



HAL
open science

Spectral cube extraction for the VLT/SPHERE IFS. Open-source pipeline with full forward modeling and improved sensitivity

M. Samland, T. D. Brandt, J. Milli, P. Delorme, A Vigan

► **To cite this version:**

M. Samland, T. D. Brandt, J. Milli, P. Delorme, A Vigan. Spectral cube extraction for the VLT/SPHERE IFS. Open-source pipeline with full forward modeling and improved sensitivity. Astronomy and Astrophysics - A&A, 2022, 668, 10.1051/0004-6361/202244587 . hal-03935716

HAL Id: hal-03935716

<https://hal.science/hal-03935716v1>

Submitted on 20 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Spectral cube extraction for the VLT/SPHERE IFS

Open-source pipeline with full forward modeling and improved sensitivity[★]

M. Samland¹ , T. D. Brandt² , J. Milli³ , P. Delorme³ , and A. Vigan⁴ 

¹ Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany
e-mail: samland@mpa.de

² Department of Physics, University of California, Santa Barbara, CA 93106, USA

³ Univ. Grenoble Alpes, CNRS, IPAG, 38000 Grenoble, France

⁴ Aix-Marseille Univ., CNRS, CNES, LAM, Marseille, France

Received 24 July 2022 / Accepted 16 September 2022

ABSTRACT

We present a new open-source data-reduction pipeline to reconstruct spectral data cubes from raw SPHERE integral-field spectrograph (IFS) data. The pipeline is written in Python and based on the pipeline that was developed for the CHARIS IFS. It introduces several improvements to SPHERE data analysis that ultimately produce significant improvements in postprocessing sensitivity. We first used new data to measure SPHERE lenslet point spread functions (PSFs) at the four laser calibration wavelengths. These lenslet PSFs enabled us to forward-model SPHERE data, to extract spectra using a least-squares fit, and to remove spectral crosstalk using the measured lenslet PSFs. Our approach also reduces the number of required interpolations, both spectral and spatial, and can preserve the original hexagonal lenslet geometry in the SPHERE IFS. In the case of least-squares extraction, no interpolation of the data is performed. We demonstrate this new pipeline on the directly imaged exoplanet 51 Eri b and on observations of the hot white dwarf companion to HD 2133. The extracted spectrum of HD 2133B matches theoretical models, demonstrating spectrophotometric calibration that is good to a few percent. Postprocessing on two 51 Eri b data sets demonstrates a median improvement in sensitivity of 80 and 30% for the 2015 and 2017 data, respectively, compared to the use of cubes reconstructed by the SPHERE Data Center. The largest improvements are seen for poorer observing conditions. The new SPHERE pipeline takes less than three minutes to produce a data cube on a modern laptop, making it practical to reprocess all SPHERE IFS data.

Key words. planets and satellites: detection – methods: data analysis – techniques: imaging spectroscopy – techniques: high angular resolution

1. Introduction

In recent years, a new generation of integral-field spectrographs (IFSs) on 8-m class telescopes have been developed especially with the purpose of aiding the direct detection and characterization of extrasolar planets. Integral-field spectrographs allow us to obtain spectral information across the field of view (FoV) of the instrument, hence providing spatially resolved images at a wide range of wavelengths at the same time.

Integral-field spectrographs have proven to be a powerful asset in the toolbox of high-contrast imaging; not only can the spectral information be used to characterize the properties of any astrophysical object detected, it can also be used to exploit the different chromatic behavior of diffraction speckles from a real astrophysical source to suppress noise and improve detection contrast (e.g., Racine et al. 1999; Sparks & Ford 2002; Hoeijmakers et al. 2018).

The recent generation of IFS instruments combines the high angular resolution of the telescope with extreme adaptive optics (XAO) to provide near diffraction-limited images of stars and their surroundings (e.g., Guyon 2005; Jovanovic et al. 2015). The difference in brightness between the star and potential planets necessitates the use of further technologies such as coronagraphs

(e.g., Soummer 2005; Guyon et al. 2005; Snik et al. 2012) to suppress the light from the central star. Dedicated high-contrast imaging techniques are also needed to distinguish between the unavoidable residual starlight halo and genuine astrophysical signals in the direct vicinity of the star, such as angular differential imaging (ADI; Marois et al. 2006) and spectral differential imaging (SDI; Racine et al. 1999).

Currently, four such dedicated (XAO) high-contrast imaging instruments are equipped with an IFS: GPI on Gemini South (soon Gemini North, Macintosh et al. 2014), ALES at the Large Binocular Telescope (LBT; Skemer et al. 2015), SPHERE at the Very Large Telescope (VLT; Beuzit et al. 2019), and CHARIS on the Subaru telescope (Groff et al. 2015). All of these instruments use a similar technology that is based on lenslet arrays to map spatial elements onto small spatially separated spots that will be dispersed onto the detector as distinct spectra. Although these instruments differ in implementation details and design (e.g., wavelength range, resolution, and size of the FoV), they are sufficiently similar for a suitably flexible software package to be adapted to the peculiarities of each individual instrument.

The reduction and processing of IFS data is inherently complex. After investing a large part of their budgets in the initial development and building of an instrument, individual teams therefore often struggle to provide the manpower for continuous support and improvement of the software that is used for

[★] The pipeline and documentation are publicly available at <https://github.com/PrincetonUniversity/charis-dep>

data reduction. This can be remedied by providing the scientific community access to source codes and encourage community involvement in their further development and in improving their reliability. In this paper we present our effort to take the first step in this direction by adapting the CHARIS data-reduction pipeline (Brandt et al. 2017, henceforth B17), an open-source package written in Python and Cython available on Github, to be able to process SPHERE IFS data into scientifically usable products. The largest difference of the SPHERE IFS compared to the other IFS instruments is its hexagonal lenslet geometry, known as BIGRE (Antichi et al. 2009). This will receive special attention in this work because this fact is not widely known to end-users of the data and has implications for future improvements of data-reduction algorithms and planet detection algorithms.

The pipeline described in this paper offers various improvements over the original SPHERE pipeline in how the spectral extraction is performed, and it addresses known shortcomings. The paper is structured as follows. Section 2 gives an overview of the design properties and peculiarities of the SPHERE IFS. It also describes the current ESO SPHERE pipeline. Section 3 shows the detailed steps we performed to obtain the necessary calibration data, and the static calibration products that are incorporated into the pipeline. Section 4 describes the spectral extraction methods we implemented, the process that transforms raw data into the spectral image cubes. Section 5 compares the results obtained with our new pipeline to results obtained with the official pipeline. Section 6 details what the active user of the software needs to know to use the pipeline, and it provides a brief overview of the computational performance. Finally, we discuss our findings and conclude in Sect. 7.

2. Overview of the SPHERE IFS and ESO pipeline

The SPHERE IFS is part of the SPHERE instrument (Beuzit et al. 2019), which is located on the stable Nasmyth platform of the VLT, which is operated by the European Southern Observatory (ESO) as part of a system of three dedicated high-contrast imaging instruments: ZIMPOL (visible light polarimetry; Schmid et al. 2018), IRDIS (dual-band imager; Dohlen et al. 2008), and IFS (Claudi et al. 2008; Mesa et al. 2015). These instruments are located behind a common path and infrastructure (CPI) module that feeds the subsystems with a highly stabilized, AO-corrected beam. The SAXO XAO system for SPHERE achieves Strehl ratios better than 90% (Petit et al. 2012) in the near-infrared (NIR). In this paper we focus exclusively on the extraction and analysis of data cubes from the IFS subsystem. Table 1 lists the basic properties of the SPHERE IFS.

The SPHERE IFS has been used to discover and characterize exoplanets and brown dwarfs (e.g., Chauvin et al. 2017b, 2018; Keppler et al. 2018; Samland et al. 2017; Milli et al. 2017a; Müller et al. 2018; Cheetham et al. 2019; Maire et al. 2020). It can also be used to study extended objects such as debris disks (e.g., Boccaletti et al. 2015; Milli et al. 2017b) and asteroids (Hanuš et al. 2017). Several large surveys have been performed using the instrument (e.g., Chauvin et al. 2017a; Bohn et al. 2020; Janson et al. 2021; Bonavita et al. 2022), and a considerable archive of hundreds of observation sequences exists. It is being used to study the demographics of young giant planets in the SHINE survey (e.g., Vigan et al. 2021).

The SPHERE IFS images a $1''.73 \times 1''.73$ patch of sky onto a lenslet array. Each hexagonal lenslet focuses the light incident on it to a point in the plane of a detector. A prism residing between

Table 1. Basic properties of the SPHERE IFS.

| Parameter | SPHERE-IFS |
|-------------------------------|-------------------------------------|
| Detector | 2048 × 2048 H2RG |
| No. of lenslets | ≈152 × 152 |
| Lenslet lattice configuration | hexagonal |
| Lenslet area | 370 mas ² ^(a) |
| Field of view | 1''.73 × 1''.73 |
| Wavelength coverage | 0.95–1.66 μm |
| Microspectrum length | ~39 pixel ^(b) |
| Separation between spectra | ~5 pixel ^(b) |
| $R = \lambda/\delta\lambda$ | ~55; ~35 ^(b) |
| Available modes | Y–J; Y–H |

Notes. ^(a)Area of hexagonal lenslet on sky. Resampled by the pipeline to $\sim(7.4 \text{ mas})^2$ per square pixel. The side length t of the hexagon is 7.4 mas, and the area $A = 3\sqrt{3}t^2/2$. ^(b)Microspectrum length, spacing, and R may vary slightly depending on definitions and may vary slightly across the FoV.

the lenslet array and the detector disperses each point into an about 39-pixel-long low-resolution spectrum. These spectra may cover either the Y–J spectral range at a spectral resolution $R = \lambda/\delta\lambda \approx 55$ (henceforth referred to as YJ-mode) or the Y–H spectral range at a lower $R \approx 35$ (henceforth referred to as YH-mode)¹.

The SPHERE IFS has an internal calibration unit. This unit can uniformly illuminate the detector to create a detector flat-field, or it can uniformly illuminate the lenslet array to create a lenslet flat. The lenslet flat appears on the detector as a grid of dispersed microspectra. The calibration unit also takes wavelength calibration data by illuminating the lenslets with lasers of known wavelength. All of these data are taken during the day, both for calibration purposes and to monitor the health of the instrument. Dark and background images are likewise obtained by closing shutters at different positions (in front of the IFS module for darks, or in front of the SPHERE entrance shutter for instrument backgrounds) during daytime. This saves observing time during the night for actual science data. The calibration unit consists of (a) a series of lamps, both narrow and broad band, for taking detector flat-fields and flat-fields taken through the integral field unit (IFU), and (b) a series of monochromatic lasers for wavelength calibration. For SPHERE, the ‘wavecal’ calibration template that is routinely executed takes images with three (YJ-mode) or four (YH-mode) lasers simultaneously (see left panel, Fig. 1).

The first step in the analysis of SPHERE IFS data is to take the 2D array of microspectra and reconstruct a 3D data cube: an image at each sampled wavelength. Subsequent analysis steps, usually referred to as postprocessing, subtract starlight to search for faint exoplanets, calibrate and extract spectra, and calculate contrasts and sensitivities. The ESO data-reduction and handling (DRH) pipeline (Pavlov et al. 2008) is designed to perform all steps that are required to reconstruct data cubes from SPHERE IFS raw data. This includes everything from thermal background subtraction and flat-field correction to the more intricate steps of wavelength calibration and extraction of the spectra into usable

¹ The official designations of the modes are IRDIFS and IRDIFS_EXT. In these modes, IRDIS dual-band imaging data are taken at longer wavelengths at the same time as the IFS data. As IRDIS is not the topic of this paper, we use a more intuitive designation for the IFS modes: YJ- and YH-mode.

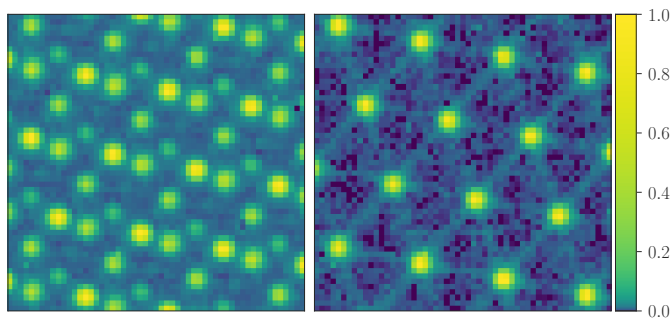


Fig. 1. Normalized monochromatic laser flat-fields in logarithmic scale. The dispersion direction is from bottom to top. *Left:* standard wavelength calibration image for the extended (YH) mode using all four lasers simultaneously. This is provided by ESO as part of the standard calibration sequence and is used by the ESO pipeline. *Right:* similar image, but with only the 1123.7 nm laser on. Images at this and three other wavelengths were taken and used to construct our lenslet PSF library (Sect. 3.4).

spectral image cubes. The SPHERE Data Center² (DC, Delorme et al. 2017), as well as a third-party Python wrapper (Vigan 2020), both combine the DRH with additional functions and improvements. These include a more efficient correction for bad pixels, a correction for spectral crosstalk caused by overlapping lenslet spectra, an additional step to correct erroneous wavelength calibration, and a correction for anamorphism in the instrument.

