



Universiteit
Leiden
The Netherlands

Impacts of high-contrast image processing on atmospheric retrievals

Nasedkin, E.; Mollière, P.; Wang, J.; Cantalloube, F.; Kreidberg, L.; Pueyo, L.; ... ; Vigan, A.

Citation

Nasedkin, E., Mollière, P., Wang, J., Cantalloube, F., Kreidberg, L., Pueyo, L., ... Vigan, A. (2023). Impacts of high-contrast image processing on atmospheric retrievals. *Astronomy And Astrophysics*, 678. doi:10.1051/0004-6361/202346585





Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3717452>

Note: To cite this publication please use the final published version (if applicable).

Impacts of high-contrast image processing on atmospheric retrievals

E. Nasedkin¹ , P. Mollière¹, J. Wang² , F. Cantalloube³, L. Kreidberg¹, L. Pueyo⁴, T. Stolker⁵ , and A. Vigan³ 

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

² Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA

³ Aix Marseille Univ., CNRS, CNES, LAM, 13000 Marseille, France

⁴ Space Telescope Science Institute (STScI), Baltimore, MD 21218, USA

⁵ Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, The Netherlands

Received 3 April 2023 / Accepted 2 August 2023

ABSTRACT

Many post-processing algorithms have been developed in order to better separate the signal of a companion from the bright light of the host star, but the effect of such algorithms on the shape of exoplanet spectra extracted from integral field spectrograph data is poorly understood. The resulting spectra are affected by noise that is correlated in wavelength space due to both optical and data processing effects. Within the framework of Bayesian atmospheric retrievals, we aim to understand how these correlations and other systematic effects impact the inferred physical parameters. We consider three algorithms (KLIP, PynPoint, and ANDROMEDA), optimising the choice of algorithmic parameters using a series of injection tests on archival SPHERE and GPI data of the HR 8799 system. The wavelength-dependent covariance matrix was calculated to provide a measure of instrumental and algorithmic systematics. We perform atmospheric retrievals using petitRADTRANS on optimally extracted spectra to measure how these data processing systematics influence the retrieved parameter distributions. The choice of data processing algorithm and parameters significantly impact the accuracy of retrieval results, with the mean posterior parameter bias ranging from 1 to 3 σ from the true input parameters. Including the full covariance matrix in the likelihood improves the accuracy of the inferred parameters, and cannot be accounted for using ad hoc scaling parameters in the retrieval framework. Using the Bayesian information criterion and other statistical measures as heuristic goodness-of-fit metrics, the retrievals including the full covariance matrix are favoured when compared to using only the diagonal elements.

Key words. planets and satellites: atmospheres – methods: data analysis

1. Introduction

The field of high contrast imaging (HCI) has advanced dramatically over the last two decades. From the first detection of 2M 1207 b (Chauvin et al. 2005) to the ongoing large surveys such as GPIES (Nielsen et al. 2019) and SHINE (Desidera et al. 2021; Langlois et al. 2021; Vigan et al. 2021), we have seen improvements in instrumentation, adaptive optics and data processing that have led to the discovery of numerous new exoplanets. Such surveys have established the rarity of giant, widely separated companions, finding that < 10% of high-mass stars have planetary mass companions between 10 and 100 AU. However, much of this work has remained focussed on the detection of new companions at higher contrast ratios and smaller angular separations. The spectroscopic characterisation of known planets has seen less dedicated effort – to date there has not been a uniform survey of known objects to present a homogeneous sample of spectroscopic measurements. This can lead to systematic discrepancies between measurements made with different instruments, and challenges in fitting datasets with different spectral resolutions (Xuan et al. 2022). Such biases may impact the conclusions made from population studies, such as the exploration of the C/O ratios of the directly imaged planet population in Hoch et al. (2023). Individual characterisation efforts have nevertheless led to intriguing findings: measurements of water and carbon monoxide

abundances in the HR 8799 planets (Konopacky et al. 2013; Lavie et al. 2017; Wang et al. 2020), precise constraints on the C/O ratio and metallicity of β Pictoris b (GRAVITY Collaboration 2020), measurements of isotope ratios (Zhang et al. 2021) and the detection of a dusty envelope around PDS 70 b and c (Wang et al. 2021; Benisty et al. 2021).

This characterisation work remains challenging. Extensive post-processing is required to extract the faint signal of the target. Even in the most careful analysis systematic biases usually remain from both instrumental and processing effects. Integral field spectrograph (IFS) measurements introduce correlated noise as a function of wavelength due to pixel cross talk, interpolation effects and imperfect adaptive optics correction (speckles). Greco & Brandt (2016) provide a method for empirically estimating the correlation from IFS data. They demonstrated that accounting for such correlations is necessary when analysing exoplanet atmospheres, and failing to do so leads to biased and overconfident posterior distributions on measured parameters. Efforts such as the Exoplanet Imaging Data Challenge (Cantalloube et al. 2020b) explored the detection abilities of a suite of HCI algorithms but a systematic algorithmic comparison for spectral characterisation has not yet been performed.

Most post-processing techniques are based on Angular Differential Imaging (ADI; Marois et al. 2006, 2008b), where the telescope is pupil stabilised and the field is allowed to rotate. This

Table 1. Epochs of SPHERE ([1] Zurlo et al. (2016) 60.A-9249(C)) and GPI ([2] Greenbaum et al. (2018) GS-2015B-Q-500-1394) data used.

Instrument	Date	Band	$\lambda/\Delta\lambda$	Field rotation [°]	Med. seeing (as)	DIT (s)	NEXP	Ref.
SPHERE	2014-08-12	<i>YJH</i>	29	29.65 15.37	0.87	100 60	32 48	[1]
GPI	2016-09-19	<i>H</i>	45	20.93	0.97	60	60	[2]

provides differential motion of the planet over the course of the observations and allows for the removal of stellar speckles by deroating and stacking the resulting images. Ongoing development of this method has been largely driven by the goals of increasing sensitivity at small angular separations. To this end, different algorithms have been developed to maximise the information available in imaging datasets, leveraging spatial and spectral information in order to separate the faint planet signal from the bright host star. Kiefer et al. (2021) explored how different approaches impact the signal-to-noise (S/N) of IFS observations, but did not examine the impact of the processing on the extracted spectral shape.

The use of atmospheric retrievals to study directly imaged planets is relatively new, with only a small but growing selection of targets being subject to such an analysis (e.g. Lee et al. 2013; Lavie et al. 2017; Mollière et al. 2020; GRAVITY Collaboration 2020; Brown-Sevilla et al. 2023; Whiteford et al. 2023). While the effects of systematics are well understood for transmission spectroscopy using HST (Ih & Kempton 2021), with significant efforts extending this to *James Webb* Space Telescope (JWST; Barstow et al. 2015; Rocchetto et al. 2016; Lacy & Burrows 2020), the impact of systematic uncertainties in ground-based high-contrast data on atmospheric retrievals has not been thoroughly explored. Even in the era of JWST, understanding systematics is critical to interpreting model fits to data. Ground-based observations will remain a key component of this understanding due to their higher spectral and angular resolution that cannot yet be achieved from space.

In this work, we explore the systematic effects introduced through high-contrast data processing on the retrieval of atmospheric parameters. The details of our example datasets used are described in Sect. 2. Section 3 outlines our methods, exploring the different algorithm tested in Sect. 3.1, together with the measurement and interpretation of the covariance matrix in Sects. 3.2 and 3.4. We determined the optimal parameters for spectral extraction through the injection and recovery of synthetic companions into the data in Sect. 4. The results of our retrieval comparisons are described in Sect. 5, while the implications and limitations of these results are discussed in Sect. 6.

2. Observations

While the first goal of our study is to demonstrate the effects that post-processing algorithms can have on inferred atmospheric parameters for general high-contrast spectroscopy, we still had to select demonstration datasets. We chose GPI and SPHERE observations of the well-known four-planet system in HR 8799 (Marois et al. 2008b, 2010), where discrepancies between GPI and SPHERE datasets, covering the same wavelength range, had already been noted (Lavie et al. 2017; Mollière et al. 2020). HR 8799 has seen extensive photometric and spectroscopic observing campaigns, (e.g. Konopacky et al. 2013; Zurlo et al. 2016; Lavie et al. 2017; Greenbaum et al. 2018; GRAVITY Collaboration 2019; Mollière et al. 2020;

Wang et al. 2020, 2023; Ruffio et al. 2021). The importance of this system, together with the abundance of high contrast data from multiple instruments make it an ideal object of study for our purposes. As a benchmark target, the companions have luminosity and spectra typical of this class of low surface gravity object and are representative of the current directly imaged exoplanet population.

SPHERE. The SPHERE data were taken during the commissioning run of the SPHERE instrument (Beuzit et al. 2008, 2019) in 2014, and were originally presented in Zurlo et al. (2016). It remains the best *YJH* band spectrum of HR 8799 to date in terms of signal-to-noise and spectral resolution. IFS frames in the *YJH* band were taken with a series of both 60 s and 100 s integrations, using pupil-stabilised observations to allow for ADI post-processing. Total field rotations of 15.37° and 29.65° were observed for the 60 s data cube and for the 100 s data cube, respectively. To compensate for the difference in exposure time, we multiply each 60 s exposure by a factor of 100/60, in order to process the data as a whole. We rereduced the SPHERE data using the pipeline described in Vigan (2020): details of which are described in Appendix A.

GPI. The GPI (Macintosh et al. 2014) observations of HR8799 were originally published in Greenbaum et al. (2018) and were taken on 17 November 2013, 18 November 2013, and 19 September 2016 for the *K1*, *K2*, and *H* bands respectively. As with the SPHERE data, the telescope was pupil-stabilised to take advantage of ADI post-processing. These were reduced using the standard GPI reduction pipeline (version 1.4.0). The median seeing of the observations was 0".97; the observing conditions are more thoroughly described in Ingraham et al. (2014). While data were taken in the *H*, *K1*, and *K2* bands of GPI, we only considered the *H* band observations due to the low S/N of the *K*-band observations. The observations from both GPI and SPHERE are summarised in Table 1.

2.1. Data preprocessing

In order to reduce the systematic variation between the datasets, we first rereduced the data with up-to-date pipelines. For both the SPHERE and GPI datasets, we then preprocess the IFS cubes using the Vortex Image Processing (VIP) library in order to select the optimal frames for further ADI processing. The `cube_detect_badfr_correlation` function computes the similarity between each frame and a reference frame in order to identify frames that are outliers when compared to the rest of the sequence. We choose the frame which maximises the mean similarity of all frames as the reference frame, and remove the most different 12% of frames from each the SPHERE and GPI datasets. Such variation in the data is typically due to changing observing conditions, introducing effects into the data such as the low-wind effect Milli et al. (2018) or the wind-driven halo Cantalloube et al. (2020a) This threshold is sufficient to remove frames which are significantly outlying and visually show differences when

Table 2. Stellar properties of HR 8799 A.

HR 8799		
Parameter	Value	Note
α (J2000)	23 ^h 07 ^m 28.7157 ^s \pm 0.0685 ^s	[1]
δ (J2000)	+21° 08′ 03.3021″ \pm 0.0799″	[1]
μ_α (mas yr ⁻¹)	108.301 \pm 0.168	[1]
μ_δ (mas yr ⁻¹)	-49.480 \pm 0.152	[1]
$\bar{\omega}$ (mas)	24.2175 \pm 0.0881	[1]
d (pc)	41.2925 \pm 0.1502	[1]
RV (km s ⁻¹)	-12.60 \pm 1.4	[2]
Spectral type	F0+VkA5mA5 C	[3]
T_{eff} (K)	7200 \pm 50	[4]
$\log g$ (cgs)	3.0 \pm 0.25	[4]
[Fe/H] (dex)	0.0 \pm 0.2	[4]
R_\star (R_\odot)	1.496 \pm 0.0054	[4]
L_\star (L_\odot)	5.230 \pm 0.0498	[4]
C/O	0.54 ^{+0.12} _{-0.09}	[5]

Notes. [1] [Gaia Collaboration \(2018\)](#). [2] [Gontcharov \(2006\)](#). [3] [Gray et al. \(2003\)](#). [4] BT-NextGen best-fit to photometry ([Hauschildt et al. 1999](#)). [5] [Wang et al. \(2020\)](#). HR 8799 is a λ Boötis star, for further discussion see [Mollière et al. \(2022\)](#).

compared to a typical frame. This leaves 69 ADI frames for the SPHERE dataset, and 51 for the GPI H -band dataset.

2.2. Stellar model for flux extraction

In order to obtain the absolute flux of the companions we use a model of the stellar spectrum to flux-calibrate the contrast measurements. HR8799 is an F0+VkA5mA5 C star ([Gray et al. 2003](#)) located 41.3 ± 0.2 pc ([Gaia Collaboration 2018](#)). Stellar photometry of HR 8799 from WISE and 2MASS is used to fit model stellar spectrum ([Cutri et al. 2021, 2003](#)). We exclude data points beyond $5 \mu\text{m}$ so that the fit is not impacted by the infrared excess from the debris disk ([Su et al. 2009; Faramaz et al. 2021](#)). Using the `species` package ([Stolker et al. 2020](#)), we fitted a BT-Nextgen model to the photometry within our wavelength range of interest. The best-fit model has parameters of $T_{\text{eff}} = 7200$ K, $\log g = 3.0$ and $[\text{Fe}/\text{H}] = 0.0$, slightly cooler than the models used in previous studies ([Zurlo et al. 2016; Greenbaum et al. 2018](#)). The full set of stellar parameters is listed in Table 2. This spectrum is normalised to a 10 pc distance. The model is convolved to the instrumental spectral resolutions and binned to the instrumental wavelength channels to allow for spectrophotometric calibration of contrast measurements.

In order place measurements of planet properties in context it is also necessary to understand the properties of the host star. [Wang et al. \(2020\)](#) used HARPS observations to directly measure the C and O abundances of the star, finding a C/O ratio of $0.54^{+0.12}_{-0.09}$. HR 8799 is a λ Boötis star, known to be depleted in iron ([Gray & Corbally 2002](#)). Consistent with this, the authors fit Fe I and Fe II lines, finding a metallicity of $[\text{Fe}/\text{H}] = -0.52 \pm 0.08$. Both of the carbon and oxygen abundances were measured to be consistent with solar composition, suggesting that the iron metallicity is not representative of the bulk stellar composition, and that our BT-NextGen is still applicable. At the low spectral resolution considered in this study the metallicity does not significantly impact the SED of the star and variations in its measurement will not affect the calculation of the results, though

will ultimately impact the context – and thus interpretation – of planetary metallicity measurements.

3. ADI data processing

The atmospheric properties of directly imaged exoplanets are presently accessible only through their thermal emission. For directly imaged planets, these spectra are usually obtained through low to moderate resolution IFS instruments equipped with coronagraphic optics. IFS data is complex, with a large array of systematic and random noise effects imprinted onto the data. Cross talk between neighbouring pixels due to optical effects ([Antichi et al. 2009; Larkin et al. 2014](#)) and scattered light can introduce correlations in wavelength space. Once the data has been reduced from raw detector frames to data cubes, quasi-static stellar speckles – light of the host star scattered by the telescope optics – is the dominant noise source ([Marois et al. 2005, 2008a](#)). ADI processing is used to remove the stellar PSF and speckle noise, taking advantage of the stability of the PSF over time ([Marois et al. 2006](#)). ADI exploits the rotation of the planet through the frame, which produces a signal that is different from the stellar speckles, which remain fixed in position. By derotating and stacking the images, the residual speckles following post-processing are averaged out, while the planet signal is enhanced. The stability assumption is not without flaws, as the PSF varies due to thermal variation in the telescope, short and long-term atmospheric changes and more ([Milli et al. 2016](#)), but in practice it is robust enough to allow for planet detection. Obtaining an exoplanet spectrum is generally achieved by applying an ADI algorithm to each spectral channel on a 4D cube of IFS data. Modern ADI processing is more sophisticated than simply derotating and stacking the images, but the algorithms generally fall into three broad categories:

Speckle subtraction methods attempt to directly subtract the residual stellar speckles from each frame of the image cube. The planet signal is then measured either through aperture photometry or through fitting a model of the PSF to the signal and minimising the residuals. This is the most commonly used method, and includes algorithms such as (Template) Locally Optimised Combination of Images (LOCI and TLOCI, [Lafrenière et al. 2007; Maire et al. 2012; Marois et al. 2014b](#)), low rank plus sparse decomposition (LLSG, [Gomez Gonzalez et al. 2016](#)) and various implementations of principal component analysis (PCA) based methods, including Karhunen-Loève Image Projection (KLIP, [Soummer et al. 2012](#)) and Standardised Trajectory Intensity Mean (STIM, [Pairet et al. 2019](#)). Such PCA-based methods construct an ordered library of principal components of the data: low orders describe the most important components of the stellar PSF, while higher orders describe high frequency noise. By building a library to describe the host star PSF, it can be more effectively subtracted from each frame before stacking the images, improving the S/N of the companion.

