



Use the 4S (Signal-Safe Speckle Subtraction): Explainable Machine Learning Reveals the Giant Exoplanet AF Lep b in High-contrast Imaging Data from 2011

Markus J. Bonse^{1,2} , Timothy D. Gebhard^{1,2,3} , Felix A. Dannert^{1,4} , Olivier Absil⁵ , Faustine Cantalloube⁶ ,
Valentin Christiaens^{5,7} , Gabriele Cugno^{1,8} , Emily O. Garvin¹ , Jean Hayoz¹ , Markus Kasper⁹ , Elisabeth Matthews¹⁰ ,
Bernhard Schölkopf^{2,11} , and Sascha P. Quanz^{1,4,12}

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

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

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

⁴ National Center of Competence in Research PlanetS, Switzerland

⁵ University of Liège, STAR Institute, Allée du Six Août 19C, 4000 Liège, Belgium

⁶ Université Grenoble Alpes, CNRS, IPAG, 38000 Grenoble, France

⁷ KU Leuven, Institute for Astronomy, Celestijnenlaan 200D, Leuven, Belgium

⁸ University of Michigan, Department of Astronomy, Ann Arbor, MI 48109, USA

⁹ European Southern Observatory, Garching bei München, Germany

¹⁰ Max Planck Institute for Astronomy, Königstuhl 17, 69117 Heidelberg, Germany

¹¹ ETH Zurich, Department of Computer Science, Universitätsstrasse 6, 8092 Zurich, Switzerland

¹² ETH Zurich, Department of Earth and Planetary Sciences, Sonneggstrasse 5, 8092 Zurich, Switzerland

Received 2024 June 3; revised 2024 November 29; accepted 2025 January 15; published 2025 March 5

Abstract

The main challenge of exoplanet high-contrast imaging (HCI) is to separate the signal of exoplanets from their host stars, which are many orders of magnitude brighter. HCI for ground-based observations is further exacerbated by speckle noise originating from perturbations in Earth's atmosphere and imperfections in the telescope optics. Various data postprocessing techniques are used to remove this speckle noise and reveal the faint planet signal. Often, however, a significant part of the planet signal is accidentally subtracted together with the noise. In the present work, we use explainable machine learning to investigate the reason for the loss of the planet signal for one of the most used postprocessing methods: principal component analysis (PCA). We find that PCA learns the shape of the telescope point-spread function for high numbers of PCA components. This representation of the noise captures not only the speckle noise but also the characteristic shape of the planet signal. Building on these insights, we develop a new postprocessing method (4S) that constrains the noise model to minimize this signal loss. We apply our model to 11 archival HCI data sets from the Very Large Telescope NACO instrument in the L' band and find that our model consistently outperforms PCA. The improvement is largest at close separations to the star ($\leq 4\lambda/D$), providing up to 1.5 mag deeper contrast. This enhancement enables us to detect the exoplanet AF Lep b in data from 2011, 11 yr before its subsequent discovery. We present updated orbital parameters for this object.

Unified Astronomy Thesaurus concepts: Direct imaging (387); Astronomy data reduction (1861); Exoplanets (498); Interdisciplinary astronomy (804); High angular resolution (2167); Astronomy image processing (2306)

1. Introduction

The primary challenge of exoplanet high-contrast imaging (HCI) is to separate and detect the light of faint exoplanets from their host stars, which are many orders of magnitude brighter. During an observation, the light of the star gets scattered in the image plane, creating a halo of highly structured speckle noise superimposed on the planet signal. Speckles tend to mimic the expected shape of the planet, which makes it difficult to distinguish them from a real planet signal. There are two types of speckle noise (J. R. Males et al. 2021): quasi-static speckles, caused by nanometer-scale imperfections in the telescope optics, and atmospheric speckles, resulting from the turbulence of Earth's atmosphere. Over the past decade, considerable efforts have been made to push the limits of HCI observations and reduce speckle noise. In this context, one key aspect is the development of faster and more accurate extreme adaptive optics (XAO) systems. Examples of such systems are the

Gemini Planet Imager (GPI), the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE), the Subaru Coronagraphic Extreme Adaptive Optics (SCExAO), and MagAO-X (B. Macintosh et al. 2014; N. Jovanovic et al. 2015; J. R. Males et al. 2018; J. L. Beuzit et al. 2019). XAO uses new detectors, advanced wave-front sensors, and deformable mirrors with higher actuator counts, along with faster computing hardware and control software. For HCI, they are often coupled with different types of coronagraphs to efficiently suppress the stellar point-spread function (PSF; M. A. Kenworthy et al. 2007; D. Mawet et al. 2009; F. Snik et al. 2012; M. N'Diaye et al. 2015; G. P. P. L. Otten et al. 2017).

Equally crucial for achieving deeper contrast and for removing speckle noise is the development of postprocessing techniques (see F. Cantalloube et al. 2021, for an overview). Postprocessing techniques are combined with dedicated observing strategies like angular differential imaging (ADI; C. Marois et al. 2006), spectral differential imaging (SDI; R. Racine et al. 1999; W. B. Sparks & H. C. Ford 2002), reference star differential imaging (RDI; D. Mawet et al. 2012; G. Ruane et al. 2019), or polarimetric differential imaging (PDI; J. R. Kuhn et al. 2001; S. P. Quanz et al. 2011), which are designed to



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

enable planet signals to be distinguished from speckle noise. A major breakthrough for the processing of HCI data was the introduction of principal component analysis (PCA/KLIP; A. Amara & S. P. Quanz 2012; R. Soummer et al. 2012) to estimate the correlated speckle noise. Since its adoption in HCI, PCA has formed the basis for the analysis of many HCI surveys (E. L. Nielsen et al. 2019; R. Launhardt et al. 2020; M. Janson et al. 2021; M. Langlois et al. 2021; G. Cugno et al. 2023) and allowed the detection of several new planets (B. Macintosh et al. 2015; G. Chauvin et al. 2017).

However, after more than a decade of HCI observations, only a relatively small number of objects have been found (see T. Currie et al. 2023 for an overview and a table of substellar companions), and the results of several surveys have consistently shown that super-Jovian planets with masses between $5M_J$ and $13M_J$ at large separations (10–100 au) are extremely rare (between $\sim 0.3\%$ and 2.2% depending on stellar type; see A. Vigan et al. 2021). Only recently, taking advantage of the latest generation of instruments, has it been possible to search for less massive planets at closer separations. While the sample size is still small, new planets have been detected in this regime (e.g., 51 Eri b, B. Macintosh et al. 2015; AF Lep b, R. J. De Rosa et al. 2023; K. Franson et al. 2023; D. Mesa et al. 2023), hinting at an elevated occurrence rate toward smaller masses and separations. On the other hand, radial velocity (RV) studies have found that the occurrence rate of giant planets peaks around $3.6^{+2.0}_{-1.8}$ au, decreasing closer to the star and farther out (B. J. Fulton et al. 2021). However, the decrease beyond 5–8 au has not been confirmed and is still the subject of ongoing research (A. M. Lagrange et al. 2023). It is therefore crucial to push down the limits of HCI observations close to the star to bridge the gap and create an overlap between RV and HCI surveys. This will provide empirical access to an important parameter space for exoplanet population and formation models and allow for direct imaging follow-up of exoplanets detected by RV surveys.

Unfortunately, close to the star, it becomes increasingly difficult to suppress speckle noise using data postprocessing techniques. If we use methods like PCA, a significant fraction of the planet signal is accidentally subtracted together with the noise. While at large separations the loss of planet signal is moderate, close to the star more than 95% of the planet signal might get lost (see T. D. Gebhard et al. 2022; M. J. Bonse et al. 2023). This effect is caused by the fact that the speckle noise and the signal of the planet share the same morphology. Early work on the LOCI postprocessing method attempted to overcome this problem by using different data masking strategies (C. Marois et al. 2010; R. Soummer et al. 2011). Later, forward modeling approaches were introduced, leading to a better understanding of the effect, now known as planet oversubtraction and self-subtraction (L. Pueyo 2016). Despite these efforts, the loss of planet signal remains a major bottleneck in the postprocessing of HCI data.

Contributions. In this paper, we gain a deep understanding of the reason for the loss of planet signal in PCA and related methods and use it to develop a new postprocessing algorithm. To this end, we revisit the mathematical background of PCA and give an overview of more recent postprocessing methods in Section 2. By using methods from the field of explainable machine learning, we gain additional insights about planet oversubtraction and self-subtraction (Section 3), which we use to develop our new postprocessing algorithm 4S (Section 4).

Unlike previous works that mask the data, 4S uses a mask applied to the model parameters, as well as a new loss function and different regularization strategies. A detailed analysis of the contrast performance of 4S on archival data from the Very Large Telescope (VLT) NACO instrument is given in Section 5. We find that 4S reaches significantly deeper contrast at close separations to the star and learns a model of the noise that matches our understanding of the speckle noise. Using 4S, we can reveal the faint exoplanet AF Lep b in archival VLT-NACO data from 2011 (Section 6), which is not detectable with PCA. In Sections 7 and 8, we outline directions for future work and provide concluding remarks.

Reproducibility. Along with this paper, we are releasing our algorithm 4S as a Python package called `fours`. The code is available on [GitHub](#) and comes with a documentation page on [ReadtheDocs](#). All raw data (M. J. Bonse 2024b)¹³ and intermediate results (M. J. Bonse 2024a)¹⁴ to reproduce our plots are available on Zenodo.