The DRH pipeline begins by subtracting backgrounds and flat-fielding the detector. Backgrounds are significant because most of the instrument is not cryogenically cooled. The DRH then uses the dispersed lenslet flat-field image to identify the illuminated pixels of all microspectra and creates a mask for all of them in a step called *specpos*. The information about the position of each lenslet and the pixels associated with each spectrum is used in further steps. The DRH then fits for a wavelength solution (*wavelcal*), a mapping between lenslet, wavelength, and position on the detector, using images taken with lasers of known wavelength. This mapping is required in order to extract the spectrum of each lenslet. The last calibration required by the DRH pipeline is a map of relative lenslet transmissions, called the instrument flat-field. This is measured using a spectral image cube extracted from a uniform illumination of the lenslet array and is used to correct the flux of each lenslet in the spectral extraction routine.

The DRH pipeline uses aperture photometry to extract the lenslet spectra. In this approach, the flux is summed across a predefined aperture in the direction perpendicular to the dispersion. The total flux in this direction is then assigned to the wavelength of this pixel row of the spectrum. Aperture photometry involves tradeoffs. A larger aperture includes more of the flux, but adds additional read noise and photon noise from the background. Because of the relatively close spacing of the microspectra in SPHERE, a larger aperture also includes more flux from neighboring spectra at other wavelengths (spectral crosstalk). These shortcomings may be overcome by optimal extraction (Horne 1986), as long as the line-spread function perpendicular to the dispersion direction (the instrumental profile) is known. This has not previously been measured for SPHERE,

preventing an implementation of more sophisticated spectral extraction approaches.

The SPHERE instrument uses hexagonal lenslets that are arranged in a honeycomb pattern (for an example of this). The DRH resamples the microspectra onto a rectilinear grid (see also Fig. 8). This requires an interpolation and sacrifices some of the information in the raw data. Previous generations of high-contrast instruments lacked spectral information and took data on rectilinear grids (the arrangement of pixels on nearly all detectors). The interpolation step taken by the DRH enables established postprocessing routines to work on SPHERE data without accounting for a new geometry. The use of aperture photometry of the SPHERE DRH to extract microspectra and its resampling of the data cubes onto a rectilinear grid suggest that a new extraction approach could result in higher-quality data cubes. Berdeu et al. (2020) suggested an alternative regularized least-squares approach to reduce SPHERE IFS data, called PIC, and saw a reduction in data artifacts and improvements in final postprocessed contrast. In the following sections, we describe an independent data-reduction pipeline based on that of the CHARIS IFS (Groff et al. 2016; B17). This approach revisits every aspect of the data reduction and also takes advantage of new calibration data.

3. Calibrations

Calibration data for the SPHERE IFS consist of uniformly illuminated images of the detector for flat-fielding, dark images for background subtraction, and wavelength calibration frames. The latter consist of the lenslet array uniformly illuminated by three or four lasers of known wavelength fed through an integrating sphere. Dark images have significant counts because most of SPHERE is not cryogenic.

Darks in SPHERE IFS images consist of two components: an undispersed background from scattered light and emission after the lenslet array, and a dispersed background that is imaged by the lenslet array into microspectra. Sky images include both backgrounds, while dark images include only the undispersed background. Our approach uses only the dark (thermal instrument background) images and the detector flat-fields. From these, we aim to construct a detector flat-field, a bad-pixel mask, and an undispersed background template. All calibration files necessary to reduce SPHERE IFS data are packaged with our pipeline, except for the laser flat-field needed to adjust the wavelength solution for each observing night. We include these master calibrations rather than adopting nightly calibrations for two reasons. First, each individual calibration contains its own realization of read noise and photon noise. If nightly calibration frames are used, the same realization of noise is added to each science frame, potentially doubling the read and background noise. Master calibrations avoid this drawback, thereby providing better results as long as the calibrations themselves are stable. In the following sections, we demonstrate the requisite stability of both the background and the flat-fields. Our second reason for supplying master calibrations is that calibration frames occasionally fail, for instance, due to an incorrect shutter position or an uneven illumination of the detector. The use of curated master frames avoids the possibility of bad darks and/or flats from user error or hardware failure.

For this work, we have taken additional SPHERE IFS calibration data: images with only one laser illuminating the lenslets at a time. This enabled us to clearly resolve the lenslet point spread functions (PSFs) and to extract the oversampled lenslet

² The SPHERE DC performs data reduction on request and also processes all SPHERE public data to make them available publicly. More information is available at <https://sphere.osug.fr/spip.php?rubrique16>

PSF in Sect. 3.4. For YJ-mode, the laser wavelengths are 987.7 nm, 1123.7 nm, and 1309.4 nm. In YH-mode, we have one additional laser at 1545.1 nm. Figure 1 shows the difference between using all four lasers at the same time in YH-mode and using an individual laser. Both calibration images have been dark and flat corrected.

3.1. Bad-pixel mask

Unfortunately, the detector in the SPHERE IFS has a significant number of bad pixels in clusters scattered throughout its FoV comprising $\approx 4.8\%$ of all pixels. The number of bad pixels further significantly increased after the pandemic shutdown in April 2020, during which the instrument was warmed up, to $\approx 5.0\%$. Two master bad-pixel masks and master flat-fields for before and after the shutdown are therefore included with the pipeline to reflect the changes in the detector, which is otherwise very stable over time.

The first step of our new data-processing approach is to identify and construct a binary bad-pixel mask, which we did using the dark images and the detector flat-field images. We first constructed a master dark from 100 (pre-shutdown) and 50 (post-shutdown) randomly selected dark calibration files, each containing ten exposures at an exposure time of 11 seconds. The master dark frames are a median combination of these 1000 or 500 dark frames, respectively. We then masked all pixels whose values exceed a 3.5σ deviation in an iterative FWHM = 4 pixel median filter kernel. Finally, we constructed a master flat-field as described below in detail. We also masked pixels that differed significantly from the average value in a flat-field, along with pixels whose values fluctuated across flat-field images. To do this, we masked all pixels that are below 90% or more than 110% of the median pixel counts or have a temporal standard deviation larger than 3%.

Our approach identified a total of 167 386 bad pixels from the dark frames and 31 639 bad pixels from the flat-field images pre-shutdown, with numbers increasing by several thousand pixels post-shutdown. We took the combined dark and flat mask as our final bad-pixel mask in both cases. Our pipeline does not perform a bad-pixel interpolation because both spectral extraction methods, the optimal extraction (see Sect. 4.2) and the least-squares extraction approaches (see Sect. 4.3), make it easy to exclude all bad pixels from the extraction process altogether. This means that interpolated bad-pixel data will not bias uncertainties due to newly introduced correlations. The SPHERE DRH pipeline, however, interpolates bad pixels based on all their neighbors, which is not ideal in the case of microspectra extraction. For this reason, the SPHERE Data Center (DC) reductions (Delorme et al. 2017), and the `vlt-sphere` Python pipeline (Vigan 2020), which implements a wrapper for the SPHERE DRH ESO pipeline and adds additional improvements, both implement an approach that linearly interpolates bad pixels using only the nearest good neighbor in the dispersion direction.

3.2. Detector flat-field

The SPHERE IFS contains an internal calibration lamp that can uniformly illuminate the detector to correct for pixel-to-pixel variations in sensitivity. We constructed the detector flat-field using two different exposure time settings. We randomly selected pairs of files of short (3 s) and long (11 s) exposures taken within one hour of each other (100 pairs before the pandemic shutdown; 50 pairs after the pandemic shutdown). Each file contains ten exposures, which we median combined. We took the difference

image between corresponding long and short exposures to reduce nonsensitivity related effects and subsequently normalized them by their median while excluding already identified bad pixels. We discarded 17 visibly discrepant flat-fields pre-shutdown (none post-shutdown) that had a very uneven illumination of the detector, for example. We performed no further pre-selection of frames. The remaining frames were median combined to obtain the respective master flat. We distribute the two flat-fields (before and after the shutdown) with the pipeline. The detector flat-field is stable with time at a level $\lesssim 0.35\%$. We do not divide our image by this detector flat, but rather use the flat later when extracting spectra: we multiply the PSFlet model counts by the flat-field rather than dividing the data by the flat-field.

The left panel of Fig. 2 shows the master detector flat-field distributed with our pipeline. The flat is normalized to a median intensity of one. The flat-field is not uniform, but it shows little variation with time. We characterized this variation using principal component analysis (PCA): the right panel of Fig. 2 shows one principal component of an unranked empirical PCA decomposition (Bailey 2012) of thousands of individual flat-field images taken at different times. This image highlights the structures that normally are not visible in individual exposures, providing a visual overview of the existing substructures in the flat, such as small dust on optical elements. The vertical stripes show the slightly different sensitivity of different readout channels of the H2RG detector.

The lenslet flat-field corrects for the variable transmission of lenslets in the lenslet array. This requires a different approach from the detector flat-fielding described in this section. Instead, we used images with SPHERE illuminated by a uniform white-light source (integrating sphere). This produces microspectra that must be extracted to produce the lenslet flat-field. We therefore defer discussion of the lenslet flat-field to Sect. 4.4 after we detail our approach to spectral extraction.

3.3. Undispersed background template

Because the SPHERE IFS is not cryogenic (except for the detector), background levels are high. Much of this background is undispersed and illuminates the detector relatively uniformly. We wish to remove this background before extracting the spectrum of each lenslet. All frames obtained with SPHERE are already corrected for additive bias by the read-out electronics.

Standard SPHERE observing sequences take a small number of sky frames after the science frames. These sky frames may be used to remove both the undispersed background and the dispersed background (from the sky itself). Unfortunately, using a sky frame for this purpose adds its noise to the science frame. If both the sky frame and the science frame are of the same exposure time, this step increases the minimum photon noise by a factor of $\sqrt{2}$. Sky frames taken after the coronagraphic exposure also display noticeable persistence, which introduces additional noise.

In an effort to overcome this limit, we have constructed a master template for the undispersed background. We seek a background without any dispersed light (e.g., from the sky), therefore we began with internal instrument background frames that do not contain sky. We took 234 randomly selected background frames taken between 2015 and 2021 and performed PCA on them. The number of frames was limited by the 64 GB of RAM on our computer and the need to perform PCA. These frames come with the instrument in YH-mode. We repeated the exercise with YJ backgrounds and found that YH-derived calibrations fit YJ data just as well as YJ-derived calibrations do. We used about 98%

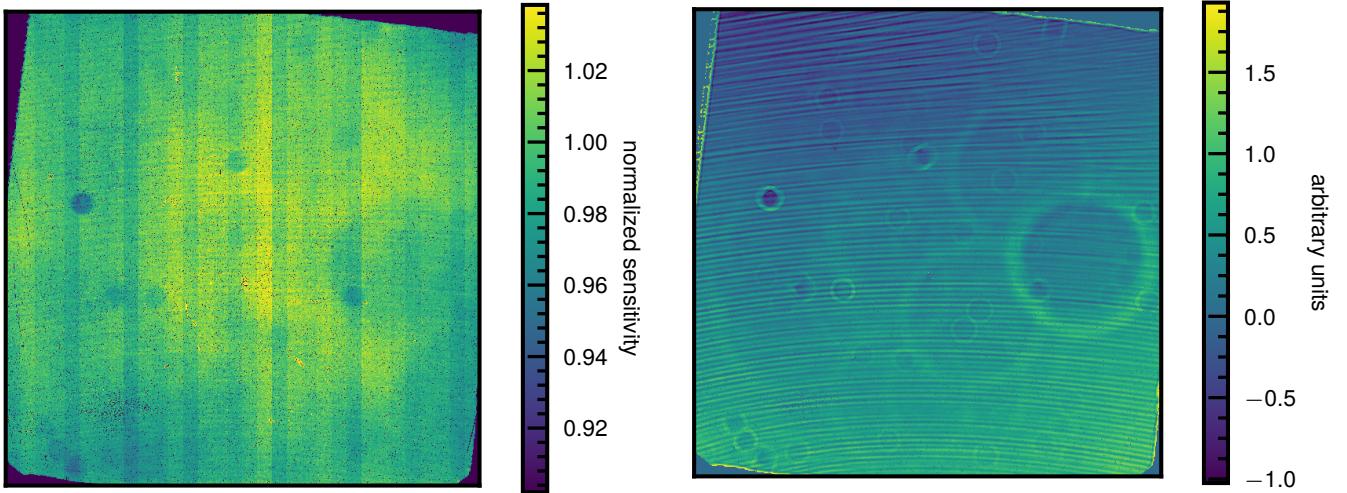


Fig. 2. Overview of structures visible in the detector flat-field. *Left*: master flat-field generated by averaging hundreds of individual flat-field frames. *Right*: one component of an unranked empirical principal component decomposition of the temporal variance of the flat-fields into five basis vectors. This component highlights minute details of the flat-field variability. Details that are normally not clearly visible become apparent due to the changes in illumination by the internal calibration lamp affecting diffraction patterns over time. This includes the fringe-like horizontal stripes, as well as dust diffraction by various optical surfaces.

