type PSF subtractions generally reach up to ~50% deeper final contrasts than our L-type subtraction of the same target (Figure 8), illustrating that more correlated wavelength slices close to the target image wavelength are being used to improve contrast. This is consistent with the findings of point 1
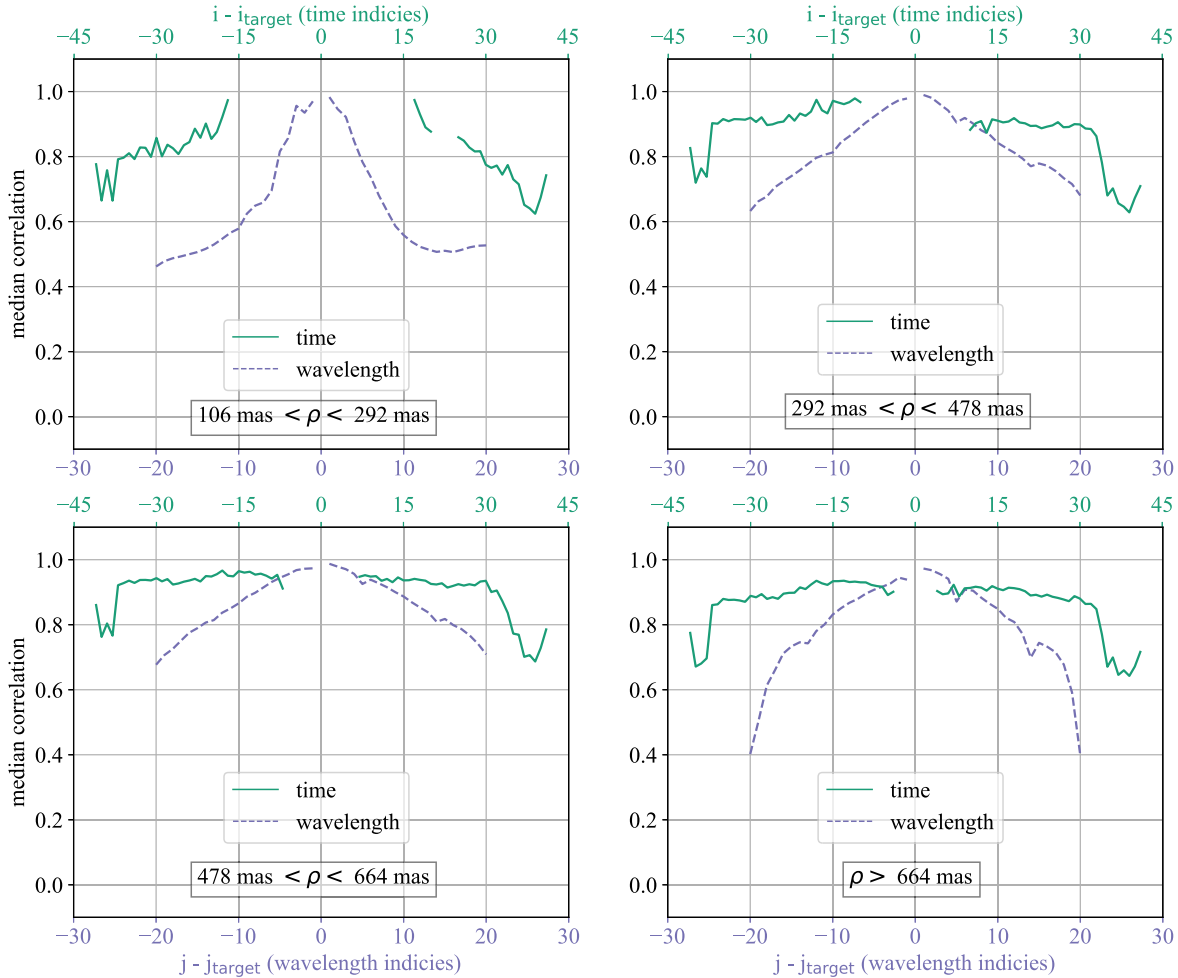
kappa And

Figure 7. Median correlation with time and wavelength for the κ And data set from Currie et al. (2018a), illustrating that on a more stable data set chromaticity is generally also much better compared to Figure 6.

above (i.e., Figure 6); together, these results illustrate that improved contrast from SDI processing of our targets comes from selecting reference slices closer in wavelength to the target image where chromaticity is less significant.

3. We present a modified framework for exoplanet algorithmic throughput correction using FM that is designed for the CHARIS broadband mode (Section 3.2). As a result, simulated planets are recovered in good agreement with their input values even when close to the 5σ noise floor (Figure 5).