Inverse methods such as ANDROMEDA ([Mugnier et al. 2009; Cantalloube et al. 2015](#)), PACO ([Flasseur et al. 2018](#)) and TRAP ([Samland et al. 2021](#)) use likelihood minimisation to directly estimate the position and contrast of proposed signal at each point in the field. To do this, a parameterised forward model of the companion signature is fit to the data, and the parameters are optimised through a likelihood minimisation process. This yields a statistical interpretation of the residuals, and provides confidence region estimates that provide a metric for detection significance, under varying assumptions of the noise properties of the data.

Finally, supervised machine learning methods (Gomez Gonzalez et al. 2018; Hou Yip et al. 2019; Gebhard et al. 2022) are trained on large sets of data with injected targets and learn how to identify the presence of a companion in an image. These methods typically only produce a binary maps where a planet is either detected or not, and do not measure the strength of the planet signal. Thus these methods have not generally been used for exoplanet characterisation.

3.1. Post-processing algorithms

We chose to compare three widely used ADI techniques in order to determine the impact of such post-processing on the spectral shape and noise properties of the extracted exoplanet spectrum. In order to compare a diverse range of techniques we chose to use KLIP and PynPoint, which are different flavours of PCA-based speckle subtraction methods, and Andromeda, which is an inverse method. As our goal is to understand the impact of systematic effects, we chose these algorithms for their broad community use, typifying the effects likely present in existing work. A more complete examination of the diversity of algorithms, including spectral differential imaging (SDI) and ADI+SDI algorithms will be explored using a larger set of data in a forthcoming publication based on Phase 2 of the Exoplanet Imaging Data Challenge.

In order to assess our choice of algorithm, we compared extractions of known injected spectra at different positions and contrasts in order to optimise the parameter selection for extracting the true spectrum. In this section we present the specific steps we took to reduce SPHERE and GPI datasets of the HR 8799 system using each of these algorithms. While a wide range of parameters were explored, Table 3 summarises the parameter choices used in this analysis for each algorithm.

3.1.1. KLIP

KLIP is a PCA-based speckle subtraction algorithm, described in Soummer et al. (2012); Wang et al. (2015) and Pueyo (2016). A Karhunen–Loève transform of an optimised combination of reference images is used to define the basis of eigenimages, onto which the science frames are projected. Often this set of reference images is derived from the science observations, but in principle can be any representative measurements of the PSF. Mathematically, this is equivalent to building the basis of principal components. This projection is subtracted from the science frames in order to produce the final residual image. A forward model of the PSF is then injected in order to measure the position and contrast of a detected companion.

Our choice of KLIP parameters is guided by Pueyo (2016) and Greenbaum et al. (2018). For this study, we use KLIP in ADI mode. Comparison tests showed that the full ADI+SDI mode provided modest increases in S/N at low contrasts, but the overall shape of the spectrum remained similar. We set a region around the proposed location of the planet extending 13 pixels radially in each direction, and 18° on either side of the planet. The flux overlap parameter is used to set the aggressiveness of the subtraction, using a value of 0.1. Fixing these parameters may result in sub-optimal spectral extraction, particularly at very small separations where the rotational movement of the planet through the frame is small. However, we are primarily concerned with the overall trends in the spectral extractions and noise properties across different tools, and do not attempt to fine tune each algorithm for each individual injections.

We use the pyKLIP astrometric measurement tools to compute the location of each target within the field of view, which

Table 3. Parameters used for each of the algorithms considered.

Parameter	Value
pyKLIP	
nPC	1–25
flux_overlap	0.1
highpass	True
maxnumbasis	150
mode	ADI
PynPoint	
nPC	1–15
merit	Gaussian
aperture	4 px
tolerance	0.01
cent_size	0"12
offset	2 px
ANDROMEDA	
filtering_frac	0.35, 0.30
min_sep	0.45 λ/d , 0.25 λ/d
width	0.8 λ/d , 1.2 λ/d
iwa	2.0 λ/d , 1.0 λ/d
owa	60/S, 45/S
opt_method	lsq

Notes. Parameters that were not varied were set based on previously reported values in literature (Pueyo 2016; Greenbaum et al. 2018; Cantalloube et al. 2015). For ANDROMEDA, the first column of parameters was used for the SPHERE data, and the second column for the GPI data. The oversampling parameter S is defined in Eq. (1). Further information about each of these parameters is available in the documentation of each package.

is used to provide our initial estimate for the planet position for each of the algorithms we consider. The extracted spectrum is highly sensitive to the inferred companion position, and so we use the KLIP astrometry as the location for all three algorithms. Pueyo (2016) outlines the procedure to extract the spectrum from KLIP processed data using the forward model extraction tool. For each target at each wavelength, a forward model is generated from the unsaturated PSF obtained during the observation. KLIP processing is then applied to subtract the stellar PSF and measure the contrast of the companion. This is converted into a flux measurement using the BT-NextGen model of the host star spectrum from Sect. 2.2.

3.1.2. PynPoint

PynPoint is a Python package designed for high contrast imaging data processing (Amara & Quanz 2012; Stolker et al. 2019). The standard PSF subtraction method used in the package is based on full-frame PCA. We process each wavelength channel of the IFS data independently, filtering for bad pixels and running ADI-PCA on each stack of images. In contrast to KLIP, which builds a model of principal components in a local region near the planet, PynPoint builds its PC library from the full available field of view. The central 0"12 of each frame is masked out, due to the large residuals close to the host star.

Following the PSF subtraction, a PSF model with negative flux is injected at the position of the planet of interest, which is known from previously computed KLIP astrometry. The PSF model for the planet is simply the stellar PSF, which is either

derived from satellite spots (for GPI data) or from unocculted observations of the host star (for SPHERE data). The position and magnitude of the negative planet are iteratively fit to the data using a simplex minimisation routine to minimise the χ^2 between the PSF model and the data. The minimisation is considered within an aperture with a radius of 4 pixels around the proposed location of the planet. The iteration continues until a tolerance of 0.01 is reached for both the planet position and contrast in magnitude units. We allow the planet position to vary by up to 3 pixels (`offset`) from the initial estimate from `pyKLIP` astrometry. This produces a best-fit value of the position and contrast-magnitude of the planet. While we allowed the number of principal components used to vary from 1 to 25, we found that the extraction quality degraded substantially after 15 components, which sets the upper bound we present in this work. This is then converted from magnitude to contrast, and multiplied by the BT-NextGen stellar model of Sect. 2.2 to find the absolute flux of the planet.

3.1.3. Andromeda

ANDROMEDA (ANgular DiffeRential OptiMal Exoplanet Detection Algorithm) is a maximum likelihood estimation algorithm for ADI data, and estimates the position and flux of point sources within the field of view (Mugnier et al. 2009; Cantalloube et al. 2015). We run the VIP implementation of the algorithm on each wavelength channel independently, and combine the extracted contrast and standard deviation to build the planet spectrum. ANDROMEDA begins by high-pass filtering the data to remove large spatial scale structure from each data frame. This step induces signal loss, and we chose a value of 0.3 for the filtering fraction parameter, leading to a $\sim 20\%$ energy loss as in Fig. 1 of Cantalloube et al. (2015). We calculate the oversampling parameter for each wavelength channel to ensure the sampling is constant across wavelength, and additionally use this parameter to determine the outer working angle, which is provided in λ/d . The oversampling parameter, S is defined to be

$$S = \frac{\lambda}{2\sigma_{\text{px}}D} \quad (1)$$

for pixel scale σ_{px} , telescope diameter D and wavelength λ .

Internally to ANDROMEDA pairs of images are chosen such that they are as close together in time as possible to preserve speckle self-similarity, while still ensuring movement of the proposed companion in order to avoid self-subtraction. This is done on an annular basis, as the motion of the planet depends on the separation from the host star. A scaling factor γ is fit using a least squares method to ensure that the mean of the intensity distribution of both images in the pair is equal. Using the assumption that the residual noise is white and Gaussian, ANDROMEDA then can perform a likelihood test to identify the presence of a companion, by minimising the difference between the residuals and a model of the companion signal.

Among the outputs of this algorithm are a contrast map, where each pixel represents the contrast of the planet if it was centred on that pixel, and a standard deviation map, specifying the uncertainty associated with each contrast estimate. This is different from the output of a speckle-subtraction algorithm, where the flux of an object must be estimated through aperture photometry or via fitting a PSF model to the residuals. To extract the spectrum, we sum the S/N map along the wavelength axis, and identify the maximum S/N pixel in a 10 pixel box around the known position of the planet. We then use this location to measure the contrast and standard deviation as a function of wavelength.

3.2. Spectral covariance estimation

Both high contrast imaging and IFS observations present challenges when deriving robust uncertainty estimates, as correlations are naturally present in the data. Due to aberrations in the telescope optics, imperfect correction for atmospheric turbulence from the adaptive optics systems, and imperfect stellar PSF subtraction, speckles from the stellar PSF are the dominant noise source for AO assisted, high-contrast datasets (Marois et al. 2006). These speckles move radially as a function of wavelength, scaling with the size of the stellar PSF. This induces a correlation between wavelength channels, as a speckle will take several channels of movement to pass over a pixel at a fixed separation. Crosstalk – light from a single lenslet in the lenslet array diffracting into neighbouring channels – will also couple these channels. Finally, as noted in Ruffio et al. (2021), additional correlation can be introduced through the interpolation of the 4D (λ, t, x, y) spectral cube during reconstruction from the detector images. This interpolation to a fixed wavelength grid guarantees the correlation of the noise in the IFS cubes, as noise in neighbouring detector pixels will be interpolated to build the IFS spaxels.

Greco & Brandt (2016) demonstrate the necessity of accounting for these correlated errors when retrieving physical properties from IFS data. If these correlations are not accounted for, they find that the retrieved confidence intervals are both artificially small and unreliable, often excluding the true parameter values at $> 95\%$ confidence. This was reinforced by Ih & Kempton (2021), where they explored the impact of correlated noise on atmospheric retrievals for transiting planets, finding that the assumption of non-correlated noise leads to biased posteriors and overfitting of the data.

3.2.1. Measuring noise correlation

Greco & Brandt (2016) introduce a procedure for empirically measuring the correlation in IFS datasets, and demonstrated the importance of including the full covariance matrix when fitting IFS spectra. In this work we extend their method by measuring the spectrum of injected planets and the resulting covariance, as opposed to the parameterised noise instance used in their work. This allows us to explore how the noise properties vary across instruments and over different post-processing methods. For each PSF-subtracted dataset we compute the average correlation within a 6 pixel wide annulus centred at the separation of the companion of interest. As in their work, we find that the correlation matrix does not depend strongly on the width of this annulus. The companion itself is masked out, leaving only residual noise. Such an annulus is chosen in order to maintain consistent noise properties in the sample of pixels: in general the noise varies more strongly with radius than with position angle. Work such as Gebhard et al. (2022) explores choosing more a more representative sample of pixels to describe the noise at the location of the planet, but such methods are computationally expensive, and we see little azimuthal asymmetry in the residuals shown in Fig. 1.

Within the annulus, we compute the elements of the correlation matrix, ψ_{ij} as

$$\psi_{ij} = \frac{\langle I_i I_j \rangle}{\sqrt{\langle I_i^2 \rangle \langle I_j^2 \rangle}} = \frac{C_{ij}}{\sqrt{C_{ii} C_{jj}}}, \quad (2)$$

where $\langle I_i \rangle$ is the mean pixel intensity in the i th spectral channel, and C_{ij} is the covariance between the two channels.

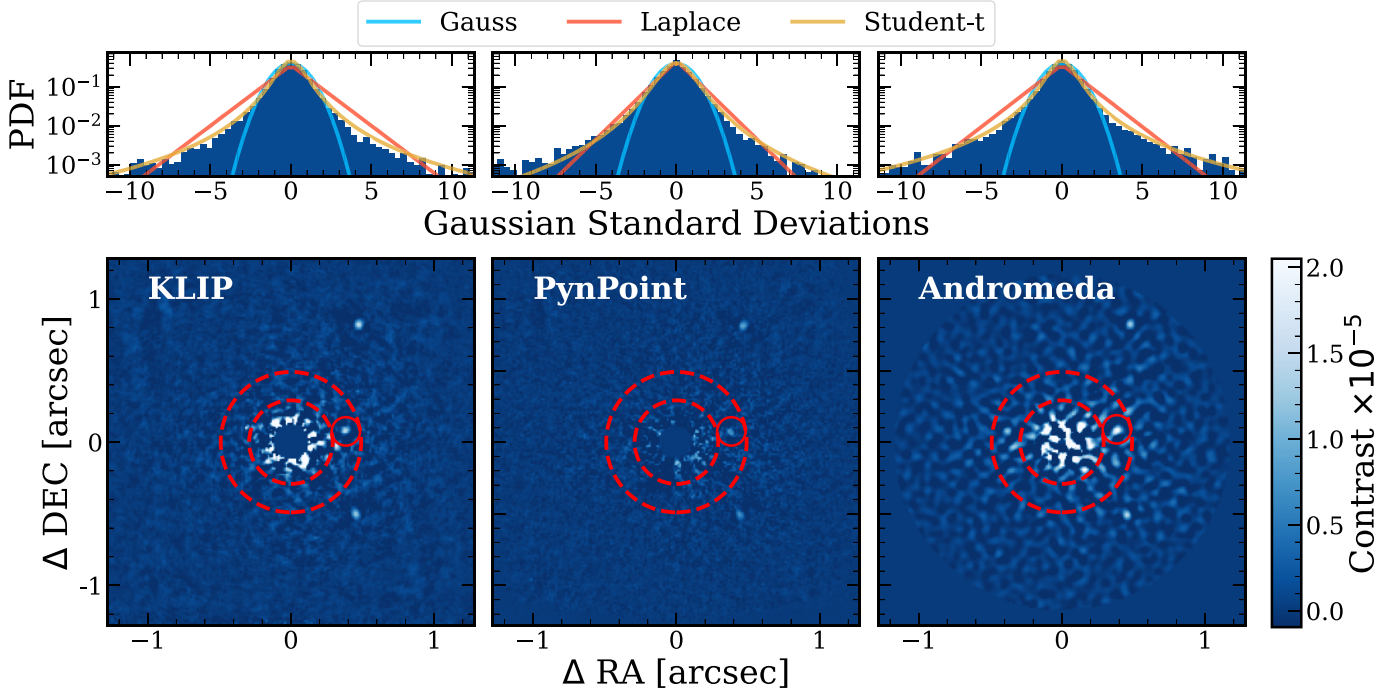


Fig. 1. Single wavelength channel from the GPI H -band data of the HR 8799 system, post-processed with KLIP (left), PynPoint (centre) and Andromeda (right). The solid red circle denotes the position of HR 8799 e as computed using the pyKLIP astrometry module. Marked in red dashed lines are the regions which are compared when finding the mean correlation between different wavelength channels. At the top of the figure, histograms of the residuals are plotted in units of σ . The light blue line is a Gaussian fit with the width defined as the Gaussian standard deviation of the residual frame below.

3.2.2. Estimating uncertainties

In order to compute the covariance matrix from the correlation matrix, we must know the diagonal, or uncorrelated elements of the covariance matrix. Several methods of measuring the photometric uncertainty were considered. We estimate the uncorrelated error in each wavelength channel by combining the photometric uncertainty of the stellar PSF, σ_{star} , with the residual noise at the location of the planet, σ_{residual} . We include the stellar uncertainty because near the edges of the bands in which spectra are observed, the filter transmission drops and atmospheric absorption increases, resulting in an increase in the uncertainty on the host star photometry. To measure the uncertainty on the stellar photometry we measure the standard deviation of the background in an annulus far from the stellar PSF in each wavelength channel, and use this to calculate the signal to noise. This represents an optimistic estimate of the stellar uncertainty, as we are unable to monitor photometric variability due to atmospheric conditions over the course of the observation, which represents the dominant source of uncertainty for the stellar photometry. To measure the uncertainty on the planet photometry we take the standard deviation of the residuals in an annulus at the separation of the companion, masking out the planet itself.

The histograms of Fig. 1 show that the assumption of Gaussian errors across the entire frame is inconsistent with the noise, and would underestimate the tails of the distribution. Pairet et al. (2019) demonstrate that a Laplacian provides a better fit to the tails of the residual distribution than a Gaussian, while Mawet et al. (2014) shows that the residuals tend to follow a Student-t distribution. We find that a Student-t distribution best matches the full frame residuals. However, as the likelihood function for a general Student-t distribution is not analytic, and a Gaussian distribution accurately captures the residuals to 2.15σ , we continued

to follow the standard practice of defining uncertainties as the Gaussian standard deviation. Taking KLIP as an example, the best fit Student-t distribution has 1.75 DoF, a mean of 2.23×10^{-8} and a width $t = 5.29 \times 10^7$. At the point where the Gaussian distribution intersects the Student-t distribution, 89% of the residuals are enclosed, compared to 97% if the residuals were Gaussian distributed. We also note that, relative to the speckle subtraction algorithms, ANDROMEDA shows an excess of 10σ outliers in the residuals, leading to difficulties in distinguishing between true positives and false positives, consistent with the findings of Cantalloube et al. (2020b). The long tails of these distribution add additional noise to each frame and need to be accounted for the detection of planet candidates in order to avoid false positives. However, for a known companion where we are concerned with inferring physical parameters to within $1-2\sigma$ confidence intervals, accounting for the 90% of the noise that is contained within the Gaussian fit to the residuals is sufficient for defining the uncertainties.