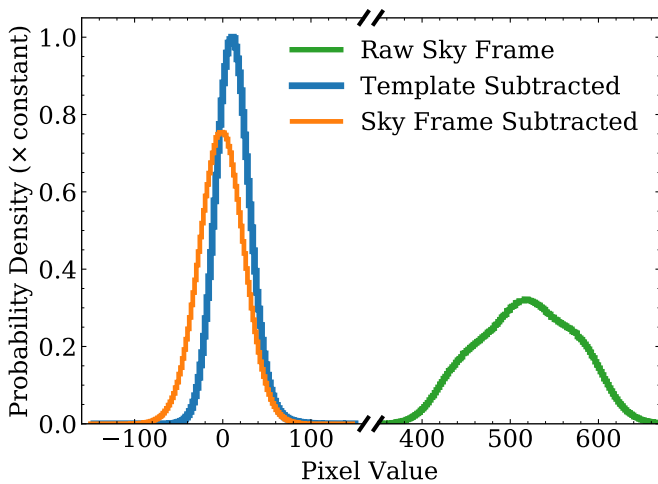


Fig. 3. Distribution of pixel values in a raw 64-s sky frame (medium green histogram) after subtraction of another sky frame (thin orange histogram) and after subtraction of a background template (thick blue line).

of the pixels to perform PCA, masking the 1% with the highest variance of the background frames and those with less than half or more than twice the average background level.

A mean plus three principal components offers excellent performance and accounts for slightly more than half of the mean-subtracted variance in the background frames. Additional principal components offer little improvement. We fit a linear combination of these four images to the four corner regions of the detector that are unilluminated by lenslets. In this way, we can fit our background templates to science frames and sky frames and preserve any light that is dispersed by the lenslet array and avoid the addition of read noise and photon noise.

Figure 3 shows the distribution of pixel values in a 64-s sky frame before and after background subtraction. Subtracting another sky frame produces residuals that are centered about zero (the thin orange histogram) but adds read noise and photon noise. A different sky exposure may also contain varying levels

of persistence, based on how brightly the coronagraphic images were exposed, which adds additional noise. Subtracting the best-fit four component template produces a narrower distribution (thick blue histogram) because a negligible amount of additional noise is added. This distribution is centered on a positive pixel value because it does not subtract the dispersed sky background and contains detector persistence that is intentionally not included in our model.

Our SPHERE IFS pipeline implements the template subtraction described here for the undispersed background. This approach minimizes the noise in the resulting data cubes. A dispersed sky background does appear as a constant, uniform background in the extracted data cubes. However, this is typically indistinguishable from the halo of the stellar PSF and is easily removed by postprocessing algorithms. Only when a library of PSFs is used that were taken on different nights might this introduce a small bias for extended sources if the dispersed background is significantly different. In this case, a constant may be fit to model the background. Constructing a reference library of sky background frames is not feasible, as these are generally taken after the coronagraphic observations and show an imprint of the stellar halo due to detector persistence. This is another general issue with the normal sky subtraction approach that is typically performed, which adds a constant oversubtraction bias to all frames.

3.4. Oversampled PSFlet templates

Spectral extraction requires at least a wavelength solution: a mapping between wavelength, lenslet, and location on the detector. Many extraction techniques also require knowledge of the monochromatic PSF of each lenslet. Optimal extraction requires the profile of this PSF (integrated along the dispersion direction), while least-squares extraction techniques require the full lenslet PSF henceforth, also referred to as PSFlet for convenience. The standard SPHERE calibration images include either three or four wavelengths simultaneously, creating an image like that shown in the left panel of Fig. 1. The simultaneous presence of multiple wavelengths is not a major problem for deriving a

wavelength solution, but it makes it extremely difficult to extract the monochromatic lenslet PSF.

In order to measure the monochromatic lenslet PSFs, we have obtained calibration images in which the detector was uniformly illuminated by a single laser at a time. The right panel of Fig. 1 shows an inset of one such image. The lenslet PSFs are now separated by ≈ 15 pixels, enabling their measurement out to ≈ 7 pixels from the center.

We followed the same approach as B17 in order to extract the lenslet PSF positions and to extract the oversampled PSFlets. Briefly, we began with the coordinates of the lenslets. The side-length t of the hexagons is 7.4 mas (see Table 1) and the hexagon centers can be represented in a rectilinear grid for lenslet number (i, j) as

$$\begin{aligned} x_{ij} &= t\sqrt{3}(i + 0.5(i \bmod 2)), \\ y_{ij} &= t j. \end{aligned} \quad (1)$$

This corresponds to the representation internal to the pipeline. We then assumed a wavelength-dependent 2D third-order polynomial mapping between the x_{ij} and y_{ij} lenslet coordinates from Eq. (1) and the PSFlet locations on the detector. This differs from B17 only in the coordinates of the lenslets themselves, which in this case reflect the hexagonal geometry of the SPHERE lenslets (CHARIS has square lenslets).

We lightly smoothed the monochromatic background-subtracted laser images by convolving them with a narrow Gaussian and then fit for the polynomial coefficients that maximize the peak intensity at the fit position of each lenslet on the detector. In other words, we determined the polynomial functions of the lenslet coordinates (x_{ij}, y_{ij}) that give the most accurate pixel locations of the laser spots shown in Fig. 1. These polynomials then give the mapping between lenslet coordinate and PSFlet centroid on the detector at a given wavelength.

After fixing the mapping between lenslet and PSFlet location, we used an approach similar to that of Anderson & King (2000) to derive an oversampled effective PSF. Briefly, we constructed a grid that was oversampled by a factor of 9 and placed each empirical PSFlet at the appropriate position within this grid. We then took this sparsely sampled grid, convolved with a Gaussian to smooth it, and then performed a deconvolution. This step accounts for the varying signal-to-noise ratios (S/N), where some offsets are sampled more thoroughly than others depending on the exact positions of the PSFlets. We constructed the empirical oversampled PSFlets separately in 25 regions of the detector to account for potential spatial variations of the PSFlets. We refer to B17 for additional details.

Figure 4 shows our resulting oversampled lenslet PSFs. The PSFs have a central core surrounded by six diffraction spikes; these result from the hexagonal lenslet geometry. The PSFlets are more symmetric than those of CHARIS (B17) and vary little across the FoV. This suggests a uniform focus for the SPHERE lenslets throughout the FoV. These lenslet PSFs include both optical and pixel sampling effects. Interpolating them gives an empirical PSFlet at each position. In practice, we use bilinear interpolation spatially to obtain the PSFlet for a given lenslet, and then interpolate within the oversampled PSFlet to obtain model monochromatic PSFlets. These may be used to measure the instrumental profile for optimal extraction or to reconstruct the microspectra themselves for a least-squares extraction.

The instrument profile, that is, the PSFlet profile perpendicular to the dispersion direction, is required for optimal extraction. The oversampled PSFlets make this straightforward to determine. The instrumental profile is simply the PSFlet collapsed along the dispersion direction (vertical in Fig. 4).

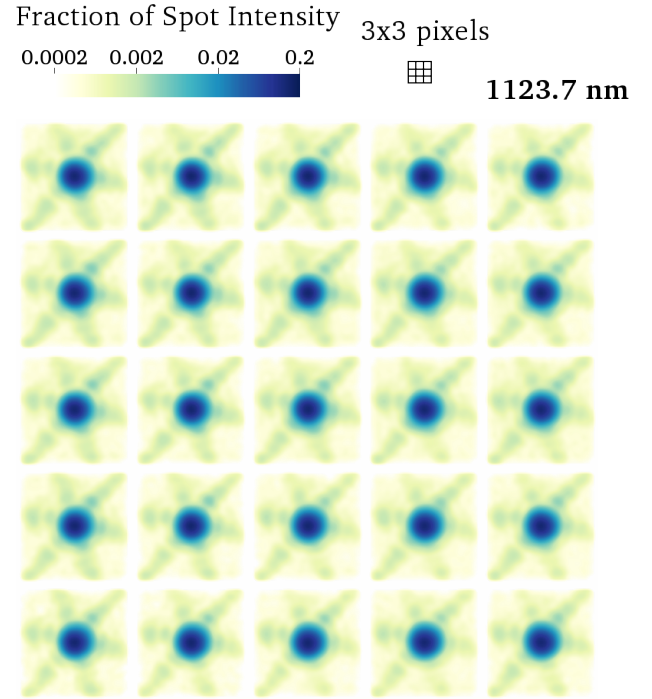


Fig. 4. Oversampled PSFlets at 1123.7 nm reconstructed from monochromatic laser images like that shown in Fig. 1 and normalized to unit intensity after pixel sampling. The detector is divided into 25 regions, each panel corresponding to the associated location on the detector: the top left PSFlet corresponds to the top left corner of the detector, etc. A 3×3 pixel grid is shown as a size comparison. Similar libraries are created at the other three calibration wavelengths and included with the pipeline. The oversampled PSFlets can be used to construct a model of the pixellated microspectra corresponding to monochromatic or broadband light imaged by the lenslet array. The six diffraction spikes result from the hexagonal lenslet geometry, while the PSFlet shapes are extremely homogeneous across the FoV and give no indications of a changing focus.

3.5. Wavelength calibration

Figure 5 shows an inset of a raw SPHERE IFS image; the image consists of microspectra arranged in the geometry of the lenslet array. One of the most important steps in the data reduction is therefore assigning each pixel membership to a specific spectrum (lenslet) and the wavelength corresponding to the position in the dispersed spectrum. We refer to these steps collectively as obtaining the wavelength solution.

Section 3.4 describes the extraction of oversampled PSFlets from a series of monochromatic laser images. The first step of extracting these PSFlets is to derive the mapping between lenslet and detector position at a given wavelength. This results in anchors to the wavelength solution at either three (in YJ-mode) or four (in YH-mode) wavelengths. We treat the full wavelength solution as a quadratic in $\log \lambda$ for YJ-mode and a cubic in YH-mode. This results in interpolation in both cases.

Figure 6 shows the measured PSFlet positions in the center of the detector at the laser calibration wavelengths for the YJ- and YH-modes. It also shows the deviations from constant spectral resolution, that is, a constant number of pixels per logarithmic unit of wavelength. These deviations reach a maximum of ≈ 0.3 pixels, or about $0.005 \mu\text{m}$, across most of the instrument spectral range. They reach ≈ 1 pixel, or about $0.02 \mu\text{m}$, at the long-wavelength extreme of the YH-mode bandpass. The lack of a calibration laser at a wavelength longer than $1.54 \mu\text{m}$

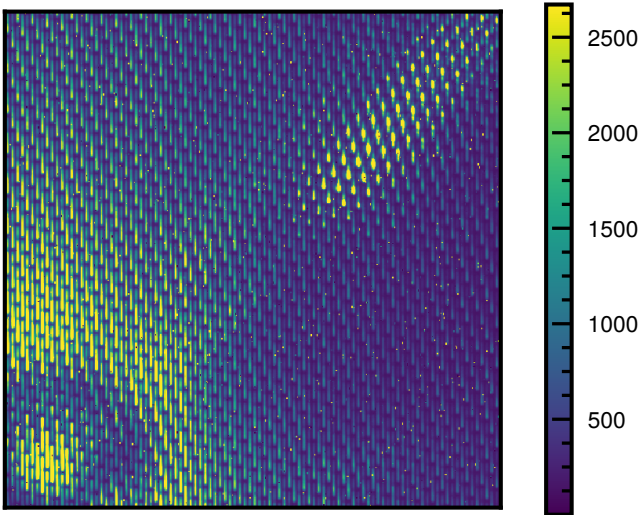


Fig. 5. Cutout of unprocessed IFS raw data of 51 Eri b in YH-mode with satellite spots induced by the deformable mirror. The dispersion is in vertical direction from longer wavelengths at the top to shorter wavelengths at the bottom. The star center is in the bottom left corner.

