

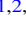







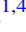




Comparing Apples with Apples: Robust Detection Limits for Exoplanet High-contrast Imaging in the Presence of Non-Gaussian Noise

Markus J. Bonse^{1,2} , Emily O. Garvin¹ , Timothy D. Gebhard^{1,2,3} , Felix A. Dannert^{1,4} , Faustine Cantalloube⁵ , Gabriele Cugno^{1,6} , Olivier Absil⁷ , Jean Hayoz^{1,4} , Julien Milli⁸ , Markus Kasper⁹ , and Sascha P. Quanz^{1,4} 

¹ETH Zurich, Institute for Particle Physics & Astrophysics, Wolfgang-Pauli-Str. 27, 8093 Zurich, Switzerland; mbonse@phys.ethz.ch

²Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, D-72076 Tübingen, Germany

³Max Planck ETH Center for Learning Systems, Max-Planck-Ring 4, D-72076 Tübingen, Germany

⁴National Center of Competence in Research Planets,¹⁰ Switzerland

⁵Aix Marseille University, CNRS, CNES, LAM, Marseille, France

⁶University of Michigan, Department of Astronomy, Ann Arbor, MI 48109, USA

⁷Université de Liège, STAR Institute, Allée du Six Août 19c, B-4000 Liège, Belgium

⁸Université Grenoble Alpes, CNRS, IPAG, F-38000 Grenoble, France

⁹European Southern Observatory, Garching bei München, Germany

Received 2022 September 20; revised 2023 March 15; accepted 2023 March 29; published 2023 July 19

Abstract

Over the past decade, hundreds of nights have been spent on the world’s largest telescopes to search for and directly detect new exoplanets using high-contrast imaging (HCI). Thereby, two scientific goals are of central interest: first, to study the characteristics of the underlying planet population and distinguish between different planet formation and evolution theories. Second, to find and characterize planets in our immediate solar neighborhood. Both goals heavily rely on the metric used to quantify planet detections and nondetections. Current standards often rely on several explicit or implicit assumptions about noise. For example, it is often assumed that the residual noise after data postprocessing is Gaussian. While being an inseparable part of the metric, these assumptions are rarely verified. This is problematic as any violation of these assumptions can lead to systematic biases. This makes it hard, if not impossible, to compare results across data sets or instruments with different noise characteristics. We revisit the fundamental question of how to quantify detection limits in HCI. We focus our analysis on the error budget resulting from violated assumptions. To this end, we propose a new metric based on bootstrapping that generalizes current standards to non-Gaussian noise. We apply our method to archival HCI data from the NACO instrument at the Very Large Telescope and derive detection limits for different types of noise. Our analysis shows that current standards tend to give detection limits that are about one magnitude too optimistic in the speckle-dominated regime. That is, HCI surveys may have excluded planets that can still exist.

Unified Astronomy Thesaurus concepts: [Astrostatistics \(1882\)](#); [Bootstrap \(1906\)](#); [Direct imaging \(387\)](#); [High angular resolution \(2167\)](#)

1. Introduction

In order to translate the information of exoplanet high-contrast imaging (HCI) observations into quantitative scientific results, an objective metric to quantify the brightness contrast between planet and host star is indispensable. In this context, two scientific questions are of particular interest: first, in case a new planet candidate is found, we want to know how confident we can be that the discovery is real. If several candidates are found during a survey, the calculated confidence can be used as a guideline for the management of follow-up time. Second, we want to know which planets we can safely rule out, i.e., we are sure that, given some confidence level, we should have seen them in our data. This second question is the subject of studies about planet occurrence rates and completeness in large surveys such as SHINE (Desidera et al. 2021; Langlois et al. 2021; Vigan et al. 2021) or GPIES (Nielsen et al. 2019). Any metric used to answer these two questions must be able to deal with the large diversity of data sets available in HCI. Over the

last years, HCI has gained increasing scientific importance as its discovery space complements other exoplanet detection and characterization techniques, such as radial velocity or transit observations. As a result, existing ground-based observatories feature, and continue to be upgraded with, dedicated HCI instruments (e.g., Kenworthy et al. 2018). Furthermore, HCI exoplanet science is driving the development of future instruments for the upcoming 30 m class telescopes (e.g., Quanz et al. 2015; Bowens et al. 2021; Brandl & Bettonvil 2021; Kasper et al. 2021). Also, the recently launched James Webb Space Telescope (JWST) is equipped with a suite of HCI modes¹¹ and a dedicated Early Release Science (ERS) program was accepted for early execution (Hinkley et al. 2022). In the long run, highly optimized space missions of large scale will be needed for the direct detection and characterization of a statistically relevant sample of temperate, terrestrial exoplanets (e.g., Gaudi et al. 2018; Quanz et al. 2022; Team 2019).

Finding a metric that provides comparable estimates for contrast across current and future instruments is challenging because any change in instrumentation, observing conditions or data postprocessing influences the characteristics of the data. Currently used standards often rely on several explicit or

¹⁰ www.nccr-planets.ch

¹¹ <https://jwst-docs.stsci.edu/methods-and-roadmaps/jwst-high-contrast-imaging>

Table 1
Standards for Contrast Quantification Used in HCI

Method and Reference	Q1	Q2	Input type	What is Noise?	SSS	Addresses	Completeness	Accounts for Non-Gaussian Noise
Estimation of noise PDF ^a	✓	✓	RI	pixel				✓ ^y
S/N ^b	✓	✓	RI	both				
<i>t</i> -Test (small sample) ^c	✓	✓	RI	apertures	✓			
Parametric P-Map ^d		✓	RI	apertures	✓		✓	
Nonparametric P-Map ^d		✓	RI	apertures	n.a.		✓	✓ ^z
STIM ^e	✓		RI	pixel				✓ ^y
ANDROMEDA ^f	✓	✓	RS	pixel	✓			
FMMF ^g	✓	✓	RS	pixel			✓	✓ ^v
ROC curves ^h			detection map and RI	apertures or pixel	n.a.			✓ ^z
Supervised ML ⁱ	✓		multiple RIs	pixel patches	n.a.			✓ ^u
RSM ^j	✓		RS	apertures or pixel	n.a.		✓	✓ ^w
This paper	✓	✓	RI	spaced pixel	✓		✓ ^x	✓

Notes. This table does not claim to be complete and is meant as an overview. We further note that the classifications used here only consider a selection of the relevant aspects. A ✓ displays that the aspect is addressed in the cited paper, not including follow-up work. Abbreviations: Q1: quantifies the uncertainty of a new detection; Q2: provides detection limits; SSS: accounts for small sample statistics; RI: residual image; RS: sequence of residuals along time; *u*: through learning typical noise pattern; *v*: estimates the distribution of the S/N over a whole survey; *w*: Gaussian and Laplacian noise; *x*: future work; *y*: estimates the noise PDF based on pixel values; and *z*: no assumption about the noise.

References.

^a Marois et al. (2008); ^b Meshkat et al. (2013); ^c Mawet et al. (2014); ^d Jensen-Clem (2017); ^e Pairet et al. (2019); ^f Cantalloube et al. (2015); ^g Ruffio et al. (2017); ^h Gomez Gonzalez (2016); ⁱ Gomez Gonzalez (2018); ^j Dahlqvist et al. (2020).

implicit assumptions about the noise. For example, it is often assumed that the noise in a postprocessed residual image is independent identically distributed (i.i.d.) and Gaussian (e.g., Cantalloube et al. 2015). While being an inseparable part of the metric, such assumptions are rarely verified. In fact, previous work has shown that the noise in ground-based HCI observations often deviates from Gaussian noise due to the presence of systematic speckle noise (compare Perrin et al. 2003; Aime & Soummer 2004; Soummer et al. 2007 for the noise statistics in raw images and Pairet et al. 2019; Dahlqvist et al. 2020 for the noise statistics of postprocessed residual images). If one computes detection limits under the assumption of Gaussian noise, but the actual noise is dominated by speckles, the results will be biased. The severity of this error depends on the extent to which the assumptions are violated. This makes it hard, or even impossible, to compare results across data sets and instruments with different noise characteristics.

In this paper, we revisit the fundamental question of “how to quantify detection limits” in HCI. Compared to previous work (Section 2), we focus our analysis on the error budget caused by violated assumptions (Section 4). For this purpose we propose a new metric based on bootstrapping, which generalizes the widely used standard by Mawet et al. (2014; Section 3) to non-Gaussian noise (Section 5). This allows us to study the systematic error of our detection uncertainty as a function of the noise characteristics. We further highlight the important difference between the signal-to-noise ratio (S/N) and the detection uncertainty. In Section 6 we demonstrate how to quantify detections with our new bootstrapping metric. For this purpose, we present a new approach which is consistent with previous work but universally applicable irrespective of the algorithm used during data postprocessing. We call this method the contrast grid. Finally, we analyze the error budget caused by non-Gaussian noise w.r.t. the detection limits in Section 7.

The code of the metrics presented in this paper is public available as a python package called `applefy` on GitHub <https://github.com/markusbonse/applefy>. A detailed documentation page on how to use the package, including many tutorials, can be found at ReadtheDocs <https://applefy.readthedocs.io/>.

2. Current Standards for Contrast Quantification

To this date, detection limits in HCI are not uniformly quantified, but are based on a wide range of standards (see Table 1 for an overview). Many of these standards are used along with particular data-reduction algorithms, which in turn are specialized for certain types of observing strategies. Most common observing strategies include angular differential imaging (ADI; Marois et al. 2006), spectral differential imaging (SDI; Racine et al. 1999; Sparks & Ford 2002), reference star differential imaging (RDI; Lafrenière et al. 2009), multi-reference star differential imaging (mRDI; Ruane et al. 2019), and polarimetric differential imaging (PDI; Kuhn et al. 2001; Quanz et al. 2011).

Data-reduction techniques can be categorized into three families (Cantalloube 2020): first, subtraction-based methods that try to model and subtract the stellar point-spread function (PSF). Famous examples are PCA/KLIP (Amara & Quanz 2012; Soummer et al. 2012), LOCI and its variations (Marois et al. 2010, 2014; Thompson & Marois 2021), LLSG (Gomez Gonzalez 2016), and the recently proposed HSR (Gebhard et al. 2022). The output of these techniques is a residual image whose values are related to flux. The quantification of the detection limits is carried out in a separate step. Second, forward-modeling techniques aim at tracking potential planetary signals during the observing sequence. Their result is not a residual, but a detection map showing the model’s belief about the presence of a planet. Contrast and detection uncertainty are defined as a match between the expected signal and the data. Forward modeling for HCI was introduced in Cantalloube et al. (2015) with the

ANDROMEDA algorithm and extended in various ways; see, for example, FMMF (Ruffio et al. 2017), PACO (Flasseur et al. 2018), and TRAP (Samland et al. 2021). Third, supervised machine-learning (ML) methods (Gomez Gonzalez 2018) try to learn the specific signatures of the signal and noise in the data. Their output is again a detection map.

Because numerous methods are used to postprocess the data, the metrics differ fundamentally in their definition of what a detection is. This starts with the research question they address: some metrics focus only on the detection problem, which is whether a potential candidate is real (Q1). Other metrics try to constrain detection limits (Q2). Yet other methods specialize solely on comparing postprocessing algorithms (e.g., ROC curves; Gomez Gonzalez 2016). Even more fundamentally, there is no uniform standard to define what is signal and what is noise. For example, some methods calculate the strength of the noise based on areas around the signal (Marois et al. 2008; Mesa et al. 2015; Otten et al. 2017; Golomb et al. 2021) while others consider noise with the same distance from the star (Cantalloube et al. 2015; Mawet et al. 2014; Jensen-Clem 2017). Other variants use the opposite angle rotated residual image (Wahhaj et al. 2013; Pairet et al. 2019) or estimates along time (Dahlqvist et al. 2020). The noise statistic is sometimes calculated directly on pixel values (Marois et al. 2008; Mesa et al. 2015; Cantalloube et al. 2015; Ruffio et al. 2017), while others average multiple pixel inside apertures (Mawet et al. 2014; Jensen-Clem 2017). Moreover, different types of statistics are used. While some authors use classical definitions of S/N (Rameau et al. 2013; Meshkat et al. 2013; Mesa et al. 2015; Uyama et al. 2017), others address the quantification using Frequentist hypothesis testing (Mawet et al. 2014; Jensen-Clem 2017) or Bayesian methods (Ruffio et al. 2018; Golomb et al. 2021).

Due to these differences, the results calculated with different standards are often not comparable. But which of the presented metrics gives the “right” scientific answer? On the one hand, an ideal metric should be universally applicable irrespective of the observing strategy or data reduction used. On the other hand, it should be robust under different data characteristics and provide reliable estimates for the achieved contrast. Current standards are limited with respect to both criteria: first, they are often specialized to be used with specific postprocessing algorithms. For example, FMMF is used after postprocessing with PCA. Second, they are reliant on fixed assumptions about the noise, making them not robust under varying conditions.

In the following Section 3 we revisit the metric presented in Mawet et al. (2014) and assess its limitations. Afterwards, we propose modifications and extensions to this approach in order to improve the robustness and universal applicability of the metric.