Thus the total uncorrelated uncertainty for the i^{th} wavelength channel is given as:

$$C_{ii} = \sigma_{i,\text{star}}^2 \left(\frac{f_{i,\text{pl}}}{f_{i,\text{star}}} \right)^2 + \sigma_{i,\text{residual}}^2. \quad (3)$$

The method described here provides an empirical estimate of the covariance of the noise after high-contrast image processing. However, as it relies on measurements of mean pixel intensities in a residual image, it is only applicable for speckle-subtraction methods. As ANDROMEDA produces an estimate of planet contrast at each pixel location, rather than residual noise following PSF subtraction, this method cannot be directly applied to the processed ANDROMEDA frames. An example of such a frame is shown in Fig. 1, where highly structured noise is visible in the frame. The

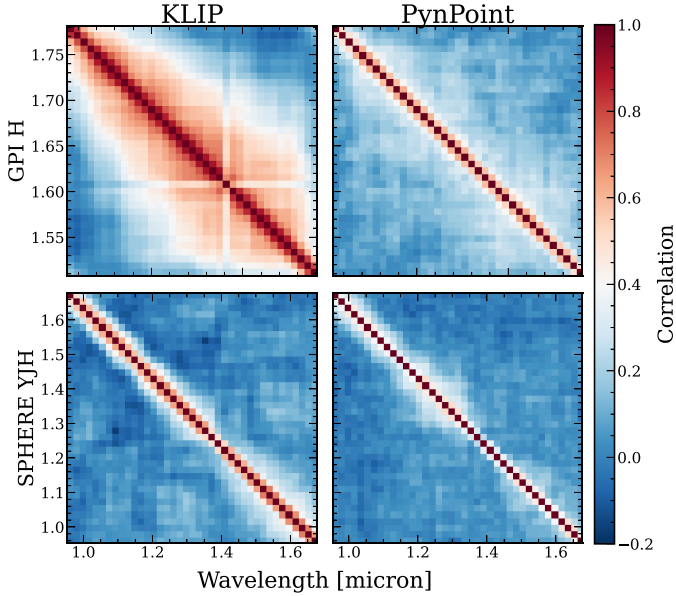


Fig. 2. Correlation matrices for each dataset for HR 8799 e, with the GPI *H*-band data shown in the top row and the SPHERE *YJH* data on the bottom. Following the processing using KLIP (left) or PynPoint (right), we calculate the correlation and covariance matrices as described in Sect. 3.2. The correlation is computed as in Eq. (2). The GPI data is more strongly correlated than the SPHERE data, particularly following the KLIP processing. The SPHERE data shows structure similar to the correlation matrix, with the correlation width following the shape of the water absorption spectrum.

noise pattern is highly correlated through wavelength space, and indeed would lead to very strong residual correlation. Rather than applying the procedure for measuring the covariance matrix for ANDROMEDA, we instead rely on the estimate of the standard deviation that is also provided by the algorithm, that is also measured during the likelihood minimisation.

Figure 2 shows the results of computing the correlation matrix for both the KLIP and PynPoint reductions. There is a strong, narrow correlation component along the diagonal with a width of around 2–3 pixels, with a weaker correlation extending out to 10 pixels in width. In the SPHERE data, the correlation decreases in the water absorption features at 1.15 μm and 1.4 μm . The KLIP data typically displays stronger correlations than the PynPoint reductions. This difference may be because the pyKLIP implementation of the KLIP algorithm uses only the most correlated frames from the PSF library to build the PSF model, which introduces an additional source of correlation in the data.

3.3. Bias correction

PCA methods tend to see increased self-subtraction as the number of principal components increases, naturally leading to poorer extractions as the number of components increases (Lagrange et al. 2010). Therefore we also considered an empirical estimate of the uncertainty by injecting and recovering a sample of planets in an annulus at the location of the planet. The standard deviation of the recovered spectra provide a measure for the uncertainty due to planet position and variation in the effectiveness of the post-processing. This can also be used as a method to correct for bias introduced by self-subtraction caused by the post-processing algorithm: by comparing the recovered spectra to the known input, a scaling factor can be computed. This can then be used to mitigate the self-subtraction induced by the PCA processing

(Marois et al. 2014a; Gerard & Marois 2016; Ruffio et al. 2017). We found that when applying such a bias correction, the χ^2 between the injected and recovered spectra was often worse than that of the nominal spectral extraction. As the noise properties are not truly azimuthally symmetric, the average bias correction does not provide a good correction for any individual planet location, particularly at the relatively faint contrasts considered in this work. Any improvement in the spectral extraction is dominated by the both the changes to the shape of the spectrum introduced by the data processing and the random noise of the measurement. Therefore we choose not to include bias correction as a step in our post processing, and do not include uncertainty from injection and recovery tests in our error estimate. The finding that bias correction can reduce the accuracy of the spectral extraction is surprising, and warrants further investigation into where this widely used technique should be applied. We leave such a study to future work.

3.4. Impact of covariance on retrieved parameters

In the previous section, we discuss how to measure a covariance matrix for IFS data. Here we explore how the covariance impacts Bayesian inference, that is, when estimating the parameters of the model used to explain the observations. Greco & Brandt (2016) demonstrated how failing to include the full covariance matrix when fitting atmospheric models to IFS data results in overly confident and biased parameter estimates. In Sect. 5.1, we expand their work to higher dimensional models using atmospheric retrievals, but first we want to pedagogically understand how the covariance is linked to the posterior probability distributions. We subsequently show how the precision of a posterior parameter estimate depends on the ratio between the length scale of the correlation in the data and the length scale over which the parameter introduces changes in the model spectra. If this ratio is larger than unity, the posterior width decreases relative to the case without correlation. If the ratio is about unity the posterior width will increase.

Consider a toy model, where the data \mathbf{y} is given by a simple sine function with period T and offset D :

$$y_i = \sin(2\pi x_i/T) + D. \quad (4)$$

In the context of atmospheric parameters, we can think of this model as the first term of a Fourier series, which can be used to describe an atmospheric spectrum to arbitrary precision if extended to a high enough order. While the variation in the spectrum due to physical parameters is more complicated than this model, we can view an offset in the toy model as a change in the overall flux, while a change in the period would be reflected in the spectral shape, such as the near infrared water features.

The period and offset were arbitrarily chosen to be 30 and 0 respectively. Assuming that this toy model describes an observed experimental setup, we construct a synthetic dataset containing a total of 300 points, with x -coordinate values from 1 to 300, that is, ten periods. We applied different noise models to these toy data, and used nested sampling as implemented in MultiNest (Feroz & Hobson 2008; Feroz et al. 2009, 2019) to retrieve the value of the period parameters T and D . Such nested sampling algorithms improve on MCMC techniques to sample large parameter spaces efficiently, gradually restricting the sampling volume to regions of high likelihood. They are also robust to multimodal posterior probability distributions, and provide an estimate for the Bayesian evidence, \mathcal{Z} , as well as the posterior distributions and maximum likelihood fit.

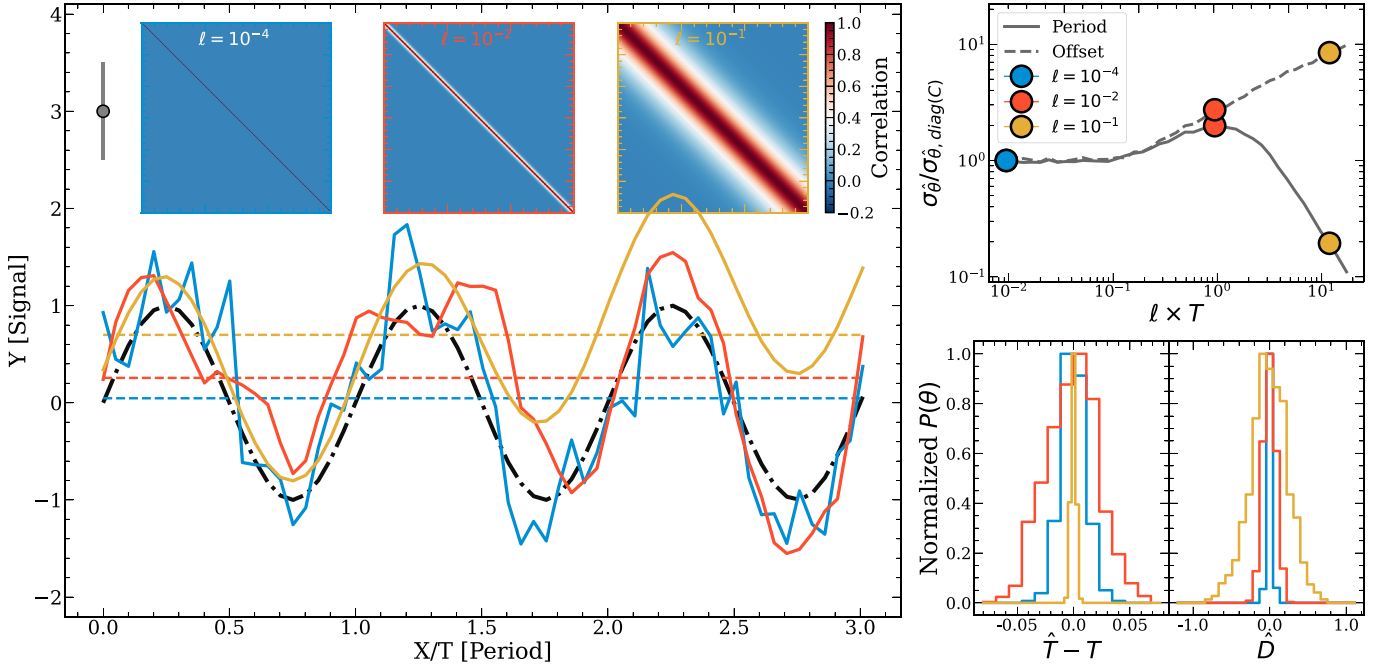


Fig. 3. Toy model for demonstrating the effect of correlated uncertainties. Left: data drawn from a toy sinusoidal model (Eq. (4)) when considering three cases of covariance in the data. The first (blue) is data drawn from a univariate Gaussian distribution with no correlations, as shown in the left inset. The second (orange) is drawn from a multivariate Gaussian distribution where the correlation length scale – as defined by the ℓ parameter of the Matérn kernel (Eq. (5)) – is less than the period of the model (centre inset). The third (yellow) is drawn from a multivariate Gaussian distribution with a correlation length scale greater than the period of the model (right inset). In the background the true input is plotted in gray. The gray datapoint indicates the 1σ error bar associated with each data point. Right: in the top panel, the posterior width of the period (solid gray line) and offset (dashed gray line) parameters scaled to the uncorrelated case as a function of the ratio between the correlation length scale and the period of the sine (so the model length scale). Marked in blue, orange and yellow are the draws plotted in the right panel. The histograms in the bottom panel are the posterior histograms for the period (left) and offset (right) for each of the highlighted cases.

We use the Matérn 3/2 kernel (Rasmussen & Williams 2005) to describe the covariance of our dataset:

$$C_{3/2}(d)_{ij} = \sigma_i \sigma_j \left(1 + \frac{\sqrt{3}d_{ij}}{\ell} \right) \exp\left(-\frac{\sqrt{3}d_{ij}}{\ell} \right). \quad (5)$$

This is a model of for the correlation between two points separated by distance d_{ij} , where we can adjust characteristic correlation length scale through the correlation length parameter, ℓ . As ℓ decreases, the correlation matrix becomes more diagonal, while as ℓ increases the data becomes more strongly correlated across broad scales. The strength of the correlation is determined by the uncertainty on each point (σ_i), which we set to a constant value of 0.5 for each data point.

In Fig. 3, we show three instances of the correlation matrix for $\ell = 10^{-4}$, 10^{-2} and 10^{-1} , ranging from uncorrelated to strongly correlated noise. Plotted in gray is the noise-less data, and the coloured lines show noise instances drawn from each of the three correlation matrices. For the diagonal case ($\ell = 10^{-4}$), we see that the data are randomly scattered around the true model. With no correlation, we see true, univariate Gaussian noise. As the correlation scale increases to $\ell = 10^{-2}$, we see that the data appear smoother, and the variations occur on larger spatial scales than in the case without any correlation. Finally, with $\ell = 10^{-1}$, we see that the data are offset from the ground truth, but do not have any small-scale scatter. This is the impact of covariance on the data: as the correlation length increases, every point is more strongly determined by the initial random draw any other point (effectively there are less points, as there are less independent measurements). In summary, we observe high frequency variation due to noise

if the correlation length is small. In this case the mean of the data, parameterised by parameter D should be accurately and precisely inferred. As the correlation length increases, the scatter of the mean D across multiple draws increases, but we see less small-scale variation, allowing a better estimate of the period, T .

We note here that the Matérn 3/2 kernel is only one model for the covariance, and is not perfectly suited for IFU data. A more robust model (such as described in Greco & Brandt 2016) would incorporate both a broad correlation term and a diagonal Gaussian term to the correlation matrix, which would introduce small-scale scatter in the data, even with a large correlation scale. Nevertheless, this is a suitable toy model to explore how changing the correlation length scale impacts parameter estimation.

To determine the impact on parameter inference, we vary the correlation matrix across a range of ℓ values from 10^{-4} to 10^0 , and use `pyMultiNest` (Buchner et al. 2014) with 400 live points to fit the true model, accounting for the covariance in the likelihood. We did not perturb the toy dataset by with an error model as defined by the covariance matrix, so we run noise-free retrievals. This is equivalent to running multiple inferences where the data are perturbed by draws from the covariance matrix and averaging over the posteriors of each inference. We set uniform priors on the period $P(\omega) = \mathcal{U}(0, 100)$ and offset $P(D) = \mathcal{U}(-10, 10)$. The results are not sensitive to the choice of number of live points ($n_{\text{live}} \gg n_{\text{param}}$) or priors. The upper right panel of Fig. 3 shows the ratio between the width of the posterior distribution for this parameter and the width of the distribution in the case of univariate Gaussian noise (i.e. no correlation). We observe that as the correlation length scale approaches the length scale of the sine function (the period) the width of the posterior increases: the

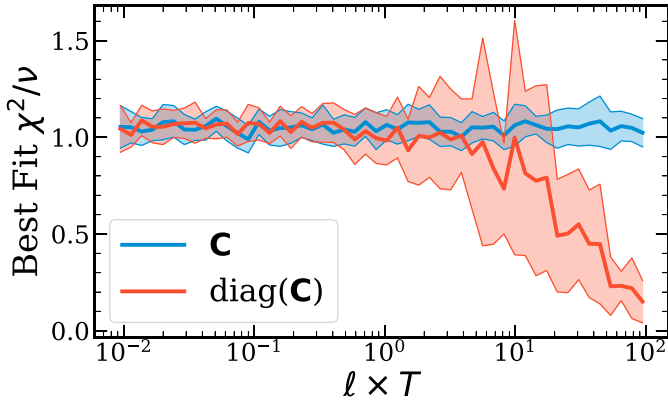


Fig. 4. Best fit χ^2/ν as a function of the ratio between the correlation length scale (proportional to $1/\ell$) and the period, T . The χ^2 was computed for fits of Eq. (4) to data perturbed by draws from the covariance matrix, varying the correlation length scale. For each ℓ , 25 Multinest retrievals were run in order to compute the uncertainty on the χ^2 , shown as the shaded region around the mean. In blue, the covariance is properly accounted for in the likelihood, while in orange only the diagonal of the covariance is used in the likelihood. In order for the reduced χ^2 to be a useful metric, the covariance must be properly accounted for.

correlation introduces variations in the data on the scale of the period, making it difficult to estimate the parameter. This is the effect described in Greco & Brandt (2016), where accounting for the covariance matrix when fitting atmospheric models increases the posterior width. However, as the correlation scale continues to increase to scales larger than the period, we see that the posterior width decreases to values lower than in the case of uncorrelated noise. As is visible in the data in the left panel, without small scale variations to introduce uncertainty in the period, it becomes easier to estimate this parameter, at the cost of increased uncertainty in the estimate of the offset parameter D .