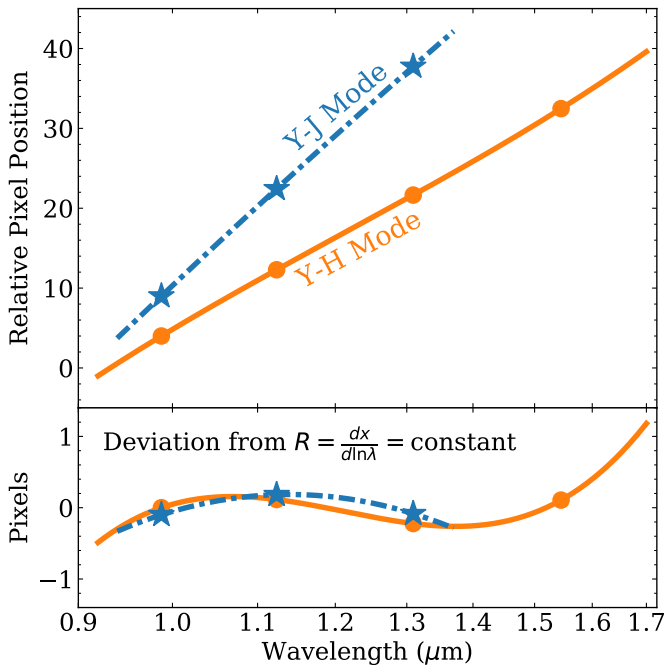


Fig. 6. Wavelength solution derived from the positions of the three (in YJ-mode) or four (in YH-mode) calibration lasers for the central lenslet. The YJ and YH lines are offset from one another for clarity. *Bottom panel:* deviations from a constant spectral resolution, i.e., a constant number of pixels per logarithmic wavelength (the straight line in the *top panel*).

means that the wavelength solution in this spectral range is an extrapolation and should be treated as uncertain.

The actual wavelength solution used by the pipeline includes the full third-order 2D polynomial fit described in Sect. 3.4. This gives a PSFlet position for each lenslet and each wavelength. Each lenslet then has either three (in YJ-mode) or four (in YH-mode) pixel-wavelength (x, λ) pairs. The unique second- or third-order polynomial $x(\ln \lambda)$ then is the wavelength solution. This part of the wavelength solution is distributed with the pipeline and treated as static: it is unchanged for all reductions of

SPHERE IFS data. The first step with a new set of raw SPHERE data is to determine the spatial shift from the static wavelength solution described above. We assume that this shift is small and can be represented by a horizontal shift, a vertical shift, and a rotation angle. For this step, we use the images taken as part of the standard SPHERE observing sequence, with either three or four calibration lasers. We first mask bad pixels in these images and then lightly smooth them with a Gaussian. We then obtain the pixel values at the locations tabulated in the static wavelength solution for all lenslets. The sum of the smoothed images at all of these locations is our figure of merit. We apply a shift in x , a shift in y , and a rotation to these positions in order to maximize this figure of merit. The best-fit shifts are small in practice: ≤ 1 pixel and $\leq 10^{-4}$ radians. These shifts are then combined with the static calibration data to describe the full wavelength solution for a given set of data. Small shifts are correlated with the temperature of the instrument due to thermal expansion and contraction. Larger shifts can occur rarely after an earthquake occurs close to Paranal.

3.6. Calibrating a raw frame

The preceding sections described the process of constructing a bad-pixel mask, a detector flat-field, a background template, a wavelength solution, and oversampled PSFlet models. All of these are static and distributed with the pipeline, except for a small shift in the wavelength solution for each set of observations. For an individual raw SPHERE IFS image, we use some of these data products before proceeding to spectral extraction. In this section, we describe these initial steps.

Our first step in reducing a raw SPHERE IFS frame is to mask bad pixels. We then subtract a background using our background template: we use least-squares fitting to the four regions in each corner of the detector that the lenslets do not illuminate (as described in Sect. 3.3). This produces a background-subtracted masked image.

We do not perform flat-fielding at this stage, but defer that step to spectral extraction. With a background-subtracted image, masked bad pixels, and a wavelength solution appropriate to a given frame, we proceed to the spectral extraction step. We defer the discussion of the lenslet flat-field to Sect. 4.4 because constructing this flat-field requires spectral extraction of white-light calibration images.

4. Extracting a data cube

In this section we present a step-by-step discussion of the process of extracting a spectral data cube from raw data using our adaptation of the CHARIS pipeline. A SPHERE IFS frame, after background removal, consists of $\approx 20\,000$ microspectra that are arranged in a grid across the detector. The core step in the data reduction is to extract each of these spectra and reconstruct the 3D (x, y, λ) data cube. The wavelength solution, described in Sect. 3.5, gives the location, or trace, of each microspectrum. Each is ≈ 39 pixels long on the detector, and the dispersion direction is very nearly parallel to the vertical dimension of the detector. The trace of each spectrum is separated by ≈ 5 pixels from its nearest neighbor.

The SPHERE IFS pipeline described here implements three approaches to extract the microspectra: boxcar extraction, optimal extraction, and least-squares extraction. The DRH (ESO) pipeline, as well as the GPI pipeline (Perrin et al. 2014), implement only boxcar extraction. In this section we summarize each

approach in turn. The approaches and algorithms we use are substantially identical to those used for CHARIS; we refer to B17 for detailed discussions.

4.1. Boxcar extraction

The first spectral extraction technique implemented for SPHERE is boxcar extraction, a technique similar to aperture photometry. In this approach, the pipeline adds all flux perpendicular to the dispersion direction out to a certain number of pixels from the trace. With a spacing of ≈ 5 pixels between traces, a typical choice would be to add the flux in a 5-pixel-wide aperture centered on a given spectrum's trace (2 pixels out in each direction). The total flux is then assigned to the wavelength of the central pixel in this box, the one along the trace. This wavelength is known from the wavelength solution. Each microspectrum is then defined on a different wavelength array: the wavelength corresponding to integer pixels along the trace of the microspectrum. The final step is then to interpolate these microspectra onto a common wavelength array.

Boxcar extraction has several important limitations. The first is that it has no way of treating bad pixels. Bad pixels must be corrected for and then treated exactly as if they were good pixels. The second limitation is that it applies the same weight to flux in the core of the trace, where the S/N is higher, as to the outskirts, where it is lower. A wider aperture captures more of the flux, but at the expense of S/N. This limitation is minor in the limit of a high S/N, but when read noise and background noise are significant, boxcar extraction entails a significant S/N penalty relative to other approaches. Finally, boxcar extraction lacks a straightforward way of treating crosstalk, where light from one microspectrum contributes flux to the aperture around a neighboring microspectrum. The DC and v1t-sphere pipelines implement a crosstalk correction before boxcar extraction.

4.2. Optimal extraction

Other spectral extraction approaches offer a superior performance to boxcar extraction if the spectral profile perpendicular to the trace is known. In Sect. 3.4 we describe the calculation of this profile using monochromatic PSFs out to ≈ 7 pixels from the trace. This captures the profile ≈ 2 pixels past the trace of the neighboring spectrum. This provides a way to implement optimal extraction (Horne 1986) and least-squares extraction (B17), and to implement a crosstalk correction by iterative modeling of the full 2D array of microspectra.

The implementation of an optimal extraction in the SPHERE IFS pipeline works similarly to boxcar extraction, but weights the flux from each pixel by the profile calculated from the PSFs. We include the pixel flat-field in this profile for each lenslet, that is, the pixel response is included in the model rather than in the data.

Each spectrum contains light from its neighbors, both along and perpendicular to the dispersion direction. This light is often called spectral crosstalk. We remove crosstalk in optimal extraction by first performing a least-squares extraction as described in the following subsection. This involves a full modeling of the 2D array of microspectra, including an attribution of light to each lenslet. We can then remove the light extending more than 2 pixels from the trace of each microspectrum before performing optimal extraction.

After optimal extraction, each microspectrum remains defined on its own unique wavelength array corresponding to the central wavelengths of each pixel along the trace. In the

same way as for the boxcar extraction, we take the final step of interpolating all microspectra onto a common wavelength array.

4.3. Least-squares extraction

The final extraction technique implemented in our SPHERE IFS pipeline is least-squares extraction, which was successfully implemented for CHARIS by B17. In this approach, we model each microspectrum in 2D. We take the model spectrum incident on a lenslet to be locally flat over short ranges of wavelength. The actual wavelength range corresponds to one spectral resolution element; we discuss the choice of extracted spectral resolution below. We then assume that the full spectrum is a linear combination of these top-hat spectra in wavelength space. Because we know the lenslet PSFs, these top-hat spectra over small ranges of wavelength correspond to PSFs in detector space that have been dispersed by one spectral resolution element. A set of coefficients for the amplitudes of these slightly dispersed PSFs corresponds to a model microspectrum on the detector. Extracting these coefficients becomes a least-squares problem of fitting the microspectra themselves. Least-squares extraction implicitly deconvolves the spectra with the line-spread function. Because the top-hat spectra are defined over the same wavelength ranges for each microspectrum, least-squares extraction does not require interpolation onto a common wavelength array.

We adopted the same procedure for crosstalk correction in least-squares extraction as for optimal extraction. This procedure was also described in B17. We first extract each microspectrum using the procedure described in the previous paragraph. We model each microspectrum using the dispersed PSFs, fitting the model microspectra over a region extending 2 pixels perpendicular to the trace in each direction. These microspectra include light from their neighbors as spectral crosstalk. Our model PSFs, however, extend 7 pixels from the trace. We compute the 2D image expected from our extracted spectra, now adding the light out to this larger extent. This full image has more photons than the actual measured image because photons from several lenslets contributed to the flux in the core of each microspectrum. For optimal extraction, we subtract the observed detector image from the full model image; this gives the overcounting of photons due to spectral crosstalk. We then subtract this residual from the actual image before performing optimal extraction. For the least-squares extraction, we fit this residual 2D array of microspectra exactly as we did for the original array of microspectra. We then subtract the residual 1D spectra from the previously extracted 1D spectra to perform the crosstalk correction.

Least-squares extraction has a number of advantages over other approaches in principle. It models the full 2D detector image, allowing it to account for (and fit) read noise patterns (B17). Least-squares extraction can naturally account for bad or missing data, using the true uncertainties on the measured flux of each pixel. It also performs a deconvolution of the microspectra with the line-spread function. This deconvolution comes at a price, however. The least-squares extracted microspectra are covariant at neighboring wavelengths, and this covariance can become very strongly negative when attempting to deconvolve to a higher spectral resolution. The SPHERE IFS pipeline extracts the microspectra at a resolution $R = 55$ for YJ-mode and $R = 35$ for YH-mode, corresponding to ≈ 2 pixels along the dispersion direction. This represents a compromise: A higher spectral resolution will enable better modeling of the detector image and return a higher-resolution data cube, but at the cost of stronger covariances between neighboring wavelengths. To

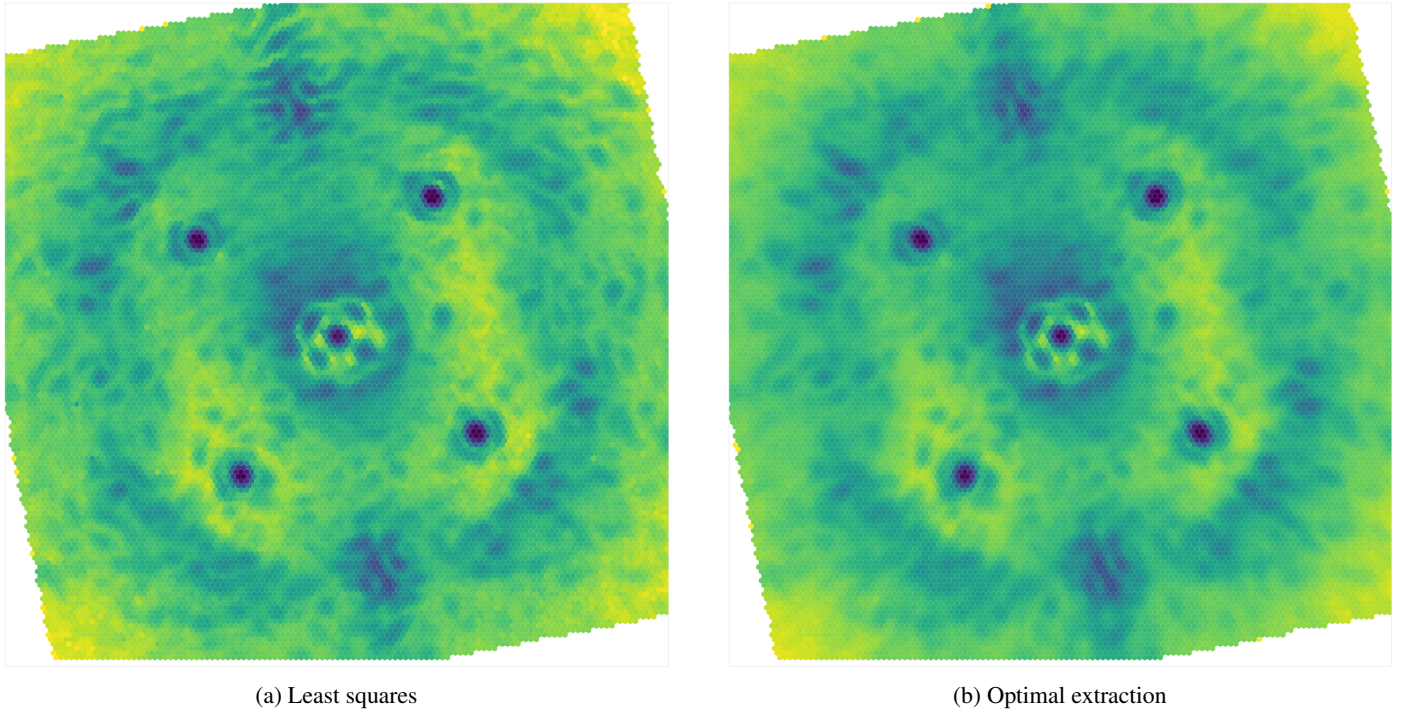


Fig. 7. Single slice through the same YH-mode data cube extracted two ways and shown in the original hexagonal geometry: (a) Least-squares extraction, and (b) optimal extraction. Least-squares extraction performs a deconvolution with the line-spread function and as such appears less smoothed in general, but may appear more noisy in the low-S/N regime for this reason.

save disk space, the pipeline does not save the full covariance matrix for each spectrum (which would multiply the file size by more than a factor of 10). Depending on the needs of the user and on the properties of a data set, either optimal extraction or least-squares extraction is likely to offer the best-quality data cubes.