3. What is a Detection?

The data postprocessing routine combines the sequence of individual observations taken over the course of one night into a single image. In this residual image potential planet candidates appear as bright spots. A priori, we do not know if a bright spot is actually a planet, and the decision of whether it is bright enough to be counted as a detection is always a balancing of the risk that it is just part of the noise. In order to quantify this risk, we use hypothesis testing. The classical approach based on the t -test (for a general background on hypothesis testing and the t -test

see Chapter 9 of Larsen & Marx 2012) was introduced by Mawet et al. (2014) and is illustrated in Figure 1. The test procedure can be split into three main steps.

In the first step, we extract values for the noise and the potential planet. It is important that the noise is taken from positions that are representative for the noise we expect at the position of the signal. A common approach is to use noise with the same separation from the star. Further, the noise and signal must be extracted in the same way, ideally independent of the resolution of the detector. This can be done by averaging pixel values inside apertures of $1\lambda/D$ diameter. The result is two samples: the noise sample $\mathcal{X} = \{\bar{X}_1, \dots, \bar{X}_n\}$ and the planet sample $\mathcal{Y} = \{\bar{Y}_1\}$. Note that the planet sample only contains a single observation.

In the next step, two competing hypotheses are formulated: first, that our observation can be explained without the presence of a planet. This is the *null hypothesis* H_0 . Second, the hypothesis that our observation is indeed attributable to the existence of a planet. This is the *alternative hypothesis* H_1 . In order to decide which of the two hypotheses to favor we need to make assumptions and explicitly formulate H_0 and H_1 . The t -test assumes that our observations are drawn independently and are identically distributed (i.i.d.) from normal distributions $\{\bar{X}_1, \dots, \bar{X}_n\} \sim \mathcal{N}(\mu_{\mathcal{X}}, \sigma_{\mathcal{X}})$, $\{\bar{Y}_1\} \sim \mathcal{N}(\mu_{\mathcal{Y}}, \sigma_{\mathcal{Y}})$ where the parameters $\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}, \sigma_{\mathcal{X}}, \sigma_{\mathcal{Y}}$ are unknown. In addition, it is assumed that the distributions of the signal and the noise have the same variance¹² $\sigma_{\mathcal{X}} = \sigma_{\mathcal{Y}} = \sigma$. In case no planet is present we would expect that \mathcal{X} and \mathcal{Y} originate from the same distribution. Thus, we can formulate $H_0: \mu_{\mathcal{X}} = \mu_{\mathcal{Y}}$. In case a planet is present we expect \mathcal{Y} to be brighter than \mathcal{X} . This gives us $H_1: \mu_{\mathcal{X}} < \mu_{\mathcal{Y}}$. A closer look at step 2 in Figure 1 shows that H_0 becomes unlikely as the distance between the noise observations (blue crosses) and planet observation (orange cross) increases and the variance of the noise sample decreases. We can quantify this effect using the test statistic of the two-sample t -test (Mawet et al. 2014)

$$T = \frac{\hat{\mu}_{\mathcal{Y}} - \hat{\mu}_{\mathcal{X}}}{\hat{\sigma}_{\mathcal{X}} \sqrt{1 + 1/n}}. \quad (1)$$


The equation is simplified for the special case in HCI where \mathcal{Y} contains only a single value. The mean estimates of the signal and noise sample are denoted as $\hat{\mu}_{\mathcal{Y}} = \bar{Y}_1$ and $\hat{\mu}_{\mathcal{X}} = \bar{X}$, respectively, while $\hat{\sigma}_{\mathcal{X}}$ is the Bessel-corrected standard deviation of \mathcal{X} . Under H_0 and if $\hat{\mu}_{\mathcal{X}}$ and $\hat{\mu}_{\mathcal{Y}}$ follow a normal distribution, T follows a student t -distribution with $\nu = n - 1$ degrees of freedom. The t -distribution does not depend on the unknown parameters $\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}, \sigma$. This allows us to compute the detection uncertainty by

$$\text{FPF} = \int_{T_{\text{obs}}}^{\infty} p(T = t|H_0) dx, \quad (2)$$

where T_{obs} denotes the value of T that we compute for our observation and $p(T = t|H_0)$ is the probability density function


¹² Note, that the value of \bar{Y}_1 is not just the pure planet signal but also contains speckle noise: $\bar{Y}_1 = S_{\text{planet}} + \bar{X}_{n+1}$, where S_{planet} is the contribution of the planet signal and \bar{X}_{n+1} is the speckle noise at the position of the planet. We can describe S_{planet} as Poisson (photon shot noise) and keep the Gaussian assumption for \bar{X}_{n+1} . The resulting PDF for \bar{Y}_1 is then the convolution of the Poisson PDF with the Gaussian PDF (Theorem 5.2.9 in Casella & Berger 2002). Thus, by assuming equal variance, we assume that the effect of photon shot noise is negligible compared to the speckle noise.

1. Observation & Measurement



$$S/N = T_{\text{obs}} = 2.28$$

2. Hypothesis & Assumptions



3. Statistical Test & Conclusion

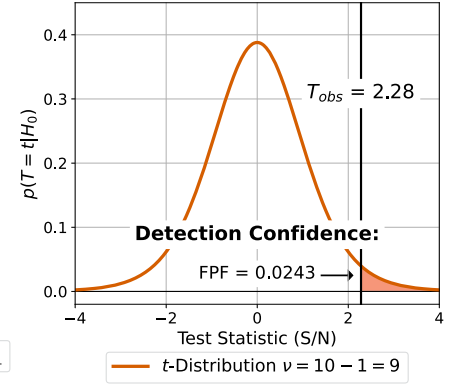


Figure 1. The three main steps of the hypothesis testing framework introduced by Mawet et al. (2014), which we generalize in this paper. Step 1 shows a typical example of a residual image obtained after data postprocessing with PCA. A previously inserted artificial companion appears as a spot slightly brighter than the remaining speckle noise. More details on the data set are given in Section 4. In order to determine if the planet is real we extract values for the noise X and signal Y by averaging pixel values inside apertures. In Step 2 we formulate two competing hypotheses H_0 (top) and H_1 (bottom) for detection and nondetection, respectively. The orange and blue crosses correspond to the signal and noise values obtained from the residual image on the left. Is the signal bright enough to show that \mathcal{X} and \mathcal{Y} differ significantly in their means? In order to answer this question we calculate the test statistic of the two-sample t -test $T_{\text{obs}} = 2.28$ (Equation (1)). Step 3: the detection uncertainty or FPF is given by the shaded red area under the t -distribution with $\nu = n - 1 = 9$ degrees of freedom.

(PDF) of the t -distribution. The false-positive fraction (FPF) gives the risk that we reject H_0 in favor of H_1 although no planet is present. In panel 3 of Figure 1 we obtain $T_{\text{obs}} = 2.28$, which corresponds to an FPF = 0.0243. That is, in 2.43% of the cases in which no planet is present we get a $T_{\text{obs}} \geq 2.28$. If the calculated FPF is below a previously defined detection threshold we treat our observation as a detection. Since small values of the FPF quickly become difficult to read, we use the following notation in this paper

$$x \sigma_N := 1 - \Phi(x) = \text{FPF}, \quad (3)$$

where Φ is the cumulative density function of the standard normal distribution. That is, we express the FPF values in terms of the quantiles of the standard normal distribution x . For example we write $5\sigma_N$ for FPF = 2.87×10^{-7} and $3\sigma_N$ for FPF = 1.25×10^{-3} . Note, σ_N should not to be confused with $\hat{\sigma}_X$.

It is crucial to distinguish between the assumed distribution of the noise (Gaussian) and the distribution of the test statistic (t -distribution). The t -distribution does not describe the nature of the noise but the effect of the sample size. For small n the values of $\hat{\mu}_X$, $\hat{\mu}_Y$, and $\hat{\sigma}_X$ are less accurate w.r.t. the true but unknown parameters μ_X , μ_Y , and σ . This uncertainty is compensated by the heavier tails of the t -distribution for small ν . For $n \rightarrow \infty$ Equation (1) converges to the classical definition of the S/N (compare, e.g., Meshkat et al. 2013) and the t -distribution converges to a normal distribution. In this limit we obtain the classical false alarm probabilities of FPF = 2.87×10^{-7} for $T = 5$ and FPF = 1.25×10^{-3} for $T = 3$. Note that the factor $1/\sqrt{1 + 1/n}$ in Equation (1) is not a correction for the classical S/N in the case of small sample sizes. It is a normalization that ensures that $p(T = t|H_0)$ follows a standard t -distribution. The small sample size at inner radii affects both the value of T_{obs} as well as the shape of the t -distribution. Any value of T_{obs} or S/N is meaningless if we do not consider the sample size and underlying assumptions of the test. For example at $2\lambda/D$ we need $T_{\text{obs}} \geq 11.2$ for a $5\sigma_N$ detection while at $10\lambda/D$ a value of $T_{\text{obs}} \geq 5.6$ is sufficient.

Therefore, any detection threshold should be specified as an FPF and not in terms of T_{obs} nor S/N.

4. Critical Assessment of Assumptions

In the following, we use the t -test described in the previous section as an example to assess critically the validity of the assumptions used in HCI. We note, that some of the metrics shown in Table 1 are based on assumptions similar to those of the t -test. For example the assumed noise distribution in forward-modeling techniques defines the maximum likelihood expression which is the basis for the calculation of the detection map. Thus, although our analysis is based on the t -test, some aspects discussed in this paper might be of general interest. Apart from the assumptions used in HCI, some implementation details can have a nonnegligible effect on the results. One example is the placement of the apertures in the residual image. We cover this aspect in Appendix A and focus our discussion on the two main assumptions of the t -test: independence and the Gaussianity of the residual noise.

Throughout the rest of this paper we compute our results based on an L' -data set of β Pictoris taken with the AGPM coronagraph (Delacroix et al. 2013) and NACO (Rousset et al. 2003) at the Very Large Telescope (VLT). The data set is the same as the one used in Mawet et al. (2014). The planet β Pictoris b was removed by insertion of a negative fake planet. More details on the data can be found in Absil et al. (2013). Our data preprocessing routine uses the `python` library PynPoint given in Stolker et al. (2019). The final stack of preprocessed frames consists of 29,681 images.

4.1. Independence

A central assumption of the t -test is that \mathcal{X} and \mathcal{Y} are drawn independently. This means that all noise values have to be uncorrelated with respect to all other noise values and the signal. Depending on the observing strategy, instrument, wavelength, and data reduction, the noise characteristics of HCI observations can be fundamentally different. For example, observations at 3–5 μm are often dominated by thermal background noise whereas observations at shorter wavelengths of 1–2.5 μm are more affected by speckles (Hunziker et al. 2018). The noise

Independence of Noise Observations

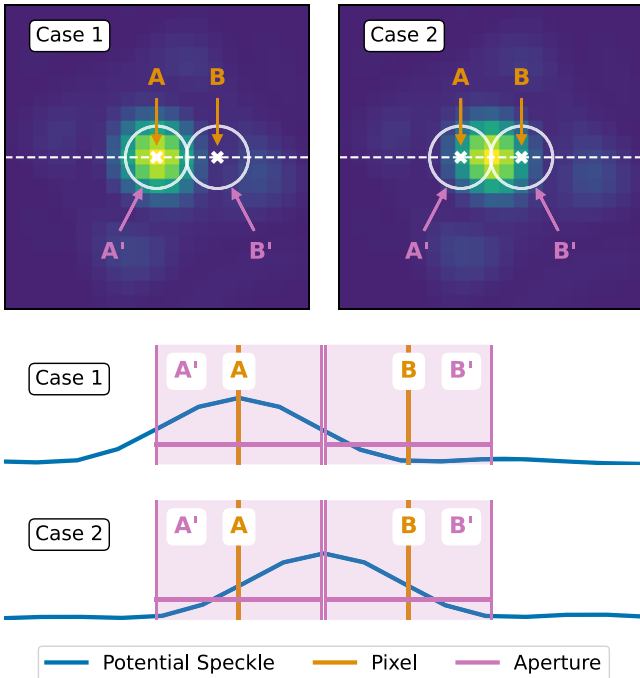


Figure 2. Violations of the independence assumption in the presence of speckle noise. The top two images illustrate the two cases described in the text. Pixels A and B are separated by 1 FWHM. We use the unsaturated PSF of the NACO data set to display the speckle. The plots below give 1D profiles of the two cases, cut along the white dashed lines. The shaded areas highlight the information contributing to the aperture averages A' and B' .

type also affects the spatial correlations in the data. While some noise sources such as photon noise are independent on a pixel-to-pixel level, the planet signal and speckle noise will follow the shape of the telescope’s PSF. If the noise is dominated by speckles and the pixel size is smaller than the width of the PSF, neighboring pixels are not independent. Some methods listed in Table 1 compute their noise statistics directly on pixel values while assuming independence. In the presence of speckle noise these methods will not provide accurate FPF estimates.