2. PCA Revisited

The concept of using PCA (or Karhunen–Loève transformation) in HCI is to estimate the stellar PSF and speckle noise to subtract it from the science data. For this purpose, PCA explores spatial covariances utilizing the time dimension and brightness of the speckles. A typical HCI data set consists of a sequence of individual exposures $\mathbf{x}_{t,\text{raw}} \in \mathbb{R}^D$, where $D = n_i \cdot n_j$ is the number of pixels per frame and $t = 1, \dots, N$ is the temporal index of the frames. If we ignore background and detector noise, each image is composed of light from the star s_t and possibly light from the planet p_t :

$$\mathbf{x}_{t,\text{raw}} = \mathbf{s}_t + \mathbf{p}_t. \quad (1)$$

The stellar light s_t outshines the planet signal p_t by several orders of magnitude. As we go through the temporal stack of science images, the PSF and speckle pattern changes. This means that the strongest brightness variations in the data are due to the light of the star and not to the planet. This data characteristic is the foundation of speckle subtraction with PCA: PCA is a linear method that tries to find a low-dimensional representation of the data with dimension $K \ll D$, which only captures the strong covariances of the data. If the planet signal is weak compared to the speckle noise, it should thus not be part of this low-dimensional representation. If we reconstruct the data from the low-dimensional space, we get an estimate of the speckle noise that can be subtracted from the science data. It is important to note that PCA only utilizes the time dimension to explore spatial covariances. It does not use the temporal evolution of the speckle noise. It is also possible to apply PCA along the time dimension as discussed by M. Samland et al. (2021) and J. D. Long et al. (2023). More details can be found in the related work in Section 2.3. Applying PCA to ADI data in HCI along the spatial dimension involves the following steps (see Figure 1):

1. *Data preparation.* The temporal average of the science data $\bar{\mathbf{x}}$ is subtracted from each individual science frame:

$$\mathbf{x}_t = \mathbf{x}_{t,\text{raw}} - \bar{\mathbf{x}}. \quad (2)$$

¹³ doi:10.5281/zenodo.11456704.

¹⁴ doi:10.5281/zenodo.11457071.

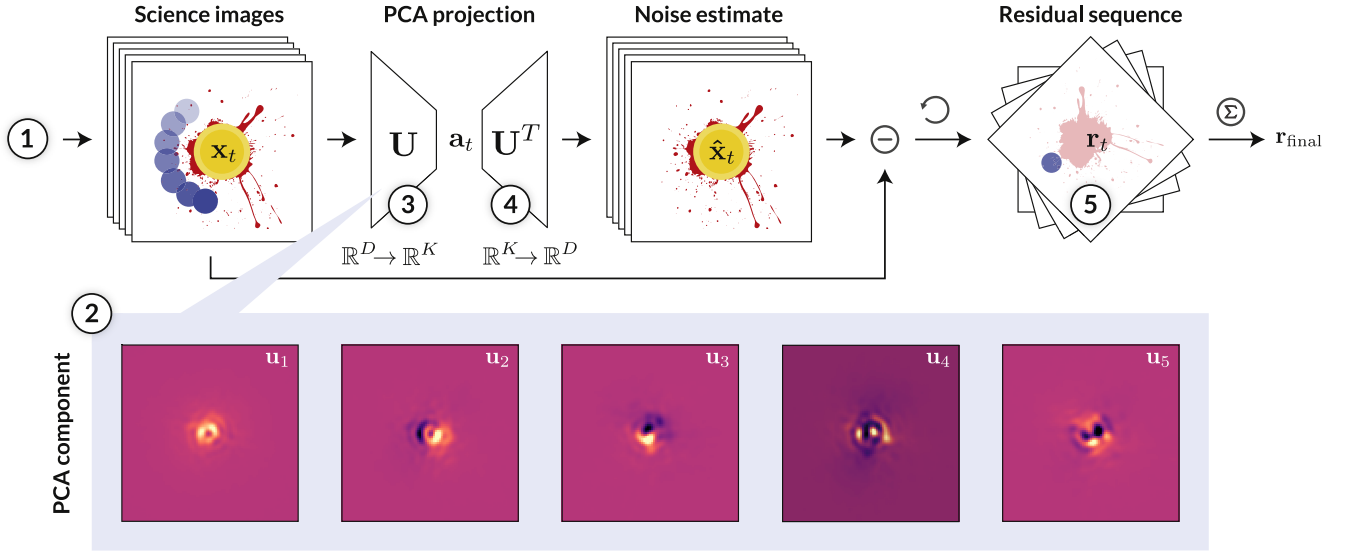


Figure 1. Top row: illustration of how PCA is applied to ADI data in HCI. The planetary signal is shown in blue, while the stellar speckle noise is represented by the red splash. Details about the steps (1) to (5) are given in the text. Bottom row: examples of the first five principal components are shown, estimated for data set #5 (HD 22049) in Table 1. Further details on the data are given in Section 5. The first components of PCA are often interpretable. For example, u_1 in the given example models the overall brightness variations of the postcoronagraphic PSF. The components u_2 and u_3 represent horizontal and vertical variations caused by small errors in the centering of the star behind the coronagraph.

The mean-subtracted 2D science frames are then flattened into one-dimensional vectors and stacked into a data matrix $X \in \mathbb{R}^{N \times D}$:

$$X = \begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_N^T \end{bmatrix} = \begin{bmatrix} x_{1,1} & x_{1,2} & \dots & x_{1,D} \\ x_{2,1} & x_{2,2} & \dots & x_{2,D} \\ \vdots & \vdots & \ddots & \vdots \\ x_{N,1} & x_{N,2} & \dots & x_{N,D} \end{bmatrix}. \quad (3)$$

Each row in X corresponds to one mean-subtracted science frame.¹⁵

2. *Calculation of the projection.* We choose the dimensionality of the low-dimensional space K and calculate the PCA basis (i.e., the component matrix $U \in \mathbb{R}^{D \times K}$), which transforms our data into the low-dimensional space:

$$U = [\mathbf{u}_1 \dots \mathbf{u}_K]. \quad (4)$$

The column vectors of U are the principal components $\mathbf{u}_1, \dots, \mathbf{u}_K$, $\forall \mathbf{u} \in \mathbb{R}^D$. Details on how to calculate these components are given in Section 2.1.

3. *Dimensionality reduction.* We project our data into the low-dimensional space $\mathbf{x}_t \rightarrow \mathbf{a}_t$, $\mathbf{a}_t \in \mathbb{R}^K$:

$$X \cdot U = \begin{bmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \\ \vdots \\ \mathbf{a}_N^T \end{bmatrix} = A \in \mathbb{R}^{N \times K}. \quad (5)$$

If the dimensionality K is chosen carefully, the low-dimensional representation A should capture the variations caused by the star only, and not the signal of the planet.

4. *Data reconstruction.* We reconstruct our data from the low-dimensional space $\mathbf{a}_t \rightarrow \hat{\mathbf{x}}_t$, $\hat{\mathbf{x}}_t \in \mathbb{R}^D$:

$$A \cdot U^T = X \cdot U U^T = \begin{bmatrix} \hat{\mathbf{x}}_1^T \\ \hat{\mathbf{x}}_2^T \\ \vdots \\ \hat{\mathbf{x}}_N^T \end{bmatrix} = \hat{X}. \quad (6)$$

Since \mathbf{a}_t mostly captures the strong covariances of the data, we can use its reconstruction to approximate the speckle noise:

$$\hat{\mathbf{x}}_t \approx \mathbf{s}_t. \quad (7)$$

5. *Calculation of the residual.* We subtract the noise estimate $\hat{\mathbf{x}}_t$ for each mean-subtracted science frame \mathbf{x}_t . The result is the so-called residual sequence \mathbf{r}_t :

$$\mathbf{r}_t = \phi_t(\mathbf{x}_t - \hat{\mathbf{x}}_t), \quad (8)$$

where ϕ_t is a de-rotation function that uses the parallactic angles to realign north in all images (ADI; C. Marois et al. 2006). The final residual image $\mathbf{r}_{\text{final}}$ can then be obtained by averaging along the time axis:¹⁶

$$\mathbf{r}_{\text{final}} = \frac{1}{N} \sum_{t=1}^N \mathbf{r}_t. \quad (9)$$

2.1. How to Find the Optimal Projection with PCA?

Depending on the type of the data set, the component matrix U can be calculated based on the science data directly (e.g., in the case of ADI; A. Amara & S. P. Quanz 2012; R. Soummer et al. 2012) or based on reference stars (RDI; G. Ruane et al. 2019). In this work, we focus on ADI, but the same concepts apply for RDI. As motivated in the previous section, we want

¹⁵ We choose the notation to be consistent with C. M. Bishop (2006) and our Python implementation. All vectors in this paper (e.g., \mathbf{x} , \mathbf{a} , and \mathbf{u}) are column vectors. This notation might deviate from the conventions used in other papers (e.g., B. Lewis et al. 2023).

¹⁶ Some implementations take the median along time instead of the mean, which can be advantageous in the case of detector artifacts such as bad pixels.




Figure 2. Variation in planet visibility as a function of principal components ($K = 20$, $K = 60$, $K = 120$, $K = 300$). The shown data set is data set #5 (HD 22049) in Table 1. More details on data preprocessing are given in Section 5. Two fake planets have been inserted into the data: one positioned at $\sim 3\lambda/D$ separation from the star with a contrast brightness of 9 mag (white circle), and one at $\sim 10\lambda/D$ with a contrast brightness of 12.5 mag (white dashed circle). The inner planet reaches the highest S/N at $\sim K = 60$ components (left panel, solid line) and the outer planet at $\sim K = 300$ components (left panel, dashed line). This example highlights the challenge of choosing the right number of principal components within a single data set.

to choose \mathbf{U} such that the strongest variations of the data, that is, the speckle noise, are captured in the low-dimensional space. This objective is equivalent to the maximum variance derivation of PCA (see, e.g., chapter 12.1 of C. M. Bishop 2006, for details).