Figure 7 shows an extracted image using the two main methods. The speckle pattern using the least-squares method is less blurry because it performs a deconvolution with the line-spread function and does not require interpolation in the spectral dimension. Both images are shown in the original hexagonal geometry in which the data are taken.

4.4. Lenslet flat-field

The lenslet flat-field is distinct from the detector flat-field we described in Sect. 3.2. The lenslet, or IFU, flat-field tells us the relative throughput of each lenslet. The SPHERE DRH pipeline uses images in which the lenslet array is uniformly illuminated by white light to construct this flat-field.

The CHARIS instrument uses the wavelength-averaged (near-)monochromatic spot amplitude patterns obtained with the supercontinuum source and a tunable narrow-band filter directly as an indicator for the lenslet transmission. Unfortunately, this is not possible for SPHERE as the lasers are coherent light sources and susceptible to fringing and spurious amplitude variations across the lenslet array. This is not an issue when we make the PSFlet library as the models are normalized, but they are not reliable indicators for absolute transmission. We therefore used the same approach as the SPHERE DRH pipeline “instrument flat-field” step, which corresponds to our lenslet flat-field, and extracted the spectra for a white lamp illuminating the IFS using the optimal extraction algorithm described above. We then

normalized by the average of all lenslets and divided the output data cube by the wavelength-averaged resulting flat-field.

4.5. Lenslet geometry

The native lenslet geometry of the SPHERE IFS is hexagonal. This differs from the similar CHARIS and GPI IFSs, which both contain square lenslet arrays. It also differs from the rectilinear geometry that is most natural to a numerical array. For this reason and for compatibility with the many software tools now available to postprocess high-contrast imaging data (e.g. Wang et al. 2015; Cantalloube et al. 2015; Gomez Gonzalez et al. 2016; Galicher et al. 2018; Dahlqvist et al. 2020; Samland et al. 2021), the SPHERE DRH interpolates data cubes spatially onto a rectilinear grid.

The SPHERE IFS pipeline that we present here offers two options. The first is to interpolate the data cubes onto the same rectilinear grid geometry as that used by the SPHERE IFS pipeline. The second option is to preserve the native hexagonal lenslet geometry, which means that each physical hexagonal lenslet corresponds to an individual hexagonal spaxel in the output image. In this image, the midpoint of each spaxel corresponds to a position on a hexagonal grid. The side-length t of the hexagons is 7.4 mas (see Table 1), and the orientation of the hexagon grid is in the so-called pointy topped configuration, as opposed to flat topped. The hexagon centers are given by Eq. (1).

The hexagonal grid is resampled onto a rectilinear grid using the Sutherland-Hodgeman algorithm, which can be used to compute the area of overlap between two polygons (in this case, the original hexagon and each pixel of the new rectilinear grid). The area of overlap between pixels of the two grids are only computed once and saved in a calibration file that is provided with the pipeline. The areas of overlap are then used to divide the flux of the original image onto the rectilinear pixel grid.

We include routines that allow a user to visualize the extracted cubes in the native hexagonal geometry without resampling, as most normal image viewers like DS9 can only view images with square pixels. A command-line script allows quick access to viewing an extracted data cube, similar to the other routines of the package. It allows the interactive viewing of an extracted image cube with a slider for wavelength.

4.6. Astrometric and spectrophotometric calibration

After the spectral cube is extracted, we extract the location of the occulted central star using the four satellite spots, which are generated using the deformable mirror for this purpose. We fit a 2D Gaussian to each spot in the resampled images and compute the point at which the lines connecting the spots on opposite side of the star cross to obtain the image center for each wavelength. For this purpose, we adapted the routines found in [Vigan \(2020\)](#). The image center positions are then fit by a polynomial of suitably high order in the wavelength dimension (second order in x -direction and third order in y -direction) to obtain the final result³. The first and last wavelength channel and the channels that are most strongly affected by atmospheric transmission (≈ 1.14 and ≈ 1.37 micron, see, e.g., the lower panels of [Fig. 9](#)) are excluded from the fit. These channels have a lower S/N due to low atmospheric transmission, and in the case of the first and last channels, low instrumental transmission. The measured position can be biased by the gradient in the transmission profile, which is significant in all of these channels. These channels, due to their inherently lower S/N, are also more prone to residual spectral crosstalk. The crosstalk modeling is discussed in [Sect. 4.3](#).

For the spectrophotometric calibration, we used the unsaturated noncoronagraphic PSF images obtained before and after the coronagraphic sequences. We accounted for the neutral density filter throughput at each wavelength and adjusted the integration time to match the coronagraphic observations. We performed aperture photometry on the PSF ($r = 3$ pix) and subtracted the sky background as determined by an annulus between $r = 15$ – 18 pix. We then scaled the flux of all PSF images to match the two PSFs that are closest in time to the coronagraphic sequence. The average of the scaled and background-subtracted PSF is used as our planet model PSF in postprocessing. We performed the same aperture photometry on the satellite spots to determine their flux and calibrate them by matching the counts of the first and last satellite spot image to the flux image closest in time, respectively. For observations that continuously use the satellite spots during the coronagraphic sequence, we have a fully flux-calibrated sequence that allows adapting the planet model for postprocessing for each frame and wavelength. For those sequences that only obtain satellite spots at the beginning and end of the coronagraphic sequence, we can still use this method to effectively move the flux calibration closer in time to the data that are used in postprocessing.

4.7. Summary of changes to the original CHARIS pipeline

The adaptation of the CHARIS pipeline [B17](#) to the SPHERE IFS required a number of changes that we summarize here. The pipeline can now extract data cubes from either instrument.

The original CHARIS pipeline was written in Python 2.7; we have modernized the entire code-base to be compatible with

³ More information about the astrometry and optical distortions of the instrument can be found in [Maire et al. \(2016\)](#).

new Python versions. To enable it to handle both SPHERE IFS and CHARIS data, we isolated all instrument-specific and hard-coded variables and moved them to a new `Instrument` module. Depending on the data provided to the pipeline, the corresponding instrument class is instantiated and changes the behavior of the pipeline according to the instrument and instrument mode of the data. The pipeline behavior then reflects the lenslet geometry, wavelength range and sampling, instrument-specific header keywords, and properties of the observatory location and average atmospheric transmission.

The construction of calibration files using the `buildcal` script is substantially identical to the original CHARIS pipeline. The appropriate calibration templates are specified by an instance of the respective instrument class. These calibration templates, including oversampled PSFs, background templates, and both detector and lenslet flat-fields, were derived as described in [Sect. 3](#). The spectral extraction step incorporates the new undispersed background-fitting routine, which currently only works for SPHERE IFS. To account for the hexagonal lenslet geometry, we modified the bad-lenslet detection routine, which cosmetically replaces outlier lenslets by its neighbors in the extracted spectral data cube, as there are six equally distant neighbors as opposed to nine unequally distant neighbors in the rectilinear geometry. Last, we implemented a routine to resample the hexagonal spaxels into a rectilinear grid and included tools to plot the spectral data cube in its native hexagonal geometry.

5. Results

In order to show the performance of our adaptation of the CHARIS pipeline for SPHERE IFS data, we extracted image cubes and compared the results with the extracted image cubes obtained from the SPHERE Data Center reduction, which is based on the ESO DRH pipeline and additional refinement tools ([Delorme et al. 2017](#)). We first briefly discuss the visual differences in the extracted images of the two pipelines and the two main extraction methods in our pipeline (optimal extraction and least-squares extraction).

We then proceed to show the performance of the pipeline by performing two tests. First, we use both pipelines to extract a high-S/N spectrum of the known white dwarf companion HD 2133B to a main-sequence star to ascertain the reliability of the spectral extraction requiring minimum postprocessing. Second, we perform the same ADI postprocessing on a high-contrast imaging sequence of the exoplanet 51 Eridani b using both pipelines. The postprocessing parameters and pipeline (TRAP, [Samland et al. 2021](#)) are kept the same in both cases. We then compare the spectra and contrast curves.

5.1. Visual comparison of extracted images

As a first step, we extracted a single wavelength slice from a data cube in the 51 Eri b SPHERE IFS data set taken on the 2015 September 25 ([Samland et al. 2017](#)) in the YH-mode. The data set was taken in below median conditions and shows a relatively strong low-wind effect ([Couder 1949](#); [Sauvage et al. 2016](#); [Milli et al. 2018](#)). [Figure 8](#) shows two wavelength slices (top: ≈ 1065 nm, bottom: ≈ 1372 nm) of the extracted image cubes using the CHARIS-based pipeline with optimal extraction in hexagonal geometry (left panel) and resampled on a rectilinear grid (middle panel), and the extraction using the DRH-based Data Center reduction (right panel). [Section 4](#) describes the extraction process in detail.

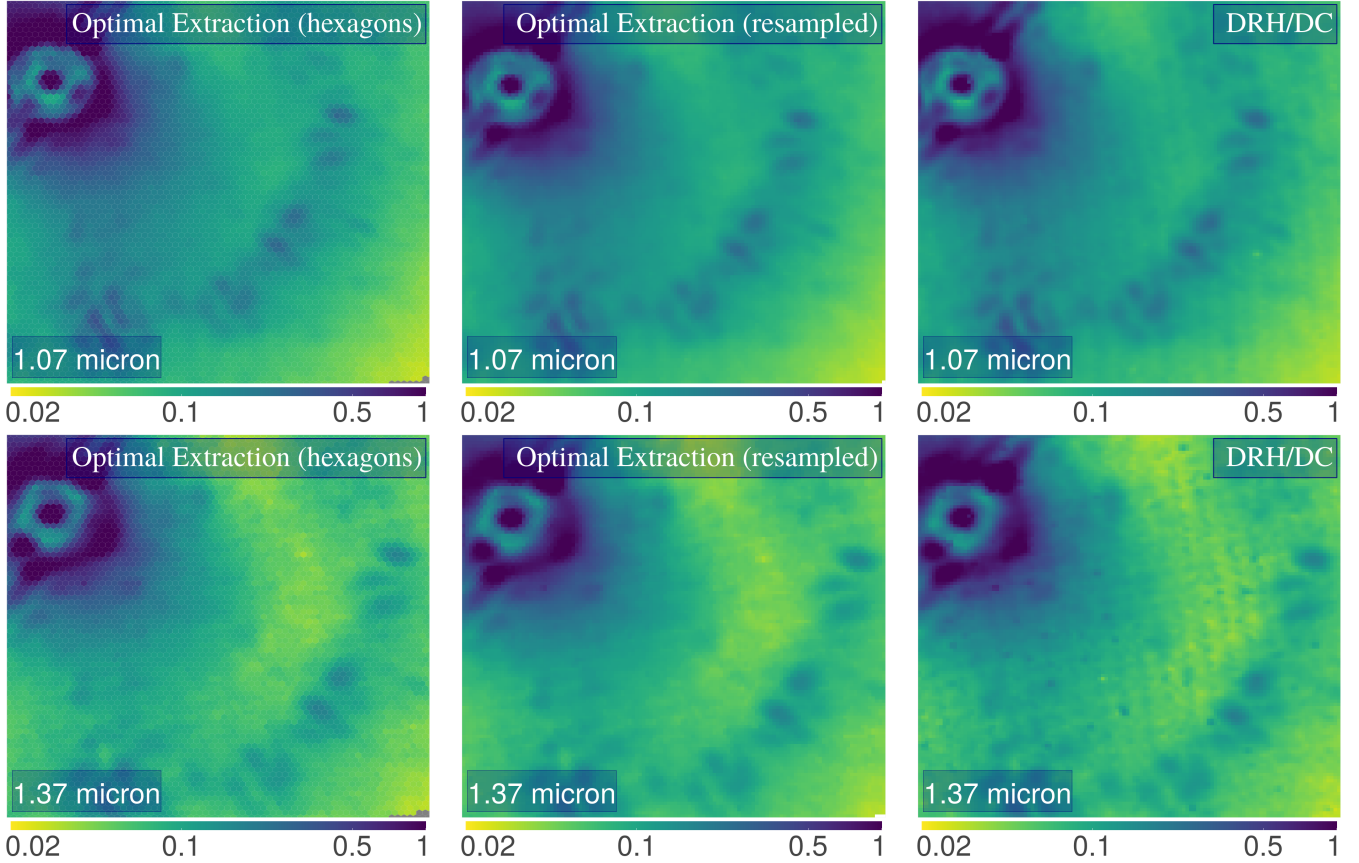


Fig. 8. Comparison of normalized image slices extracted from the 2015 51 Eri data set with our pipeline and the DRH/DC for two wavelengths: one at high S/N (*top row*) and one at low S/N (*bottom row*). *Top row*: at the peak of emission at around 1065 nm, shown in logarithmic scale. *Bottom row*: same for the channel with the lowest S/N, at around 1370 nm. *From left to right*: optimal extraction frame (YH-mode) in hexagonal geometry; the same frame resampled to rectilinear grid; DRH pipeline reduction with additional routines for crosstalk correction and wavelength correction from the SPHERE DC.