Compared to an analysis directly on the pixels, the use of apertures suggest that \mathcal{X} and \mathcal{Y} are independent. A careful investigation, however, reveals that this impression is incorrect. Let us consider two pixels A and B together with a speckle on the detector (compare Figure 2). If the two pixels are distant by less than 1 FWHM they will be partly correlated as both are influenced by the same speckle. This is the case if the speckle is located at A (case 1) or between A and B (case 2). Therefore, if we calculate aperture averages A' and B' around A and B , respectively, they will always be based on nonindependent pixel values. If we use apertures we implicitly filter the residual image with a box filter which has the shape of the aperture. Consequently, the length of the spatial correlations increases and with it the risk for a violation of the independence assumption.

We therefore propose to use the pixel values at the positions A and B directly instead of apertures for data sets in the speckle-dominated regime. The separation between A and B should be chosen according to the expected spatial correlation length in the data. As bad seeing conditions or poor adaptive optics performance can influence the shape of the PSF, we

space A and B by one FWHM and do not use the theoretical size of λ/D . The FWHM can be calculated by fitting a 2D Gaussian or Moffat to the unsaturated PSF (Stolker et al. 2019). The FWHM for the β Pictoris data set is 4.2 pixels and slightly larger than the theoretical size in λ/D (see also Jensen-Clem 2017). We note, that spacing A and B by one FWHM does not guarantee that their values are completely independent (compare case 2). But their values will be less dependent compared to the values of A' and B' . In addition to speckles other effects might influence the spatial dependencies in the data. A common example for this is the wind-driven halo discussed in Cantalloube et al. (2020). In such cases high-pass spatial filtering can be used to recover the spatial independence of the noise. The use of low-pass filters such as the Gaussian blur (Absil et al. 2013), however, exacerbates the independence problem. We further note that in reality, the spatial correlations can deviate from the shape of the PSF e.g., due to data postprocessing. An empirical analysis on this topic is given in Appendix E.

In the background-dominated regime the use of apertures might be preferable over the use of spaced pixels. Since, the photon noise occurs on a pixel-by-pixel basis it is not problematic with respect to independence. But, the signal estimate based on the brightest pixel is prone to statistical fluctuations and biases such as hot pixels. The choice whether apertures or spaced pixel are used should be taken on a case-by-case basis, which is why our `python` package `applefy` provides an implementation of both.

Future work should further investigate more sophisticated alternatives to the use of spaced pixel including explicit models of the pixel-to-pixel dependencies (compare, e.g., Golomb et al. 2021).

4.2. The Assumption of Gaussian Noise

The t -test assumes that the noise and the signal averages $\bar{\mathcal{X}}$, $\bar{\mathcal{Y}}$ follow a normal distribution. This assumption is usually justified by the use of observing strategies such as ADI (Marois et al. 2006), which average many individual observations. In this way, a temporal sequence of sufficiently i.i.d. observations will yield a residual image with normally distributed noise by virtue of the central limit theorem (CLT; Marois et al. 2008; Mawet et al. 2014). In addition, the use of data postprocessing techniques, such as a PSF subtraction with PCA, have been demonstrated to improve the normality of residuals considerably (Amara & Quanz 2012; Soummer et al. 2012; Cantalloube et al. 2015). Despite frame averaging and data postprocessing, the residual noise of real observations can still deviate from Gaussian. In Figure 3 we use Q–Q plots¹³ to compare the noise of the β Pictoris data set with normally distributed noise.

At 1 FWHM and 8 FWHM (labels 1 and 8, respectively) we notice that the noise for large values is above the diagonal line of the Q–Q plot. This implies that bright pixels appear more frequently in the data than we expect from Gaussian noise. Similar noise statistics were previously observed by Marois et al. (2008), Cantalloube et al. (2015), and Pairet et al. (2019). At 4 FWHM (labels 4a and 4p) we observe the opposite behavior: large values are found less frequently in our data compared to the normal distribution. A measure for these

¹³ Q–Q plots are a statistical tool to compare the quantiles of two distributions with each other. For a detailed explanation, see Pairet et al. (2019) and references therein.

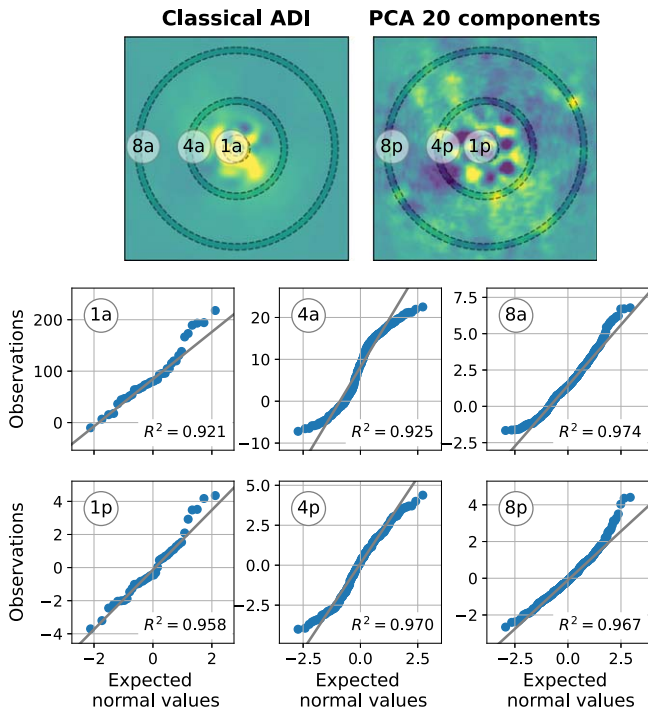


Figure 3. Deviations from a normal distribution in HCI residuals. The top two images show residual images for the β Pictoris data set: left with classical median ADI (Marois et al. 2006); right with full-frame PCA and 20 components. Below Q–Q plots are given to study the noise statistic of the pixel values inside the shaded areas at 1 FWHM, 4 FWHM, and 8 FWHM. The Q–Q plots compare the observed pixel values with normally distributed noise. A perfect match of observations and normal distribution would result in points exclusively on the gray diagonal line. We note, that the pixel values extracted from the shaded areas are not independent. The shown Q–Q plots only provide indicative evidence for the type of residual noise. But, they are not a proof for or against Gaussian distributed noise. A discussion on this topic can be found in Appendix C.

deviations is the coefficient of determination R^2 , i.e., the Pearson correlation, between the paired sample quantiles. See Pairet et al. (2019) for a definition and detailed explanation. The closer R^2 is to 1, the better the observed noise can be explained by Gaussian noise. As suggested by the results shown in Figure 3, the use of PCA partially mitigates the problem of non-Gaussian noise. Nevertheless, the noise is still not perfectly normal.

In the presence of noise which has a high probability of large values to occur, the probability that we observe a large value of T (Equation (1)) that is caused by the noise increases. Consequently, $p(T = t|H_0)$ will no longer follow a t -distribution and the interpretability of the test statistic T w.r.t. the FPF is lost. That is, we can still calculate values for T , but we no longer know with which detection uncertainty they are associated. Depending on the type of noise, different values of T_{obs} might be required to reject H_0 . This is especially problematic as noise characteristics can change from data set to data set. If we ignore potential violations of the Gaussian assumption, we under- or overestimate the detection uncertainty. For many applications outside HCI, the t -test is robust to slightly non-Gaussian data given a large sample size. As the sample size increases the average values of $\bar{\mathcal{X}}$, $\bar{\mathcal{Y}}$ will be Gaussian thanks to the CLT. In HCI, we cannot take advantage of this effect as \mathcal{Y} always contains a single observation. For small separations, where speckle noise is most important, the sample size of \mathcal{X} is also limited. How strong is the effect of

non-Gaussian noise on the detection uncertainty? Is a comparison between data sets or instruments with different noise characteristics still possible?

5. Beyond Gaussian Noise


In general, we can distinguish three concepts to deal with non-Gaussian noise: first, we can assume a different noise distribution which better describes the speckle behavior. The RSM presented in Dahlqvist et al. (2020) explores this idea, though along the time domain and not in the residual image. Second, we can try to estimate the distribution of the noise directly from the data. This idea was previously studied by Marois et al. (2008) and Pairet et al. (2019) who tried to estimate the noise PDF based on pixel values in the residuals. As mentioned before, this can be problematic since the pixels in the residuals are not independent. Further, extensive extrapolation is required in order to describe how the noise behaves in the tails of the PDF. The concept is therefore better suited for large surveys as, e.g., done by Ruffio et al. (2017). Finally, we can try to estimate detection limits by not making any assumptions about the noise. An example for this are ROC curves (Gomez Gonzalez 2016; Jensen-Clem 2017), which are often used to benchmark postprocessing algorithms.

In the following, we propose a new metric based on bootstrapping which follows the first concept and allows us to compute the detection uncertainty for any type of noise. The other two concepts are subject of future work. In order to retain the interpretability of the test statistic T , we need to find its distribution $p(T = t|H_0)$. In principle this could be done empirically. That is, we repeat our observation a lot of times on identical stars that do not host a planet and compute values for T for every observation. The probability $p(T = t|H_0)$ is then given by the frequency that a certain value of T is observed. While repeating the observation that many times is not feasible in practice, it gives the key idea of the bootstrap: we aim at finding a distribution \hat{F} that is close to the true but unknown distribution of the noise F and use it to repeat the experiment. Two main concepts can be distinguished: first, nonparametric bootstrapping, which makes no further assumptions about the noise and tries to approximate F directly from the data. Second, parametric bootstrapping, which assumes that F follows a parametric distribution with unknown parameters. In practice, if only a few noise observations are available, the estimated noise model \hat{F} of nonparametric bootstrapping might be inaccurate w.r.t. the true F . This can lead to less accurate estimates of the test statistic distribution $p(T = t|H_0)$. In this paper we study how to compute detection uncertainties using parametric bootstrapping, that is, under the assumption that the noise distribution is known. Data-driven nonparametric bootstrapping will be explored in future work.


5.1. Parametric Bootstrapping

We adopted the bootstrap test discussed in example 3.4.1 of Zoubir & Iskander (2004) and chapter 16.2 of Efron & Tibshirani (1994) to the special case of HCI where only a single value for \mathcal{Y} is available. The test is similar to the two-sample t -test explained in Section 3 except that we can assume any noise distribution. The following procedure describes the main steps of the test, which are the same irrespective of the assumed noise distribution. For demonstration purposes, we

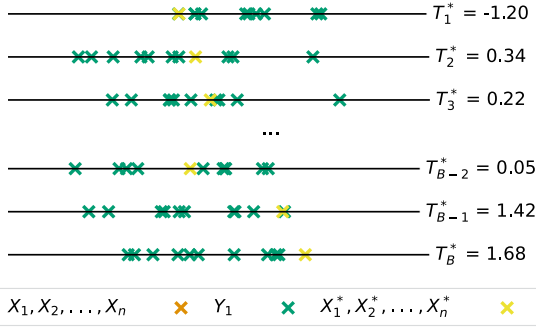
1. / 2. Measurement



3. / 4. MLE



5. / 6. Bootstrapping



7. Hypothesis & Assumption

$$p(T = t|H_0) \approx p(T^* = t|H_0)$$

7. Statistical Test & Conclusion

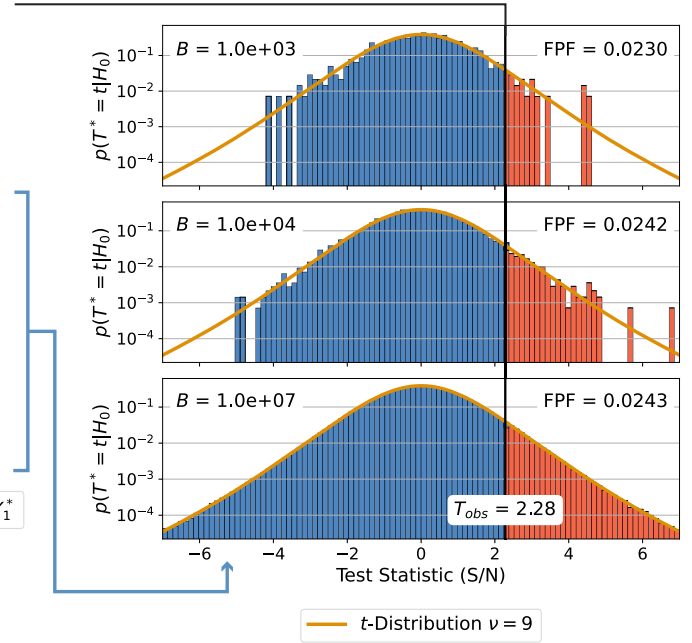


Figure 4. The parametric bootstrap test introduced in this paper. The left side of the figure illustrates the key steps of the test, exemplary for Gaussian noise. The values used are identical to those in Figure 1. The numbers of the steps match with those used in the text. On the right, three histograms are given, which show the empirical distribution of $p(T^* = t|H_0)$ for different numbers of bootstrap samples B . The actual test does not require us to compute histograms and the FPF is computed from the bootstrap samples $T_{(1)}^* \leq \dots \leq T_{(B)}^*$ directly (see Equations (4) and (5)). The results are compared with the t -distribution with $\nu = n - 1 = 9$ degrees of freedom. Compared to the t -test, bootstrapping is much more versatile, since any distribution, not necessary Gaussian, can be assumed in step 3.