Effects of ignoring covariance. It is often the case that the full covariance matrix is not used when performing atmospheric retrievals, and we wanted to explore the impact of using only the diagonal terms when fitting a model to correlated data. Figure 4 shows the best-fit reduced χ^2 as a function of the ratio between the period and the correlation length scale, as in the right panel of Fig. 3. For each ℓ , we perform an ensemble of 25 retrievals using Multinest in order to reduce the scatter and to measure the uncertainty in the χ^2 due to the variation between individual noise instances. In this case, the data are perturbed by draws from the covariance matrix, in order to test the impact of using the incorrect covariance in the likelihood when the data are correlated. We define ν as the number of data points (300) minus the number of parameters (2). This procedure is repeated using both the full covariance matrix, \mathbf{C} , in the likelihood, as well as using only the diagonal elements of the matrix – that is, we assume that the data are uncorrelated. We find that the reduced χ^2 is a useful metric if the covariance is properly accounted for. If the data are correlated and an only the uncorrelated uncertainties are used in the likelihood then the reduced χ^2 will be underestimated, and the scatter of the χ^2 increased. Often a $\chi^2/\nu < 1$ is interpreted either as overfitting of the data or overestimation of the uncertainties. However, we demonstrate here that for $n_{\text{param}} < n_{\text{data}}$ a $\chi^2/\nu < 1$ can be interpreted as an underestimation of the correlation of the data.

Applicability to atmospheric retrievals. We expect similar effects to be present in atmospheric retrievals with correlated data.

Parameters that affect model spectra on wavelength scales smaller than the correlation scale may be retrieved to higher precision than expected if the uncertainties were uncorrelated, while parameters that are sensitive to approximately the correlation length scale will have larger posterior uncertainties. As seen in Fig. 2, the correlation length can be an appreciable fraction of the total data, particularly in the case of the KLIP reduction of GPI data. Large scale correlations in the data can introduce offsets in the average flux measurement, which can lead to inconsistencies between datasets from different instruments or measured during different epochs. Correlations on moderate scales can alter the spectral shape, in turn impacting parameter estimates. For example, the surface gravity is particularly sensitive to the shape of the H -band, and changes to the shape of this band will lead to biased estimates. Thus for IFS data it is critical to account for the covariance matrix when fitting models to the data, in order to correctly capture the noise structure imprinted onto the signal.

4. Injection testing

In order to best extract a true signal, we want to optimise the data-processing parameters. However, without knowledge of the ground truth spectrum, it is unclear how these parameters should be tuned a priori. By injecting fake companions with a known spectrum, applying the post-processing, and comparing the extracted spectrum to the input we can then optimise the parameters, and use this setup to extract the true planet signal. In particular, we try to optimise the choice of the number of principal components used in PSF subtraction for KLIP and PynPoint as a function of the separation. This injection-extraction study also provides us with a metric for comparing the three algorithms described in Sect. 3.1.

Using the pyKLIP injection tool we injected companions into both the SPHERE and GPI HR 8799 datasets. The normalised stellar PSF was used both as a model for the planet PSF and to scale the bulk contrast of the injected companion. The spectrum was convolved with a Gaussian kernel to the instrumental resolving power, and binned to the instrumental wavelength grid using the `rebin_give_width` function available in `petitRADTRANS`, which accounts for non-uniform bin sizes as the number of pixels per instrumental resolution element varies with wavelength. Only a single planet was injected at a time before the data processing, which was repeated for each planet position in order to avoid potential contamination from nearby signals. We injected the companions at varying positions into both the SPHERE and GPI datasets, with a spectrum generated using `petitRADTRANS` as described in Sect. 5. These were positions representative of the known separations of the inner three companions. Planets were injected at position angles from 120° to 240° from the location of HR 8799 e in 30° increments, and between 300 and 800 mas in 100 mas steps. This process was repeated for the SPHERE YJH and the GPI H -band datasets at mean contrasts from 10^{-7} to 10^{-4} .

Once the data were prepared, we ran each of the three data processing algorithms on each injected dataset, spanning a range of algorithmic parameters. While we could not exhaustively study the effect of each parameter, we chose to focus on the impact of the number of principal components used during PSF subtraction in order to optimise the spectral extraction. Other parameters, such as the `flux_overlap` parameter in KLIP, or the filtering fraction in ANDROMEDA were set based on suggested values from previous studies (Zurlo et al. 2016; Cantalloube et al. 2015) or from qualitative examination of the post-processed data. Several parameters, such as the `tolerance` and `merit` parameters of

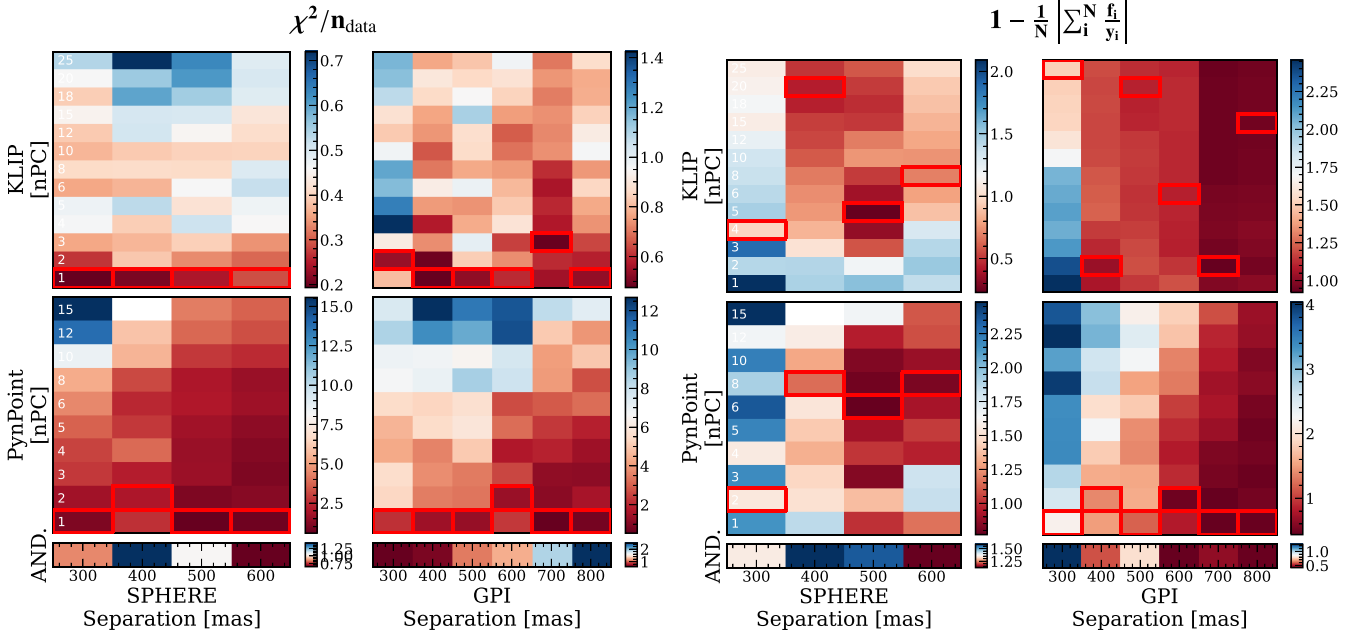


Fig. 5. Goodness-of-fit metrics mapped across separation and number of principal components. The colour scale indicates median metric value from the injections at 4 separate position angles. The range of the colour scale for each sub plot is different in order to capture the variation within a single map. Highlighted in red are the optimal extractions for each separation. Left: χ^2/n_{data} map. Right: the same as the left panel, but calculated using the relative discrepancy (e) instead of the χ^2 .

PynPoint were chosen to ensure accurate extractions within reasonable computation time. Various geometric parameters, such as the inner and outer working angles together with the width parameter in ANDROMEDA, or the subsection and annuli parameters of KLIP were set based on recommendations from the documentation¹, and ensuring that the region under consideration would contain the entirety of the planet signal, extending to at least twice the FWHM of the signal. A full table of parameter choices for each algorithm is given in Table 3.

4.1. Choice of goodness-of-fit metric

We considered several goodness-of-fit metrics with which to determine the optimal extraction, including the signal-to-noise ratio (S/N), the relative discrepancy (e) and reduced χ^2 (χ^2/n_{data}). We take the median of each metric across the five different position angles where planets were injected. We exclude spectra that are over 20σ discrepant from the input, or that display strong outliers with contrast $> 2 \times 10^{-5}$, though the results are robust to including the outlying data. Each of these metrics identified different optimal spectra.

The mean S/N always identified the spectra processed using the largest number of components as optimal. However, the resulting spectra do not correctly retrieve the shape of the input spectrum, because they typically overestimate the flux, so we did not consider this metric further.

We define the mean relative discrepancy e between a measured flux s and known input spectrum \bar{s} as

$$e = 1 - \frac{1}{N} \left| \sum_i^N \frac{s_i}{\bar{s}_i} \right|. \quad (6)$$

¹ https://pyklip.readthedocs.io/en/latest/fm_spect.html

To identify the best fit spectrum we simply find the minimum value of this function. In contrast to the χ^2 or other distance metrics, the discrepancy is invariant of the magnitude of the measured quantity, and so provides a metric to compare spectra injected at different contrasts.

The χ^2 is a standard metric for measuring the similarity of distributions, but can also favour measurements with overestimated uncertainties. The χ^2 value between the extracted spectrum s with covariance C and the known injected spectrum \bar{s} was calculated for each post-processed dataset as

$$\chi^2 = (s - \bar{s})^T C^{-1} (s - \bar{s}). \quad (7)$$

We present χ^2/n_{data} , dividing by the number of wavelength channels n_{data} to allow for a more straightforward comparison between instruments. We do not subtract the degrees of freedom from the number of data points as is typical when computing the reduced χ^2 , as principal components are not free parameters in a statistical sense, thus making the definition of degrees of freedom challenging.

4.2. Optimising spectral extractions

We present in Fig. 5 the map of χ^2/n_{data} (left) and mean relative discrepancy (right) as a function of both separation and number of principal components used, taking the median across the injections at different parallactic angles. Extracted spectra for typical injected fake companions are shown in Fig. 6. The precision and accuracy of our spectral measurements depends strongly on the separation, visible in the variation of the both metrics. There is strong position-dependent variation in the shape of the extracted spectra, including cases where the injected companion is not detected in any wavelength channel, as well as cases where the peak contrast is overestimated by a factor of 2. Examining Fig. 1, we see that beyond about 400 mas from the host star the noise properties are relatively unstructured, while

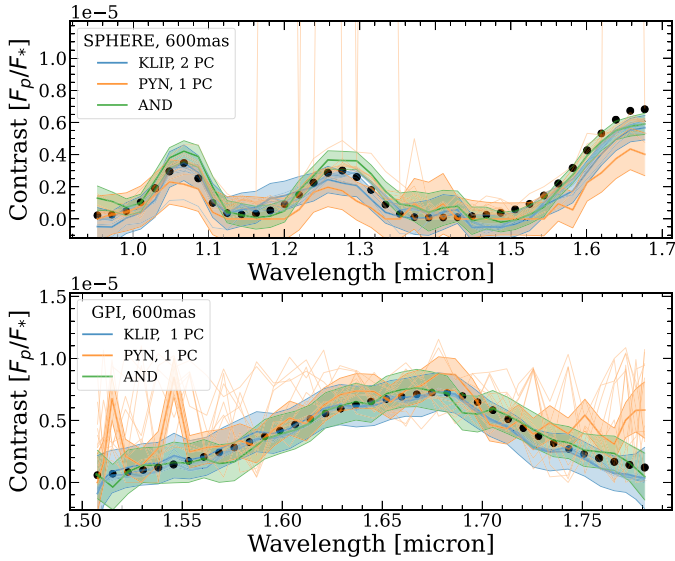


Fig. 6. Typical spectral extractions for injected planets located at separations of 600 mas. These spectra are representative of the HR 8799 planets with ($F_p/F_* \sim 2 \times 10^{-6}$). The injections into the SPHERE data are shown on the top panel, the GPI on the bottom. Each injected planet was positioned 150° from HR 8799 e. Extractions for each algorithm are plotted, with the best fit spectrum (χ^2) and 1σ error bars from the diagonal of the covariance matrix highlighted by the shaded region. The faint lines show the variation in the extractions using different numbers of principal components.

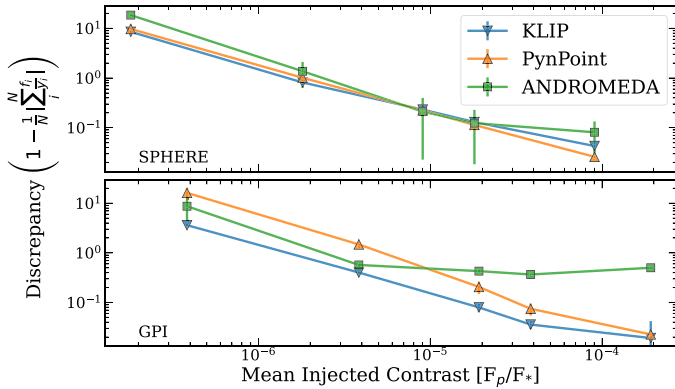


Fig. 7. Best-fit discrepancy (Eq. (6)) as a function of input contrast at 400 mas. The top panel shows the results for injections into the SPHERE data cube, while the bottom is for GPI. The injections were repeated at three position angles, and the uncertainty presented is the standard deviation of these measurements.

at 400 mas and closer the GPI data is dominated by the residual speckle noise. Such a trend is also present in the SPHERE data, though the speckle dominated regime extends out to only 300 mas. This transition in the underlying noise properties together with the greater angular displacement at wider separations results in the improved detections at wider separations. This can be disguised by the χ^2 metric, where large uncertainty estimates at small separations can result in a better χ^2/n_{data} , while the mean relative discrepancy provides a clearer trend as a function of separation. When using χ^2/n_{data} as the goodness-of-fit metric, we find that both KLIP and PynPoint favour low numbers of principal components.

Such results depend strongly on both the dataset, the reduction used, and the choice of metric. The SPHERE extractions are universally better than the GPI, largely due to the brightness of the injected spectra in the Y and J bands, the smaller inner working angle and pixel scale, and the longer integration time. PynPoint most strongly favours low numbers of principal components, across both metrics and at all separations. However, when we consider the relative discrepancy between the input and extracted spectra, and find that the number of PCs components favoured is much higher than when using the χ^2 , particularly for KLIP. These spectra may more closely match the shape of the input spectra, but may also underestimate the uncertainties, leading to them being disfavoured by the χ^2 . While the shape of the spectrum does depend on the number of components used, it is more strongly dependent on the particular location in the frame where it is injected.

Each algorithm displayed its own trends in the quality of its spectral extractions. KLIP produced smooth spectra, with systematics that were relatively consistent in shape over the full range of principal components used. However, it struggled to recover the brightest sources accurately, reproducing the over-subtraction effect described in Pueyo (2016). In contrast, PynPoint and ANDROMEDA performed worse at fainter contrasts, as demonstrated in Fig. 7. The PynPoint spectra are more dominated by random scatter than by systematic variation, reflected in the typically diagonal correlation matrices. ANDROMEDA produced some of the best overall fits, but struggled to achieve the correct flux calibration, both over- and under-estimating the flux in different cases. For the bright injection case, PynPoint consistently performed the best, producing the lowest χ^2 values for each dataset and separation. KLIP struggled to extract the brightest spectra, over-subtracting the planet signal at the red end of the SPHERE data. This effect was more severe when larger numbers of principal components were used. However, KLIP also displayed a tendency to over-estimate the flux of the signals injected into the GPI data. ANDROMEDA was able to accurately extract the high S/N SPHERE injection, but systematically underestimated the flux in the GPI data.

The best fit of each algorithm performs relatively well at extracting the true spectrum, with typical best-fit reduced χ^2/n_{data} values approaching 1. Depending on the injected position angle and separation, the χ^2 for the same algorithm at the same separation can vary by a factor of ~ 10 , with typical standard deviations on the order of 10–100 depending on the dataset and algorithm. This variation in extraction suggests that injection recovery tests to measure and correct for algorithm throughput, such as detailed in Greenbaum et al. (2018), may introduce additional biases depending on the precise positioning of the injected companions. In our reproduction of this method, we find that it does not provide better χ^2 values, and can introduce spurious wavelength-dependent signals. This variation is separation-dependent and impacts the extraction less strongly at wider separations, outside the speckle noise regime.

These results point to differences in the approach to data analysis required between detection and characterisation efforts. High numbers of principal components tend to whiten the noise and improve the detection significance, potentially allowing the discovery of fainter companions. However, this comes at the cost of reduced photometric accuracy, which is critical when attempting to recover the physical atmospheric parameters. For both speckle subtraction methods, low numbers of principal components produce the most accurate spectral extractions, though the precise number of components will depend on the brightness of the companion in question.

5. Retrieval tests

Atmospheric retrievals provide a useful, data-driven tool for exploring the properties of exoplanet atmospheres. Retrieval results are dependent on the quality of the input data and the assumptions made about both the data and the model. For our investigation below we have two primary aims:

1. Exploring the impact of high-contrast image processing on the inferred atmospheric parameters through retrievals on synthetic data.