For simplicity, let us consider the projection onto the first component \mathbf{u}_1 . We want to choose \mathbf{u}_1 such that the variance along the first axis of the low-dimensional space $\sigma^2(a_{t,1})$ is large:

$$\sigma^2(a_{t,1}) = \mathbf{u}_1^T \mathbf{C} \mathbf{u}_1, \quad (10)$$

where $\mathbf{C} \in \mathbb{R}^{D \times D}$ is the covariance matrix of the data

$$\mathbf{C} = \frac{1}{T} \mathbf{X}^T \mathbf{X}. \quad (11)$$

We add another constraint $\|\mathbf{u}_1\| = \mathbf{u}_1^T \mathbf{u}_1 = 1$ to avoid $\|\mathbf{u}_1\| \rightarrow \infty$. To enforce this constraint, we introduce a Lagrange multiplier λ_1 , which gives us the following optimization objective:

$$\arg \max_{\mathbf{u}_1} \mathbf{u}_1^T \mathbf{C} \mathbf{u}_1 + \lambda_1 (1 - \mathbf{u}_1^T \mathbf{u}_1). \quad (12)$$

If we set the derivative with respect to \mathbf{u}_1 equal to zero, we get

$$\mathbf{C} \mathbf{u}_1 = \lambda_1 \mathbf{u}_1. \quad (13)$$

This means that \mathbf{u}_1 should be chosen to be an eigenvector of \mathbf{C} . Using $\mathbf{u}_1^T \mathbf{u}_1 = 1$, we see that

$$\sigma^2(a_{t,1}) = \mathbf{u}_1^T \mathbf{C} \mathbf{u}_1 = \lambda_1. \quad (14)$$

To maximize the variance of the projection, we should therefore choose the eigenvector \mathbf{u} with the largest eigenvalue. We sort all eigenvectors of \mathbf{C} by their eigenvalues:

$$\begin{array}{ccccccc} \lambda_1 & > & \lambda_2 & > & \dots & > & \lambda_K & > & \dots & > & \lambda_D \\ \mathbf{u}_1 & & \mathbf{u}_2 & & \dots & & \mathbf{u}_K & & \dots & & \mathbf{u}_D, \end{array} \quad (15)$$

and select the first K eigenvectors with the largest eigenvalues. In practice, \mathbf{U} can be obtained by computing a singular value decomposition (SVD) of the data matrix \mathbf{X} (see, e.g., implementations in VIP and PynPoint; T. Stolker et al. 2019; V. Christiaens et al. 2023). This approach improves numerical stability and often reduces computation time.

A big advantage of PCA is that the first components of the matrix \mathbf{U} can be interpreted visually (see Figure 1). For example, the first component \mathbf{u}_1 often models the general brightness

variations of the stellar PSF. The data used to calculate the components in Figure 1 were taken with a vortex coronagraph (D. Mawet et al. 2013). Thus, the first component shows the typical donut shape of the postcoronagraphic PSF. The second and third components show the response of the coronagraphic PSF to tip-tilt errors (i.e., the first modes of the AO).

2.2. Limitations of PCA

PCA is widely used by the exoplanet imaging community and available in most HCI packages such as VIP (V. Christiaens et al. 2023), PynPoint (T. Stolker et al. 2019), or pyKLIP (J. J. Wang et al. 2015). Its appeal is mainly due to its simplicity and to the fact that it comes with only one central hyperparameter¹⁷: the number of principal components K . However, this apparent simplicity hides a complex interplay of several factors, which makes it difficult to tune K in practice. The optimal choice is highly data set specific, influenced by variables such as the instrument, observing mode, wavelength, and observing conditions. Even within a single data set, the optimal number of components can vary as a function of the separation from the star (T. Meshkat et al. 2013). This is to be expected, as low-order wave front modes, which affect the closest separations to the star, vary the most. The farther away from the center, the higher the modal order and the larger the required K . An example of this effect is shown in Figure 2. Two artificial planets are inserted into the same data set: one close to the star ($\sim 3\lambda/D$) and another farther out ($\sim 10\lambda/D$). The closer planet is best seen if 60 principal components are selected, whereas the farther one requires about 300 components to be fully visible. Notably, at 300 components, the signal of the inner planet is completely lost. This demonstrates a critical limitation of PCA: as the number of components increases, PCA not only models the noise but also starts to include parts of the planet signal, leading to what is known as planet self-subtraction and oversubtraction (L. Pueyo 2016). As discussed further in Section 5, the choice of the number of components can affect the final contrast by several magnitudes. Hence, selecting the right number of components is crucial; it can be the difference between a detection and a nondetection.

¹⁷ In practice, there are additional geometry hyperparameters. These include masking strategies and temporal exclusion criteria and depend on the implementation (see J. J. Wang et al. 2015; T. Stolker et al. 2019; V. Christiaens et al. 2023).

2.3. Beyond PCA

The loss of planet signal in PCA has encouraged the development of more advanced postprocessing techniques. Nevertheless, many of these methods either enhance PCA or incorporate its principles. Annular PCA, for instance, applies PCA in concentric annular segments around the star (O. Absil et al. 2013; C. A. Gomez Gonzalez et al. 2017). A specific exclusion criterion for ADI prevents the signal of the planet from entering the PCA basis. Similar exclusion criteria have been proposed for LOCI and its variations (D. Lafrenière et al. 2007; C. Marois et al. 2010, 2014; R. Soummer et al. 2011; Z. Wahhaj et al. 2015; W. Thompson & C. Marois 2021). B. Lewis et al. (2023) extend PCA to utilize the temporal and spatial axis simultaneously. Conceptually similar to PCA are other subtraction-based techniques that replace the bottleneck of PCA. These methods include, for example, nonnegative matrix factorization (C. A. Gomez Gonzalez et al. 2017; C. Arcidiacono & V. Simoncini 2018; B. Ren et al. 2018), low-rank plus sparse decomposition (LLSG; C. A. Gomez Gonzalez et al. 2016), or more recently the half-sibling regression (HSR; T. D. Gebhard et al. 2022). For example, LLSG adds an extra sparse term to the model with the goal of limiting the inclusion of the planet in the noise model. H. Daglayan et al. (2023) propose to change the objective function to an L1 loss to account for the heavy-tailed distribution of the speckle noise. The recently proposed ConStruct method uses nonlinear autoencoders to locally predict the speckle noise (T. N. Wolf et al. 2023). Conventional PCA as discussed in Section 2 does not explore the temporal evolution of the speckle noise. M. J. Bonse et al. (2018), M. Samland et al. (2021), and J. D. Long et al. (2023) propose methods that explicitly use the time dimension. Most of these methods discuss different exclusion criteria and masks to avoid subtraction of the planet signal.

After the speckle subtraction, the residual images are commonly analyzed with meta-techniques that quantify detections and nondetections, for example, using t -tests (D. Mawet et al. 2014), STIM maps (B. Pairet et al. 2019), or performance maps (R. Jensen-Clem et al. 2017). A recent overview of these methods can be found in M. J. Bonse et al. (2023). Beyond the analysis of a single residual, supervised machine learning techniques have been proposed (C. A. Gomez Gonzalez et al. 2018; C. Cantero et al. 2023). These methods use several PCA residuals with different numbers of components as input and train a neural network to detect the planet. The regime-switching model (RSM) map by C.-H. Dahlqvist et al. (2020) and C.-H. Dahlqvist & O. Absil (2021) replaces the temporal averaging in the last step of PCA. It allows the combination of multiple PCA outputs alongside other postprocessing techniques.

Complementary to subtraction-based techniques are inverse-problem-based methods. Instead of modeling and subtracting the noise, these methods rely on forward modeling to search for the expected signature of planets in the data. Examples are ANDROMEDA (F. Cantalloube et al. 2015), FMMF (L. Pueyo 2016; J.-B. Ruffio et al. 2017), PACO (O. Flasseur et al. 2018, 2024), and TRAP (M. Samland et al. 2021). In this context, FMMF is the direct extension of PCA to forward modeling.

Although the concepts behind these modern postprocessing techniques are diverse, they share the common idea of better using the knowledge about the shape and behavior of the planet

signal. However, despite all these efforts, PCA still represents the main benchmark in practice: for example, K. Franson et al. (2023) used it for the discovery of the exoplanet AF Lep b. Further, none of these advanced algorithms have managed to completely avoid the problem of planet signal loss, which is the main bottleneck for achieving deeper detection limits.

3. Explaining Signal Loss in PCA

In the previous section, we observed that with an increasing number of principal components, progressively more signal of the planet is lost. A common explanation for this loss is that the principal components are not orthogonal to the signal of the planet, which is why they can fit and subtract the planet (see Appendix A). However, is there a specific component responsible for this effect? While the first components of PCA can be interpreted visually (Figure 1), it is difficult to understand the higher-order modes. To better understand why we lose signal in PCA, we propose to use methods from the field of explainable artificial intelligence (XAI). With the recent boom in machine learning, there is a trend toward increasingly complex deep network architectures. Unfortunately, the decisions made by such models are hard to understand and therefore hard to trust. XAI seeks to gain a better understanding of *why* and *how* machine learning models make their decisions, thus giving us a glimpse into the black box. A large number of XAI methods have been proposed in the literature (see S. Ali et al. 2023, for a recent overview). Some of them are model specific (e.g., tailored to neural networks), while other methods are more generally applicable (e.g., LIME by M. T. Ribeiro et al. 2016). In the following paragraph, we will keep our explanation of the XAI method used in the present paper general, to provide the basis for a better understanding of other postprocessing methods beyond PCA.

3.1. Saliency Maps

A widely used approach to explain the predictions of machine learning methods is the use of input gradients. The most basic approach was proposed by K. Simonyan et al. (2013). It is often referred to as the vanilla gradient and is straightforward to apply to subtraction-based postprocessing methods in HCI.