give examples of how the test is carried out in the case of Gaussian noise (see Figure 4):

1. Extract the noise sample $\mathcal{X} = \{X_1, \dots, X_n\}$ and the signal sample $\mathcal{Y} = \{Y_1\}$ from the residual image.¹⁴ As for the t -test we assume that the noise is independent and identically distributed.
2. Compute T_{obs} by using Equation (1).
3. Assume a parametric distribution for the noise and list its unknown parameters. In the case of Gaussian noise this would be $F = \mathcal{N}(\mu_{\mathcal{X}}, \sigma)$.
4. Compute the maximum likelihood estimate of the unknown parameters. For Gaussian noise we compute the mean $\hat{\mu}_{\mathcal{X}} = \bar{X}$ and the standard deviation $\hat{\sigma} = \text{std}(\mathcal{X})$ and set $\hat{F} = \mathcal{N}(\hat{\mu}_{\mathcal{X}}, \hat{\sigma})$.
5. Repeat the experiment B times under H_0 using numerical simulations based on \hat{F} . For every repetition we resample the noise $\mathcal{X}^* = \{X_1^*, \dots, X_n^*\} \sim \hat{F}$ and the signal $\mathcal{Y}^* = \{Y_1^*\} \sim \hat{F}$ from the same distribution. It is important to draw the same number of noise observations for \mathcal{X}^* as we have in \mathcal{X} . The results are B so-called bootstrap samples

$$\begin{aligned} & (\{X_1^*, \dots, X_n^*\}, \{Y_1^*\})_1, \dots, (\{X_1^*, \dots, X_n^*\}, \{Y_1^*\})_B \\ & = (\mathcal{X}_1^*, \mathcal{Y}_1^*), \dots, (\mathcal{X}_B^*, \mathcal{Y}_B^*). \end{aligned}$$

6. For every bootstrap sample $(\mathcal{X}_i^*, \mathcal{Y}_i^*)$ compute the test statistic T_i^* by using Equation (1). Rank the results with increasing order $T_{(1)}^* \leq \dots \leq T_{(B)}^*$.

7. Under the assumption that \hat{F} is close to F the distribution of $p(T^* = t|H_0)$ will converge to $p(T = t|H_0)$ if B becomes large. The convergence depends on the maximum likelihood estimates calculated in step 4. If this estimate is inaccurate the bootstrap results will be biased. However, if the assumed noise distribution satisfies some additional properties explained in Section 5.3, we are able to obtain the exact distribution of $p(T = t|H_0)$ irrespective of how close \hat{F} is to F .

The FPF is given by the fraction of values in $\{T_1^*, \dots, T_B^*\}$ which are larger than T_{obs} :

$$\text{FPF} = \#\{T_i^* > T_{\text{obs}}\} / B. \quad (4)$$

More accurate results can be obtained by linear interpolation. For this purpose we use the sorted bootstrap results $T_{(1)}^* \leq \dots \leq T_{(B)}^*$ from the previous step and search for the two values of T^* which are adjacent to T_{obs} : $T_{(a)}^* \leq T_{\text{obs}} \leq T_{(b)}^*$, where a and b are the corresponding indices. The detection uncertainty can be computed with

$$\begin{aligned} 1 - \text{FPF} &= \frac{a}{B-1} \cdot \frac{T_{\text{obs}} - T_{(a)}^*}{T_{(b)}^* - T_{(a)}^*} \\ &+ \frac{b}{B-1} \cdot \frac{T_{(b)}^* - T_{\text{obs}}}{T_{(b)}^* - T_{(a)}^*}. \end{aligned} \quad (5)$$

As shown in Figure 4 the calculated FPF of the bootstrap test converges to $\text{FPF} = 0.0243$. This is exactly the same value that we obtained with the t -test in Section 3. This means that the parametric bootstrap is equivalent to the t -test if we assume Gaussian noise.

¹⁴ We write $\mathcal{X} = \{X_1, \dots, X_n\}$ and $\mathcal{Y} = \{Y_1\}$ instead of $\mathcal{X} = \{\bar{X}_1, \dots, \bar{X}_n\}$ and $\mathcal{Y} = \{\bar{Y}_1\}$ to emphasize that we use spaced pixels instead of apertures (see the discussion in Section 4.1).

5.2. Parametric Bootstrapping for Laplacian Noise

In contrast to the t -test, parametric bootstrapping allows us to compute the FPF for any type of noise. This can be done by swapping out the Gaussian assumption in step 3. In coronagraphic images with small static residuals, the modified Rician distribution of the speckle noise (Soummer et al. 2007; Aime & Soummer 2004) is reduced to a one-sided exponential (Fitzgerald & Graham 2006). This applies for example in the case of coronagraphic imaging with small non-common-path aberrations or after some basic data postprocessing. If we subtract two such images from each other, as done in the ANDROMEDA algorithm (Cantalloube et al. 2015), we expect the noise to follow a two-sided exponential distribution, i.e., a Laplacian. Similarly, if a PCA-model tries to subtract bright speckles, it sometimes erroneously induces negative speckles leading to a similar type of noise. Even after averaging several frames along the temporal dimension, the residual noise of some HCI data sets is often better described¹⁵ by Laplacian and not by Gaussian noise (see Pairet et al. 2019 for a detailed analysis on this topic). In order to extend the bootstrap procedure to Laplacian noise we first assume $F = \mathcal{L}(\mu_{\mathcal{X}}, b)$ with PDF

$$f(x|\mu_{\mathcal{X}}, b) = \frac{1}{2b} \exp\left(-\frac{|x - \mu_{\mathcal{X}}|}{b}\right). \quad (6)$$

The maximum likelihood parameters (step 4) for $\mu_{\mathcal{X}}$ and b are given by Norton (1984)

$$\begin{aligned} \hat{\mu}_{\mathcal{X}} &= \text{median}(\mathcal{X}), \\ \hat{b} &= \frac{1}{n} \sum_{i=1}^n |X_i - \hat{\mu}_{\mathcal{X}}|. \end{aligned} \quad (7)$$

This gives us $\hat{F} = \mathcal{L}(\hat{\mu}_{\mathcal{X}}, \hat{b})$, which we use to repeat the experiment by resampling (compare the left side of Figure 4). The remaining steps of the test are identical to those in Section 5.1.¹⁶ By changing F from a Gaussian to a Laplacian, we are able to compute the FPF under the assumption of Laplacian noise and the distribution of the test statistic $p(T = t|H_0)$ (compare the right side of Figure 4) will no longer be a t -distribution. A comparison of $p(T = t|H_0)$ under the assumption of Gaussian and Laplacian noise is given in Figure 5. If the assumed noise distribution is Gaussian, the distribution of $p(T = t|H_0)$ resulting from bootstrapping agrees with the t -distribution. This applies to separations close to the star as well as further out. If the assumed noise distribution is Laplacian, the bootstrapping converges to a distribution with even heavier tails than the t -distribution. These heavy tails allow us to correct for the high occurrence of large noise values in Laplacian noise. We note that the distributions shown in Figure 5 are not the distributions of the noise (here Gaussian

¹⁵ Often it is not possible to find sufficient evidence that the residual noise is indeed Laplacian or Gaussian. More details on the topic are given in Appendix C.

¹⁶ The test statistic in Equation (1) is not necessarily optimal for all types of noise. It might be possible to derive a new test statistic for specific situations which still gives correct FPF estimates and at the same time offers higher power. This could be done by utilizing the Neyman–Pearson lemma (Theorem 3.1.1 in Zoubir & Iskander 2004). We leave this idea open for future work and focus our analysis on the calculation of the error budget using the test statistic in Equation (1).




Figure 5. The convergence distributions of the parametric bootstrapping, shown as histograms, in comparison with the normal and t -distribution. The top two plots show $p(T = t|H_0)$ under the assumption that F is Gaussian, one at a separation of 2 FWHM i.e., $n = 11$ and one at 8 FWHM i.e., $n = 49$. The plots show the same result as the right side of Figure 4 but for different n . The bottom two plots show $p(T = t|H_0)$ under the assumption of Laplacian noise. We use $B = 10^8$ bootstrap iterations.

and Laplacian), but the distribution of the test statistic $p(T = t|H_0)$.

In step 5 of the bootstrap algorithm we sample the same number of noise observations n as in the original sample \mathcal{X} . This allows us to consider simultaneously the effects of the small sample statistics and non-Gaussian noise. The procedure can be extended to any parameterized noise distribution F for which the maximum likelihood estimates in step 4 are known.

5.3. The Importance of Pivoting

In general, if we make no further restrictions on the type of noise distribution assumed, the accuracy of the bootstrap test depends on the maximum likelihood estimates calculated in step 4 in Section 5.1. Any deviation of the estimated \hat{F} from the true F will change $p(T^* = t|H_0)$ and with it bias the FPF. However, for some specific types of noise distributions, we can overcome this limitation by taking advantage of so-called pivotal quantities (see definition 9.2.6 in Casella & Berger 2002): “a random variable $Q(\mathcal{X}, \theta)$ is a pivotal quantity if the distribution of $Q(\mathcal{X}, \theta)$ is independent of all parameters. That is, if $\mathcal{X} \sim F(x|\theta)$, then $Q(\mathcal{X}, \theta)$ has the same distribution for all values θ .”

In case of the t -test the noise and planet sample follow a normal distribution $\{X_1, \dots, X_n\} \sim \mathcal{N}(\mu_{\mathcal{X}}, \sigma)$, $\{Y_1\} \sim \mathcal{N}(\mu_{\mathcal{Y}}, \sigma)$ with unknown parameters $\theta = \{\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}, \sigma\}$. Under the null hypothesis, the test statistic T in Equation (1) is a pivotal quantity. That is, the distribution of $p(T = t|H_0)$ is always the same, irrespective of the unknown parameters $\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}, \sigma$. This means that no matter which maximum likelihood estimates $\hat{\mu}_{\mathcal{X}}$ and $\hat{\sigma}$ we calculate, the bootstrap procedure will always converge to the same distribution. For the t -test this distribution is the t -distribution. The fact that T is a pivot is crucial for the t -test: it allows us to compute the exact FPF based on the t -distribution without knowing $\mu_{\mathcal{X}}, \mu_{\mathcal{Y}}$, and σ .

In Appendix D we prove that T is a pivot not only under the assumption of Gaussian noise but for all distributions from a location-scale family. A location-scale family is characterized by two unknown parameters: a scale parameter q and a shift parameter w . Let $Z \sim \mathcal{F}$ be a random variable following a standard distribution with no unknown parameters. By shifting and scaling Z we obtain the location-scale family with random variables $X = Zq + w$. In the case of the normal distribution Z is the standard normal distribution $Z \sim \mathcal{N}(0, 1)$ with $q = \sigma$ and $w = \mu$. Also the Laplacian distribution is a location-scale family where $Z \sim \mathcal{L}(0, 1)$ with PDF

$$f(x) = \frac{1}{2} \exp(-|x|) \quad (8)$$

where $q = b$ and $w = \mu$ (compare to Equation (6)). The result that T under H_0 is a pivotal quantity for location-scale family distributions has direct implications for the bootstrap procedure explained in Section 5.1: the computation of the maximum likelihood parameters in step 4 is no longer needed. The resampling in steps 5 and 6 is independent of $\hat{\mu}_{\mathcal{X}}$ and $\hat{\sigma}$ (or \hat{b}) and always converges to the same $p(T^* = t|H_0)$. No matter how close $\hat{\mu}_{\mathcal{X}}$ and $\hat{\sigma}$ (or \hat{b}) are to the true values of $\mu_{\mathcal{X}}$ and σ (or b), we always obtain the exact FPF. The estimation of the FPF becomes a simple lookup. We can calculate $p(T^* = t|H_0)$ once and reuse it for future calculations, which drastically reduces the computation time.

The distribution of $p(T^* = t|H_0)$ depends on the sample size n , which is why we have to compute separate lookups for different separations from the star. In order to achieve accurate results for low FPFs, the number of bootstrap samples B has to be large. For example, if we aim for a FPF of 2.87×10^{-7} ($5\sigma_N$), we expect to observe one false positive every $B = 1/2.87 \times 10^{-7} = 3.48 \times 10^6$ iterations. To constrain $p(T^* = t|H_0)$ with sufficient accuracy, we need about $B = 10^8$ bootstrap samples. In our `python` package `applefy` we provide precomputed lookups within $1\text{--}20 \lambda/D$ for $B = 10^8$ and interfaces which can be used to extend the code to other noise distributions.

6. How to Quantify Detections With Parametric Bootstrapping?

Current standards for contrast quantification in HCI (compare Section 2) address two main research questions: first, the detection problem, i.e., whether a potential candidate is real (Q1), and second, the quantification of the detection limits (Q2). The following two subsections give a detailed recipe of how parametric bootstrapping can be used to address these two questions. An implementation is given in our `python` package `applefy`.