2. Characterising how correlated noise influences the fits of the synthetic data. With a known ground truth, we can explore how the use of the covariance matrix can help mitigate the impact of systematic effects introduced by the data processing.

To this end, we use a representative selection of optimised spectral extractions as described in Sect. 4. We choose to use the model injected at 600 mas, and positioned 150° rotated from HR 8799 e, combining both the SPHERE *YJH* and GPI *H*-band datasets. The best extracted spectrum as measured by the relative discrepancy were chosen as the baseline inputs to the retrievals. This represents a realistic, though challenging spectrum on which to perform atmospheric retrievals. For validation we also explored a set of retrievals on different locations and choice of extraction, finding that while the precision often varies with the S/N of the extracted spectrum, the overall trends of our results are reproducible. In contrast to Greco & Brandt (2016), these retrievals explore the full impact of IFS data processing on the spectra, as opposed to using data synthesised from a parametric estimate of the covariance. In contrast to their use of a 3-parameter BT-Sett1 model, we use an ~8 parameter forward model in order to understand the cumulative impacts of data post-processing on the inferred atmospheric parameters in the context of high-dimension on-the-fly retrievals. Such retrievals are highly flexible, and are more likely to try to fit spurious data features than more physically motivated fits from self-consistent grids.

Atmospheric model. The models we use in our atmospheric retrieval setup are computed using `petitRADTRANS` (Mollière et al. 2019), a fast, open-source radiative transfer code with which we can calculate the emission spectrum of an atmosphere². In this framework, the atmosphere of a planet is divided up into pressure bins. Temperature and chemical structures are calculated and applied to each bin, and radiative transfer using the correlated-k method for the opacities (Goody et al. 1989; Lacis & Oinas 1991) is performed to calculate the emission spectrum. The correlated-k opacities are binned from their native spectral resolving power of 1000 to a user-supplied model resolution using the `exo-k` package (Leconte 2021), improving the computation time of the retrieval. A wavelength binning of at least twice the data resolution is used for the models, in order that the binned model spectrum is Nyquist sampled. This spectral model is convolved with a Gaussian kernel with the width of the instrumental spectra resolution and then binned to the wavelength grid of the input data for the retrieval using the `rebin_giv_width` function. At the spectral resolutions considered in this work, the effects of the convolution and binning on the spectrum can dominate the spectral shape over data processing effects. For this reason we ensured that we use the same convolution and binning procedure during the spectral injections as during the retrieval. However, future work should investigate incorporating better instrumental models and wavelength dependent kernels into retrieval frameworks.

² <https://petitradtrans.readthedocs.io/>

Table 4. Priors for retrieval setup.

Parameter	Prior	Input
Guillot, free chemistry		
$\log g$	$\mathcal{U}(2.0, 5.5)$	4.0
R_{pl}	$\mathcal{N}(1.0 R_{\text{jup}}, 0.2 R_{\text{jup}})$	$1.0 R_{\text{jup}}$
T_{int}	$\mathcal{U}(300 \text{ K}, 2000 \text{ K})$	750 K
T_{equ}	$\mathcal{U}(0 \text{ K}, 300 \text{ K})$	100 K
γ	$\mathcal{N}(1, 0.2)$	0.5
$\log \kappa_{\text{IR}}$	$\mathcal{U}(-3.0, 1.0)$	-1.0
$\log X_{\text{H}_2\text{O}}$	$\mathcal{U}(-7.0, 0.0)$	-1.5
$\log X_{\text{CO}}$	$\mathcal{U}(-7.0, 0.0)$	-2.0

Notes. $\mathcal{U}(a, b)$ denotes uniform priors with bounds a and b , and $\mathcal{N}(\mu, \sigma)$ denotes a normal distribution centred at a mean μ with standard deviation σ . The final column indicates the true values of the spectrum injected into the IFS cubes.

Our baseline model uses a Guillot temperature profile (Guillot 2010) and freely retrieved chemical abundances. This profile is a simple analytical model, constructed to estimate the thermal structure of irradiated planets:

$$T_{\text{Guillot}} = \frac{3T_{\text{int}}^4}{4} \left(\frac{2}{3} + \tau \right) + \frac{3T_{\text{irr}}^4}{4} \left(\frac{2}{3} + \frac{1}{\gamma\sqrt{3}} + \left(\frac{\gamma}{\sqrt{3}} - \frac{1}{\gamma\sqrt{3}} \right) e^{-\gamma\tau\sqrt{3}} \right), \quad (8)$$

where $T_{\text{irr}} = \sqrt{2}T_{\text{equ}}$ and $\tau = P \times \kappa_{\text{IR}}/g$. T_{equ} is the standard equilibrium temperature of an irradiated body, g is the surface gravity. P is the atmospheric pressure, divided up into a total of 80 log-spaced layers from 10^3 bar to 10^{-6} bar. The remainder of the parameters are as in Guillot (2010): T_{int} is the intrinsic internal temperature of the planet, κ_{IR} is the mean infrared opacity, and γ is the ratio between the optical and infrared opacities. All of these parameters are freely retrieved, rather than being derived from the opacities in each atmospheric layer.

This provides a simple but flexible model for the P-T profile, and is the model used to generate the injected spectrum. By setting the irradiation temperature to low values the Guillot model can reproduce the general shape of typical directly imaged planet temperature profiles. Setting the equilibrium temperature T_{equ} to zero provides the limiting case of the Eddington profile (Eddington 1930). Together with the planet radius, $\log g$, and the chemical abundances this model uses a total of 8 parameters. The only included sources of line opacities are H₂O from the ExoMol data base (Chubb et al. 2021; Tennyson & Yurchenko 2012) and CO from HITEMP (Rothman et al. 2010). Both Rayleigh scattering in an H₂ and He dominated atmosphere and collisionally induced absorption between H₂-H₂ and H₂-He are included as continuum opacity sources. The priors used for all parameters in the retrieval are presented in Table 4.

Retrieval setup. `pyMultiNest` is used to generate samples and determine both the posterior parameter distributions and the Bayesian evidence of the retrieval (Buchner et al. 2014). This is a Python wrapper for the `MultiNest` sampler and likelihood integration method of Feroz & Hobson (2008). For all of the retrievals we use 4000 live points to thoroughly explore the parameter space, and a sampling efficiency of 0.8, as recommended

Table 5. Summary of retrievals run on synthetic data.

Dataset	χ^2/n_{data}	$\log \mathcal{L}$	BIC	d_m
KLIP, C	0.69	2692	-5289	2.35
KLIP, $\text{diag}(\mathbf{C})$	0.41	2656	-5218	3.29
PynPoint, C	1.72	2646	-5220	5.63
PynPoint, $\text{diag}(\mathbf{C})$	1.43	2641	-5208	9.60
Andromeda, $\text{diag}(\mathbf{C})$	0.39	2650	-5266	4.62
Gaussian, $\text{diag}(\mathbf{C})$	1.06	2651	-5267	2.21
Noise Free, $\text{diag}(\mathbf{C})$	0.06	2686	-5336	1.69

Notes. We compare best fit reduced χ^2/n_{data} , the (negative) log likelihood which includes the covariance weighting term of Eq. (9), the Bayesian Information Criterion of Eq. (10) and the Mahalanobis d_M from Eq. (11). Retrievals were performed on data processed with each algorithm, both using the full and diagonal only terms of the covariance matrix. Toy models using univariate Gaussian scatter about the input and no scatter are also included, with the uncertainties defined as $\text{diag}(\mathbf{C})$ from the KLIP data.

in the `pyMultiNest` documentation for parameter estimation. We compute negative log likelihood, the value of which is minimised in order to find the best-fit set of parameters. Across many samples, this provides a measurement of the posterior probability distribution of model parameters given the data. Under the assumption of Gaussian distributed errors, the log likelihood function takes the form of a simple χ^2 likelihood distribution. Using the covariance matrix \mathbf{C} of the data from Sect. 3.2 with elements C_{ij} , we compute the log likelihood function $\log \mathcal{L}$, which is the log-probability of measuring the observed spectrum \mathbf{S} given a forward model \mathbf{F} . A normalisation term is included which allows for a varying covariance matrix or uncertainty for each dataset and penalises samples with higher uncertainties. Thus our likelihood function is computed as:

$$-2 \log \mathcal{L} = (\mathbf{S} - \mathbf{F})^T \mathbf{C}^{-1} (\mathbf{S} - \mathbf{F}) + \log (2\pi \det(\mathbf{C})). \quad (9)$$

Goodness-of-fit metrics. In a retrieval on real exoplanet data without a ground truth value to compare to, we must turn to different metrics in order to determine which retrieval best describes the underlying spectrum. Table 5 lists χ^2 values of each best fit, as well as the minimum negative log likelihood as computed in Eq. (9) and the Bayesian Information Criterion (BIC, [Wit et al. 2012](#)). Typically when performing model comparison in a Bayesian framework, we would turn to the Bayes factor in order to reject the null hypothesis. However, when comparing the impact of different data reductions on the retrieval outcomes, the Bayes factor as computed through the nested sampling evidence estimate is insufficient, as not all of the free parameters are included in the sampling process or in the prior volume, namely the those related to the post-processing algorithms. This would bias the evidence estimate, which depends on the choice of priors and thus the overall prior volume. A full treatment would require marginalising over these algorithmic parameters, and computing a forward model of the planet signal in the IFS data. At the present time, such a joint approach is computationally infeasible. [Ruffio et al. \(2019\)](#) and [Wilcomb et al. \(2020\)](#) demonstrate that this is possible given a linear model of the starlight, the planet signal and the residuals, which can be optimised and analytically marginalised over to determine posterior distributions. However, this approach loses information on the continuum shape of the spectrum, and relies on moderate-to-high spectral resolution to infer physical quantities. The atmospheric model is also not computed on the fly,

and instead relies on a precomputed grid, limiting the parameter space available for exploration.

Therefore, instead of the Bayes factor, we rely on the BIC as a summary statistic:

$$\text{BIC} = k \log n - 2 \log \mathcal{L}_{\text{max}} \quad (10)$$

for k free parameters and n data points. This formulation allows us to account for the free parameters of the atmospheric model, as well as the parameters of the data processing, where we add one parameter for each principal component used during PSF subtraction. Unlike the Bayes factor, the BIC is only a heuristic for model comparison, and differences in the BIC cannot be treated as a metric for statistical significance. Nevertheless, models with a lower BIC can be considered more strongly favoured. As the BIC depends on the likelihood, this also means we cannot directly compare retrievals which include or neglect the off-diagonal terms of the covariance matrix. Bayes factors and the BIC estimate whether a certain forward model is favoured when compared to another one, whereas turning covariance on or off corresponds to changing the functional form of the likelihood function. It is, therefore, not a question of forward model selection. Thus no single summary statistic can determine the overall goodness-of-fit of the retrieval.

The χ^2/n_{data} statistic is useful for understanding the impact of varying the covariance and quantifying the similarity of the model to the spectrum, while the BIC is useful for heuristically evaluating the goodness-of-fit, accounting for possible over-fitting from the addition of extra parameters. In general however, we cannot directly compare the likelihood or the BIC when comparing the cases including the covariance to those using only the diagonal of the matrix with the usual motivation of model selection. Adding or neglecting the covariance does not correspond to a different forward model choice, instead it is equivalent to using a correct or incorrect functional form of the likelihood function. Therefore, assuming that the covariance is correctly measured, it is always better to include the full matrix in the likelihood in order to make statistically robust statements about the data. Thus even though the reduced χ^2 of the covariance case may be larger than that with only the diagonal, it still provides a more honest analysis of the data. It is also not surprising if the χ^2 increases if the covariance is added, since a Gaussian distribution defined by a covariance matrix with non-zero off-diagonal elements will always have a higher information content (e.g. [Rodgers 2000](#)). As discussed in Sect. 3.4, including the covariance may either increase or decrease the width of the parameter posteriors.

For comparing retrieval results that include or neglect the covariance matrix we make use of the Mahalanobis distance d_M ([Mahalanobis 1936](#)), to quantify the absolute distance between the posterior probability distributions $P(\theta | \mathbf{x})$ with means $\boldsymbol{\mu}$ and covariance S , and the true parameter values $\hat{\boldsymbol{\theta}}$:

$$d_M(\hat{\boldsymbol{\theta}}, P(\theta | \mathbf{x})) = \sqrt{(\hat{\boldsymbol{\theta}} - \boldsymbol{\mu})^T S^{-1} (\hat{\boldsymbol{\theta}} - \boldsymbol{\mu})}. \quad (11)$$

This provides a metric for the overall accuracy of the retrieval when the true input parameters are known.

Finally, for this work we use the median parameter values and associated spectra as our point of comparison as opposed to the maximum likelihood fit. We find that although the spectrum generated by the median parameter values is a worse fit (by definition) than the best fit spectrum, the median parameters are a more accurate measurement of the input parameters.

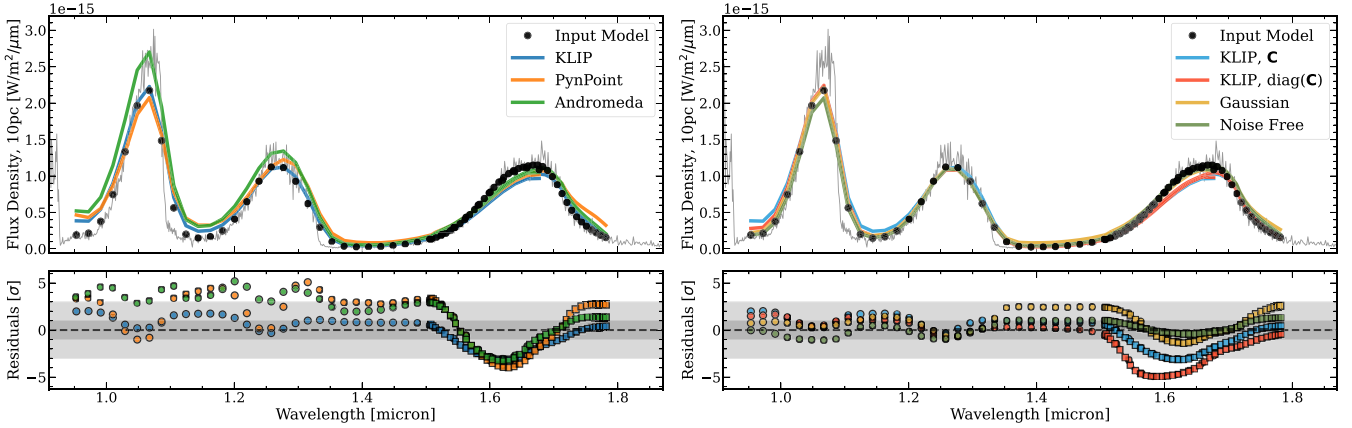


Fig. 8. Best-fit models compared to the ground truth spectrum, in order to determine how well the retrieval can account for the systematics introduced through the data analysis. The residuals are calculated by sampling the posterior distributions to generate spectra, and taking the standard deviation at each wavelength. Left: the best fit from the retrievals for each KLIP (blue), PynPoint (orange) and ANDROMEDA (green). Right: the best fits of the KLIP retrievals with (blue) and without (yellow) including the covariance matrix, as well as the retrieval on the Gaussian noise (red) and Noise-free datasets (green). These are again compared to the true input spectrum.

5.1. Outline of retrievals

We performed three main tests to answer address the central theme of this paper:

1. Comparing retrievals on spectra extracted using different post processing algorithms.
2. Comparing retrievals that either include or ignore the covariance matrix in the likelihood function.
3. Testing if a lack of correlation information can be accounted for using additional ad-hoc data-processing parameters in the retrieval.

Our primary retrieval results are summarised in Table 5. In Sect. 5.2 we compare the cases of data that has been processed with KLIP, PynPoint and ANDROMEDA, both with and without the use of the covariance matrix from Sect. 3.2. As a benchmark, we also include a retrieval using the nominal input spectrum, perturbed with draws from a Gaussian distribution, where the covariance is given by the diagonal of the KLIP covariance matrix. This represents how the data would appear without systematics from HCI data processing and without the correlations introduced by the instrument optics. We also include a retrieval using the same uncertainties as in the Gaussian case, but without scatter about the input spectrum to validate our retrieval method and choice of goodness-of-fit metrics. For the sake of brevity, we refer to these as the ‘Gaussian’ and ‘noise-free’ cases respectively. We explore the impact of incorporating the covariance matrix in the retrieval framework in Sect. 5.3, using the KLIP, Gaussian and noise-free cases. Section 5.5 explores whether we can account for ignorance of the covariance in the data by introducing scaling factors and offsets in the retrieval.

Validation. To verify the validity of our results, we also ran a series of validation retrievals to test the sensitivity of our results to the choice of datasets, priors, and models. We ran retrievals on each dataset independently, as well as with broad and tight priors. Neither dataset was able to retrieve the parameters as precisely as the combined retrievals. The posterior distributions and fits were insensitive to our choice of priors. We ran additional retrievals using a spline temperature profile, as a proxy for our model not truly matching the underlying data. With 5 spline nodes, we were able to retrieve the $\log g$, R_{pl} , T_{int} and the water mass fraction to the same precision and accuracy as using the Guillot profile

used to generate the data, thus concluding that the retrievals are flexible enough to account for some degree of imperfect model assumptions.