4. PSF subtracted sequences using only the *H*-band slices within the broadband data set generally reach final contrasts that are similar to the final contrasts of the same slices in the broadband sequences (Figure 9), validating the claim made in points 1 and 2 above.

Beyond the work presented in this paper, as long as ExAO systems remain limited by temporal stability (i.e., preventing subtraction to the photon noise limit using ADI and/or RDI post-processing, thus motivating the additional need for SDI processing), chromaticity will also limit the deepest contrasts that can be reached. Looking ahead to the future, coherent differential imaging (CDI; e.g., Guyon 2004; Bordé &

Traub 2006; Baudoz et al. 2006; Give'On et al. 2007; Serabyn et al. 2011; Sauvage et al. 2012; Jovanovic et al. 2018, and references therein) is a complimentary approach toward reaching deeper contrasts on ExAO systems. CDI can operate in a narrow bandpass (i.e., in contrast with SDI) and is therefore, in principle, not limited by chromaticity. However, detailed characterization of colder/lower mass exoplanet atmospheres will continually push the ExAO community toward reaching deeper broadband contrasts, in which case chromaticity will ultimately become a fundamental limitation to this goal. The future push toward reaching deeper contrasts will therefore also require a push toward implementing more achromatic broadband systems.
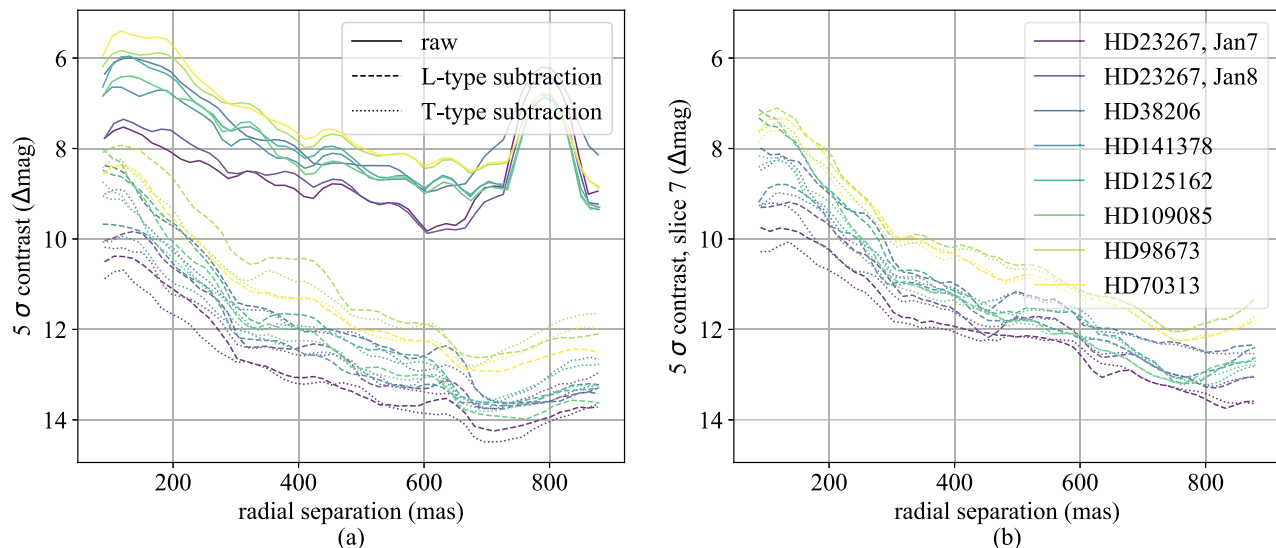
Figure 8. (a) Contrast curves for all of our targets. (b) Final contrasts at a single wavelength, slice 7 = 1.5 μm. Each target is color coded (the same in both figures); the solid, dashed, and dotted lines indicate the raw, final L-type, and final T-type contrasts, respectively.
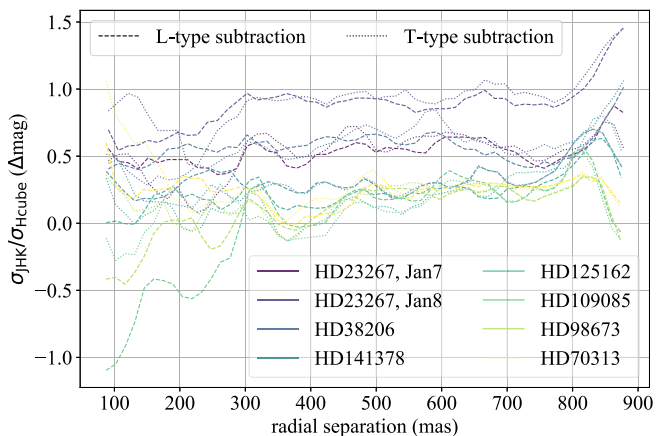

Figure 9. Difference in contrast for each target when using the full *JHK* data set vs. only the *H* band to do PSF subtraction, illustrated here for a single wavelength slice at the center of the *H* band (slice 11 in the full *JHK* cube). The color and line styles are the same as in Figure 8.