6.1. How to Compute the Uncertainty of a Detection?

In Section 4 we identified two limitations of the t -test when used in HCI: (1) violation of the independence assumption if we use apertures and (2) non-Gaussian residual noise. To overcome these limitations, we propose the following procedure to quantify whether a signal in the data is a real planet candidate:

1. Make an assumption about the type of noise present in the residual. For a conservative choice and in the case of speckle noise we recommend choosing Laplacian noise

over Gaussian noise. If the noise is from a location-scale family, use parametric bootstrapping to precompute the distribution of the test statistic under the null hypothesis $p(T^* = t|H_0)$ (see Section 5.3). In the case of Gaussian noise use the t -distribution. For surveys the use of the nonparametric bootstrapping briefly mentioned in Section 5 represents an alternative to a fixed assumption about the noise.

2. The planet sample \mathcal{Y} and noise sample \mathcal{X} need to be extracted in the same way. Since we use spaced pixels for the noise, we have to use one pixel as the planet signal to get commensurable quantities. We search for the position of the planet by fitting a 2D Gaussian. The flux integrated over the circular area of one pixel around the best-fit position gives us the planet signal \mathcal{Y} .
3. Extract noise values \mathcal{X} with the same separation from the star. To ensure that the values in \mathcal{X} are approximately independent, we do not use apertures but integrate over the circular area of one pixel spaced by one FWHM. In the background-limited regime the use of apertures might be preferable for this and the previous step.
4. Use Equation (1) to compute the test statistic T_{obs} .¹⁷
5. Translate T_{obs} into the detection uncertainty (FPF) by using Equations (2) or (5).
6. Repeat steps 3–5 for different noise positions. This can be done by rotation of the initial noise positions (compare the discussion in Appendix A and Figure 10).
7. Report the median FPF over all noise positions. The mean absolute deviation from the median can be used as a measure of the uncertainties introduced by the placement of the noise values.

6.2. How to Compute Detection Limits?

If no planet candidate is found, we can compute detection limits to constrain the existence of which planets we can confidently rule out. To this end, artificial planets are usually inserted in order to determine the minimum planet brightness still detectable (Marois et al. 2010; Morzinski et al. 2015). Hereby, the calculation of the detection limit has to be consistent with the calculation of the FPF in the previous section. That is, if we insert an artificial planet above the detection limit we expect it to be counted as a detection. Vice versa, if we insert an artificial planet below the detection limit it should not be detectable.

In reality detection limits are never a hard limit and only give the point at which approximately 50% of the planets are detected. There is no guarantee that all planets above the contrast curve will be detected and no planet below the curve will not be detected. At the position of the planet we always observe a combination of planet signal and speckle noise. If we are lucky and a faint planet falls on top of a speckle, we might be able to detect it although it is below our calculated limits. Conversely, if a negative speckle caused by the data postprocessing is added to the planet signal, it may fall below the detection threshold. This effect can be quantified through calculations of the true-positive rate (TPR) also called the power of the test. A detailed discussion about the topic is presented in Jensen-Clem (2017). In this paper we focus only on the FPF and leave power calculations for the parametric bootstrapping open for future work.

¹⁷ The assumption made in step 1 only influences the distribution of the test statistic $p(T = t|H_0)$. The test statistic itself is the same.

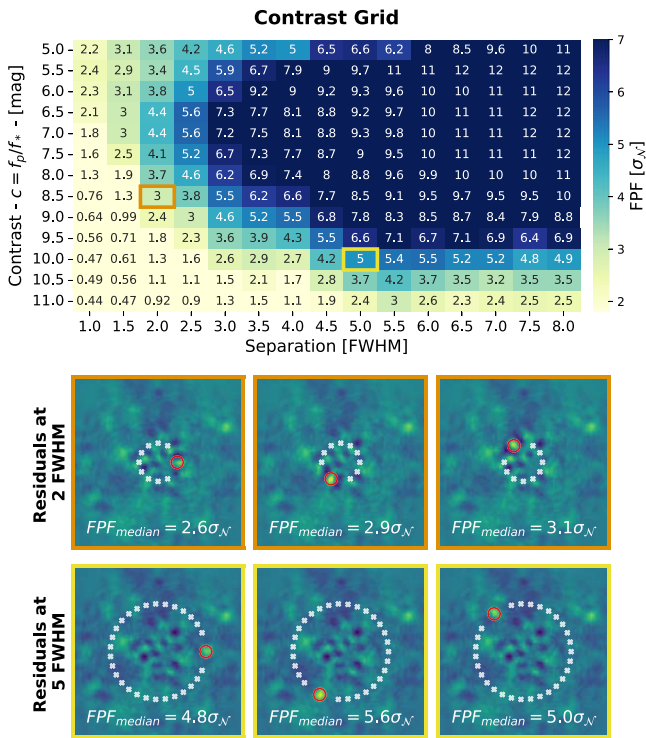


Figure 6. Detection limits for the β Pictoris data set under the assumption of Gaussian noise. The results under the assumption of Laplacian noise are presented later. The contrast grid (top panel) gives the detection uncertainty (FPF) as a function of separation and contrast. Each value in the grid is based on six fake planet residuals. Below three example residuals are shown, which were used to compute the contrast grid at the positions marked in orange and yellow. The $\text{FPF}_{\text{median}}$ given for each residual is the median FPF over different noise sample positions (see the details in Section 6.1). The mean of all six $\text{FPF}_{\text{median}}$ values gives the final FPF value of the contrast grid. An interactive version of the plot is available on the documentation page of `applefy`: <https://applefy.readthedocs.io/>.

In the following we introduce a new approach to compute detection limits which is independent of the data postprocessing. We call this approach the contrast grid. The calculation starts with the insertion of artificial planets at different separations, e.g., $s = 1, 1.5, \dots, 7.5, 8$ FWHM and planet brightnesses, e.g., $f_p/f_* = 5, 5.5, \dots, 10.5, 11$ mag into the raw data set. We use the unsaturated PSF scaled by f_p/f_* and potential attenuation due to a coronagraph as the planetary signal.

In order to account for azimuthal variations we insert six artificial planets for each separation, one every 60° . Some postprocessing algorithms like the subtraction-based half-sibling regression presented in Gebhard et al. (2022) make use of spatial correlations in the data. If we insert multiple planets at the same time these methods might learn to subtract the planet based on the movement of other planets in the data. Therefore, we insert only one planet at a time. For every inserted fake planet we run the data postprocessing, in our case PCA, and save the residual images. Next, we estimate the FPF for each fake planet using the procedure explained in Section 6.1. Depending on the type of noise we can choose between the t -test or parametric bootstrapping. The contrast grid, shown in Figure 6, combines all FPF estimates into a single 2D grid. The $5\sigma_N$ and $3\sigma_N$ contrast curves can be obtained by thresholding and interpolation of the grid. As shown in the figure, no fake planet within 2 FWHM exceeds the $5\sigma_N$ detection threshold, no matter how bright it is. The

reason for this is planet self-subtraction: at close separations a large fraction of the planetary signal is subtracted by the PCA noise model. If the planet gets brighter a progressively larger fraction of the planet flux is subtracted. In extreme cases an increase in planet brightness does not lead to a stronger signal in the residual. A detailed discussion of the effect is given in Appendix B.

The number of PCA components used during data postprocessing affects the achieved contrast. While a higher number of PCA components leads to a stronger reduction of the noise, it also causes more loss of planetary signal. For a given data set fewer components give higher contrast at close separations and more components provide better results for large separations. We recommend computing contrast curves for a range of PCA components and report the overall best (Xuan et al. 2018). Our `python` package `applefy` provides the code needed to automate and parallelize these computations.

Under additional assumptions and for some data postprocessing algorithms it is possible to compute the contrast curve analytically. This way the computation time can be reduced significantly. Current implementations in HCI packages like VIP (Gomez Gonzalez 2017) or PynPoint (Stolker et al. 2019) make use of this alternative. Also `applefy` allows one to compute analytical contrast curves. A detailed discussion of this approach and how it can be used with parametric bootstrapping is presented in Appendix B.

7. Results and Discussion

The parametric bootstrapping presented in Section 5 quantifies planet detections for different types of residual noise. This way we can calculate the FPF of a potential planet candidate or, conversely, fix the FPF in order to determine detection limits. But, how accurate is the calculated FPF for non-Gaussian noise? Can we afford to make the wrong assumption?

To answer these questions we run four Monte Carlo simulations considering Gaussian noise with ($\mu = 5, \sigma = 5$), Laplacian noise with ($\mu = 5, b = 5$) at 2 FWHM and 8 FWHM separation from the star. The separation in this simulation only influences the number of available noise values. Each Monte Carlo simulation is based on 10^9 experiments for which we sample n values from the respective noise distribution representing the planet signal under H_0 and the $n - 1$ noise values available at 2 FWHM ($n = 11$) and 8 FWHM ($n = 49$) separation. We run three tests to calculate the FPF for each experiment of the Monte Carlo simulation: the t -test, parametric bootstrapping under the assumption of Gaussian noise, and parametric bootstrapping under the assumption of Laplacian noise. For the parametric bootstrapping we use precomputed lookups as discussed in Section 5.3 with $B = 10^8$.

We set different detection thresholds and count the number of experiments for which the tests are fooled by the noise; that is, the calculated FPF is smaller than the chosen threshold. This way we obtain the true FPF of the tests for a given detection threshold. A comparison of the true FPF with the FPF promised by the tests is shown in Figure 7.

We observe that if the simulated noise is Laplacian the obtained confidence for the t -test is substantially lower than the desired confidence. For example at 2 FWHM, a chosen confidence of $5\sigma_N = 2.867 \times 10^{-7}$ FPF results in $3.8\sigma_N = 7.235 \times 10^{-5}$ FPF. This means that we

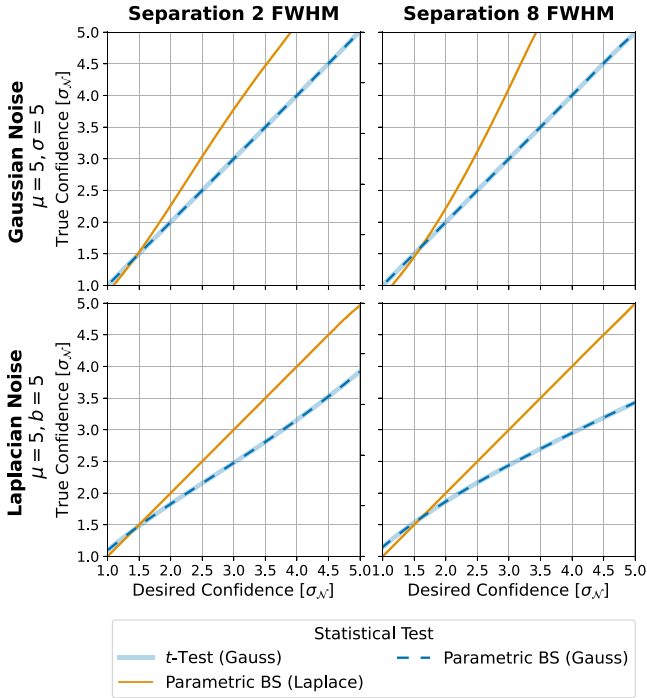


Figure 7. Effect of different noise distributions on the detection uncertainty. The plots give the relationship between the selected detection threshold and the actually observed FPF for the t -test and the parametric bootstrap tests presented in this paper. Any deviation from the diagonal corresponds to an over- or underestimation of the FPF. The results are based on a Monte Carlo simulation described in the text.

underestimate the number of false positives by a factor of 250. Unlike the t -test, the Laplacian parametric bootstrapping computes the correct FPF. If we suspect that our residual noise distribution is better described by a Laplacian we should choose the parametric bootstrap over the t -test. If the noise is Gaussian the t -test is well calibrated and provides accurate FPF results. The Laplacian parametric bootstrapping, however, overestimates the FPF and the results become too pessimistic. This means that accurate knowledge about the underlying noise distribution is critical to determine the FPF. The parametric bootstrapping under the assumption of Gaussian noise is consistent with the t -test. This result demonstrates that bootstrapping is able to account for non-Gaussian noise and at the same time for the small sample statistics.

Required S/N. For many applications, such as the computation of contrast curves, we are not interested in the FPF of a potential planet. Instead, we want to fix the FPF to constrain which value of the test statistic T_{obs} (S/N) is needed to achieve the desired confidence (FPF). For Gaussian noise this can be done by solving Equation (2) for T_{obs} . For Laplacian noise we use the procedure explained in Appendix B. The results for different separations are summarized in Figure 8. As shown in the plot a significantly larger value of T_{obs} is required under Laplacian noise. This is because the Laplacian distribution has heavier tails and the occurrence rate of large noise values is higher. If we aim for a detection confidence of $5\sigma_N$ the signal of the planet needs to be more than two times brighter compared to the limits of Gaussian noise. The effect is important irrespective of the separation from the star. This means that even if we have a large sample size, we are not robust to non-Gaussian noise. This is due to the noise at the position of the planet and the fact that the sample of the signal




Figure 8. Relationship between the value of the test statistic T_{obs} and the detection uncertainty specified as $5\sigma_N = 2.87 \times 10^{-7}$ FPF and $3\sigma_N = 1.35 \times 10^{-3}$ FPF. The t -distribution shown in blue converges toward a Gaussian for large separations (black dashed line) while the Laplacian parametric bootstrap remains heavier-tailed.

contains only one observation (see the discussion in Section 3). For a detection threshold of $3\sigma_N$ the difference between Gaussian and Laplacian noise is less important but still not negligible. At large separations the t -test converges to the classical 5σ or 3σ limits. Both tests account for the effect of small sample statistics at separations close to the star.