Figure 8 shows the overall behavior of the different extraction approaches. In high-S/N channels, our pipeline extractions are visually similar to the DRH/DC extraction. The largest difference appears in the lowest-S/N channel. In the case of DRH, a stripe pattern is visible from the spectral extraction. This pattern gradually becomes more dominant as the S/N decreases, but is present in all channels at low amplitude. This pattern is not visible in our pipeline reductions. In data sets with very high detector counts ($>20\,000$), the crosstalk correction in all reductions creates a negative feature at the location of the brightest speckles at shorter wavelengths in the lowest S/N channels, especially at the location of satellite spots. This might indicate that the satellite spots, if strongly exposed, may not be perfectly reliable.

5.2. Spectrum of the white dwarf companion HD 2133B

Our first quantitative test of the new cube extraction pipeline uses the white dwarf HD 2113B. The white dwarf has an effective temperature of $\approx 28\,000$ K (Burleigh et al. 1997) and orbits the F7 star HD 2133 at a separation of about $0''.68$. It is bright enough to be easily visible in individual exposures and wavelength channels of the coronagraphic sequence. Hot white dwarfs have relatively simple spectra on their Rayleigh-Jeans tails in the near-infrared, making them good candidates for understanding spectrophotometric calibrations.

We extracted data cubes of HD 2133B using the DRH pipeline, our CHARIS-based pipeline with optimal extraction

and least-squares extraction. In all cases, we resampled the data onto a rectilinear grid to facilitate postprocessing. Because the white dwarf is located at the adaptive-optics correction radius, we still performed a simple postprocessing step to reduce speckle noise while at the same time extracting the contrast of the white dwarf relative to its main-sequence binary component. To do this, we ran the TRAP temporal detrending algorithm without any principle components, which means that only an offset was fit together with the forward model of the PSF for each pixel that was affected by signal of the white dwarf companion. This corresponds to a temporal version of classical ADI (Marois et al. 2006), except that we simultaneously fit a temporal forward model and an offset to each affected pixel using the known PSF instead of subtracting a median image from each frame. We performed this parameter-free extraction of the white dwarf contrast for all three reductions (optimal and least-squares extraction and DRH/DC).

We converted the measured contrast into a flux using a BT-NextGen model for the star. Based on comparing different combinations of model parameters for both the star and white dwarf taken from the literature with our data, we determined that the stellar parameters of $T_{\text{eff}} = 6400$ K, $\log g = 4.3$, solar metallicity, and a white dwarf model of $T_{\text{eff}} = 26\,000$ K, $\log g = 8.0$, provide a close fit to the data. The normalized fluxes were compared both to a simple blackbody of 26 000 K and to a theoretical white dwarf spectrum from the Levenhagen et al. (2017) grid. Figure 9 shows the extracted white dwarf spectra compared to

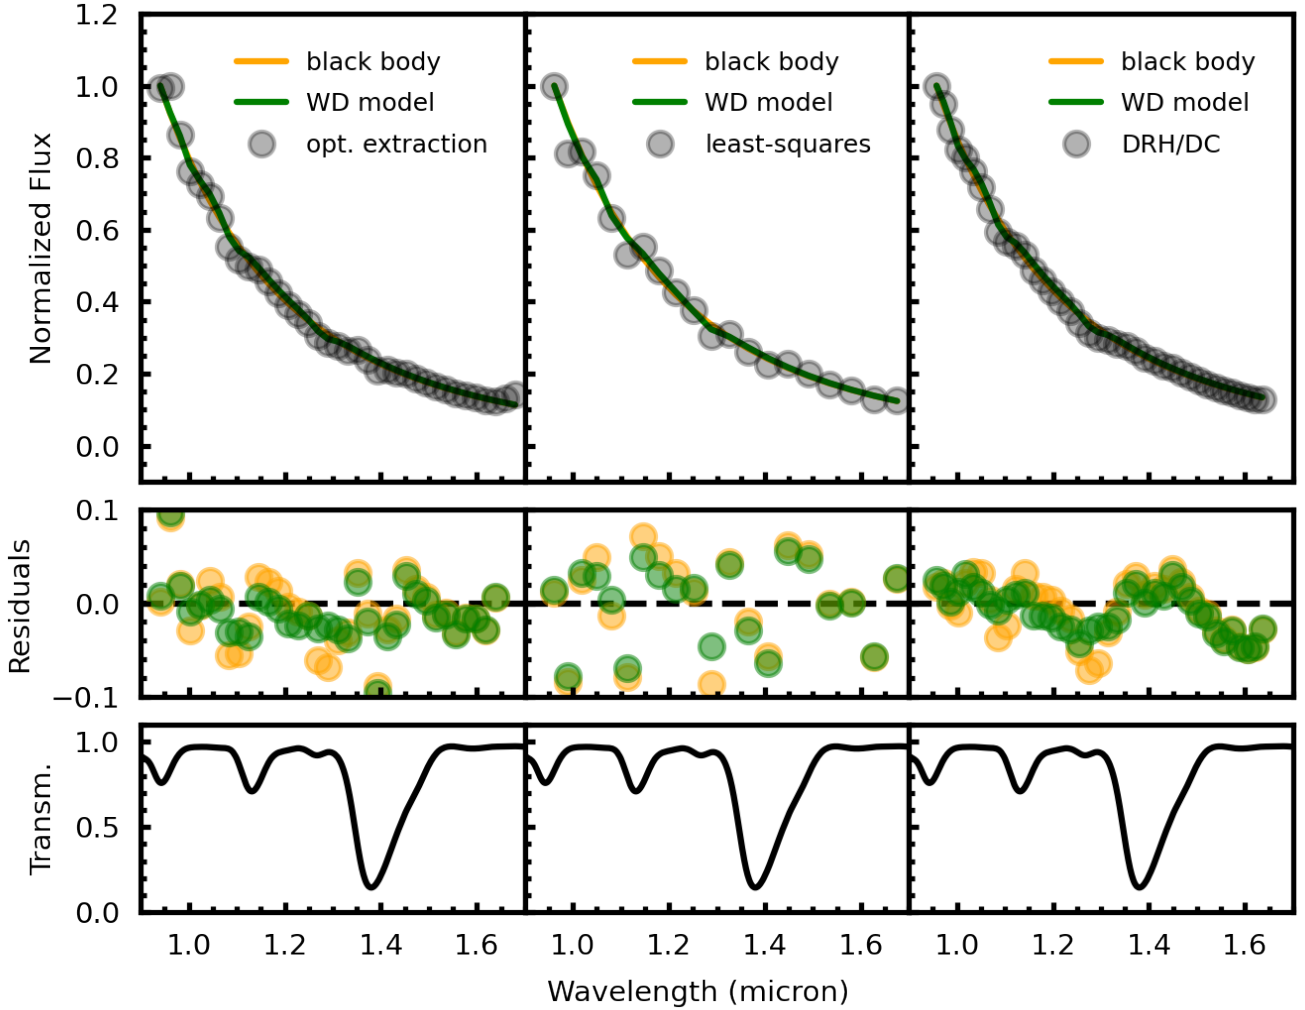


Fig. 9. Extracted spectrum of the white dwarf HD 2133 B compared to a blackbody and a white dwarf model spectrum of the same temperature, 26 000 K. The extracted spectrum has a high S/N (>200) and follows the spectrum to within $\sim 5\%$.

models. As the S/N is very high (>200), the statistical uncertainties of the extracted contrast are small (not shown because they are smaller than the symbols); the residuals are dominated by other effects. We also show the average atmospheric transmission profile obtained for Cerro Paranal Observatory using SKYCALC (Noll et al. 2012) to provide context for small deviations from the expected spectral shape.

The two pipelines and the different extraction methods are all consistent with the model at a level of $\lesssim 5\%$. As expected, the white dwarf model provides a slightly better match to the data than the blackbody because it includes the main absorption features of the white dwarf spectrum. We computed two metrics for the three extractions: first the median absolute deviation of the residuals from the white dwarf model,

$$\text{MAD} = \text{median}(|x|), \quad (2)$$

where $|x|$ are the absolute values of the model residuals normalized to the model flux. Second, we computed the correlation between pairs of residual values at adjacent wavelengths. Assuming that x has N elements, we computed the neighbor correlation as

$$\psi = \left(\sum_{i=1}^{N-1} (x_i x_{i+1}) + x_N \cdot x_1 \right) \left(\sum_{i=1}^N x_i^2 \right)^{-1}, \quad (3)$$

Table 2. White dwarf model residual metrics.

| Extraction | MAD | Correlation ψ |
|--------------------------|------|--------------------|
| Optimal extraction | 1.9% | 0.48 |
| Least-squares extraction | 3.7% | -0.05 |
| DRH/DC | 1.7% | 0.84 |

where the last term uses periodic boundary conditions to match the number of terms in the numerator and denominator. Table 2 summarizes the results of these two metrics from Eqs. (2) and (3) for the three extractions.

As our pipeline performs fewer interpolations (e.g., no post-hoc wavelength scaling, no anamorphism correction), the residuals are less spectrally correlated than in the DRH pipeline, and as expected from B17, they are slightly anticorrelated for least-squares extraction. In the case of least-squares extraction, which effectively deconvolves the images with the line-spread function, the spectrum appears to scatter more (seen in the higher MAD value). However, this is an artifact that arises because the spectrum is less strongly smoothed. The residual scatter of the optimal extraction and DRH/DC reduction is comparable (difference of 0.2%). All reductions show a high reliability of the spectral extraction.