Let us consider a single science image \mathbf{x}_t at time t . Given a speckle estimation method S , we calculate the speckle noise for this particular image:

$$S(\mathbf{x}_t) = \hat{\mathbf{x}}_t. \quad (16)$$

For each pixel position in the noise estimate, we can use input gradients to calculate so-called *saliency maps*. Saliency maps are heatmaps that highlight areas in the input image \mathbf{x}_t that contribute the most to the noise estimate at the selected pixel position. Given a pixel position (i, j) , we calculate the derivative of S with respect to the input image \mathbf{x} at the point \mathbf{x}_t :

$$\mathbf{w}_l = \frac{\partial S_l}{\partial \mathbf{x}} \bigg|_{\mathbf{x}_t}, \quad (17)$$

where $l = i \cdot n_j + j = 1, \dots, D$ gives the index for the selected pixel position. The magnitude of the derivative \mathbf{w} points to pixels in \mathbf{x} that cause the largest difference on the noise estimate at the position (i, j) . For each pixel position in the noise estimate, we get one saliency map M_l by first calculating

the input gradient using Equation (17), then rearranging the one-dimensional vector \mathbf{w} back into a two-dimensional image, and finally taking the absolute value:

$$M_l = |\mathbf{w}_l|. \quad (18)$$

Saliency maps can be computed for any speckle estimation method S for which we can calculate the gradient in Equation (17). For nonlinear methods such as ConStruct (T. N. Wolf et al. 2023) or SODINN (C. A. Gomez Gonzalez et al. 2018), more advanced methods such as SmoothGrad (D. Smilkov et al. 2017) or integrated gradients (M. Sundararajan et al. 2017) are likely to give better results than vanilla gradients.

3.2. Right for the Wrong Reasons

In the case of PCA, the noise estimation method S is just the projection onto the low-dimensional space and its reconstruction:

$$S_{\text{PCA}}(\mathbf{x}_t) = \mathbf{x}_t^T \mathbf{U} \mathbf{U}^T = \hat{\mathbf{x}}_t. \quad (19)$$

Given the number of principal components K , we compute \mathbf{U} using the procedure explained in Section 2.1. Assuming that \mathbf{U} is fixed, the saliency map of PCA is then given by

$$M_{l,\text{PCA}} = |\mathbf{U} \mathbf{U}^T|_{:,l}. \quad (20)$$

Each pixel value in the noise estimate of PCA is a weighted sum of the pixel values in the science frame. The weights of this sum are the column vectors of the projection matrix $\mathbf{U} \mathbf{U}^T$. This means that *each column* of $\mathbf{U} \mathbf{U}^T$ gives us an explanation for which information is used to estimate the noise at *one output pixel position*. To get the saliency map for the pixel position l , we have to select the l th column vector of $\mathbf{U} \mathbf{U}^T$ and take its absolute value. Note that since PCA is a linear model, the saliency map for a given position is the same for all science frames along time.

Figure 3 illustrates the concept and provides examples of saliency maps for PCA based on data set #4 (HD 22049; see Section 5 for more details on the data set). To better understand the loss of planet signal in PCA, let us assume that the planet signal in the science frame (Figure 3, top left panel) is at the position marked by the blue circle. If parts of the planet signal are lost in the subtraction step, it means that some signal leaks into the noise estimate at the same position (top right panel, orange plus sign). We can use Equation (20) to compute the saliency map for this position (Figure 3, middle left panel). It turns out that most of the information used to estimate the noise at the planet’s position is only taken from the input image at the planet’s position. This means that PCA just copies the planet signal from the input image to the noise estimate. The same pattern can be observed for other pixel positions that do not contain the planet signal (Figure 3, middle right panel). Even if the planet’s position changes as a function of time (e.g., in the case of ADI), it is still subtracted.

It is interesting to note that the shape of the information used corresponds to the shape of the unsaturated PSF. We know that, to first order, the speckles follow the shape of the PSF (i.e., discarding wavelength-smearing effects). This means that PCA explains the speckle noise by its local morphology. However, since both the speckle noise and the planet signal follow the shape of the PSF, the planet signal is partially subtracted. By

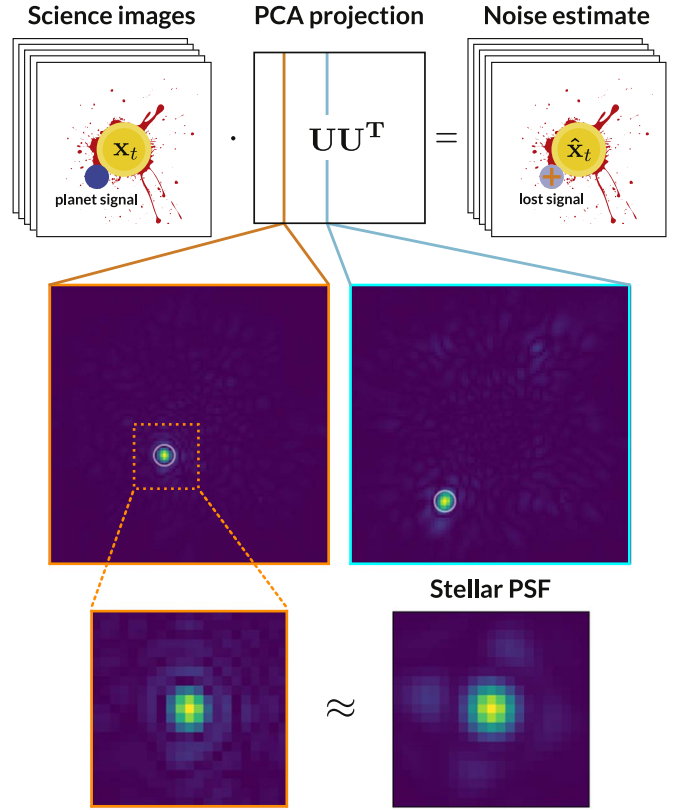


Figure 3. Illustration of how to compute saliency maps for PCA. For each pixel position in the noise estimate, we can compute one saliency map. Each saliency map highlights the areas in the science frame that are used to obtain the noise estimate at the respective position. Since the noise model of PCA is linear, every pixel in the noise estimate is a weighted sum of the pixels in the input image. The saliency map shows the absolute value of these weights. Two examples of such saliency maps are shown in the middle row (300 principal components). The example used in the text is marked in orange. For large numbers of principal components, the saliency map converges to the shape of the PSF. A gallery of saliency maps for different separations and numbers of principal components is given in Figure 17. The saliency maps should not be confused with the component vectors $\mathbf{u}_1, \dots, \mathbf{u}_K$ shown in Figure 1.

using the shape of the PSF, PCA provides a good estimate of the noise (correct answer). However, the explanation is not limited to the noise but also explains the planet signal. PCA gives the right answers for the noise, but for the wrong reasons.

It should be noted that this problem occurs regardless of whether the planet is actually in the data. That is, even if we build the PCA basis using data from a different star (RDI), we still get a model that can subtract the planet signal. As the number of components increases, the projection $\mathbf{U} \mathbf{U}^T$ converges to a de facto identity by learning the shape of the PSF. This convergence prevents us from learning more complex noise patterns.

Quantification of the effect. The convergence rate of PCA to the shape of the PSF is very data set and instrument specific and depends on whether the noise is dominated by speckles. For reasons explained later in Section 5, we confine our analysis to the VLT-NACO data taken in the L' filter, leaving research on other instruments for future work.

To quantify how many components are needed to converge to the shape of the PSF, we compute $\mathbf{U} \mathbf{U}^T$ for different numbers of principal components $K \in [1, \dots, 500]$. We calculate saliency maps for pixel positions at $1.5\lambda/D$, $4\lambda/D$, and $8\lambda/D$ and estimate the *match* between the saliency map and the




Figure 4. Convergence of the saliency map to the shape of the PSF as a function of the number of principal components used. At closer separations from the star, the convergence is faster, resulting in a greater loss of the planet signal with a smaller number of components.

unsaturated PSF template. More precisely, we compute a scalar product between the local saliency map and the unsaturated PSF and normalize it by the maximum value over all separations and components. The result of this experiment is shown in Figure 4. Examples of saliency maps used for this analysis are shown in Figure 17. For separations close to the star, PCA converges much faster to the shape of the PSF than farther out. This observation is in line with the fake planet experiment shown in Figure 2. While the signal of the fake planet close to the star is lost at 300 components, the planet farther out is still clearly visible. The speckle intensity gradually decreases with separation from the star. As PCA models the strongest covariances first, the convergence to the shape of the PSF is faster close to the star. As a result, the risk of losing parts of the planet signal is also higher close to the star.

3.3. Misguided Loss Function

The origin of the planet signal loss described in the previous section occurs regardless of whether a planet signal is present in the data used to compute the PCA basis. It is therefore often referred to as planet oversubtraction. In the case of ADI data sets, in addition to oversubtraction, the planet signal can also directly affect the PCA basis. This second effect is known as planet self-subtraction and can lead to even greater signal loss. In Section 2.1, we explained how to derive PCA by maximizing the variance in the low-dimensional space. This derivation is well motivated for HCI, since the speckles are usually much brighter and more variable than the planet signal. An alternative derivation of PCA is the minimum reconstruction error formulation. We want to find a linear dimensionality reduction that minimizes the following reconstruction loss:

$$\mathcal{L}_{\text{PCA}} = \frac{1}{N} \sum_{t=1}^N \left\| \underbrace{x_t - \hat{x}_t}_{s_t + p_t - \hat{x}_t} \right\|^2. \quad (21)$$

That is, we want our reconstruction \hat{x}_t to be close to our science data x_t . Both the maximum variance derivation and the minimum reconstruction error formulation yield the same algorithm (see chapter 12.1.2 in C. M. Bishop 2006). In the




Figure 5. Retained flux after postprocessing with PCA as a function of planet brightness and separation from the star. The grid was computed using `applefy`. Each value in the grid is the result of three fake planet experiments with planets inserted at different position angles. A significant fraction of the planet signal is lost, especially close to the star.

context of HCI, minimizing Equation (21) can be problematic, as x_t contains both speckle noise and planet signal. This means that \mathcal{L}_{PCA} can only become small if the noise estimate \hat{x}_t also fits the signal of the planet p_t . For large numbers of principal components and for bright companions, PCA tends to include the companion signal in its basis.

To quantify this effect, we compute a throughput grid using the Python package `applefy` (M. J. Bonse et al. 2023). A throughput grid illustrates the attenuation of the planet signal due to the data postprocessing. The grid is computed by inserting fake companions with different known brightnesses and distances to the star and then measuring the fraction of the flux that remains after PSF subtraction. The result of this experiment is shown in Figure 5. We observe that for faint companions the retained flux is only a function of the separation and does not depend on the brightness of the companion. This gradual loss of planet signal closer to the star is caused by the convergence of PCA to the shape of the PSF (see Section 3.2). For bright companions, however, PCA begins to include parts of the signal in its basis, causing an even greater loss of planet signal. The brightness at which planets are affected depends on the number of principal components. If more components are used, fainter companions start to be affected.