Detection limits. In order to investigate the effect of non-Gaussian noise on the detection limits, we compute contrast curves under the assumption of Gaussian and Laplacian noise. For this purpose, we compare the results of the classical t -test with those of the Laplacian parametric bootstrap test. We follow the procedure discussed in Appendix B to calculate our contrast curves. We overcome the limitations discussed in Section 4, i.e., we use spaced pixels instead of apertures to guarantee independent noise observations and take the median contrast over several different noise positions. The results for the β Pictoris data set are shown in Figure 9.

For separations >2 FWHM the contrast of the t -test is about one magnitude deeper compared to the contrast of the parametric bootstrap. This difference corresponds to a factor of 2.5 in planetary brightness. At separations close to the star (1–2 FWHM) the difference between the two limits is smaller. At these separations the effect of small sample statistics becomes relevant. The t -test accounts for this effect through the heavier tails of the t -distribution. These heavier tails also partially mitigate the problem caused by non-Gaussian noise. If the true noise is heavy tailed, the contrast curve of the t -test is too optimistic. In such situations direct imaging surveys might have ruled out regions of the parameter space where we might still find planets. Note, non-Gaussian noise has a systematic effect on the results. That is, the error does not average out if




Figure 9. Comparison of the detection limits for the β Pictoris data sets, once under the assumption of Gaussian (blue line) and once for Laplacian (orange line) residual noise. The solid line is the median contrast over 360 different placements of the noise values in the residuals. The shaded area gives the mean absolute deviation from the median. We used 30 PCA components for all separations.

several data sets are combined within a survey. If the data sets within one survey are all affected by heavy-tailed speckle noise the whole survey is biased. In reality, the true distribution of the noise is influenced by many factors and therefore often unknown. Hence, we cannot decide which of the two detection limits is actually correct. Based on the related work discussed in Section 4.2 we would expect the noise in the speckle-dominated regime, i.e., close to the star, to be better described by a Laplacian. At these separations, the true contrast is likely closer to the contrast curve of the parametric bootstrapping under the assumption of Laplacian noise. At larger separations, the noise becomes more Gaussian and the results of the t -test or Gaussian parametric bootstrapping are likely accurate. It is important to note that we cannot prove if the noise is sufficiently normal. A discussion about this problem is given in Appendix C. Thus, as long as no additional knowledge about the noise is available, we have to accept that our contrast curves are potentially inaccurate by about one magnitude.

8. Summary and Conclusions

In this paper we presented a new framework to quantify exoplanet detections in HCI. Through the use of parametric bootstrapping, our new metric is able to estimate robust detection limits for any type of residual noise. Our metric, however, assumes that the noise distribution is known. A comparison of the detection limits under the assumption of Gaussian and Laplacian noise revealed that commonly used metrics, such as the t -test, might be too optimistic in case of speckle-dominated observations. The occurrence rate of large noise values is higher in case of Laplacian noise compared to Gaussian noise. This results in a higher risk to obtain a high S/N value originating from the noise. For example, the risk that the noise produces a false detection at 2 FWHM distance from the star with a S/N of 5 is about 250 times higher under Laplacian than under Gaussian noise. Therefore, the S/N should not be considered as a direct measure for the detection

uncertainty. Only if we take the sample size and the correct noise distribution into account the S/N becomes interpretable.

The link between the detection limit and the noise characteristics makes a fair comparison between HCI observations difficult. This is especially the case if the noise distributions differ between the data sets we want to compare. An example for this is the development of new postprocessing algorithms. If we compare algorithms that produce residuals with different types of noise, the comparison of the methods is likely biased. The same applies for any comparison between observations taken under different circumstances: e.g., ground-based versus space-based observations, different instruments or observing strategies such as ADI and RDI. If we want to compare inhomogeneous data, we have to take possible biases arising from statistics into account. We recommend computing detection limits under different assumptions to set optimistic and conservative bounds for the achieved contrast. Studies on residual noise using for example Q-Q plots can provide valuable insights about the noise. But they can never prove that the noise is sufficiently normal to use a t -test. Apart from graphical tools, quantitative tests such as the Shapiro-Wilk test can only reject normality. They can never prove that the data are actually normal.

Future work should seek to understand better the speckle statistics of HCI residuals. If we know the true distribution of noise as a function of observing conditions, instrument, and data postprocessing, we could use the bootstrapping algorithm presented in this paper to determine the true contrast. Alternatively, we can use nonparametric methods such as nonparametric bootstrapping to estimate the statistics directly from the data. Future work should further investigate the origin of non-Gaussian noise and develop new methods to account for the spatial dependencies of speckle noise.

Acknowledgments

We thank the anonymous referee for a critical and constructive review of the original manuscript which helped improve the quality of the paper significantly.

This work was supported by an ETH Zurich Research Grant. M.J.B. and S.P.Q. gratefully acknowledge the financial support from ETH Zurich. Parts of this work has been carried out within the framework of the National Centre of Competence in Research PlanetS supported by the Swiss National Science Foundation (SNSF) under grants 51NF40_182901 and 51NF40_205606. Parts of this work were supported by the SNSF via grant No. 200020_200300. S.P.Q., E.O.G., F.A.D., and J.H. acknowledge the financial support of the SNSF. O.A. acknowledges funding from F.R.S.-FNRS, and from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 819155). G.C. thanks the Swiss National Science Foundation for financial support under grant No. P500PT_206785. Author contributions: M.J.B. carried out the main analyses and wrote the manuscript. He further programmed the publicly available the python package `applefy` and wrote the documentation page on ReadTheDocs. E.O.G. proposed the initial idea to explore parametric bootstrapping and pivoting to address the detection problem. T.D.G. contributed to the programming of `applefy`. T.D.G. and G.C. helped to debug the alpha version of the code. All authors discussed the results and commented on the manuscript.

G.C. thanks the Swiss National Science Foundation for financial support under grant number P500PT 206785.

Appendix A Aperture Placement

The first step of the hypothesis test illustrated in Figure 1 is the extraction of the signal and the noise. The position of the signal can be estimated by maximizing the flux inside the signal aperture. The positions for the noise are commonly selected depending on the final signal position. This choice, however, is arbitrary and any other arrangement of nonoverlapping apertures should yield similar results. In order to check if T_{obs} is invariant to the positioning of the apertures we insert fake planets into the β Pictoris data set and perform a PCA-based PSF subtraction. We insert one fake planet at a separation of 2 FWHM and one at 3 FWHM. We measure the flux of the inserted fake planets \bar{Y}_i in each residual image and extract values for the noise $\bar{X}_1, \dots, \bar{X}_n$ by using apertures with the same separation from the star. At 2 FWHM we have $n = 12$ and at 3 FWHM we have $n = 18$ apertures available. Apertures close to the planet signal are excluded in order to account for self-subtraction artifacts next to the planet. This reduces n by about 2 apertures. Next, we compute T_{obs} and the FPF as explained in Section 3. We repeat the calculation several times for slightly rotated aperture positions. In total we rotate all apertures by 60° , which corresponds to a displacement by $60/360 \times n = 2$ apertures at 2 FWHM and 3 apertures at 3 FWHM. The results are summarized in Figure 10.

As shown in the plot the computed FPF changes significantly depending on the rotation of the aperture positions $\Delta\theta$. Especially at 2 FWHM the values of T_{obs} vary by a factor of 2. The corresponding confidence levels span from $3.4\sigma_N$ to $5.0\sigma_N$. At 3 FWHM the effect is smaller but still not negligible. The position of the apertures influences whether speckles end up between or inside the apertures. If a bright speckle is located

between two apertures, its flux will be split and $\hat{\sigma}_\chi$ decreases. Vice versa, if the speckle is in the center of an aperture, $\hat{\sigma}_\chi$ will be large. At small separations only a few apertures are available and a single speckle can already have a large effect on T_{obs} .

We recommend estimating T_{obs} several times for slightly rotated aperture positions. Instead of a 60° total rotation a rotation by $360/(2\pi r)$ degrees is sufficient for a displacement by one aperture (one oscillation; compare Figure 10). The results can be summarized by reporting the median value of T_{obs} over all rotations and the median absolute deviation (MAD) as a measure for the influence of the aperture placement. The minimum and maximum value of T_{obs} give the worst- and best-case results, respectively. The same procedure should be used if spaced pixels are used instead of apertures (see the discussion in Section 4.1).

Appendix B Analytical Contrast Curves

The calculation of one contrast grid presented in Section 6.2 requires one to process several hundred data sets with inserted fake planets. This computation can be very time consuming. Under some mild assumptions, it is possible to reduce this computation time considerably. This is the commonly used approach of how contrast curves are calculated in packages like `PynPoint` (Stolker et al. 2019) or `VIP` (Gomez Gonzalez 2017). We start by rearranging Equation (1)

$$Y_1 = T_{\text{obs}} \cdot \hat{\sigma}_\chi \sqrt{1 + \frac{1}{n}} + \hat{\mu}_\chi. \quad (\text{B1})$$

The value of Y_1 is the brightness of the signal required in the residual image to reach T_{obs} . Given a detection threshold specified as an FPF we can derive the T_{obs} needed to be counted as a detection. For Gaussian noise this can be done by solving Equation (2) for T_{obs} where $p(T = t|H_0)$ is given by the t -distribution. For non-Gaussian noise, larger values of T might be required to reach the same FPF (see Figure 8). We reuse the sorted bootstrap results $T_{(1)}^* \leq \dots \leq T_{(B)}^*$ discussed in step 6 of Section 5.1, and estimate T_{obs} by linear interpolation

$$T_{\text{obs}} = T_{([a])}^*(a - [a]) + T_{([a+1])}^*([a+1] - a), \quad (\text{B2})$$

where $a = (B-1)(1 - \text{FPF})$ gives the index of the two bootstrap results $T_{([a])}^*$ and $T_{([a+1])}^*$, which are closest to the required FPF. This step is the inverse of the linear interpolation explained in step 7 of Section 5.1. Due to planet over- and self-subtraction during the data postprocessing, the flux of the planet in the residual image is attenuated. We can describe this effect by

$$Y_1 = f_p \cdot \kappa(f_p, s) + X_{n+1}, \quad (\text{B3})$$

where X_{n+1} is the speckle noise at the position of the planet and $\kappa(f_p, s) \in [0, 1]$ is the throughput accounting for the attenuation of the data postprocessing and potential coronagraphs. The separation is denoted as s . The throughput can be computed using the following procedure:

1. As for the contrast grid in Section 6.2, insert artificial planets at different separations s with different contrasts $c = f_p/f_*$ into the raw data. Run the data postprocessing to compute their residuals. In order to account for

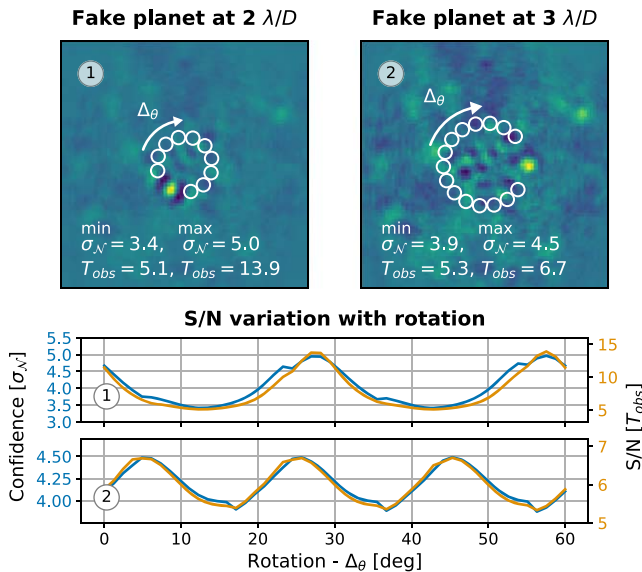


Figure 10. The effect of the aperture placement on the detection uncertainty. The top two images show the residuals with the inserted fake planets at 2 and 3 FWHM. The white circles indicate the initial positions of the apertures used to extract the noise. Below we show the computed test statistic T_{obs} (orange) and detection uncertainty (FPF; blue) for different rotations of the initial noise positions. The difference between the values of T_{obs} and the FPF are due to the number of apertures changing by ± 1 during the rotation.




Figure 11. The top two panels show example residuals for the β Pictoris data set with and without fake planets. The bottom left image gives their difference. The bottom right image shows the result of the same experiment but for a fake planet closer to the star and with 50 PCA components. For 20 PCA components the presence of the planet only affects the values of the residual image in the direct neighborhood of the signal. For 50 PCA components the whole image changes. This is potentially problematic as discussed in the text.

azimuthal variations we insert six planets, one at a time, for each separation and contrast.