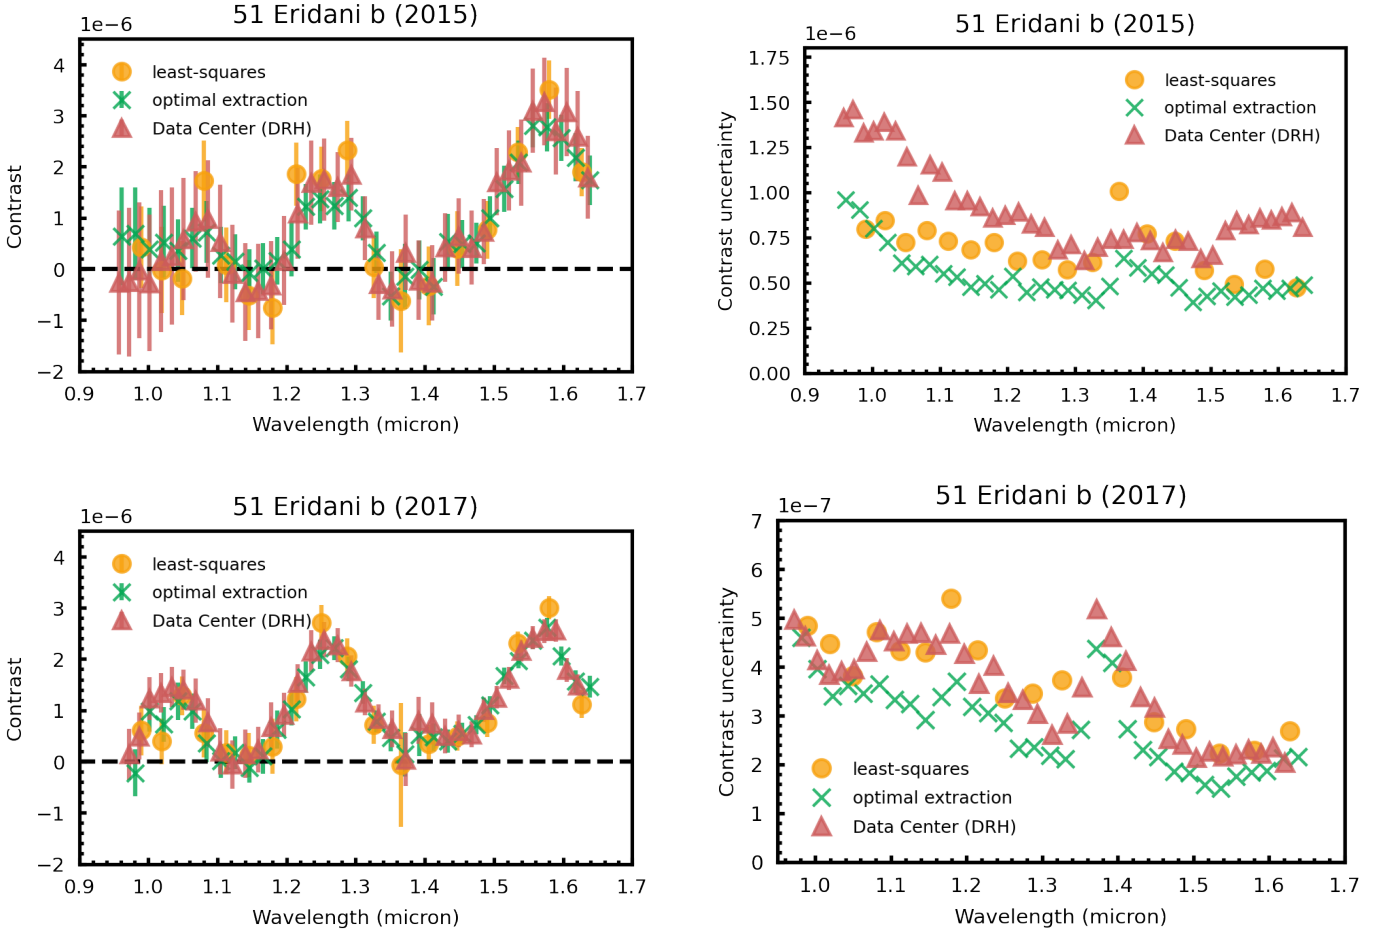


Fig. 10. Spectra extracted from 51 Eridani b data using different spectral cube extraction methods but the same postprocessing. *Top row:* data from 2015 taken in challenging conditions with a strong wind-driven halo, and the *bottom row* shows data from 2017 taken under excellent conditions. *Left panels:* extracted planet spectra, and the *right panels* show only the uncertainties of the extracted spectrum for better clarity.

We note that there are kinks in the residuals that largely correspond to the spectral channels that are most strongly affected by a change in atmospheric transmission. These are an indication that taking the atmospheric transmission in the model spectra into account becomes important in low spectral resolution spectroscopy. This is an independent problem that is not related to the spectral extraction of the IFS data itself, however.

5.3. Spectrum of 51 Eridani b

As a second test of our pipeline, we extracted and postprocessed data cubes of the directly imaged exoplanet 51 Eri b (Macintosh et al. 2015). We used two coronagraphic datasets of 51 Eri obtained by SPHERE IFS: the data from 2015 (Samland et al. 2017), and newer data from 2017 (Brown-Sevilla et al. 2022). We extracted data cubes using our pipeline, and for comparison, we used the extracted image cubes provided by the SPHERE Data Center using the ESO DRH pipeline. The DC by default temporally bins the data with a factor that avoids significant smearing in the IFS FoV. In these data sets, this results in 60 s exposures, using a binning of four for 2015 and two in 2017. We then extracted the planetary spectrum using the TRAP (Samland et al. 2021) postprocessing algorithm with the same uniform default settings for all data sets.

Figure 10 shows the resulting spectra and uncertainties for the 2015 and 2017 data sets and for each of the three cube-extraction approaches. The upper panels show the

spectrum obtained from the 2015 data, the lower panels correspond to the 2017 data. The left panels show the spectra, and the right panels show the statistical uncertainties from the extraction for better visibility because the spectra overlap strongly.

The statistical uncertainties in both data sets are smallest after postprocessing when the optimal extraction method of our pipeline is used. The least-squares method also reduces uncertainties in the extraction compared to DRH/DC in the 2015 data, but not as much in the cleaner 2017 data. However, it should be noted that least-squares performs a deconvolution with the line-spread function, therefore the spectral correlations are reduced and the spectral sampling is sparser, at the cost of larger scatter (less smoothed-out images). The direct comparison of the 1D uncertainties may therefore be misleading. On the other hand, optimal extraction is directly comparable to the DRH reduction, and there is a consistent reduction in the statistical uncertainties by a factor of $1.80^{+0.70}_{-0.67}$ for the 2015 data and $1.31^{+0.39}_{-0.37}$ for the 2017 data. These numbers refer to the median and 16–84 percentile range.

In Fig. 11 we show the corresponding contrast curves for these reductions. The contrast is given for different spectral templates (left: flat contrast, middle: L-type, and right: T-type). The least-squares and optimal extraction method of our pipeline provide comparable results and show a corresponding similar improvement compared to the DRH/DC reduction.

The contrast obtained with postprocessing of our pipeline-reduced images is consistently improved with respect to the

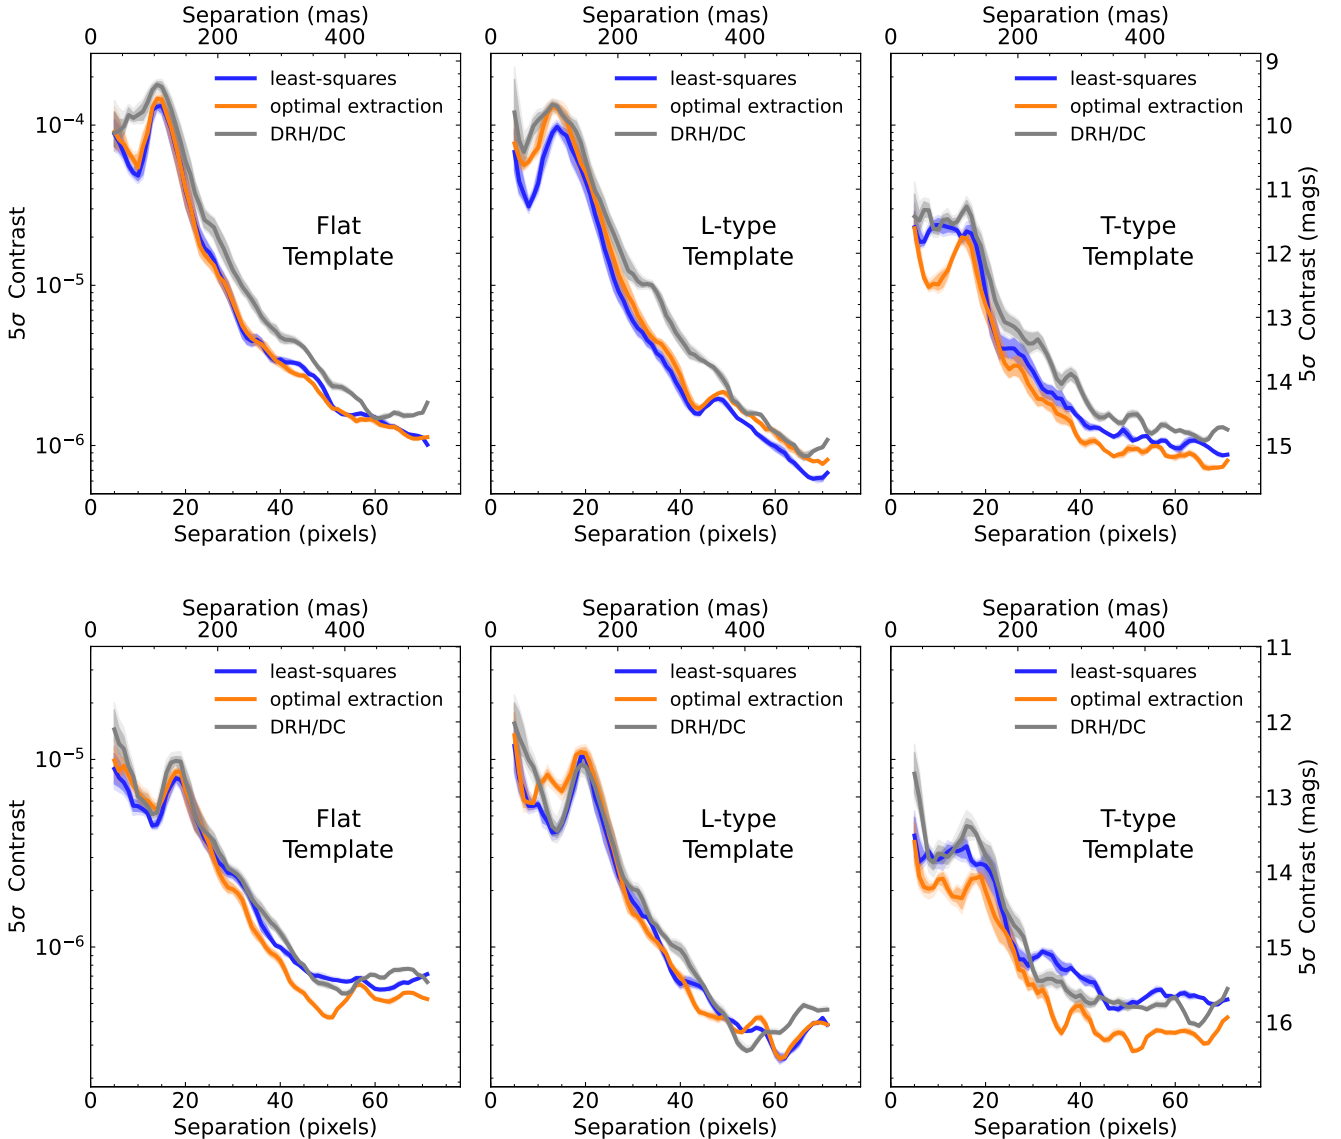


Fig. 11. Contrast curves for 51 Eri obtained from using TRAP on image cubes extracted using our pipeline and the DRH/DC pipeline. *Upper row:* results for the 2015 data, and the *lower row* shows the results for the better 2017 data. The individual panels show different spectral templates of planetary companions that were used to combine the contrast curves over the IFS wavelength channels: flat contrast, which corresponds to a mean contrast over all channels, and L-type and T-type spectral contrast templates.

DRH/DC. The strongest improvement is noticeable in the 2015 data using optimal extraction as was already seen in the extraction of the spectrum. The spectral template used to combine the wavelength channels does not impact these results significantly. As expected, we are most sensitive to T-type objects because they contain the most distinct spectral features compared to residual speckles.

6. Usage and performance

The adapted CHARIS pipeline is publicly available on Github. It has been modularized such that all instrument-specific calibrations and parameters are moved into a class structure, such that it can easily be maintained and extended for future instruments if necessary.

The pipeline functions can either be called from Python or a command-line tool can be used to extract image cubes

using a parameter file. The latter was the main way of interfacing with the original CHARIS pipeline. Calling the Python module directly is recommended for constructing an optimized workflow, however.

As with the previous CHARIS-pipeline, it only requires two steps to produce results: (1) ‘buildcal’ generates the calibration files, the wavelength calibration and (oversampled) PSFlet models, from a monochromatic flat-field for a specific night containing the laser spots, and (2) ‘extractcube’ extracts an image cube from raw image input and the calibration files. The usage of the command-line tools is analogous to the previous version (B17) for the calibration step. For the extraction step, the only difference is that three more parameters can be specified in the [Extract] category of the config-file: (1) `crosstalkscale` (float), which provides a continuous tuning of the crosstalk correction. The default of 1.0 provides a full crosstalk correction (see 4.2), and 0 provides no crosstalk correction. (2) `dcxtalkcorrection` (boolean), if set *True* tells the

pipeline to use the same crosstalk correction as implemented in the Data Center and [Vigan \(2020\)](#) instead of our method, which is based on the fitted PSFlet model. This alternative crosstalk correction smooths the data with a modified Moffat-kernel and subtracts the smoothed image to remove large-scale structures from the unextracted raw images. (3) `individual_dits` (boolean), if set `True` every individual exposure (detector integration, dit) in the data cube is extracted separately. Otherwise, the cube is averaged and then extracted. We strongly recommend extracting all exposures for obtaining the highest temporal resolution, but this parameter can be set to `False` for testing purposes in order to speed up the extraction. For SPHERE IFS data, if `bgsb` is `True` and no background images are provided, the pipeline uses the new principal component based fitting of the thermal background to perform the background subtraction. When data are directly reduced from Python, the `extractcube.getcube`-function can be called. We refer to the corresponding documentation for details of the required input parameters. An example script will be provided.

For non-coronagraphic PSF images taken with SPHERE IFS, the extraction may fail when `fitshift` is set to `True`. The reason for this is that in the case of SPHERE an overall subpixel shift of the spectra is determined from all lenslets. However, for non-coronagraphic PSF images only a very small fraction of lenslets are exposed to the light of the star making the determination unreliable and in some cases fail.