5.2. Impact of algorithm selection on retrievals

The largest variation in the extracted spectrum is due to the choice of post-processing algorithm, so our first aim is to explore how these differences in the data lead to differences in the inferred parameters. In Fig. 8, we compare the best fit results from each of the retrievals run on data processed using KLIP, PynPoint and ANDROMEDA to the ground truth spectrum injected into the IFS data. All three processing tools provide reasonable fits to the input spectrum, and share trends in the shape of their residuals, though KLIP provides the overall best reproduction of the input spectrum. The retrieved spectra tend to fit the input better at higher flux values, where the S/N is greater. Figure 9 shows the posterior distributions of most parameters; regardless of the retrievals setup the absolute uncertainties on all of the retrieved parameters are large. This highlights the importance having high S/N inputs to obtain precise constraints, as well as broad wavelength coverage to have sensitivity to a wide range of parameters. Not included in the plot are T_{equ} , γ and the CO mass fraction, none of which are constrained in any retrieval. For widely separated planets, T_{equ} is small and has little impact on the shape of the pressure temperature profile. As such, the Guillot profile is effectively reduced to the Eddington term, which does not depend on γ . Finally, there are no strong CO features present in the wavelength range considered in the injected planets. Thus we do not expect any of these parameters to impact the spectrum enough to be constrained by this retrieval. We neglect these unconstrained parameters when calculating the distance d_M .

KLIP performs the best of the three algorithms; accurately fitting the spectrum and retrieving the input parameters, measured from the χ^2/n_{data} and d_M respectively, as presented in Table 5. PynPoint performs somewhat worse again, though the GPI data suffers from two outlying data points, and the measured uncertainties are generally smaller than for the KLIP or ANDROMEDA extractions. The physical interpretation of the PynPoint is significantly different than that of KLIP or ANDROMEDA: the inferred mass from the median $\log g$ and planet radius is more than a factor of 10 smaller when using the

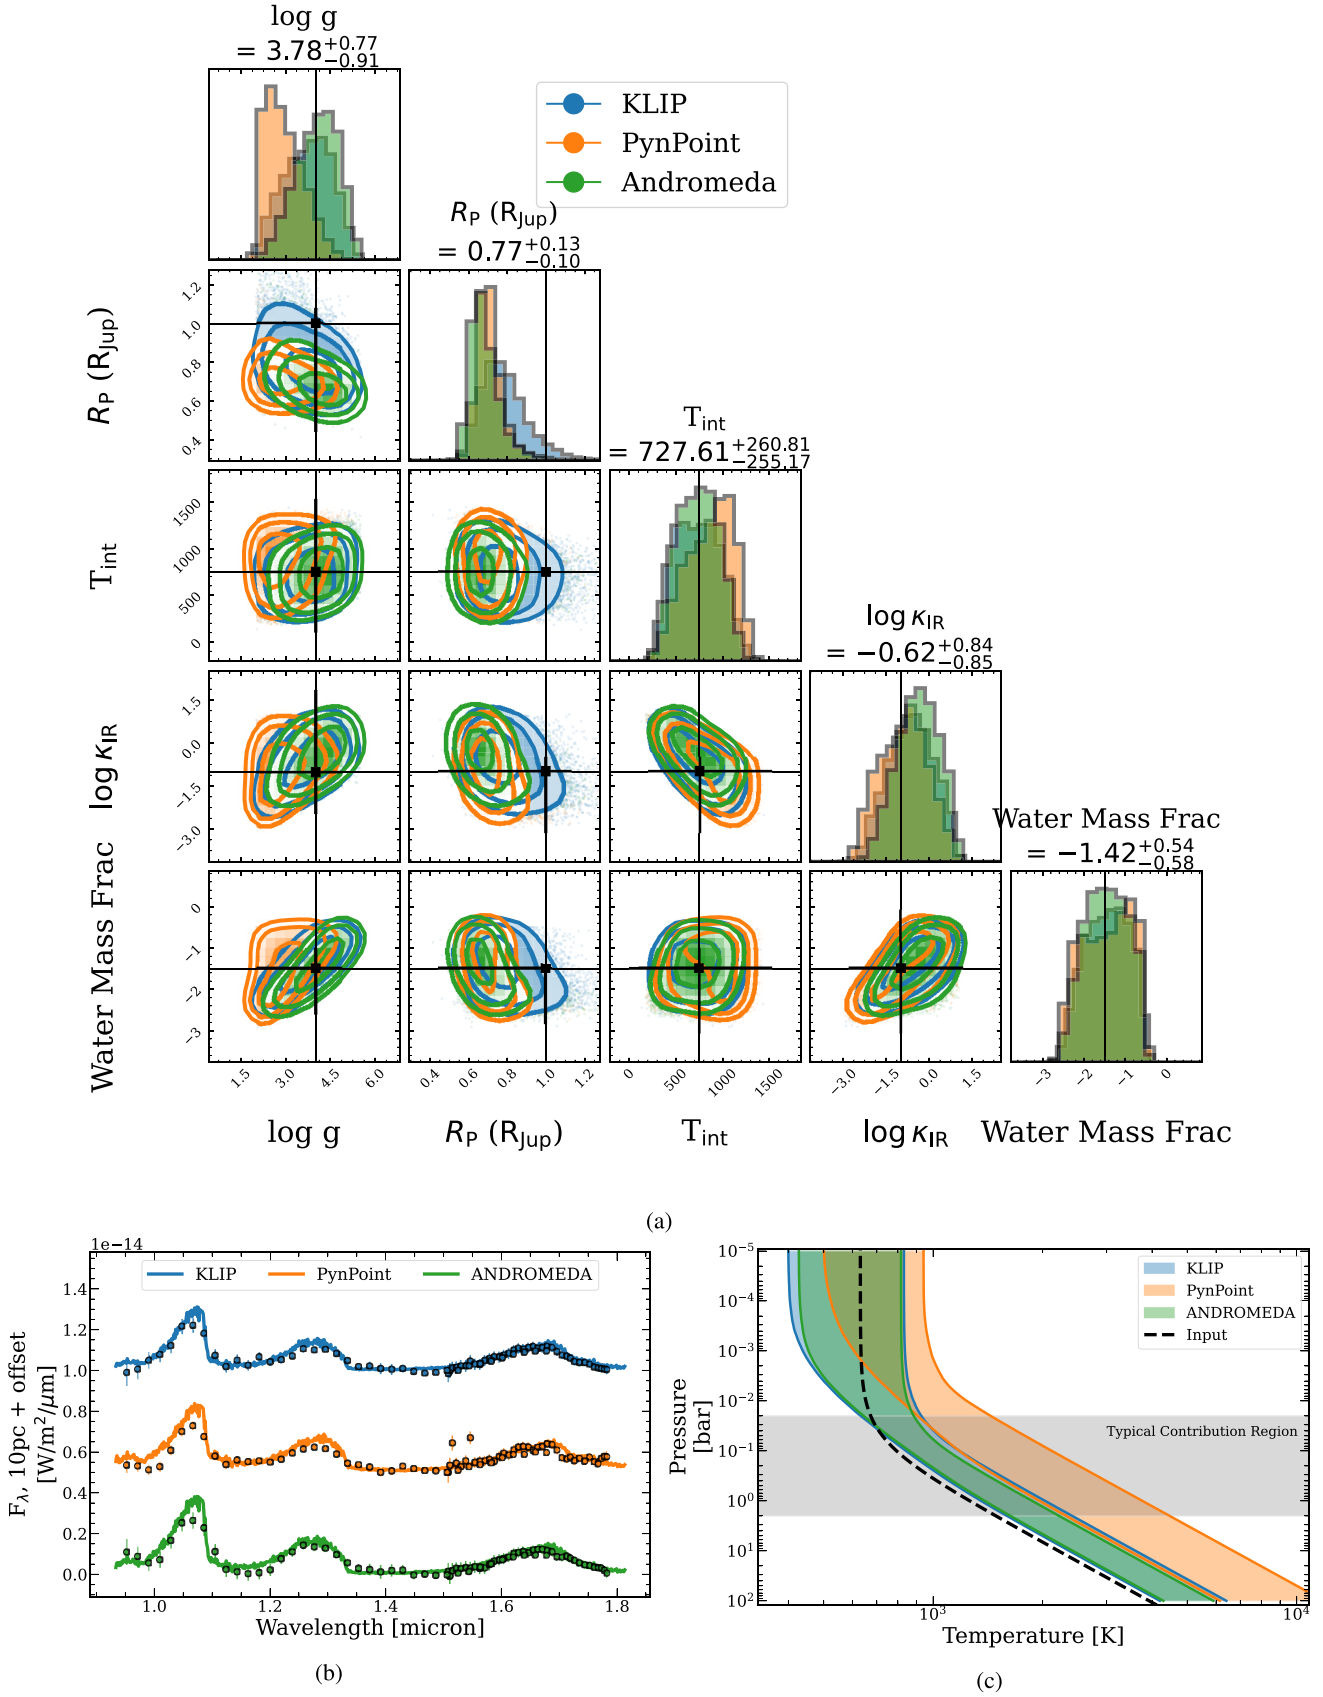


Fig. 9. Results of retrievals comparing each of the three data processing algorithms. (a) Posterior distributions for retrievals from each data processing algorithm. Contours are plotted for 2D Gaussian 1, 2 and 3 σ levels, corresponding to 36%, 86% and 99% confidence intervals, and the ground truth value is marked in black. The text labels correspond to the KLIP retrieval. (b) Median-fit spectrum from retrievals on each algorithm, with the covariance (other than for ANDROMEDA). (c) P - T Profiles for each retrieval. The shaded region indicates the 68% confidence region for the retrieved profile.

PynPoint parameters. This highlights the need for feedback between modelling and data analysis, as well as for comparison both between different data analyses and different models. This strongly impacts the measurement of $\log g$, which is sensitive to the shape of the H -band. Finally ANDROMEDA fails to reproduce the input spectrum well, but recovers the input parameters more accurately than PynPoint, though confidently excluding the true planet radius.

We find that for all of the retrieval setups, the median parameter values provide a better estimation of the true input parameters than the single maximum likelihood fit. For KLIP, we find that the median internal temperature estimate of 724 ± 260 K accurately, if imprecisely measure the true value of 750 K. However, the best fit value of 498 K is strongly biased from the true value, as are the remaining parameters. We therefore continue using only the median parameter estimates, rather than the maximum likelihood fits.

The goodness-of-fit metrics provide a mechanism to select between the different retrievals. We find that all of the metrics favour the KLIP retrieval, though noting that the reduced χ^2 for ANDROMEDA is smaller, but does not account for the covariance. Based on the variation in the BIC and the Mahalanobis distance, the effect that the algorithm choice has on the retrieved parameters is significant. Interpreting the Mahalanobis distance as standard deviations from the truth, ANDROMEDA and PynPoint are 2.3 and 3.3 standard deviations less accurate than KLIP respectively. The trend of the BIC follows that of d_M , favouring KLIP, followed by Andromeda and then PynPoint. As the ground truth is not generally known, this reinforces the use of the BIC or a similar metric (such as the Bayes factor) as a robust metric for selecting between models, even when the data is also varied.

None of the retrievals retrieve the true input pressure temperature profile to within 1σ , as is evident from Fig. 9c. However, in the region where the emission contribution is located, the retrieved PT profiles share a similar slope to the true input, at slightly higher temperatures. Such discrepancies highlight the importance of broad wavelength coverage in atmospheric retrievals, where the spectrum can probe different pressure, and thus temperature, layers of the atmosphere.

Effects of principal component optimisation. We also compare how the number of principal components used in the data processing impacts the retrieval results. This effect is more apparent at lower S/N so for this particular case we choose an injection at 400 mas, extracting the spectra with KLIP. Figure 10 highlights the impact that the choice of the number of PCs has on the precision of the posterior distributions. We compared the optimal extraction (6PCs for GPI and 8 for SPHERE) to an extraction using 25 PCs for each dataset. While both extractions retrieve the input parameters with similar accuracy, the optimised extraction is significantly more precise in its measurement of the planet radius.

5.3. Impacts of including the covariance in retrievals

While the choice of algorithm produces most of the difference in the spectral shape, we also consider how including the covariance into the log-likelihood calculation of the retrieval impacts the retrieved spectrum and inferred parameters. To understand this, we compare a KLIP retrieval with and without the use of the covariance matrices, and how these results compare to the Gaussian and noise-free cases.

The right panel of Fig. 8 shows each of these datasets compared to the true input model, while Fig. 11 shows the posterior

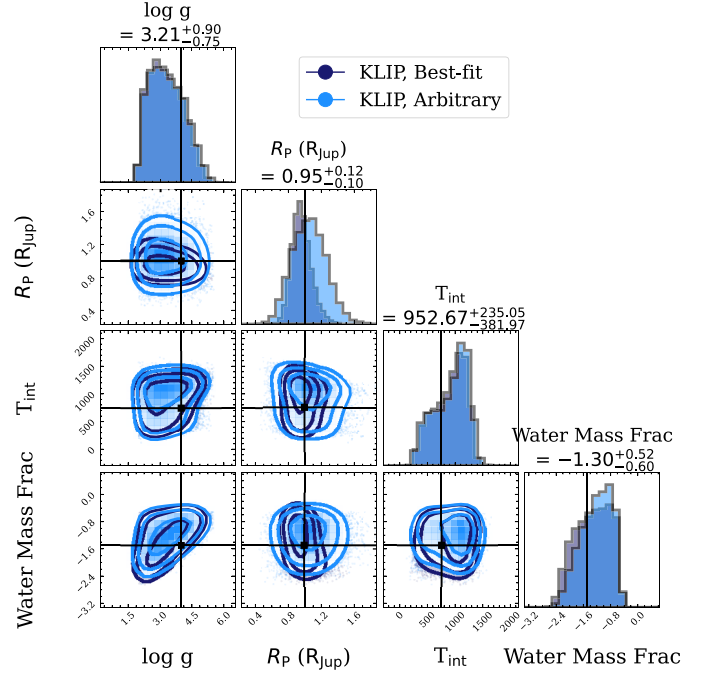


Fig. 10. Corner plot comparing the retrieved parameter distributions for two different KLIP reductions. In light blue, the input spectrum was optimised using the relative discrepancy metric (Fig. 5), while in dark blue an arbitrary extraction was chosen for each of the SPHERE and GPI datasets, reflecting a non-optimal parameter selection.

distributions, best fits and PT profiles. We find that our retrievals reproduce the results of Greco & Brandt (2016): incorporating the covariance matrix improves the accuracy of the retrieval, at the cost of lower precision. Including the covariance matrix reduced the posterior bias for all parameters. However, the $\text{diag}(C)$ case was able to more accurately retrieve the input PT profile, even though the parameter distributions were more discrepant from the true input parameters.

Quantitatively, including the covariance improved the d_M by 1 compared to the $\text{diag}(C)$, approaching the Gaussian measurement of 2.21. As in Sect. 3.4, the χ^2/n_{data} is underestimated. This is also reflected by the BIC, which favours the inclusion of the covariance matrix over both the $\text{diag}(C)$ case and the Gaussian case. In the KLIP extraction of the GPI data, there is rather broad covariance, which is easier to fit, as demonstrated by the toy model of 3.4.

Within the pressure range probed, all of the retrievals measure similar slopes, and correct temperatures to within 2σ , with the noise-free and $\text{diag}(C)$ cases most accurately retrieving the input profile. The correlation between atmospheric parameters of T_{int} and κ_{IR} in Fig. 11a shows the difficulty in inferring atmospheric properties, and explains the inability of the retrievals to perfectly infer the temperature structure. For the KLIP data, the inclusion of the covariance matrix in the log-likelihood improves the accuracy of the constraint on both parameters.

We repeated this experiment using the PynPoint extractions. Consistent with the KLIP results, we find that including the covariance improves the accuracy of the retrieval, at a marginal cost to the precision of the retrieved parameters.

5.4. Relation between measurement and posterior precision

With the nominal uncertainties, the noise-free case can reproduce the input data to within 1σ across the entire wavelength range.