4. 4S: Signal-Safe Speckle Subtraction

In the previous section, we identified two main reasons for the loss of planet signal in PCA: first, the convergence of the noise model to the telescope's PSF, and second, a problematic loss function that accidentally leads the noise estimate to fit the planet signal. Based on these findings, we propose a novel postprocessing method that adopts the main steps of PCA but overcomes its limitations. We have named this method 4S, short for Signal-Safe Speckle Subtraction. 4S is characterized by three innovations: first, a linear noise model with a right reason constraint; second, a new loss function that is invariant to the planet signal; and third, a domain-knowledge-informed regularization.

4.1. The Right Reason Constraint

In Section 3.2, we found that the noise model of PCA converges to a de facto identity that reproduces the shape of the telescope's PSF. This is problematic, as both the speckle noise

and the planet signal follow this shape. Ideally, we would like our noise model to specifically recognize the speckle noise, that is, we want to find explanations for the speckle noise that are unique to the noise and do not describe the planet signal. In the context of XAI, A. S. Ross et al. (2017) proposed to add a right reason penalty to guide the predictions of neural networks using domain knowledge. The domain knowledge is represented as a mask that penalizes wrong explanations. The method has been successfully applied in practice (P. Schramowski et al. 2020). Inspired by this work, we propose a right reason constraint to prevent the noise model of 4S from learning the shape of the telescope’s PSF.

Like PCA (see Equation (19)), 4S uses a linear model, but the projection matrix UU^T is replaced by a square matrix $B \in \mathbb{R}^{D \times D}$:

$$4S(x_i) = x_i^T B = \hat{x}_i. \quad (22)$$

It is important to note that B is much larger than the component matrix of PCA, $U \in \mathbb{R}^{D \times K}$. For PCA, each column of the projection matrix UU^T gives us an explanation of which information of the input frame x_i is used to estimate the noise at one pixel position in \hat{x}_i . Similarly, the columns of B contain the explanations of 4S. Following the idea of A. S. Ross et al. (2017), we want to limit the information 4S can use from the input frames to predict the speckle noise. We achieve this goal by designing a right reason mask that sets the values along and around the diagonal of B to zero. For each pixel position l , we have one parameter vector $b_l \in \mathbb{R}^D$ as a column of B that is masked by one right reason mask $m_l \in \mathbb{R}^D$:

$$B_{:,l} = b_l \circ m_l, \quad (23)$$

where \circ is the element-wise multiplication. In total we have D right reason masks, one for every column of B . Each mask m_l is designed such that values around the position $l = i \cdot n_j + j$ are set to zero:¹⁸

$$m_l(l') = \begin{cases} 0, & \text{if } \sqrt{(i - i')^2 + (j - j')^2} \leq d, \\ 1, & \text{otherwise} \end{cases} \quad (24)$$

where d sets the radius of the mask and $l' = i' \cdot n_j + j' = 1, \dots, D$ gives the pixel position within the l th mask (see bottom left box of Figure 6). We choose d relative to the size of the PSF: $d = 0.75 \cdot \text{FWHM}$. This mask size ensures that the core of the PSF is protected from the de facto identity, while local information about the speckle noise is still preserved. A larger mask size allows us to protect even more signal, but it also weakens the noise model. The value of $d = 0.75 \cdot \text{FWHM}$ is a compromise for which we have obtained good results in practice.

The right reason mask should not be confused with the exclusion criteria used by other postprocessing techniques (e.g., D. Lafrenière et al. 2007; Z. Wahhaj et al. 2015; C. Marois et al. 2010; R. Soummer et al. 2011; O. Absil et al. 2013; C. A. Gomez Gonzalez et al. 2017). While the masks of these techniques also aim to reduce the loss of planet signal, the way they work is fundamentally different. The masks proposed in the literature exclude PSF-sized regions in the science data to

prevent the planet signal from leaking into the noise estimate. They often have a temporal dependency to track potential trajectories of point sources as a function of field rotation. 4S, on the other hand, masks not the science data but the parameters of the noise model during the optimization. The mask does not block the signal of the planet, but instead forces the model to explore other areas of the input image that explain the speckle noise. The right reason mask has no time dependency. In Appendix A we compare the linear noise models of LOCI, PCA, and 4S and discuss how their masking strategies differ.

Our approach is similar to the postprocessing methods based on the HSR (M. Samland et al. 2021; T. D. Gebhard et al. 2022). However, while T. D. Gebhard et al. (2022) select areas in the input images that promise to contain useful information about the speckle noise, we use the entire science frame but exclude areas that correspond to the shape of the telescope’s PSF near the de facto identity.

4.2. Signal Invariant Loss Function

A second limitation of PCA is its loss function (see Equation (21)), which can only become small if the noise estimate includes the planet signal. This is problematic in the case of ADI, for which the planet signal is part of the data set used to compute the PCA basis. Different types of exclusion regions (D. Lafrenière et al. 2007; C. Marois et al. 2010; F. Cantalloube et al. 2015; W. Thompson & C. Marois 2021) or explicit planet models (M. Samland et al. 2021) have been proposed in the literature to protect the planet’s signal from becoming part of the noise estimate. These methods, however, often come with additional hyperparameters that are hard to optimize in practice. Therefore, instead of minimizing the reconstruction loss, as done in Equation (21), we propose to minimize the variance along the temporal axis in the de-rotated residual sequence:

$$\mathcal{L}_{4S} = \frac{1}{N} \sum_{i=1}^N \|r_i - \bar{r}\|^2, \quad \bar{r} = \frac{1}{N} \sum_{i=1}^N r_i. \quad (25)$$

An illustration of why \mathcal{L}_{4S} is sensitive to the speckle noise but does not affect the planet signal is shown in the bottom right box of Figure 6. In the de-rotated residual sequence r_i , the signal of the planet p_i is always at the same position and has approximately constant brightness. Consequently, it only contributes to the temporal mean \bar{r} but does not influence the temporal variance. Since the temporal mean \bar{r} is subtracted in Equation (25), the loss function becomes invariant to the presence of the planet signal. Quasi-static and atmospheric speckles, however, vary with time, which is why they affect the loss function (red dashed lines in Figure 6). Note that the noise model of 4S is still applied to the science images before de-rotation (see Equation (22)). Only the loss function is evaluated on the de-rotated residual sequence.¹⁹ Details on how to minimize this loss function are given in Section 4.4.

¹⁸ Note: the mask does not have to be binary, which is beneficial for data sets with low resolution. Our implementation makes use of circular apertures based on the `photutils` package.

¹⁹ We do not account for changing conditions during the night. However, if information about the variability of the planet signal as a function of time is available (e.g., for data taken with a vAPP coronagraph such as in G. P. L. Otten et al. 2017), one could normalize for these variations as part of the preprocessing.

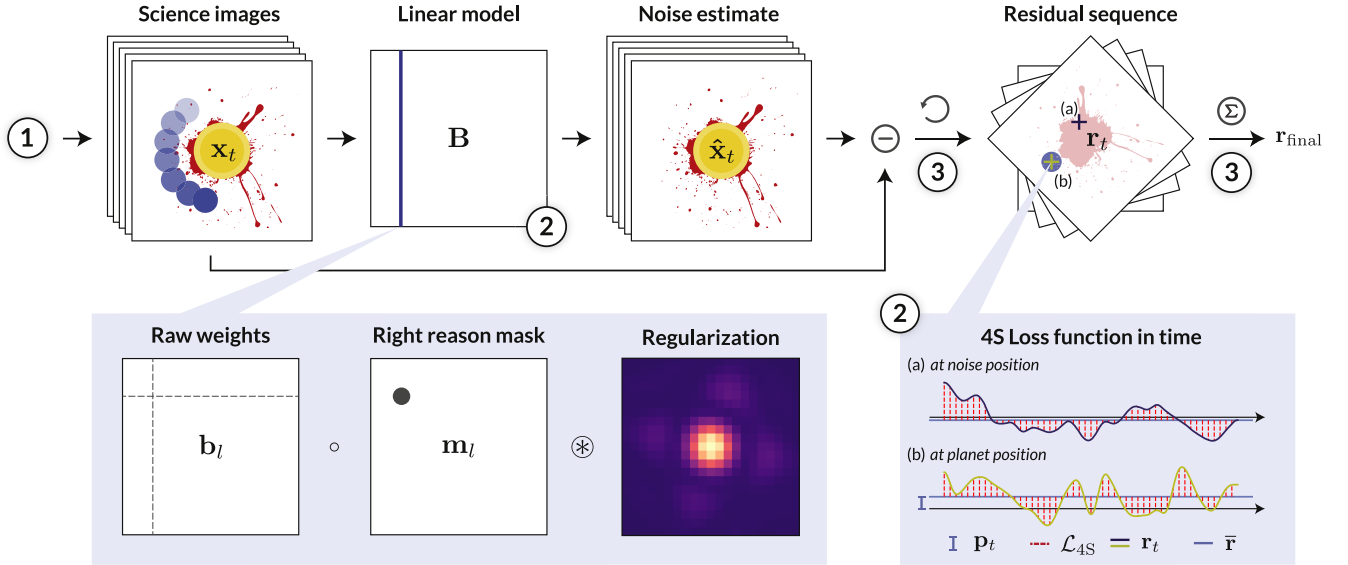


Figure 6. Illustration of our new postprocessing technique 4S. To overcome the limitations of PCA, 4S makes use of three innovations: (1) a right reason mask to prevent oversubtraction (bottom left box), (2) a loss function that is invariant to the signal of the planet (bottom right box), and (3) a regularization based on domain knowledge to prevent overfitting. The main steps ① to ③ are explained in detail in Section 4.4.