2. Compute one residual image without fake planets.
3. Subtract the planet-free residual from every fake planet residual and estimate the flux at the position of the fake planet. Thanks to the linearity of PCA this gives us $Y_1 - X_{n+1}$. Since our statistic is based on pixels spaced by one FWHM, we integrate the flux within an area of one pixel around the position of the planet (compare Section 6.1). We note that this step is only valid for sufficiently faint planets which do not affect the PCA component matrix. It further does not hold for all existing postprocessing techniques.
4. Use Equation (B3) to compute the throughput.

Examples for residuals with and without fake planets are shown in Figure 11. The throughput is summarized in Figure 12. It depends on the brightness of the inserted fake planet f_p as well as its separations from the star. For faint planets, the PCA basis is not changed by the presence of the planet and the throughput converges. Bright planets, on the other hand, influence the PCA basis and cause additional signal loss. Planets which are close to the detection limit are usually faint. That is, under the assumption that these planets do not affect the PCA basis we can simplify $\kappa(f_p, s) = \kappa(s)$ and use the convergence throughput. That is, we only need to compute the last row of Figure 12. The limit for the planet to star contrast can be calculated by

$$c = \frac{T_{\text{obs}} \cdot \hat{\sigma}_\chi \sqrt{1 + 1/n} + \hat{\mu}_\chi - X_{n+1}}{\kappa(s) \cdot f_*}. \quad (\text{B4})$$

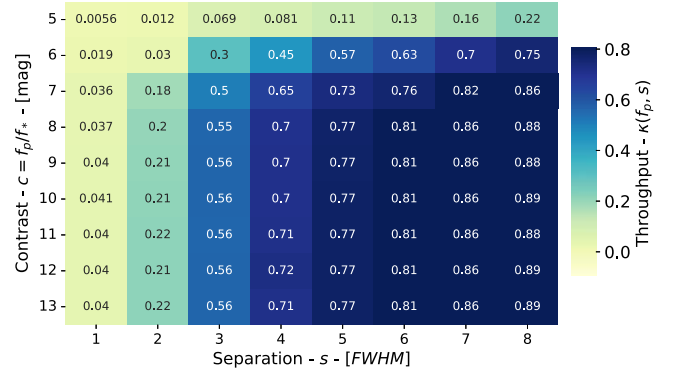


Figure 12. Throughput of the β Pictoris data set for 30 PCA components. Every value in the plot is the average of six experiments with artificial planets inserted at six different azimuthal positions.




Figure 13. Comparison of the analytical contrast curve (based on the approximation that the throughput is only a function of separation) with the complete contrast grid presented in the main paper (Section 6.2). The orange line gives the contrast grid thresholded at $5\sigma_N$. The results are based on the assumption of Gaussian noise.

The values of X_{n+1} , $\hat{\sigma}_\chi$, and $\hat{\mu}_\chi$ are based on noise observations from the planet-free residual. They are again dependent on where the noise gets extracted. We propose the following procedure to account for this effect:

1. Estimate f_* on the unsaturated PSF by integration of the flux within an area of one pixel around the star.
2. Choose a detection threshold, for example $5\sigma_N = 2.87 \times 10^{-7}\text{FPF}$.
3. For each separation extract noise values spaced by 1 FWHM from the planet-free residual. Use one value as X_{n+1} and the rest as \mathcal{X} .
4. Make an assumption about the noise to constrain T_{obs} . For Gaussian noise solve Equation (2) for T_{obs} . For non-Gaussian noise use Equation (B2).
5. Use Equation (B4) to compute the contrast.

6. Repeat steps 3–5 with different noise positions (compare Figure 10). Report the median contrast over all repetitions. The mean absolute deviation from the median can be used as a measure for the uncertainty introduced by the placement of the noise positions.

The given procedure assumes that the planet signal has no effect on the noise sample \mathcal{X} . As shown in the bottom right plot in Figure 11 this is not necessary the case for a large number of PCA components. Under such circumstances, the calculation of a complete contrast grid is favorable. A comparison of the analytical contrast curve with the results of the contrast grid is shown in Figure 13.

As shown in the figure, both methods are consistent. In other words, the detection uncertainty of the artificial planets inserted for the contrast grid agrees with the contrast curve. At small separations, the contrast curve reaches a regime where the PCA basis is changed by the planet. For these separations no $5\sigma_N$ contrast exist, meaning that no planet, no matter how bright, will ever give a $5\sigma_N$ detection. This effect can only be identified with the contrast grid.

Appendix C Testing for Non-Gaussian Noise

The noise distribution of HCI residuals can vary within and between observations. To confirm that the noise is sufficiently normal for the t -test to be applicable, various tests are carried out prior to the calculation of the detection limits. For example Mawet et al. (2014) and Otten et al. (2017) compute histograms based on pixel values within annuli of $1\lambda/D$ width. A visual inspection of histograms can provide first insights about the general distribution shape. A direct comparison with a normal distribution, however, is difficult. This is especially the case for the tails of the distribution, i.e., the occurrence rate of bright events. These events are of special importance for the calculation of the detection limits. Q–Q plots are a good alternative to histograms as they allow one to compare the noise with the normal distribution directly. In this way, insights about a possible over- or underestimation of the contrast can be obtained.

The Shapiro–Wilk test (Shapiro & Wilk 1965) can be used as a quantitative measure for normality. The test can be used to reject the null hypothesis that a sample was drawn from a normal distribution. For the test to be applicable, it is assumed that the values of the sample are independent. Due to the spatial size of the speckles, neighboring pixels in the HCI residuals are never independent. Thus, the test has to be calculated on pixels which are at least 1 FWHM apart. Results computed on nonindependent values (see, e.g., Absil et al. 2013) are likely biased. The output of the test is a p -value corresponding to the remaining risk that the sample is compatible with the null hypothesis. For example, if we observe $p = 0.01$, we know that 1% of the observations under Gaussian noise have a value of the test statistic equal or more extreme than ours. By setting a threshold, we can specify the accepted error by which Gaussian noise is misclassified as non-Gaussian noise. Although the Shapiro–Wilk test has shown superior performance compared to other tests (Razali & Wah 2011), it lacks power in detecting non-Gaussian noise at close separations to the star. We study this effect in Figure 14.

The highest TPR is reached for the threshold $p < 0.2$ at 8 FWHM and is just about 60%. The results get worse for close

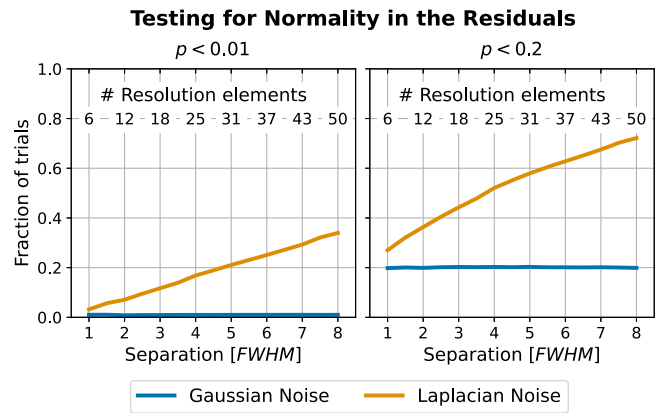


Figure 14. Sensitivity and false-positive rate (FPR) of the Shapiro–Wilk test as a function of separation from the star. For each separation, 10^5 samples are drawn, each containing as many values as independent resolution elements are available. We sample from a Gaussian distribution to measure the FPR (blue) and from a Laplacian distribution to measure the sensitivity (orange). The left and right plots show the results for two different thresholds of $p < 0.01$ and $p < 0.2$, respectively.

separations. This is problematic as speckle noise is expected to be the dominant noise source close to the star. While sensitivity is one problem, preliminary testing comes with a “logical problem” (Scholz & Stephens 1987): “Because insufficient evidence exists to reject normality, normality will be considered true” (Rochon et al. 2012). If we cannot reject the null hypothesis, we must not conclude that our data set is normal, i.e., combined with the low sensitivity for small sample sizes, we can neither prove nor disprove normality. As concluded in Rochon et al. (2012) “support for the assumption of normality must come from extra-data sources” (Easterling & Anderson 1978).

Appendix D Pivot for the Laplace Distribution

Let Z be a random variable following a standard distribution with no unknown parameters (e.g., a standard Gaussian or Laplacian). The random variables $X = Zq_X + w_X \sim \mathcal{F}(q_X, w_X)$ and $Y = Zq_Y + w_Y \sim \mathcal{F}(q_Y, w_Y)$ with $w_X, w_Y \in \mathbb{R}$ and $q_X, q_Y \in \mathbb{R}^+$ form a location-scale family of Z . In the case of a normal distribution $Z \sim \mathcal{N}(0, 1)$, with $q = \sigma$ and $t = \mu$, and $Z \sim \mathcal{L}(0, 1)$ with $q = b$ and $w = \mu$ for the Laplacian distribution.

We consider a two-sample data set $\mathcal{X} = X_1, \dots, X_n \sim \mathcal{F}(q_X, w_X)$ and $\mathcal{Y} = Y_1, \dots, Y_m \sim \mathcal{F}(q_Y, w_Y)$ under the assumption of equal scale $q_X = q_Y = q$. We want to test the null hypothesis $H_0: w_X - w_Y = 0$ against the alternative $H_1: w_Y - w_X > 0$ using the test statistic of the two-sample t -test (see Equation (1) for the special case of $m = 1$ used in HCI)

$$T = \frac{\hat{\mu}_Y - \hat{\mu}_X}{\hat{\sigma}_{X,Y} \sqrt{1/n + 1/m}}, \quad (\text{D1})$$

where $\hat{\mu}_X$ and $\hat{\mu}_Y$ are the sample averages

$$\hat{\mu}_X = \bar{X} = \frac{1}{n} \sum_{i=1}^n X_i, \quad \hat{\mu}_Y = \bar{Y} = \frac{1}{m} \sum_{j=1}^m Y_j, \quad (\text{D2})$$

and $\hat{\sigma}_{\mathcal{X},\mathcal{Y}}$ is the pooled standard deviation of the two samples

$$\hat{\sigma}_{\mathcal{X},\mathcal{Y}}^2 = \frac{(n-1)\hat{\sigma}_{\mathcal{X}}^2 + (m-1)\hat{\sigma}_{\mathcal{Y}}^2}{n+m-2}, \quad (\text{D3})$$

with

$$\hat{\sigma}_{\mathcal{X}}^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2, \quad (\text{D4})$$

$$\hat{\sigma}_{\mathcal{Y}}^2 = \frac{1}{m-1} \sum_{j=1}^m (Y_j - \bar{Y})^2. \quad (\text{D5})$$

Under H_0 the test statistic T follows a distribution that is independent of the parameters $w_{\mathcal{X}}$, $w_{\mathcal{Y}}$, q , i.e., it is a pivot. The following proof is inspired by exercise 9.9 in Casella & Berger (2002).

Proof:

$$\begin{aligned} \hat{\sigma}_{\mathcal{X}}^2 &= \frac{1}{n-1} \sum_{i=1}^n (Z_i q_{\mathcal{X}} + w_{\mathcal{X}} - \bar{Z} q_{\mathcal{X}} - w_{\mathcal{X}})^2 \\ &= \frac{q_{\mathcal{X}}^2}{n-1} \sum_{i=1}^n (Z_i - \bar{Z})^2 \\ &= q_{\mathcal{X}}^2 \hat{\sigma}_{Z_n}^2. \end{aligned}$$

The same holds for $\hat{\sigma}_{\mathcal{Y}}^2 = q_{\mathcal{Y}}^2 \hat{\sigma}_{Z_m}^2$. With $q_{\mathcal{X}} = q_{\mathcal{Y}} = q$ it follows

$$\begin{aligned} \hat{\sigma}_{\mathcal{X},\mathcal{Y}}^2 &= \frac{q^2(n-1)\hat{\sigma}_{Z_n}^2 + q^2(m-1)\hat{\sigma}_{Z_m}^2}{n+m-2} \\ &= q^2 \hat{\sigma}_{Z_n Z_m}^2. \end{aligned}$$

We rewrite the test statistic T

$$\begin{aligned} T &= \frac{\frac{1}{n} \sum_{i=1}^n X_i - \frac{1}{m} \sum_{j=1}^m Y_j}{\hat{\sigma}_{\mathcal{X},\mathcal{Y}} \sqrt{(1/n + 1/m)}} \\ &= \frac{\frac{1}{n} \sum_{i=1}^n (Z_i q + w_{\mathcal{X}}) - \frac{1}{m} \sum_{j=1}^m (Z_j q + w_{\mathcal{Y}})}{\hat{\sigma}_{\mathcal{X},\mathcal{Y}} \sqrt{(1/n + 1/m)}} \\ &= \frac{q \left(\frac{1}{n} \sum_{i=1}^n Z_i - \frac{1}{m} \sum_{j=1}^m Z_j \right)}{q \hat{\sigma}_{Z_n Z_m} \sqrt{(1/n + 1/m)}} + \frac{w_{\mathcal{X}} - w_{\mathcal{Y}}}{q \hat{\sigma}_{Z_n Z_m} \sqrt{(1/n + 1/m)}}. \end{aligned}$$