Appendix
Additional Correlation Analysis

Additional correlation analysis informed our choice of values for a correlation cut, as discussed in Section 3.2. Still using the HD 23267 January 7 data set, frame 17, slice 10 as an example (as in Figure 2), Figure 10 illustrates the trends in correlation of selected reference images at either the same time or same wavelength (i.e., $i = 17$ or $j = 10$, respectively). The top two panels show the correlation trends of the 15 most correlated images as defined in Section 3.2 ($N_{ref} = 15$) at large separations as a function of time (a) or wavelength (b) at large separations, using an aggressiveness of 0.5, and using an assumed L0 input spectrum. Uncorrelated frames can be identified as vertical streaks in these images (e.g., the water bands at $j = 5$ and $j = 14$, as in Figure 2). The chosen correlation cut value at this separation should reject these uncorrelated frames from the sequence in order to prevent unnecessary noise amplification in the least-squares algorithm (Marois et al. 2010). Because we require that all 15 reference images lie above our defined correlation cut value, the bottom panels of Figure 10 show a one-dimensional slice of the top two panels: the correlation of reference image number 14 (i.e., the fifteenth most correlated selectable reference, using a zero-index counting system) as a function of time (c) and wavelength (d) for both spectral types and all separations. Similarly, at all separations and spectral types we would like to discard the frames where there are "dips" in these plots in order to optimally reject uncorrelated images and avoid propagating them through the least-squares algorithm. Figures 10(c) and (d) also suggest that larger separations may be less correlated overall and that a different correlation cut value will be needed for each separation. Note that the low correlation at larger separations may be a result of wavefront decorrelation at those spatial frequencies and/or increased background noise because the speckles at those separations are detected at a lower S/N. At a given separation, using a correlation cut value that is too high will reject too many images at that separation and return only a minimal number of subtracted frames to the final collapsed cube; the optimal correlation cut value should balance rejecting too many images with propagating too much noise through the least-squares algorithm.

The extension of Figures 10(c) and (d) is to calculate the median behavior over a full sequence and for all targets. Thus, for a single target at a given separation, in Figure 11 we compute the median and robust standard deviation for the fifteenth most correlated reference image over all target images (time and wavelength) and input spectral types. Figure 11 illustrates that, for most targets, the fifteenth most correlated
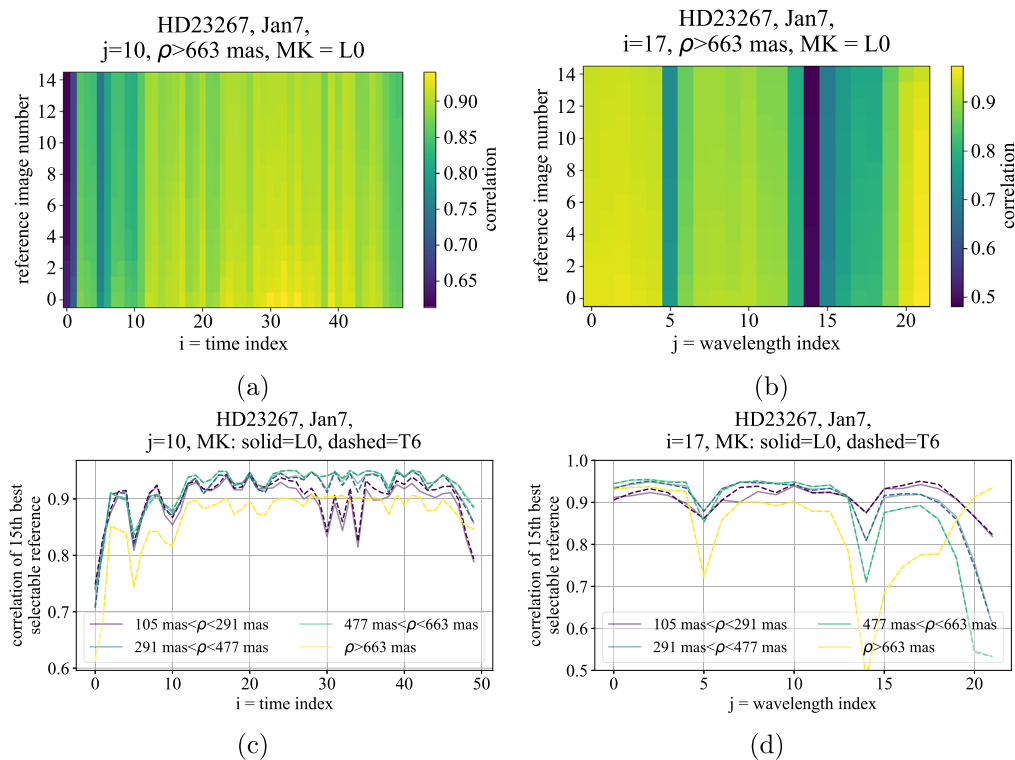
Figure 10. (a) and (b): the 15 most correlated selectable reference images (i.e., reference image numbers 0 through 14, using a zero-index counting system) at a single wavelength slice as a function of time and at a single timestamp as a function of wavelength, respectively, both for large separations and assuming an L0 template selection criterion. (c) and (d): the fifteenth most correlated selectable reference image at all separations and for two input spectral types (c) at a single wavelength as a function of time and (d) at a single time as a function of wavelength (i.e., reference image number 14 as a function of time and wavelength).
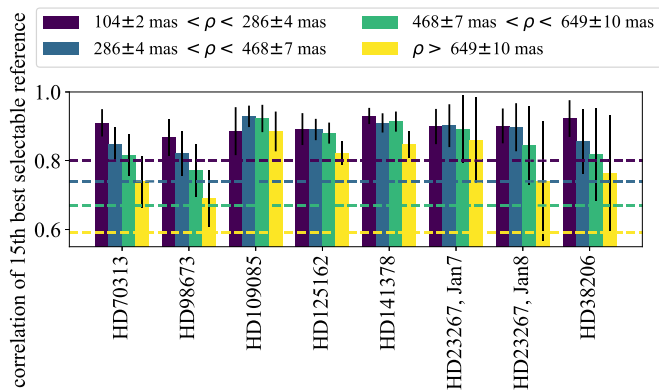