4.3. Regularization

So far, the number of free model parameters of 4S is huge: given a frame resolution of 100×100 pixels, \mathbf{B} contains about 100 million parameters. To mitigate the risk of overfitting, we use two types of regularization to reduce the expressiveness of the model. First, we use a classical L2 penalty term on the model parameters (similar to Tikhonov/ridge regression; see chapter 3.1.4 of C. M. Bishop 2006), which we add to the loss function:

$$\mathcal{L}_{4S} = \underbrace{\frac{1}{N} \sum_{t=1}^N \|\mathbf{r}_t - \bar{\mathbf{r}}\|^2}_{\text{temporal variance}} + \underbrace{\lambda \cdot \sum_{l=1}^D \|\mathbf{b}_l\|^2}_{\text{L2 penalty}}, \quad (26)$$

where λ is the only hyperparameter of 4S. Larger values of λ force the parameters \mathbf{b} to be closer to zero, reducing the expressiveness of the model. This means that a smaller value of λ results in a stronger reduction of the speckle noise while coming at the risk of overfitting (i.e., memorizing the data). In contrast to PCA, we notice that the best choice of λ does not depend on the separation from the star. Instead, it depends on the temporal resolution and the number of pixels in the science frames. If only a few frames are available in time, a stronger regularization (i.e., larger values of λ) is needed and the reduction of speckle noise is less effective. However, a detailed analysis of how to tune λ is beyond the scope of this paper.

In addition to the L2 regularization, we also convolve the parameters of our model with the unsaturated PSF. For this purpose we replace Equation (23) by

$$\mathbf{B}_{:,l} = (\mathbf{b}_l \circ \mathbf{m}_l) \circledast \text{PSF}. \quad (27)$$

The convolution ties spatially adjacent parameters together, reducing the effective number of model parameters. From our domain knowledge, we know that speckles approximately follow the shape of the PSF. All features in the data that are smaller are not due to the speckle noise but are caused, for example, by detector noise. By convolving the parameters \mathbf{b}

with the PSF, we prevent our model from overfitting pixel-level artifacts. The convolution is an integral part of the optimization, and the parameters \mathbf{b} are optimized taking into account both the right reason mask and the convolution. As explained in step ③ in the next section, we use automatic differentiation to back-propagate through the convolution.

4.4. Main Steps of 4S and Optimization

Applying 4S to ADI data in HCI consists of the following steps, which we illustrate in Figure 6:

1. *Data preparation.* Similar to PCA, 4S works on normalized data. We first subtract the temporal mean $\bar{\mathbf{x}}$ and divide by the standard deviation along the time axis:

$$\mathbf{x}_t = (\mathbf{x}_{t,\text{raw}} - \bar{\mathbf{x}}) / \text{std}(\mathbf{x}_{t,\text{raw}}). \quad (28)$$

The additional division by $\text{std}(\mathbf{x}_{t,\text{raw}})$ is required to ensure that the L2 regularization affects all areas in the image equally. The 2D science frames are flattened into one-dimensional vectors and stacked into the data matrix $\mathbf{X} \in \mathbb{R}^{N \times D}$ (see the PCA procedure in Section 2).

2. *Initialization of the model.* We initialize all parameters \mathbf{b}_l with zeros and create the right reason masks \mathbf{m}_l for all output pixel positions $l = 1, \dots, D$.
3. *Optimization of the model parameters.* We optimize the parameters of the noise model \mathbf{b}_l to minimize the loss function in Equation (26) using the Python package PyTorch (A. Paszke et al. 2019a). We chose the L-BFGS optimizer (D. C. Liu & J. Nocedal 1989), which converged faster for our application than other optimizers such as SGD with momentum (I. Sutskever et al. 2013) or Adam (D. P. Kingma & J. Ba 2017). The L-BFGS optimizer is a quasi-Newton method that iteratively minimizes the loss function. The optimization is carried out in an end-to-end fashion taking into account the right reason mask, the convolution with the PSF (see Equation (27)), and the field rotation due to ADI. We use the automatic differentiation framework (see A. G. Baydin et al. 2018, for a tutorial on the topic) of

Pytorch to compute gradients of the loss function \mathcal{L}_{4S} with respect to the parameters \mathbf{b} . Since the loss function in Equation (26) is defined with respect to the de-rotated residuals sequence, this step requires us to back-propagate through the image rotation ϕ_r . For this purpose, we take advantage of sampler grids originally developed for spatial transformer networks (M. Jaderberg et al. 2016). The L-BFGS optimizer takes the first-order gradients and approximates second-order derivatives to update the parameters \mathbf{b}_l . The iterative optimization is finished once the loss plateaus. This is usually the case after 100–500 iterations depending on the data set and the value of λ . Larger values of λ (stronger regularization) lead to faster convergence compared to smaller values of λ . We find that the use of minibatches, as often used in neural networks, leads to significantly worse results. This is probably due to the values of \bar{r} , which contain the signal of the planet and become noisy when using minibatches. Therefore, the entire data set is used in each step of the optimization.

4. *Calculation of the residual.* After the optimization is completed, we calculate the residual sequence using Equation (8). As for PCA, the final residual is obtained by averaging along the time dimension (Equation (9)).

We note that the subtraction of the temporal average in step ① is not free from planet self-subtraction. More advanced ideas to overcome this step should be explored in future work. Our Python implementation takes advantage of modern NVIDIA A100 and H100 GPUs to speed up the optimization. Depending on the spatial and temporal resolution of the data set, the optimization converges in about 10–30 minutes using one GPU. For large data sets that do not fit into GPU memory, gradient accumulation can be used. It is possible to calculate results for a range of $\lambda_1 > \lambda_2 > \lambda_3 > \dots$ without significant computational overheads. For this we first optimize the parameters starting with the strongest regularization λ_1 . For the second setup λ_2 , we repeat steps ② to ④. But instead of initializing the model parameters \mathbf{b}_l with zeros in step ②, we use the \mathbf{b}_l computed for λ_1 . This strategy significantly speeds up the convergence of the optimization for λ_2 .

5. Quantitative Analysis

Comparing the contrast performance of postprocessing techniques can be difficult for two main reasons. First, different postprocessing methods follow fundamentally different concepts. PCA and 4S are both subtraction-based techniques. They return a residual image, which is analyzed in a second step with statistical methods like t -tests or STIM maps (D. Mawet et al. 2014; B. Pairet et al. 2019). Methods like FMMF, PACO, or TRAP (J.-B. Ruffio et al. 2017; O. Flasseur et al. 2018; M. Samland et al. 2021) belong to the family of inverse-problem approaches. They produce detection maps that show the models’ belief in the presence of a planet. The detection process for subtraction-based and inverse-problem approaches relies on very different statistics. Therefore, their results are not directly comparable. In addition, the residual noise characteristics are method specific and usually unknown. This makes a fair comparison even more difficult, especially if we want to compare different methods given the same false-positive fraction (FPF). For this reason, methods should only be

compared within the same category (see, e.g., F. Cantalloube et al. 2021).

Second, the contrast performance is influenced by the instrument and the observing mode. Depending on the wavelength, different noise sources dominate, such that the best choice of postprocessing method is likely to be instrument dependent. To keep our analysis consistent, we decided to focus on archival data from the VLT-NACO instrument in the L' band ($\lambda = 3.8 \mu\text{m}$). The exposure times of these data are typically very short, which gives us a high temporal resolution. This is crucial for 4S to minimize the risk of overfitting. It also allows us to correct not only for quasi-static speckles but also for atmospheric speckles on shorter timescales. To ensure the robustness of our results under changing observing conditions and targets, we base our analysis on 11 archival data sets (see Table 1). This guarantees that any observed improvements are not random, but consistent across multiple data sets. All targets are bright ($L' < 4 \text{ mag}$) and were chosen such that no obvious companions or disk signals are present. We ensure that each data set has at least 30 minutes of observations with seeing no worse than $1''.5$. The final list is a random selection of data sets in the archive that satisfy these conditions.

We process all our data sets with the state-of-the-art pipeline `PynPoint`. The preprocessing steps include (1) a simple dark and flat calibration, (2) an interpolation of bad pixels, (3) a simple background subtraction using sky frames, (4) a centering of the star using cross-correlations, and (5) a removal of bad frames (e.g., frames with an open AO loop or bad centering of the star behind the coronagraph). To keep the computation time manageable, we combine every five frames of the science sequence and crop the frames to a field of view of radius $1''.2$ (equivalent to $12\lambda/D$). Even after temporal binning, we still have between 2400 and 13,000 frames along the time dimension. The final preprocessed data cubes are publicly available on Zenodo (M. J. Bonse 2024b).

5.1. Right for the Right Reasons

In Section 3.2, we used saliency maps to better understand which information is used by PCA to estimate the noise at a given position. This approach revealed that PCA converges to the telescope’s PSF for high numbers of components, causing significant signal loss. 4S is designed to overcome this limitation and forces the model to explore other explanations for the speckle noise using a right reason constraint. But is the noise estimate of 4S right for the right reasons?