Under H_0 we have $w_{\mathcal{X}} = w_{\mathcal{Y}}$. It follows

$$T = \frac{\bar{Z}_n - \bar{Z}_m}{\hat{\sigma}_{Z_n Z_m} \sqrt{(1/n + 1/m)}}. \quad (\text{D6})$$

All quantities \bar{Z}_n , \bar{Z}_m , $\hat{\sigma}_{Z_n Z_m}$ are independent of $w_{\mathcal{X}}$, $w_{\mathcal{Y}}$, q . \square

Appendix E Speckle Noise Correlation

A key assumption of the hypothesis tests discussed in this paper is that the observations in \mathcal{Y} and \mathcal{X} are independent. This assumption is fundamental for both the t -test as well as the parametric bootstrapping. In Section 4, we assert that pixel values spaced by 1 FWHM are approximately independent. This, however, is only the case if the actual correlations in the data follow the shape of the unsaturated PSF. In the following, we compute the spatial correlations of the noise in the β Pictoris data set processed with PCA (Amara & Quanz 2012; Soummer et al. 2012) and cADI (Marois et al. 2006):

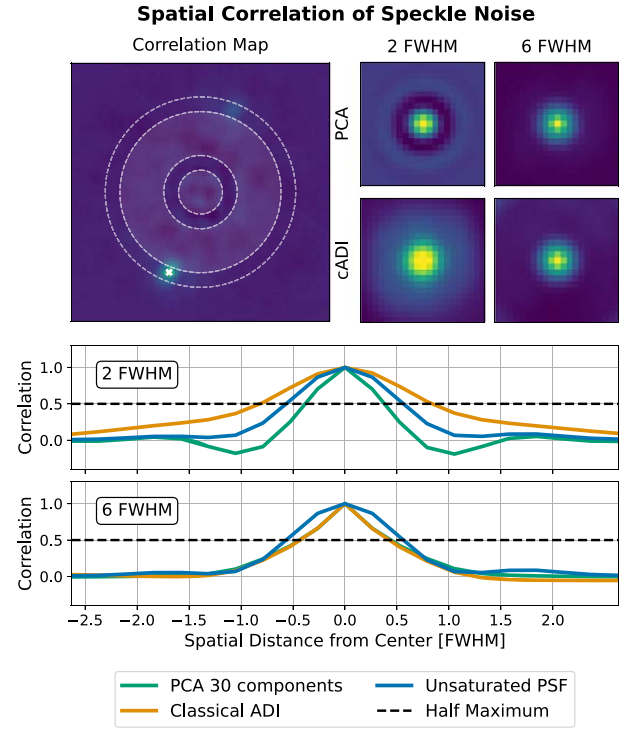


Figure 15. Estimates of the spatial noise correlations of the β Pictoris data set. The image at the top left is an example of a correlation map. It shows how much the pixel marked with the white cross is correlated w.r.t. all other pixels in the image. We compute one correlation map for every pixel in the residual image. The top-right images show the local averages of the correlation maps at 2 FWHM and 6 FWHM for cADI and PCA. The regions used to compute the averages are marked in the correlation map. The bottom two plots give the 1D profile of the local correlation maps (top-right plots) in comparison to the unsaturated PSF of the star.



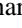

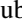






1. We start the analysis with the preprocessed stack of images in which the star is located in the center of the frames. For cADI we subtract the median frame from every image in the stack. For PCA we subtract the first 30 PCA components. We derotate the frames according to their parallactic angles. The result is a sequence of residual images, which we average to obtain the final residual image.
2. We compute the correlation of every pixel along the temporal domain w.r.t. all other pixels. For this we use the sequence of derotated residuals after PSF subtraction. This way we obtain a single 2D correlation map for each pixel.
3. We crop the correlation maps such that the pixels to which they correspond are in the center of the cropped map. The resulting maps show the local correlations for each pixel in the residual frame.
4. We average the cropped correlation maps for all pixels within an annulus at 2 FWHM and 6 FWHM distances from the star.

The results and a comparison with the unsaturated PSF of the star is shown in Figure 15.

We observe that the true correlations in the data do not always follow the shape of the PSF. Further, they are influenced by the separation from the star as well as the data postprocessing used. For cADI the FWHM of the local correlation map at 2 FWHM is about 60% larger compared to the PSF FWHM. Similar results were previously observed

by Golomb et al. (2021) for data from the Gemini Planet Imager. If the spatial correlation length is large, a spacing of 1 PSF FWHM might not be sufficient to ensure that the noise observations in \mathcal{X} are independent. At 6 FWHM the FWHM of the local correlations is slightly smaller than the PSF FWHM. If we use PCA, the local correlations at 2 FWHM change while they stay the same at 6 FWHM. Future work should further investigate this behavior to understand better the spatial correlations of speckle noise in real data. A possible path for this could be bootstrapping for dependent observations (Zoubir & Iskander 2004).

ORCID iDs

Markus J. Bonse  <https://orcid.org/0000-0003-2202-1745>
 Emily O. Garvin  <https://orcid.org/0000-0003-2530-9330>
 Timothy D. Gebhard  <https://orcid.org/0000-0001-9310-8579>
 Felix A. Dannert  <https://orcid.org/0000-0002-5476-2663>
 Faustine Cantalloube  <https://orcid.org/0000-0002-3968-3780>
 Gabriele Cugno  <https://orcid.org/0000-0001-7255-3251>
 Olivier Absil  <https://orcid.org/0000-0002-4006-6237>
 Jean Hayoz  <https://orcid.org/0000-0003-3768-5712>
 Julien Milli  <https://orcid.org/0000-0001-9325-2511>
 Markus Kasper  <https://orcid.org/0000-0002-8425-6606>
 Sascha P. Quanz  <https://orcid.org/0000-0003-3829-7412>

References

- Absil, O., Milli, J., Mawet, D., et al. 2013, *A&A*, 559, L12
 Aime, C., & Soummer, R. 2004, *ApJL*, 612, L85
 Amara, A., & Quanz, S. P. 2012, *MNRAS*, 427, 948
 Bowens, R., Meyer, M. R., Delacroix, C., et al. 2021, *A&A*, 653, A8
 Brandl, B., Bettonvil, F., van Boekel, R., et al. 2021, *Msngr*, 182, 22
 Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. 2015, *A&A*, 582, A89
 Cantalloube, F., Farley, O., Milli, J., et al. 2020, *A&A*, 638, A98
 Cantalloube, F., Gomez-Gonzalez, C., Absil, O., et al. 2020, *Proc. SPIE*, 11448, 114485A
 Casella, G., & Berger, R. L. 2002, *Statistical Inference* (2nd edn; Pacific Grove, CA: Thomson Learning)
 Dahlqvist, C.-H., Cantalloube, F., & Absil, O. 2020, *A&A*, 633, A95
 Delacroix, C., Absil, O., Forsberg, P., et al. 2013, *A&A*, 553, A98
 Desidera, S., Chauvin, G., Bonavita, M., et al. 2021, *A&A*, 651, A70
 Easterling, R. G., & Anderson, H. E. 1978, *JSCS*, 8, 1
 Efron, B., & Tibshirani, R. 1994, *An Introduction to the Bootstrap* (New York: CRC Press)
 Fitzgerald, M. P., & Graham, J. R. 2006, *ApJ*, 637, 541
 Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2018, *A&A*, 618, A138
 Gaudi, B. S., Seager, S., Mennesson, B., et al. 2018, arXiv:1809.09674
 Gebhard, T. D., Bonse, M. J., Quanz, S. P., & Schölkopf, B. 2022, *A&A*, 666, A9
 Golomb, J., Rocha, G., Meshkat, T., et al. 2021, *AJ*, 162, 304
 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
 Gomez Gonzalez, C. A., Wertz, O., Absil, O., et al. 2017, *AJ*, 154, 7
 Hinkley, S., Carter, A. L., Ray, S., et al. 2022, *PASP*, 134, 095003
 Hunziker, S., Quanz, S. P., Amara, A., & Meyer, M. R. 2018, *A&A*, 611, A23
 Jensen-Clem, R., Mawet, D., Gomez Gonzalez, C. A., et al. 2017, *AJ*, 155, 19
 Kasper, M., Urna, N. C., Pathak, P., et al. 2021, *Msngr*, 182, 38
 Kenworthy, M. A., Snik, F., Keller, C. U., et al. 2018, *Proc. SPIE*, 10702, 1070246
 Kuhn, J. R., Potter, D., & Parise, B. 2001, *ApJL*, 553, L189
 Lafrenière, D., Marois, C., Doyon, R., & Barman, T. 2009, *ApJL*, 694, L148
 Langlois, M., Gratton, R., Lagrange, A.-M., et al. 2021, *A&A*, 651, A71
 Larsen, R. J., & Marx, M. L. 2012, *An Introduction to Mathematical Statistics and Its Applications* (5th edn.; Boston, MA: Prentice Hall)
 Marois, C., Correia, C., Galicher, R., et al. 2014, *Proc. SPIE*, 9148, 91480U
 Marois, C., Lafreniere, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, 641, 556
 Marois, C., Lafreniere, D., Macintosh, B., & Doyon, R. 2008, *ApJ*, 673, 647
 Marois, C., Macintosh, B., & Véran, J.-P. 2010, *Proc. SPIE*, 7736, 77361J
 Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, 792, 97
 Mesa, D., Gratton, R., Zurlo, A., et al. 2015, *A&A*, 576, A121
 Meshkat, T., Bailey, V., Rameau, J., et al. 2013, *ApJL*, 775, L40
 Morzinski, K. M., Males, J. R., Skemer, A. J., et al. 2015, *ApJ*, 815, 108
 Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019, *AJ*, 158, 13
 Norton, R. M. 1984, *Am. Stat.*, 38, 135
 Otten, G. P. P. L., Snik, F., Kenworthy, M. A., et al. 2017, *ApJ*, 834, 175
 Pairet, B., Cantalloube, F., Gomez Gonzalez, C. A., Absil, O., & Jacques, L. 2019, *MNRAS*, 487, 2262
 Perrin, M. D., Sivaramakrishnan, A., Makidon, R. B., Oppenheimer, B. R., & Graham, J. R. 2003, *ApJ*, 596, 702
 Quanz, S. P., Crossfield, I., Meyer, M. R., Schmalzl, E., & Held, J. 2015, *IJAsB*, 14, 279
 Quanz, S. P., Schmid, H. M., Geissler, K., et al. 2011, *ApJ*, 738, 23
 Quanz, S. P., Ottiger, M., Fontanet, E., et al. 2022, *A&A*, 664, A21
 Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, *PASP*, 111, 587
 Rameau, J., Chauvin, G., Lagrange, A.-M., et al. 2013, *ApJL*, 772, L15
 Razali, N. M., & Wah, Y. B. 2011, *JSMA*, 2, 21
 Rochon, J., Gondan, M., & Kieser, M. 2012, *BMC Med. Res. Methodol.*, 12, 81
 Rousset, G., Lacombe, F., Puget, P., et al. 2003, *Proc. SPIE*, 4839, 140
 Ruane, G., Ngo, H., Mawet, D., et al. 2019, *AJ*, 157, 118
 Ruffio, J.-B., Macintosh, B., Wang, J. J., et al. 2017, *ApJ*, 842, 14
 Ruffio, J.-B., Mawet, D., Czekala, I., et al. 2018, *AJ*, 156, 196
 Samland, M., Bouwman, J., Hogg, D. W., et al. 2021, *A&A*, 646, A24
 Scholz, F. W., & Stephens, M. A. 1987, *JASA*, 82, 918
 Shapiro, S., & Wilk, M. 1965, *Biometrika*, 52, 591
 Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007, *ApJ*, 669, 642
 Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJL*, 755, L28
 Sparks, W. B., & Ford, H. C. 2002, *ApJ*, 578, 543
 Stolter, T., Bonse, M. J., Quanz, S. P., et al. 2019, *A&A*, 621, A59
 The LUVVOIR Team 2019, arXiv:1912.06219
 Thompson, W., & Marois, C. 2021, *AJ*, 161, 236
 Uyama, T., Hashimoto, J., Kuzuhara, M., et al. 2017, *AJ*, 153, 106
 Vigan, A., Fontanive, C., Meyer, M., et al. 2021, *A&A*, 651, A72
 Wahhaj, Z., Liu, M. C., Biller, B. A., et al. 2013, *ApJ*, 779, 80
 Wahhaj, Z., Cieza, L. A., Mawet, D., et al. 2015, *A&A*, 581, A24
 Xuan, W. J., Mawet, D., Ngo, H., et al. 2018, *AJ*, 156, 156
 Zoubir, A. M., & Iskander, D. R. 2004, *Bootstrap Techniques for Signal Processing* (Cambridge: Cambridge Univ. Press)