Figure 11. Median correlation of the fifteenth most correlated, selectable reference image and corresponding robust standard deviation over all times, wavelengths, and input spectral types, at each separation for each target. Our adopted correlation cut values are illustrated by the horizontal dashed lines, color coded to each separation.

image decreases in correlation with separation, motivating the use of different correlation cut values at each separation. The horizontal dashed lines in Figure 11 illustrate our chosen correlation cut values at each separation. We set these values to generally lie near the bottom of the corresponding 1σ error bars over all targets. Using this definition to define the correlation cut values will balance the rejection of bad frames relative to the median correlation value without discarding too many frames from the sequence. In general, the optimal correlation cut value will vary for each least-squares subtraction and is degenerate with the SVD cutoff and number of reference images (Marois et al. 2010); this will ultimately require a more

optimized approach (e.g., Gerard & Marois 2016), but is beyond the scope of this paper.

ORCID iDs

Benjamin L. Gerard ⓘ https://orcid.org/0000-0003-3978-9195
Christian Marois ⓘ https://orcid.org/0000-0002-4164-4182
Thayne Currie ⓘ https://orcid.org/0000-0002-7405-3119
Timothy Brandt ⓘ https://orcid.org/0000-0003-2630-8073
Jeffrey K. Chilcote ⓘ https://orcid.org/0000-0001-6305-7272
Zachary H. Draper ⓘ https://orcid.org/0000-0002-1834-3496
Tyler Groff ⓘ https://orcid.org/0000-0001-5978-3247
Olivier Guyon ⓘ https://orcid.org/0000-0002-1097-9908
Gillian R. Knapp ⓘ https://orcid.org/0000-0002-9259-1164
Tomoyuki Kudo ⓘ https://orcid.org/0000-0002-9294-1793
Jungmi Kwon ⓘ https://orcid.org/0000-0003-2815-7774
Julien Lozi ⓘ https://orcid.org/0000-0002-3047-1845
Frantz Martinache ⓘ https://orcid.org/0000-0003-1180-4138
Michael McElwain ⓘ https://orcid.org/0000-0003-0241-8956
Motohide Tamura ⓘ https://orcid.org/0000-0002-6510-0681
Taichi Uyama ⓘ https://orcid.org/0000-0002-6879-3030