Similar to PCA, we can compute saliency maps for 4S that are just the column vectors of our noise model:

$$M_{l,4S} = |\mathbf{B}|_{:,l}. \quad (29)$$

We run 4S on data sets #4 (HD 22049) and #7 (HD 40136) with a regularization strength $\lambda = 10^3$ and compute saliency maps for three different positions. The two data sets were selected randomly, with about 3 yr between observations. The output saliency maps are shown in Figure 7. We notice that 4S no longer learns the shape of the telescope’s PSF, but instead explores other areas in the science frames. The most dominant features in the saliency maps of 4S are point symmetries. This is a strong indication that 4S indeed learns the characteristic behavior of speckles, which are known to exhibit symmetric or antisymmetric correlations (E. E. Bloemhof et al. 2001; M. D. Perrin et al. 2003;

Table 1
VLT-NACO Data Sets Used for Performance Evaluation

No.	Target	Program ID	Observing Date	r_0 (arcsec)	Field Rotation (deg)	DIT (s)	Total (s)	No. Frames	AGPM
1	HD 2262	199.C-0065 (C)(a)	2017-11-01	0.56	73.7	0.35	5043	12711	✓
2	HD 7570	1101.C-0092 (C)(c)	2018-11-27	0.65	68.9	0.35	6895	14515	✓
3	HD 11171	1101.C-0092 (C)(c)	2018-11-28	0.63	108.7	0.35	7350	18572	✓
4	HD 22049	096.C-0679 (A)(c)	2015-12-17	0.57	90.1	0.08	4608	48300	✓
5	HD 22049	199.C-0065 (C)(a)	2017-10-30	0.61	104.7	0.08	5304	57966	✓
6	HD 38678	084.C-0396 (A)(b)	2009-11-27	1.47	39.2	0.2	2460	11185	
7	HD 40136	1101.C-0092(C)(c)	2018-11-29	0.87	114.7	0.3	7500	20265	✓
8	HD 115892	1101.C-0092 (E)(c)	2019-05-23	0.48	95.2	0.35	6895	17401	✓
9	HD 169022	1101.C-0092 (E)(c)	2019-05-20	1.28	113.8	0.35	6195	9935	✓
10	HD 177724	091.C-0654 (A)(c)	2013-08-19	0.65	24.9	0.2	3200	12720	
11	HD 209952	089.C-0149 (A)(b)	2012-07-14	1.12	40.4	0.2	2880	12647	
12	HD 35850	088.C-0085(A)(b)	2011-10-21	1.1	69.6	0.2	3120	13809	

Notes. Data sets #1 through #11 are used for the quantitative analysis; data set #12 (AF Lep) is used for the scientific demonstration in Section 6. The number of frames provided is the number of frames after removing bad frames. All data sets are preprocessed with `PynPoint` and publicly available on Zenodo (M. J. Bonse 2024b).

References: (a) R. Launhardt et al. (2020); (b) J. Rameau et al. (2013); (c): no reference or unpublished.

E. E. Bloemhof 2007; E. N. Ribak & S. Gladysz 2008). Another pattern, characterized by four points of equal distance from the star, can be seen in the lower left saliency map of Figure 7. This pattern is due to waffle mode, a distortion of the wave front that is not visible to the wave front sensor. It is an artifact in the data that we want to remove in the data postprocessing. Due to the width of the L' -band filter, the speckles are slightly elongated. 4S finds this characteristic of the speckle noise and uses information on a line pointing toward the center of the frame (bottom two panels of Figure 7).

The two data sets, #4 (HD 22049) and #7 (HD 40136), were acquired with the same observing mode. However, there is a significant gap of 3 yr between the two observations. Despite this time difference, it is remarkable to observe the similarity in the saliency maps produced by 4S. This suggests that 4S can learn signatures of the noise that are linked to the instrument and not the individual observation. This result opens up several possibilities for future work, which we discuss in Section 7.

5.2. Fake Planet Experiments

To reliably calculate the contrast performance of 4S, we perform extensive fake planet experiments. We use contrast grids implemented in the `python` package `applefy` (M. J. Bonse et al. 2023). Contrast grids systematically calculate the signal-to-noise ratio (S/N) of inserted fake companions as a function of distance from the star and planet brightness. We insert fake planets every 0.5 mag for contrasts [5, 15] mag and for separations $[1.5, 12.5]\lambda/D$ every $0.5\lambda/D$. For each combination of planet brightness and separation, three planets are inserted at different position angles. We insert only one fake planet at a time, resulting in 945 fake planet experiments per data set. Each data set with an inserted fake planet is processed with PCA and 4S to obtain a residual image. Afterward, the S/N of the companion is calculated using the method proposed by D. Mawet et al. (2014). Since our limits are in the speckle-dominated regime, we use

spaced pixels instead of apertures to extract the photometry, as suggested by M. J. Bonse et al. (2023). The final contrast curves are obtained by thresholding the contrast grid using all 945 fake planet experiments.

The results of PCA depend strongly on the number of components used. Therefore, 33 different setups with $K \in [1500]$ were calculated for each fake planet experiment. For 4S, the regularization strength was set to $\lambda \in [10^2, 10^3, 10^4, 10^5]$. We compute contrast grids for each data set (see Table 1) and all algorithm setups. In total, we compute more than 300,000 PCA and more than 30,000 4S residuals. This experiment required substantial compute and was thus realized on a compute cluster, allowing us to use up to 200 NVIDIA A100/H100 GPUs simultaneously. To further speed up the computations, we reimplemented the `PynPoint` version of PCA using `PyTorch` to enable GPU computations. Our new GPU implementation of PCA uses the fast SVD from N. Halko et al. (2011), resulting in a 50-fold speedup on one NVIDIA H100 GPU compared to 64 CPUs. The code for our new implementation is available along with our `Python` package `fours`. Example residuals and contrast curves for one of our data sets, #5 (HD 22049), are shown in Figure 8. As shown in the plot, 4S reaches significantly deeper contrast compared to PCA, especially at separations close to the star. The two close fake companions at $1.5\lambda/D$ and $2.5\lambda/D$ are undetectable with PCA, regardless of the number of principal components. For separations $< 2.5\lambda/D$, we reach about 1.4 mag deeper contrast with 4S compared to PCA. Even at larger separations, a decent improvement of ~ 0.3 mag can be achieved. Although 4S is designed to prevent planet self-subtraction, we still observe negative *wings* next to the planet's signal at close separations. These negative wings are caused by the subtraction of the temporal average in the first step of 4S. If we were able to subtract the temporal mean without the planet signal, the negative wings would disappear and even better results could be achieved. This path should be explored in future work within the context of RDI.

To achieve the deepest contrast of PCA, we need to carefully tune the number of components for each separation. While 16 components give the best results at $\sim 3\lambda/D$, 400 components are required at $\sim 12\lambda/D$. For 4S, on the other hand, $\lambda = 100$

4S Saliency Map

#4 (HD 22049) - Dec 2015 #7 (HD 40136) - Nov 2018




Figure 7. Saliency maps for 4S, computed for data sets #4 (HD 22049; left column) and #7 (HD 40136; right column). The saliency maps show which information from the science frames x_i is used by 4S to estimate the noise at the position marked by the white plus sign. In contrast to PCA (see Figure 3), 4S successfully learns the point-symmetric signatures of the speckle noise. The red star marks the center of the frame.

consistently outperforms all PCA setups. Since tuning the number of principal components is difficult in practice and could suffer from human bias, this is a clear advantage of 4S.

We repeat the same experiment shown in Figure 8 for all data sets summarized in Table 1. Depending on the data set, we select the best $\lambda \in [10^2, 10^3, 10^4, 10^5]$ for 4S and compare it to the *optimal* PCA limits. For data sets #10 and #11 we had to exclude the result for the innermost 2.5λ , due to strong saturation from the host star. The relative improvement of 4S over PCA accumulated over all data sets is shown in Figure 9.

The results are consistent with the improvement achieved on data set #5, with about 0.5 mag variation depending on the data set. The improvement is largest at close separations to the star. We know that the loss of the planet signal in PCA is at its maximum at these separations (see Section 3). 4S successfully mitigates this limitation, allowing it to learn a better representation of the speckle noise.

The improvement of 4S over PCA is probably larger in practice, mainly for three reasons:

1. If we go through the stack of PCA residuals for different numbers of components, the speckle pattern changes considerably. For 4S the speckle pattern in the residual images is very static for different values of λ . Since we try 33 different setups for PCA, the chance of a speckle co-aligning with the inserted fake companion increases. This artificially increases the S/N in the PCA residuals.
2. It is likely that the optimal PCA limit is not reached in practice. This is because it is difficult to find the best number of principal components. For data set #4 (HD 22049), R. Launhardt et al. (2020) report a contrast of 10.5 mag at $0''.5$ and 11.6 mag at $0''.75$. Our optimal PCA limits reach down to 11.5 mag at $0''.5$ and 12.3 mag at $0''.75$, which is almost 1 mag deeper. This difference is likely due to the large number of setups tried in our PCA analysis. In addition, better preprocessing of the data with PynPoint could explain part of this improvement.
3. The noise distribution of the PCA residuals often deviates from Gaussian noise (B. Pairet et al. 2019; M. J. Bonse et al. 2023). This is especially true for small numbers of principal components. Heavy-tailed non-Gaussian residual noise increases the risk of false positives. In Appendix C, we study the residual noise distribution of 4S and PCA. We find that the noise of 4S is not perfectly Gaussian either, but it is more Gaussian than the residual noise of PCA. This means that the chance of observing an $S/N > 5$ that is caused by noise is greater for PCA than for 4S. However, since the true noise distribution is unknown for both methods, the limits cannot be calculated at the same FPF level.

Over all data sets, the median contrast performance of 4S is consistently better than limits reached with *optimal* PCA. This improvement opens a new discovery space on archival data and for future observations.

6. Scientific Demonstration

To demonstrate that 4S is not only capable of improving the contrast performance for artificial companions, we have searched the ESO archive for companions that have been found with VLT-SPHERE but not with VLT-NACO. VLT-SPHERE has found several new companions, but unfortunately most of them have never been observed with VLT-NACO.

AF Leporis is a young (~ 28 Myr based on moving group age estimate; C. P. M. Bell et al. 2015) F8V-type star (R. O. Gray et al. 2006) belonging to the β Pictoris moving group (S. Zúñiga-Fernández et al. 2021). Since AF Lep presents a significant astrometric Hipparcos–Gaia acceleration, it was observed by multiple direct imaging surveys targeted at accelerating stars. Three direct detections of the companion AF Lep b obtained with VLT-SPHERE and Keck-NIRC2 were claimed almost simultaneously in 2023 (R. J. De Rosa et al. 2023; K. Franson et al. 2023; D. Mesa et al. 2023). While these observations firmly place AF Lep b as a faint L-type dwarf in the planetary regime, estimations of its dynamical mass via the astrometric acceleration are conflicting between instruments and with age–luminosity model-based estimates. While the latter disagreement is often observed in recent analyses (A. Cheetham et al. 2018; G. M. Brandt et al. 2021a; T. L. Tobin et al. 2024), the former is likely due to the very short time baseline on which the orbit of AF Lep b could be sampled (2021 December to 2023 February; see Figure 11).