Performance. The ‘buildcal’ step is by far the most time-consuming part of a reduction. On a typical laptop and a single core, this can take about 90 min when the ‘upsample’ option is switched on, which is recommended for best results. However, this step is only required once for any particular observation night, and the resulting calibration can be used for any of the targets that are observed on that date and can be further sped up by using multiple cores for an approximately linear decrease in computational time with the number of cores.

The extraction of a single exposure from a cube of raw images on one core takes less than three minutes on the highest setting, that is, using the oversampled calibrations and iterative fitting of the PSFlet models to the data to subtract crosstalk. This includes all the processing steps, including the identification and correction of bad hexagons based on an iterative comparison of each hexagon with its neighbours, and the subsequent resampling of the data on a rectilinear grid.

7. Discussion and conclusion

We presented a new open-source data-reduction pipeline for the SPHERE IFS instrument. Our pipeline is a candidate replacement for the current modified DRH pipeline that is the basis of all IFS data reduction. It makes use of new calibration data that were specifically obtained for the purpose of improving the IFS data reduction. The pipeline is straightforward to use for anyone with a laptop and a current Python distribution, but can also make use of heavy multiprocessing if a large amount of data is processed. We plan to include our pipeline into the SPHERE Data Center (DC) in the future to facilitate access for the community even more. Because of the significant improvement in data quality provided by our pipeline over the current version of the SPHERE-DC pipeline, the SPHERE-DC ([Delorme et al. 2017](#)) has started adapting our pipeline to its infrastructure. It will then gradually reprocess and publicly release all available public IFS data over the next few years. Our pipeline was also

recently used in [Franson et al. \(2022\)](#) as the primary reduction for a newly discovered companion.

The public DRH has several important and known limitations that prevent its use for the most advanced science, such as limited accuracy in the wavelength calibration (up to 20 nm discrepancy), a grid pattern that is especially visible in low-S/N channels, and the complicated structure of the reduction recipes, which requires expert knowledge. A user wishing to postprocess SPHERE data must therefore either contact the Data Center or use a Python wrapper for the DRH pipeline containing known fixes for some of the issues ([Vigan 2020](#)), such as correcting the wavelength calibration based on measurements of the position of the satellite spots. This poses other dangers, however; the transmission profile of the atmosphere will shift the effective wavelength measured for the satellite spot position in channels that show water or methane absorption. The grid pattern visible in the DRH reduction is particularly detrimental in the study of extended low-S/N structures such as disks, and for obtaining reliable measurements in the water- or methane-affected spectral range of exoplanets and brown dwarfs. It also causes problems for sparse aperture masking data by introducing noise at specific spatial frequencies. The extracted data cubes are also often binned temporally and always contain multiple interpolations, both spectral and spatial.

The SPHERE IFS pipeline presented here, in contrast, inherently avoids these issues. Our pipeline introduces a number of improvements over the current Data Center product based on the ESO pipeline. Our pipeline improvements are enabled by new calibration data: monochromatic images of the individual calibration lasers taken through an integrating sphere. These enable us to extract and reconstruct the monochromatic lenslet PSFs at each of the four calibration wavelengths and over all regions of the detector.

The monochromatic lenslet PSFs enable fundamental improvements and new approaches to the cube reconstruction process. We can reconstruct the 2D shape of the spectra on the detector using a forward-modeling approach using our high-quality lenslet PSFs measurements. The model can be fit with a least-squares extraction that allows us to avoid interpolating the data spectrally. This forward model is also used to correct the crosstalk in a self-consistent fashion that does not rely on removing large-scale structures with a smoothing kernel. In addition to a least-squares extraction, we also enable optimal extraction using the measured lenslet PSF profiles, which is a straight improvement compared to the aperture photometry approach implemented in the DRH. We implemented a PCA-based approach to thermal background subtraction that makes taking sky frames unnecessary. It avoids introduction of additional read and photon noise, and persistence noise from sky frames taken after coronagraphic images.

Our pipeline allows the preservation of the original hexagon geometry of the images, which avoids the most aggressive spatial interpolation step in the data-reduction process. This opens the way for many possible improvements in postprocessing. However, in order to use traditional image-based postprocessing algorithms on hexagonal images, it becomes necessary to implement algorithms for registering and aligning (and potentially rotating) hexagon images. Furthermore, algorithms that work on image regions need to select these regions based on the real hexagon position instead of 2D array indices, which is often how the selection is implemented. For TRAP ([Samland et al. 2021](#)), only this second step is necessary because the algorithm does not require the images to be aligned. The companion PSF forward model, however, still has to be shifted to the correct position.

This might be simplified by using an analytic model trained on the unsaturated PSF image, thereby avoiding this issue entirely. An algorithm such as TRAP that does not require images to be aligned because it works directly on the time-series of pixels or spaxels can work on the completely uninterpolated spaxels, thereby reducing noise correlations significantly. This will further improve the attainable contrast, especially because our pipeline provides uncertainties for each spaxel of the extracted image that can be used as input for postprocessing. The necessary changes have not yet been implemented into TRAP, but this is planned for the future. In general, the architecture of TRAP is such that the required changes are relatively minor (passing the position of each hexagon to the reduction area and reference area selection routine and adjusting the forward model of the companion PSF to the lenslet location in each frame of the ADI sequence).

Our SPHERE IFS pipeline is easy to use as it has only two steps: First, the calibration data are build from the `wavecal` file, a monochromatic measurement of the lenslet PSF using three to four lasers depending on the wavelength coverage. This step only requires one calibration file, all other calibrations are included with the pipeline. Second, the calibration data are used to extract the spectral cubes. In conclusion, we presented a new pipeline that addresses multiple known issues, introduces several significant improvements, is accessible, and is easy to use, with much potential for further improvement of postprocessing in the future.

Acknowledgements. M.S. acknowledges support from the European Research Council under the Horizon 2020 Framework Program via the ERC Advanced Grant Origins 83 24 28. This work has made use of the SPHERE Data Centre, jointly operated by OSUG/IPAG (Grenoble), PYTHEAS/LAM/CESAM (Marseille), OCA/Lagrange (Nice), Observatoire de Paris/LESIA (Paris), and Observatoire de Lyon. This work is supported by the French National Research Agency in the framework of the Investissements d’Avenir program (ANR-15-IDEX-02), through the funding of the “Origin of Life” project of the Grenoble-Alpes University. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (COBREX; grant agreement n° 885593).

References

- Anderson, J., & King, I. R. 2000, *PASP*, **112**, 1360
- Antichi, J., Dohlen, K., Gratton, R. G., et al. 2009, *ApJ*, **695**, 1042
- Bailey, S. 2012, *PASP*, **124**, 1015
- Berdeu, A., Soulez, F., Denis, L., Langlois, M., & Thiébaud, É. 2020, *A&A*, **635**, A90
- Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, **631**, A155
- Boccaletti, A., Thalmann, C., Lagrange, A.-M., et al. 2015, *Nature*, **526**, 230
- Bohn, A. J., Kenworthy, M. A., Ginski, C., et al. 2020, *MNRAS*, **492**, 431
- Bonavita, M., Fontanive, C., Gratton, R., et al. 2022, *MNRAS*, **513**, 5588
- Brandt, T. D., Rizzo, M., Groff, T., et al. 2017, *J. Astron. Teles. Instrum. Syst.*, **3**, 048002
- Brown-Sevilla, S. B., Maire, A.-L., & Mollière, P. 2022, *A&A*, submitted [arXiv:2211.14330]
- Burleigh, M. R., Barstow, M. A., & Fleming, T. A. 1997, *MNRAS*, **287**, 381
- Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. 2015, *A&A*, **582**, A89
- Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017a, in *SF2A-2017: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, eds. C. Reylé, P. Di Matteo, F. Herpin, E. Lagadec, A. Lançon, Z. Meliani, & F. Royer, 331
- Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017b, *A&A*, **605**, L9
- Chauvin, G., Gratton, R., Bonnefoy, M., et al. 2018, *A&A*, **617**, A76
- Cheetham, A. C., Samland, M., Brems, S. S., et al. 2019, *A&A*, **622**, A80
- Claudi, R. U., Turatto, M., Gratton, R. G., et al. 2008, *Proc. SPIE*, **7014**, 70143E
- Couder, A. 1949, *L’Astronomie*, **63**, 253
- Dahlqvist, C. H., Cantalloube, F., & Absil, O. 2020, *A&A*, **633**, A95
- Delorme, P., Meunier, N., Albert, D., et al. 2017, in *SF2A-2017: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, eds. C. Reylé, P. Di Matteo, F. Herpin, E. Lagadec, A. Lançon, Z. Meliani, & F. Royer, 347
- Dohlen, K., Langlois, M., Saisse, M., et al. 2008, *Proc. SPIE*, **7014**, 70143L
- Franson, K., Bowler, B. P., & Bonavita, M. 2022, *ApJ*, submitted [arXiv:2211.09840]
- Galicher, R., Boccaletti, A., Mesa, D., et al. 2018, *A&A*, **615**, A92
- Gomez Gonzalez, C. A., Wertz, O., Christiaens, V., Absil, O., & Mawet, D. 2016, *Astrophysics Source Code Library* [record ascl:1603.003]
- Groff, T. D., Kasdin, N. J., Limbach, M. A., et al. 2015, *Proc. SPIE*, **9605**, 96051C
- Groff, T. D., Chilcote, J., Kasdin, N. J., et al. 2016, *Proc. SPIE*, **9908**, 99080O
- Guyon, O. 2005, *ApJ*, **629**, 592
- Guyon, O., Pluzhnik, E., Woodruff, R., et al. 2005, *AAS Meeting Abstracts*, **207**, 191.04
- Hanuš, J., Marchis, F., Viikinkoski, M., Yang, B., & Kaasalainen, M. 2017, *A&A*, **599**, A36
- Hoeijmakers, H. J., Schwarz, H., Snellen, I. A. G., et al. 2018, *A&A*, **617**, A144
- Horne, K. 1986, *PASP*, **98**, 609
- Janson, M., Squicciarini, V., Delorme, P., et al. 2021, *A&A*, **646**, A164
- Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, **127**, 890
- Keppler, M., Benisty, M., Müller, A., et al. 2018, *A&A*, **617**, A44
- Levenhagen, R. S., Diaz, M. P., Coelho, P. R. T., & Hubeny, I. 2017, *ApJS*, **231**, 1
- Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *Proc. Natl. Acad. Sci.*, **111**, 12661
- Macintosh, B., Graham, J. R., Barman, T., et al. 2015, *Science*, **350**, 64
- Maire, A.-L., Langlois, M., Dohlen, K., et al. 2016, *SPIE Conf. Ser.*, **9908**, 990834
- Maire, A. L., Lafudino, J. L., Desidera, S., et al. 2020, *A&A*, **633**, L2
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, **641**, 556
- Mesa, D., Gratton, R., Zurlo, A., et al. 2015, *A&A*, **576**, A121
- Milli, J., Hibon, P., Christiaens, V., et al. 2017a, *A&A*, **597**, L2
- Milli, J., Vigan, A., Mouillet, D., et al. 2017b, *A&A*, **599**, A108
- Milli, J., Kasper, M., Bourget, P., et al. 2018, *SPIE Conf. Ser.*, **10703**, 107032A
- Müller, A., Keppler, M., Henning, T., et al. 2018, *A&A*, **617**, L2
- Noll, S., Kausch, W., Barden, M., et al. 2012, *A&A*, **543**, A92
- Pavlov, A., Möller-Nilsson, O., Feldt, M., et al. 2008, *Proc. SPIE*, **7019**, 701939
- Perrin, M. D., Maire, J., Ingraham, P., et al. 2014, *SPIE Conf. Ser.*, **9147**, 91473J
- Petit, C., Sauvage, J. F., Sevin, A., et al. 2012, *SPIE Conf. Ser.*, **8447**, 84471Z
- Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, *PASP*, **111**, 587
- Samland, M., Mollière, P., Bonnefoy, M., et al. 2017, *A&A*, **603**, A57
- Samland, M., Bouwman, J., Hogg, D. W., et al. 2021, *A&A*, **646**, A24
- Sauvage, J.-F., Fusco, T., Petit, C., et al. 2016, *J. Astron. Teles. Instrum. Syst.*, **2**, 025003
- Schmid, H. M., Bazzon, A., Roelfsema, R., et al. 2018, *A&A*, **619**, A9
- Skemer, A. J., Hinz, P., Montoya, M., et al. 2015, *SPIE Conf. Ser.*, **9605**, 96051D
- Snik, F., Otten, G., Kenworthy, M., et al. 2012, *SPIE Conf. Ser.*, **8450**, 84500M
- Soummer, R. 2005, *ApJ*, **618**, L161
- Sparks, W. B., & Ford, H. C. 2002, *ApJ*, **578**, 543
- Vigan, A. 2020, *Astrophysics Source Code Library* [record ascl:2009.002]
- Vigan, A., Fontanive, C., Meyer, M., et al. 2021, *A&A*, **651**, A72
- Wang, J. J., Ruffio, J.-B., De Rosa, R. J., et al. 2015, *Astrophysics Source Code Library* [record ascl:1506.001]