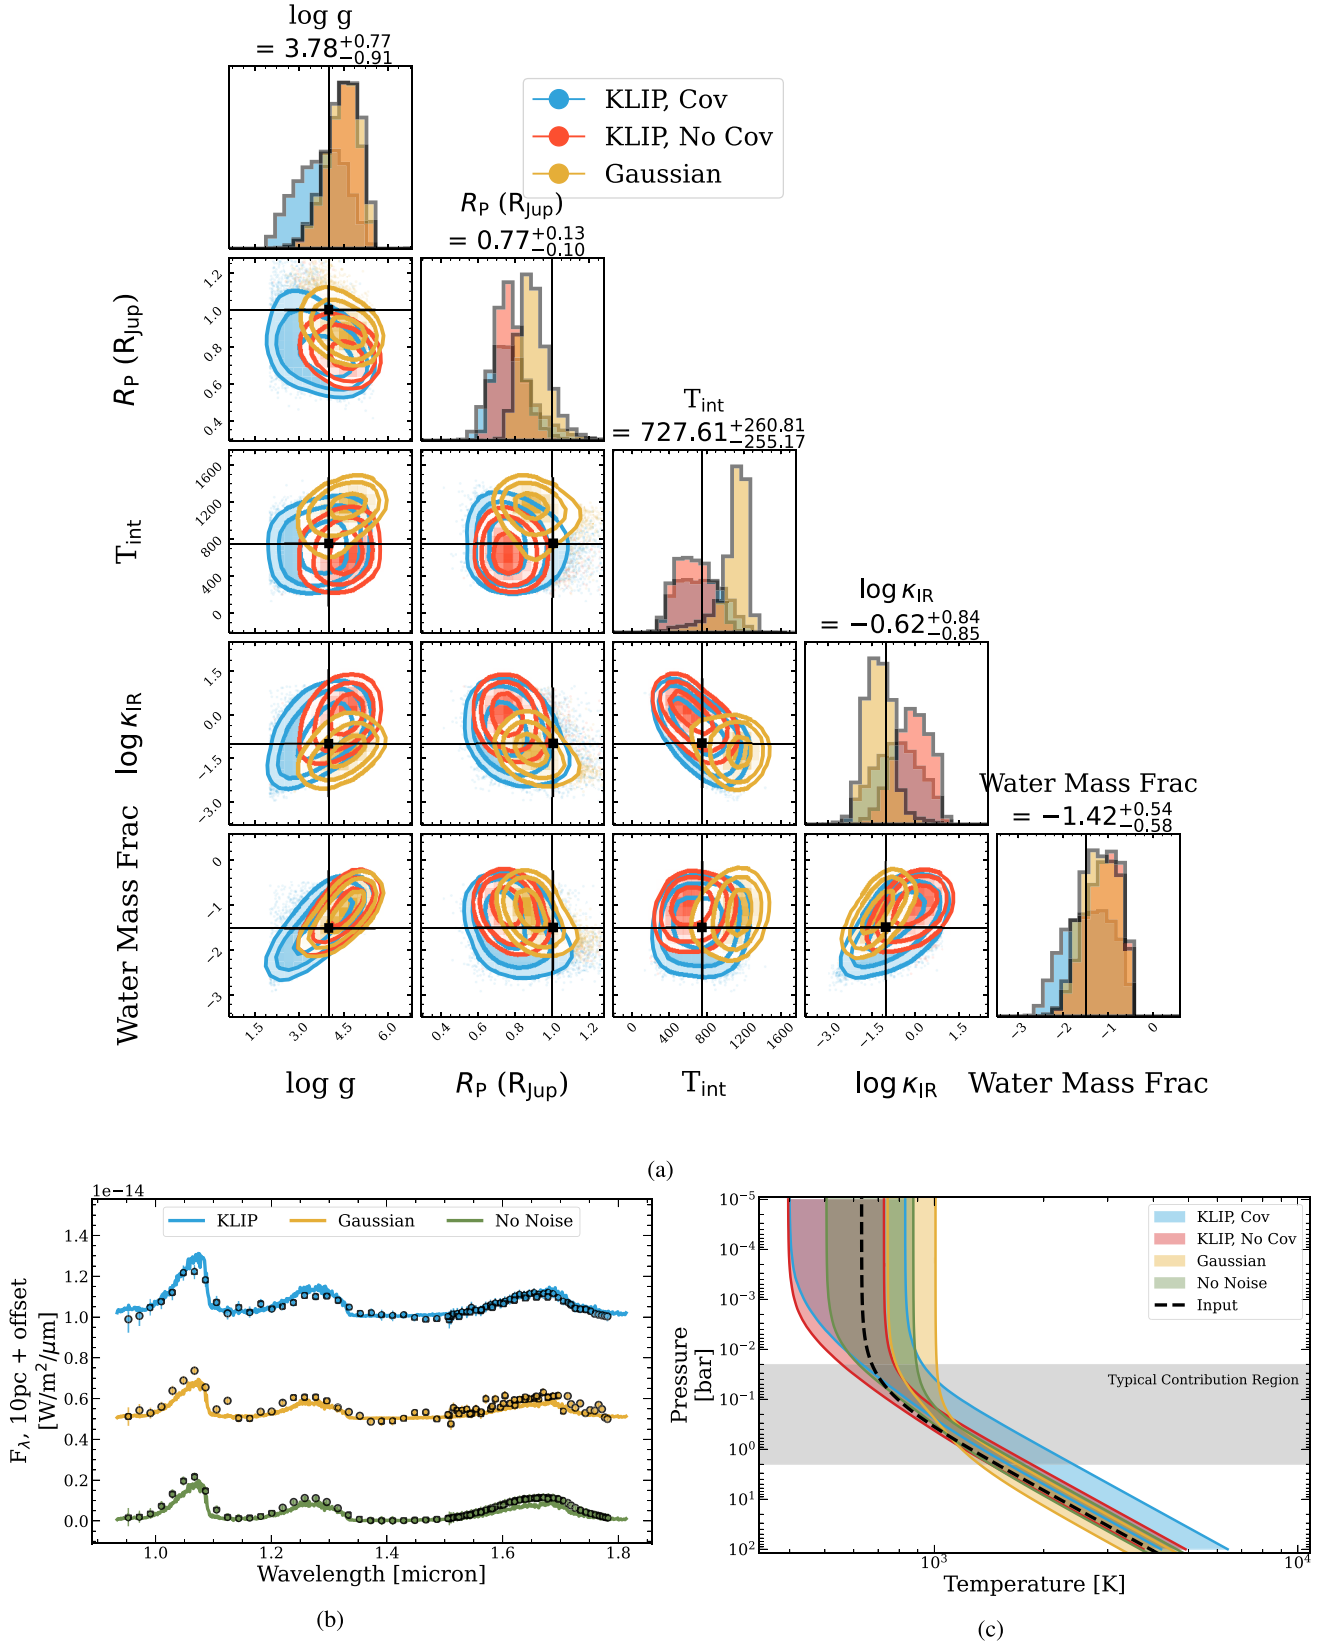


Fig. 11. Results of retrievals using KLIP, comparing each the cases of computing the likelihood using the full covariance matrix (blue), the diagonal elements only (yellow) and using truly Gaussian scattered data (red). (a) Posterior distributions for retrievals from KLIP. Contours are plotted for 2D Gaussian 1, 2 and 3σ levels, corresponding to 36%, 86% and 99% volume regions, and the ground truth value is marked in black. The text labels correspond to the retrieval including the covariance matrix. (b) Median spectrum for the KLIP retrieval with covariance, as well as for the Gaussian (yellow) and noise free cases (green). (c) P - T profiles for KLIP, Gaussian, and noise-free retrievals. The shaded region indicates the 68% confidence region for the retrieved profile.

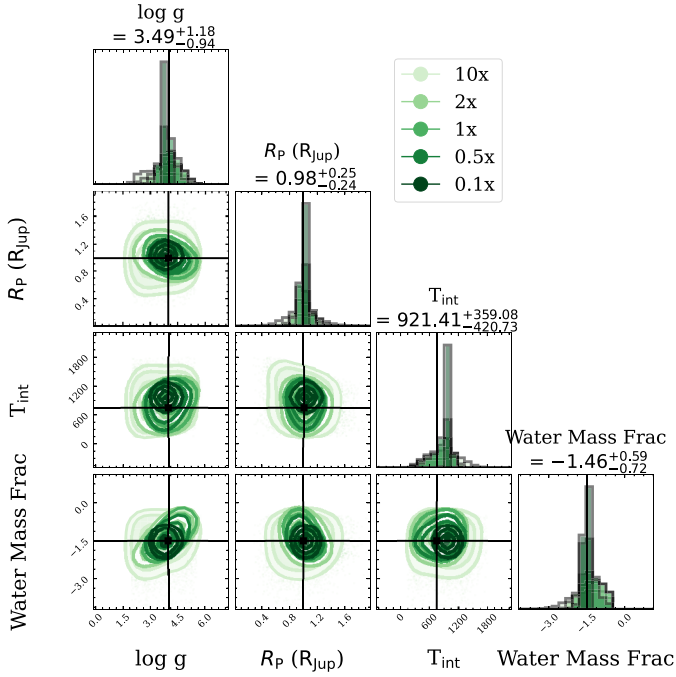


Fig. 12. Posterior distributions for retrievals in the ‘Noise-free’ case. The nominal uncertainties are from $\text{diag}(\mathbf{C})$ of the KLIP extraction, and were scaled by factors from 0.1 to 10. The titles list the uncertainties for the 10x case.

However, at the level of precision of these measurements none of the retrievals are able to put strong constraints on any of the measured parameters, though this would be improved with more precise spectroscopic measurements. Using the noise-free case, we explored how the measurement precision affects the posterior precision by scaling the uncertainties by factors from 0.1 to 10, shown in Fig. 12.

In all cases, the parameters are accurately retrieved, and the posterior precision increases as the uncertainties decrease. In the nominal case, the mean S/N per channel is 5 in the SPHERE wavelength range and 3 in the GPI-*H* band. With this precision and wavelength coverage, even in the optimistic case of no scatter in the data the internal temperature can only be constrained to within ± 223 K. This improves to ± 26 K if the S/N is improved by a factor of 10: while low S/N may be sufficient for accurate retrievals, high S/N is required to precisely measure the physical parameters. This finding complements Fig. 10, highlighting that the main impact of the principal component optimisation is to improve the precision of the posterior distributions, as the choice of PCs impacts the precision of the spectroscopic measurement. Our choice of metrics are also validated, as the noise-free case is favoured by every metric (when comparing to diagonal only cases).

5.5. Mitigating covariance impacts with nuisance parameters

One challenge to the interpretation is the suggestion of overfitting by the reduced χ^2 values. A $\chi^2/n_{\text{data}} < 1$ suggests overfitting, although it can also be interpreted as overestimated uncertainties, or underestimated correlation (which effectively translates into overestimated uncertainties as well). The retrievals both with and without covariance on the KLIP dataset both have $\chi^2/n_{\text{data}} < 1$. As the number of parameters is much lower than the number of data points and is identical to the ground truth model, this suggests that the uncertainty of the KLIP data is overestimated.

Table 6. Summary statistics for KLIP retrievals, including retrieved parameters to account for systematic biases.

Dataset	χ^2/n_{data}	$\log \mathcal{L}$	$\log_{10} \mathcal{Z}$	BIC	d_M
KLIP, C	0.69	2692	1166	-5289	2.35
KLIP, $\text{diag}(\mathbf{C})$	0.41	2656	1149	-5218	3.29
Offset GPI	0.42	2657	1147	-5214	3.29
Offset SPH	0.36	2659	1148	-5218	3.40
Scale GPI	0.51	2666	1152	-5233	4.48
Scale SPH	0.56	2659	1149	-5218	3.93
Scale GPI Err.	1.07	2684	1158	-5268	4.52
Scale SPH Err.	0.70	2660	1150	-5220	4.73
Scale Both.	1.15	2684	1158	-5263	4.54
10^b Both.	0.43	2567	1148	-5210	3.41
GPI Only, C	0.51	1320	571	-2596	0.62
SPH Only, C	1.00	1373	593	-2698	2.75

Notes. $\log_{10} \mathcal{Z}$ is the Bayesian evidence, the difference of which is the Bayes factor between two models. ‘Scale’ indicates a multiplicative factor applied to the specified dataset, while ‘Offset’ indicates an additive term.

In this section we explore the use of various parameterisations to account for systematics in the data, such as including offsets or scaling factors in the retrievals. As all of these comparisons use the same data, and additional parameters are properly included in the prior volume, we can now use the Bayes factor from Table 6 to quantitatively select the best model. We refer to Table 2 of Benneke & Seager (2013) for our interpretation of the Bayes factor: $\log \Delta \mathcal{Z}_{H_2, H_1} > 10$ is strong evidence in favour of model H_2 over H_1 .

Beginning with offsets, we fix one dataset and allow the other dataset to float, with a uniform prior of $\mathcal{U}(-10^{-14} \text{W m}^2/\mu\text{m}, 10^{-14} \text{W m}^2/\mu\text{m})$. We find that while allowing for offsets may marginally improve the fit to the data ($\chi^2_{\text{SPH, offset}}/n_{\text{data}} < \chi^2_{\text{KLIP, diag}(\mathbf{C})}/n_{\text{data}}$), it does not improve the accuracy or precision of the posteriors ($d_M > d_{M, \text{KLIP, diag}(\mathbf{C})}$), and is not favoured by the Bayes factor or BIC.

Next, we multiply one dataset and its corresponding uncertainties by a scaling factor ($\mathcal{U}(0.5, 2.0)$), fixing the remaining dataset. We find that a scaling factor of 0.73 ± 0.05 for the GPI dataset is somewhat favoured by the Bayes factor ($\Delta \log_{10} \mathcal{Z} = 3$), though the posterior precision and accuracy is somewhat reduced. Scaling the SPHERE data did not significantly improve the fit, and is not favoured by the Bayes factor.

This result is emphasised when we scale only the uncertainties for each of the datasets. We ran three retrievals: scaling the uncertainties of each dataset and fixing the other, or scaling both datasets with independent scaling factors simultaneously. To avoid hitting the prior boundaries, the scaling factor is given a uniform prior of $\mathcal{U}(0.05, 2.0)$. We find that a scaling factor of 0.28 ± 0.04 is favoured ($\Delta \log_{10} \mathcal{Z} = 9$) for the GPI dataset in both retrievals where the uncertainties are allowed to float, while allowing the SPHERE uncertainties to float does not change the fit. This implies that the either the KLIP uncertainties for the GPI dataset are underestimated, or that the correlation is not correctly accounted for, as described in Sect. 3.4. Allowing the uncertainties to float improves the fit compared to the retrieval using only the diagonal components of the covariance matrix, but is still disfavoured compared to the retrieval using the full matrix.

Line et al. (2015) introduces a different parameterisation to scale the uncertainties and reduce overfitting. Using a parameter

b , the uncertainty on the i^{th} wavelength bin is inflated as:

$$s_i^2 = \sigma_i^2 + 10^b, \quad (12)$$

where σ_i is the uncertainty on that bin. This will only allow for an increase in the size of the error bars, and allows us to account for model uncertainties and missing physics, rolling the additional uncertainty into the marginalised posterior parameter distributions. The prior range on b is set from -36 to -26 : this encompasses the suggested range from [Line et al. \(2015\)](#) such that $0.01 \times \min(\sigma^2) < 10^b < 100 \times \max(\sigma^2)$. We retrieve b independently for each dataset used in the retrieval. Using this formalism, we find that the b parameterised retrieval is disfavoured compared to the KLIP retrievals, both with and without the use of the covariance matrix. Likewise, the d_M measured for this case (3.41) suggests a marginal decrease in accuracy relative to the baseline retrievals. This suggests that the extra parameters used to fit the b parameter are not justified to help resolve the problem of overfitting of the spectrum. As this formalism can only inflate the uncertainties, it is unsurprising that it cannot correct for the overestimated GPI uncertainties as shown by the scaling factor retrievals and the small reduced χ^2 values.

We conclude that scaling factors and offsets are inadequate for accounting for systematic offsets in the data due to the data processing and correlated noise. While allowing the uncertainties to float was marginally favoured by the Bayes factor compared to the $\text{diag}(\mathbf{C})$ KLIP retrieval, it was still strongly disfavoured compared to the retrieval using the full covariance matrix. Alternative methods, such as Gaussian process regression ([Wang et al. 2021](#); [Xuan et al. 2022](#)) may be able to overcome these limitations and allow for the characterisation of systematics in a Bayesian framework.

5.6. Limitations

This work reflects many of the best practices used in both data analysis and atmospheric retrievals, but remains an optimistic assessment of our ability to infer both accurate and precise physical parameters. Additional sources of bias are inevitably present in the data, such as the differences in spectra arising from different reduction pipelines as shown in [Appendix A](#) and from the process of building the 3D cubes from the 2D detector frames. Exoplanet data relies heavily on precise photometry, yet the host star which is used as a calibration source is obscured behind the coronagraph during the observations, making temporal monitoring of the PSF challenging. Finally, we use the same model for both injection and as the basis of the retrieval, ultimately ignoring many key physical processes present in real exoplanets. Even with these limitations, we remain optimistic about the prospects for retrievals to characterise directly imaged exoplanets, particularly in the era of high precision, broad wavelength spectroscopy as enabled by VLT/GRAVITY, JWST, and the ELTs.