References

Baudoz, P., Boccaletti, A., Baudrand, J., & Rouan, D. 2006, in IAU Coll. 200, Direct Imaging of Exoplanets: Science & Techniques, ed. C. Aime & F. Vakili (Cambridge: Cambridge Univ. Press), 553
Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2008, Proc. SPIE, 7014, 701418
Bordé, P. J., & Traub, W. A. 2006, ApJ, 638, 488
Brandt, T. D., Rizzo, M., Groff, T., et al. 2017, JATIS, 3, 048002
Burgasser, A. J., Liu, M. C., Ireland, M. J., Cruz, K. L., & Dupuy, T. J. 2008, ApJ, 681, 579
Burgasser, A. J., McElwain, M. W., Kirkpatrick, J. D., et al. 2004, AJ, 127, 2856

Currie, T., Brandt, T. D., Uyama, T., et al. 2018a, AJ, 156, 291
Currie, T., Debes, J., Rodigas, T. J., et al. 2012, ApJ, 760, L32
Currie, T., Kasdin, N. J., Groff, T. D., et al. 2018b, PASP, 130, 044505
Draper, Z. H., Marois, C., Wolff, S., et al. 2014, Proc. SPIE, 9147, 91474Z
Gerard, B. 2017, S17A-TE261: Search for exoplanets scultping debris disks with SCExAO+CHARIS, Gemini Proposal: Subaru Exchange Time, NOAO Proposal ID 2017A-0261
Gerard, B. L., & Marois, C. 2016, Proc. SPIE, 9909, 990958
Give'On, A., Belikov, R., Shaklan, S., & Kasdin, J. 2007, OExpr, 15, 12338
Groff, T., Chilcote, J., Brandt, T., et al. 2017, Proc. SPIE, 10400, 1040016
Groff, T. D., Chilcote, J., Kasdin, N. J., et al. 2016, Proc. SPIE, 9908, 99080O
Guyon, O. 2004, ApJ, 615, 562
Guyon, O. 2005, ApJ, 629, 592
Hinkley, S., Oppenheimer, B. R., Soummer, R., et al. 2007, ApJ, 654, 633
Hodapp, K. 2018, UH Proposal: Subaru Time, UH-07B
Johnson-Groh, M. 2016, Master's Thesis, Univ. Victoria
Jovanovic, N., Absil, O., Baudoz, P., et al. 2018, Proc. SPIE, 10703, 107031U
Jovanovic, N., Guyon, O., Martinache, F., et al. 2015a, ApJL, 813, L24
Jovanovic, N., Martinache, F., Guyon, O., et al. 2015b, PASP, 127, 890
Kennedy, G. M., & Wyatt, M. C. 2014, MNRAS, 444, 3164
Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, ApJ, 660, 770
Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, PNAS, 111, 12661
Madurowicz, A., Macintosh, B. A., Ruffio, J.-B., et al. 2018, Proc. SPIE, 10703, 107036E

Marois, C., Correia, C., Galicher, R., et al. 2014, Proc. SPIE, 9148, 91480U
Marois, C., Doyon, R., Racine, R., & Nadeau, D. 2000, PASP, 112, 91
Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006a, ApJ, 641, 556
Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2006b, ApJ, 647, 612
Marois, C., Macintosh, B., Soummer, R., Poyneer, L., & Bauman, B. 2008, Proc. SPIE, 7015, 70151T
Marois, C., Macintosh, B., & Véran, J.-P. 2010, Proc. SPIE, 7736, 77361J
Milli, J., Banas, T., Mouillet, D., et al. 2016, Proc. SPIE, 9909, 99094Z
Minowa, Y., Hayano, Y., Oya, S., et al. 2010, Proc. SPIE, 7736, 77363N
Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019, AJ, 158, 13
Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, PASP, 111, 587
Rich, E. A., Wisniewski, J. P., Currie, T., et al. 2019, ApJ, 875, 38
Rosenthal, E. D., Gurwell, M. A., & Ho, P. T. P. 1996, Natur, 384, 243
Sauvage, J.-F., Mugnier, L., Paul, B., & Villecroze, R. 2012, OptL, 37, 4808
Serabyn, E., Wallace, J. K., & Mawet, D. 2011, ApOpt, 50, 5453
Sivaramakrishnan, A., & Oppenheimer, B. R. 2006, ApJ, 647, 620
Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJL, 755, L28
Sparks, W. B., & Ford, H. C. 2002, ApJ, 578, 543
Vigan, A., Moutou, C., Langlois, M., et al. 2010, MNRAS, 407, 71
Wang, J. J., Ruffio, J.-B., De Rosa, R. J., et al. 2015, pyKLIP: PSF Subtraction for Exoplanets and Disks, Astrophysics Source Code Library, ascl:1506.001
Wyatt, M. C. 2008, ARA&A, 46, 339