6. Summary and conclusions

Based on our comparison of high contrast imaging algorithms, it is clear that systematic variations are a more significant contribution to uncertainty than random errors for directly imaged exoplanet spectra. Such variations often lead to spectral correlation of the data, and knowledge of the length scale and strength of said correlation is crucial to accurate interpretation of the data. We used the methods of [Greco & Brandt \(2016\)](#) to compute covariance matrices for IFS data, and demonstrated that correlations in the data can both increase or decrease the posterior

width of model parameters, depending on whether the parameter is sensitive to wavelength scales greater or less than the correlation scale. Using injection testing, we optimise our choice of algorithm parameters. We find that using only the S/N as a metric to determine the quality of spectral extraction does not produce optimal extractions. Instead, data processing parameters should be tuned using injection testing, with careful consideration of what goodness-of-fit metric should be used. Using the mean relative discrepancy, we optimised the number of principal components used in PSF subtraction in order to optimally extract the companion spectrum. Of course the number of principal components used in PSF subtraction is not the only source of systematic biases during spectral extraction: the precision of the astrometric solution, choices in processing both the science frames and the unsaturated PSF frames and the details of parameter choice all introduce biases on a similar level to the number of principal components, and must be independently optimised.

Each algorithm considered performed best under different conditions: the contrast, separation, observing conditions and data volume all impact which algorithm produces the optimal extraction. Care must be taken as the parameter choices that lead to the most sensitivity in order to detect companions are often different than the parameters required to robustly extract the spectrum in order to characterise the planets. During independent comparisons such choices led to statistically significant differences in both the shape and overall flux calibration of the extracted spectra. Without a priori knowledge of the spectrum, it is therefore necessary to compare multiple independent measurements in order to determine the underlying spectral shape.

By performing atmospheric retrievals on data processed using different algorithms we show that the variation between different data reductions is larger than the statistical posterior uncertainty. Model choice is highly dependent on data quality and quantity, and Bayesian comparisons should be performed to determine whether model complexity is suitable given the data. The ideal solution is to fully understand and correct for systematics during the data processing, with broad data coverage and high spectral resolution. When comparing models, statistical tools such as the BIC or the Bayes factor should be used with care, and only when the free parameters are fully incorporated into the retrieval process to account for their impact on the prior volume and posterior distribution. In such a Bayesian framework, we find that the median parameter values are accurate measurements of the true input parameters, but that the maximum likelihood values are often strongly biased. Even using the median values, the difference in retrieved parameters from different data processing tools is significant, and can lead to dramatically different astrophysical interpretations. Retrievals should be performed on multiple data reductions to ensure that the retrieved parameters are robust to such variation.

We used the Mahalanobis distance, d_M to measure the distance between the true input parameters and the posterior distributions. Using this metric, we found that accounting for the covariance in the likelihood function of a retrieval framework can help mitigate correlations in the data, but not entirely resolve them. Compared to using only the diagonal terms using the full matrix will reduce the bias in all parameters, at the cost of slightly decreased precision, reproducing the results of [Greco & Brandt \(2016\)](#). When the data are correlated including the covariance is necessary to make meaningful statistical statements about models fitting the data. In all cases, including the full covariance matrix leads to improved accuracy of the inferred planet parameters. When testing the use of scaling factors and offsets to try to correct for these systematic biases we found worse results than when

relying on the covariance matrix. Thus we recommend that the covariance matrix always be published along with IFS data of exoplanets.

The systematic biases of the spectral extractions fundamentally limit the accuracy with which we can understand exoplanet atmospheres, with effects that can be much more significant than the statistical posterior precision. One solution to the issues discussed in this work is to acquire higher quality data. Nevertheless, it will remain important to measure the covariance and to understand systematic effects imparted by data processing.

Acknowledgements. We would like to thank the anonymous referee for their insightful and detailed report, which substantially improved the quality of this manuscript. In addition we'd like to thank Alice Zurlo and Alex Greenbaum for supplying the data used in this work. Software used: petitRADTRANS, pyKLIP, PynPoint, VIP-HCI, species, pyMultiNest, Python, numpy, astropy, phot_utils, matplotlib, scicomap.

References

- Amara, A., & Quanz, S. P. 2012, *MNRAS*, **427**, 948
- Antichi, J., Dohlen, K., Gratton, R. G., et al. 2009, *ApJ*, **695**, 1042
- Barstow, J. K., Aigrain, S., Irwin, P. G. J., Kendrew, S., & Fletcher, L. N. 2015, *MNRAS*, **448**, 2546
- Benisty, M., Bae, J., Facchini, S., et al. 2021, *ApJ*, **916**, L2
- Benneke, B., & Seager, S. 2013, *ApJ*, **778**, 153
- Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2008, *SPIE Conf. Ser.*, **7014**, 701418
- Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, **631**, A155
- Brown-Sevilla, S. B., Maire, A. L., Mollière, P., et al. 2023, *A&A*, **673**, A98
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
- Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. 2015, *A&A*, **582**, A89
- Cantalloube, F., Farley, O. J. D., Milli, J., et al. 2020a, *A&A*, **638**, A98
- Cantalloube, F., Gomez-Gonzalez, C., Absil, O., et al. 2020b, *SPIE Conf. Ser.*, **11448**, 114485A
- Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2005, *A&A*, **438**, L25
- Chubb, K. L., Rocchetto, M., Yurchenko, S. N., et al. 2021, *A&A*, **646**, A21
- Cutri, R. M., Skrutskie, M. F., van Dyk, S., et al. 2003, *VizieR Online Data Catalog*: II/246
- Cutri, R. M., Wright, E. L., Conrow, T., et al. 2021, *VizieR Online Data Catalog*: II/328
- Desidera, S., Chauvin, G., Bonavita, M., et al. 2021, *A&A*, **651**, A70
- Eddington, A. S. 1930, *MNRAS*, **90**, 279
- Faramaz, V., Marino, S., Booth, M., et al. 2021, *AJ*, **161**, 271
- Feroz, F., & Hobson, M. P. 2008, *MNRAS*, **384**, 449
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
- Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2019, *Open J. Astrophys.*, **2**, 10
- Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2018, *A&A*, **618**, A138
- Gaia Collaboration (Brown, A. G. A., et al.) 2018, *A&A*, **616**, A1
- Gebhard, T. D., Bonse, M. J., Quanz, S. P., & Schölkopf, B. 2022, *A&A*, **666**, A9
- Gerard, B. L., & Marois, C. 2016, *SPIE Conf. Ser.*, **9909**, 990958
- Gomez Gonzalez, C. A., Absil, O., Absil, P. A., et al. 2016, *A&A*, **589**, A54
- Gomez Gonzalez, C. A., Absil, O., & Van Droogenbroeck, M. 2018, *A&A*, **613**, A71
- Gontcharov, G. A. 2006, *Astron. Lett.*, **32**, 759
- Goody, R., West, R., Chen, L., & Crisp, D. 1989, *J. Quant. Spec. Radiat. Transf.*, **42**, 539
- GRAVITY Collaboration (Lacour, S., et al.) 2019, *A&A*, **623**, A11
- GRAVITY Collaboration (Nowak, M., et al.) 2020, *A&A*, **633**, A110
- Gray, R. O., & Corbally, C. J. 2002, *AJ*, **124**, 989
- Gray, R. O., Corbally, C. J., Garrison, R. F., McFadden, M. T., & Robinson, P. E. 2003, *AJ*, **126**, 2048
- Greco, J. P., & Brandt, T. D. 2016, *ApJ*, **833**, 134
- Greenbaum, A. Z., Pueyo, L., Ruffio, J.-B., et al. 2018, *AJ*, **155**, 226
- Guillot, T. 2010, *A&A*, **520**, A27
- Hauschildt, P. H., Allard, F., & Baron, E. 1999, *ApJ*, **512**, 377
- Hoch, K. K. W., Konopacky, Q. M., Theissen, C. A., et al. 2023, *AJ*, **166**, 85
- Hou Yip, K., Nikolaou, N., Coronica, P., et al. 2019, in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2019*, Würzburg, Germany, September 16–20, 2019, Proceedings, Part III (Berlin, Heidelberg: Springer-Verlag), 322
- Ih, J., & Kempton, E. M. R. 2021, *AJ*, **162**, 237
- Ingraham, P., Marley, M. S., Saumon, D., et al. 2014, *ApJ*, **794**, L15
- Kiefer, S., Bohn, A. J., Quanz, S. P., Kenworthy, M., & Stolker, T. 2021, *A&A*, **652**, A33
- Konopacky, Q. M., Barman, T. S., Macintosh, B. A., & Marois, C. 2013, *Science*, **339**, 1398
- Lacis, A. A., & Oinas, V. 1991, *J. Geophys. Res.*, **96**, 9027
- Lacy, B. I., & Burrows, A. 2020, *ApJ*, **905**, 131
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, **660**, 770
- Lagrange, A. M., Bonnefoy, M., Chauvin, G., et al. 2010, *Science*, **329**, 57
- Langlois, M., Gratton, R., Lagrange, A. M., et al. 2021, *A&A*, **651**, A71
- Larkin, J. E., Chilcote, J. K., Aliado, T., et al. 2014, *SPIE Conf. Ser.*, **9147**, 91471K
- Lavie, B., Mendonça, J. M., Mordasini, C., et al. 2017, *AJ*, **154**, 91
- Leconte, J. 2021, *A&A*, **645**, A20
- Lee, J.-M., Heng, K., & Irwin, P. G. J. 2013, *ApJ*, **778**, 97
- Line, M. R., Teske, J., Burningham, B., Fortney, J. J., & Marley, M. S. 2015, *ApJ*, **807**, 183
- Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *Proc. Natl. Acad. Sci. U.S.A.*, **111**, 12661
- Mahalanobis, P. 1936, *Proc. Natl. Inst. Sci. India*, **2**, 49
- Maire, J., Gagné, J., Lafrenière, D., et al. 2012, *SPIE Conf. Ser.*, **8447**, 844760
- Marois, C., Doyon, R., Racine, R., et al. 2005, *JRASC*, **99**, 130
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, **641**, 556
- Marois, C., Lafrenière, D., Macintosh, B., & Doyon, R. 2008a, *ApJ*, **673**, 647
- Marois, C., Macintosh, B., Barman, T., et al. 2008b, *Science*, **322**, 1348
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, *Nature*, **468**, 1080
- Marois, C., Correia, C., Galicher, R., et al. 2014a, *SPIE Conf. Ser.*, **9148**, 91480U
- Marois, C., Correia, C., Véran, J.-P., & Currie, T. 2014b, in *Exploring the Formation and Evolution of Planetary Systems*, 299, eds. M. Booth, B. C. Matthews, & J. R. Graham, 48
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, **792**, 97
- Milli, J., Banas, T., Mouillet, D., et al. 2016, in *Adaptive Optics Systems V*, (SPIE), **99094Z**
- Milli, J., Kasper, M., Bourget, P., et al. 2018, in *Adaptive Optics Systems VI*, (SPIE), **107032A**
- Mollière, P., Wardenier, J. P., van Boekel, R., et al. 2019, *A&A*, **627**, A67
- Mollière, P., Stolker, T., Lacour, S., et al. 2020, *A&A*, **640**, A131
- Mollière, P., Molyarova, T., Bitsch, B., et al. 2022, *ApJ*, **934**, 74
- Mugnier, L. M., Cornia, A., Sauvage, J.-F., et al. 2009, *J. Opt. Soc. Am. A*, **26**, 1326
- Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019, *AJ*, **158**, 13
- Pairet, B., Cantalloube, F., Gomez Gonzalez, C. A., Absil, O., & Jacques, L. 2019, *MNRAS*, **487**, 2262
- Pueyo, L. 2016, *ApJ*, **824**, 117
- Rasmussen, C. E., & Williams, C. K. I. 2005, *Gaussian Processes for Machine Learning* (The MIT Press)
- Rocchetto, M., Waldmann, I. P., Venot, O., Lagage, P. O., & Tinetti, G. 2016, *ApJ*, **833**, 120
- Rodgers, C. D. 2000, *Inverse Methods for Atmospheric Sounding: Theory and Practice* (World Scientific Publishing Co. Pte. Ltd)
- Rothman, L. S., Gordon, I. E., Barber, R. J., et al. 2010, *J. Quant. Spec. Radiat. Transf.*, **111**, 2139
- Ruffio, J.-B., Macintosh, B., Wang, J. J., et al. 2017, *ApJ*, **842**, 14
- Ruffio, J.-B., Macintosh, B., Konopacky, Q. M., et al. 2019, *AJ*, **158**, 200
- Ruffio, J.-B., Konopacky, Q. M., Barman, T., et al. 2021, *AJ*, **162**, 290
- Samlund, M., Bouwman, J., Hogg, D. W., et al. 2021, *A&A*, **646**, A24
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJ*, **755**, L28
- Stolker, T., Bonse, M. J., Quanz, S. P., et al. 2019, *A&A*, **621**, A59
- Stolker, T., Quanz, S. P., Todorov, K. O., et al. 2020, *A&A*, **635**, A182
- Su, K. Y. L., Rieke, G. H., Stapelfeldt, K. R., et al. 2009, *ApJ*, **705**, 314
- Tennyson, J., & Yurchenko, S. N. 2012, *MNRAS*, **425**, 21
- 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](#)]
- Wang, J., Wang, J. J., Ma, B., et al. 2020, *AJ*, **160**, 150
- Wang, J. J., Vigan, A., Lacour, S., et al. 2021, *AJ*, **161**, 148
- Wang, J., Wang, J. J., Ruffio, J.-B., et al. 2023, *AJ*, **165**, 4
- Whiteford, N., Glasse, A., Chubb, K. L., et al. 2023, *MNRAS*, **525**, 1375
- Wilcomb, K. K., Konopacky, Q. M., Barman, T. S., et al. 2020, *AJ*, **160**, 207
- Wit, E., Heuvel, E. v. d., & Romeijn, J.-W. 2012, *Statistica Neerlandica*, **66**, 217
- Xuan, J. W., Wang, J., Ruffio, J.-B., et al. 2022, *ApJ*, **937**, 54
- Zhang, Y., Snellen, I. A. G., Bohn, A. J., et al. 2021, *Nature*, **595**, 370
- Zurlo, A., Vigan, A., Galicher, R., et al. 2016, *A&A*, **587**, A57

Appendix A: SPHERE Data Reduction

We begin the process of extracting planetary spectra by applying a range of pre-processing steps, hereafter referred to as the data reduction stage. Dark frames, detector flats and IFS flats are subtracted from the data. The spectral positions of each slice and the wavelengths are calibrated. Bad pixels and cross talk are corrected, and the science frames are background subtracted. ND filter transmission profiles are applied to stellar flux observations taken at the beginning and end of the ADI sequence observations. To centre the science frames, the satellite spots are used where available. If available, satellite spots are used to calibrate the wavelength centre of each channel. These frames have anamorphism corrected and are shifted to a common centre, and are output as a set of files with dimensions of (x, y, λ) .

For SPHERE data, we compare the results of using the standard SPHERE Data Center pipeline for the reduction to that of [Vigan \(2020\)](#). Recent updates to the pipeline have shown discrepancies in the wavelength solution for YJH data, however as this is only applicable to data taken in the satellite spot mode. Lacking satellite spots, we instead rely on the standard ESO wavelength solution, and remain unaffected by the changes.

As it is challenging to inject fake companions into the raw detector frames for an IFS, we instead use the extracted spectrum of HR 8799 e for our comparison. The raw frames were processed using both pipelines, and post-processed using KLIP in order to extract the planetary spectrum.

Figure A.1 shows the results of this extraction. While qualitatively similar, the wavelength solutions are different between the pipelines, and both the location and amplitude of the 1.3 μm feature disagree. Finally, there is significant discrepancy between the shapes of the spectra in the H band. Without a clear metric for selecting a reduction pipeline, we choose to use the most up-to-date implementation of [Vigan \(2020\)](#) due to its ease of use and Python-based interface.

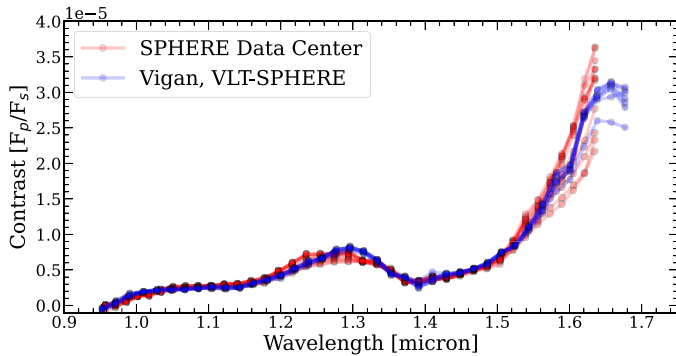


Fig. A.1: For the SPHERE data we compare the standard SPHERE Data Center data reduction (red) to the VLT-SPHERE pipeline described in [Vigan \(2020\)](#) (blue). As planets cannot be injected into the raw data, we compare the spectrum of HR 8799 e as extracted with KLIP, where each measurement in the figure represents a different number of principal components used in the spectra extraction.