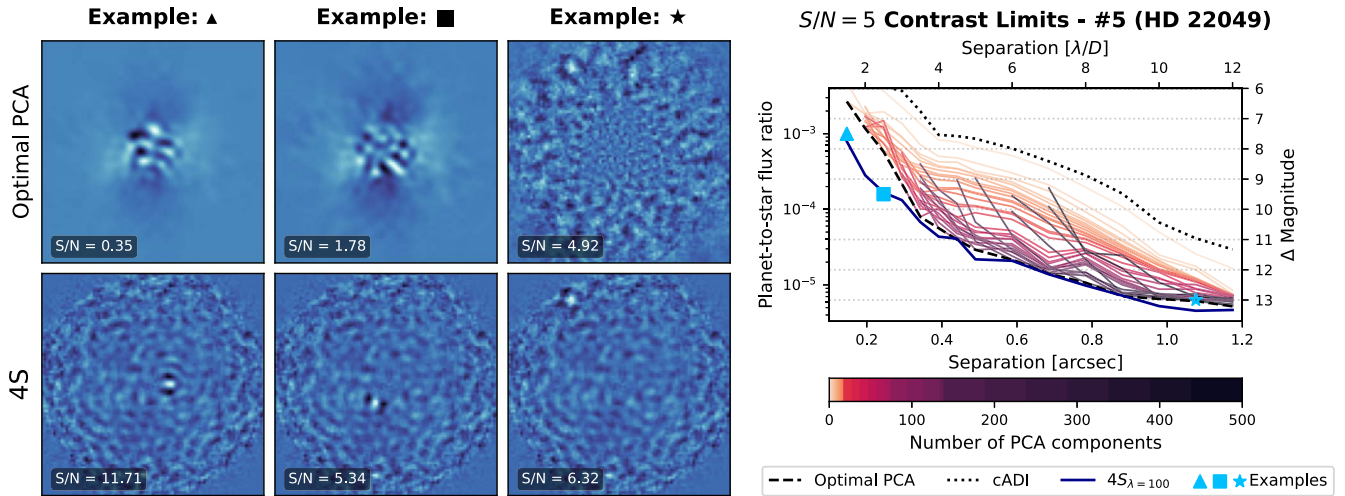


Figure 8. Results of our fake planet experiments for data set #5 (HD 22049; see Table 1). On the left, example residuals are shown for fake planets at ▲ = $(1.5\lambda/D, 7.5 \text{ mag})$, ■ = $(2.5\lambda/D, 9.5 \text{ mag})$ and ★ = $(11 \lambda/D, 13 \text{ mag})$. In the case of PCA, we tune the number of principal components to get the deepest contrast (optimal PCA). For 4S, we use the same setup with $\lambda = 100$ for all separations. An interactive version of the plot is available on our [documentation page](#). The plot on the right shows contrast curves of 4S and PCA with different numbers of principal components. Each contrast curve was obtained by computing a contrast grid using the Python package *applefy* and is based on 945 fake planet experiments. The dotted black line gives the cADI contrast, which is equivalent to a simple subtraction of the temporal median. Note that for large numbers of components and small separations, none of the inserted fake planets reach an $S/N > 5$.




Figure 9. Relative improvement in contrast of 4S over PCA, accumulated over the 11 data sets shown in Table 1. The blue line shows the median improvement, the gray boxes give the interquartile range (IQR), and the whiskers extend to points that lie within 1.5 IQR. The points marked as circles are outliers.

K. Franson et al. (2023) discovered the planet in the L' band at a separation of $\rho_{2023.090342} = 342 \pm 8 \text{ mas}$ and at a contrast of $\Delta L' = 9.94 \pm 0.14 \text{ mag}$. This is precisely the regime where we expect an improvement with 4S over PCA. The star AF Lep was imaged with VLT-NACO in 2011 at L' band using the normal imaging mode (see Table 1). The data set was taken in medium observing conditions (seeing $1''.1$) but provides 69.6° of field rotation. We process this archival data set with PynPoint (T. Stolker et al. 2019) using the same preprocessing routine explained in Section 5 and then subtract the stellar PSF using PCA and 4S. Our 4S reduction reveals the giant exoplanet AF Lep b on the opposite side of the star, 11 yr before its discovery in 2022 (see Figure 10). The signal of the planet is also present in the residuals of PCA, but it is indistinguishable from the speckle noise.

This precovery of AF Lep b presents a unique opportunity to extend the time baseline of its orbit sampling from about 1 yr to

more than 11 yr. Furthermore, it enables an independent confirmation of the existing L' -band magnitude estimate (K. Franson et al. 2023).

Astrometry and Photometry. Because of the attenuation of the planet signal due to the postprocessing, we cannot extract the photometry and astrometry from the residual images directly. Note that, due to the normalization of the data in the first step of 4S and to the finite size of the right reason mask, 4S still loses some planet signal. In the case of PCA, it is common practice to insert a negative fake planet into the raw data at the position of the detection. The photometry and astrometry, as well as their uncertainties, are then calculated using Markov Chain Monte Carlo (MCMC; see, e.g., T. Stolker et al. 2019). Each step in the MCMC requires a complete recalculation of the PCA basis and the residual. A full MCMC often relies on several thousand rereducations, which becomes computationally expensive. For 4S, a reduction of the AF Lep data set takes about 10 minutes, resulting in a total computation time of about 1 month. It is likely possible to significantly reduce the computation time required for an MCMC with 4S. This could be achieved by avoiding the recalculation of the model matrix B for each iteration. Instead, B could be calculated once and then updated continuously during the MCMC optimization. However, to verify the accuracy of such an approach, further tests are needed that are beyond the scope of this paper. Further details for future work in this direction are discussed in Section 7. Instead, we decided to estimate the astrometry and photometry based on the PCA residuals using the well-established and tested MCMC code presented in T. Stolker et al. (2019). For this, we first calculated PCA residuals for $K \in [1, 200]$ every 10 components. The peak $S/N = 4.7$ is reached at 65 components.²⁰ We fix $K=65$ and use the MCMC procedure in PynPoint to estimate the astrometry and photometry. As in T. Stolker et al. (2019), we first use a

²⁰ It is important to note that the number of principal components that gives the highest S/N can only be found given both the exact separation and the position angle. If only the separation is known, the optimal number of components may vary by a few components as a function of the position angle.

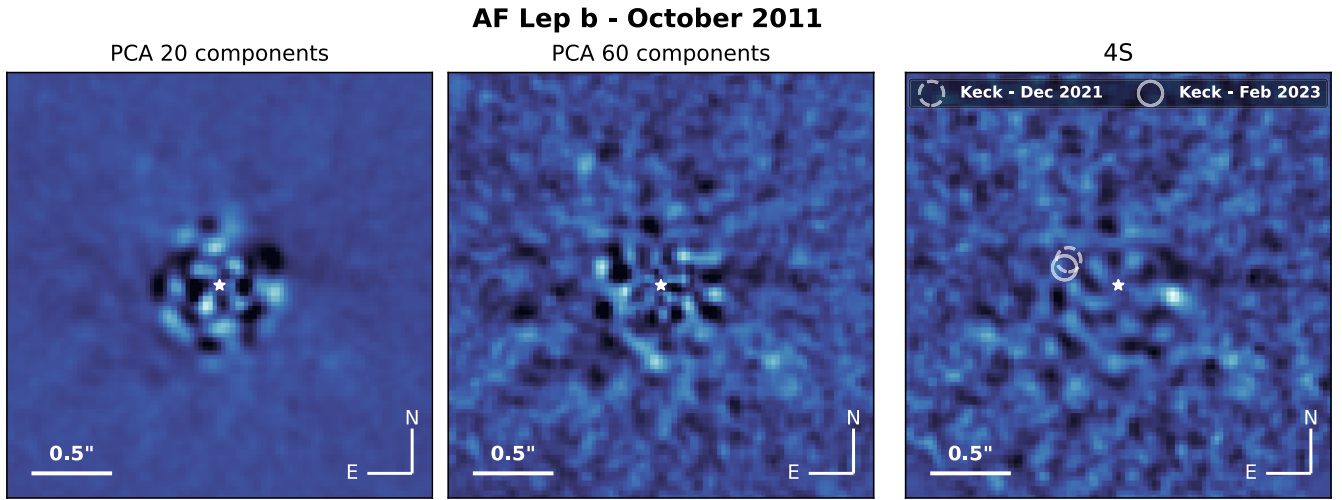


Figure 10. Residual images of the AF Lep data set (#12 in Table 1) taken with VLT-NACO in 2011. The left and middle panels show the residuals obtained with PCA with $K = 20$ and $K = 60$ components, respectively. The right panel shows the residual of 4S, which reveals the giant exoplanet AF Lep b at $S/N = 6.8$, 11 yr before its discovery. For better visibility, we filter the residual images with a Gaussian kernel ($\sigma = 0.8$ pixels). The S/N is calculated on the raw residuals before applying the Gaussian filter. An interactive version of the plot for different setups of PCA and 4S is available in our [online documentation](#). A gallery of all PCA and 4S residuals is given in Figure 18.

simplex minimization to find good initialization parameters. Afterward, we run the MCMC for 500 steps using 100 walkers. The minimization criterion is based on the Hessian (T. Stolker et al. 2019; V. Christiaens et al. 2023). Based on the injection of many fake planets, the MCMC procedure allows us to retrieve both the astrometry and photometry of the companion, as well as the corresponding error bars.

Our results are summarized in Table 2. The estimated photometry of $\Delta L' = 10.03^{+0.13}_{-0.12}$ mag is in perfect agreement (i.e., within one standard deviation) with the results of K. Franson et al. (2023), who reported $\Delta L' = 9.94 \pm 0.14$ mag. In Appendix B, we discuss an additional fake planet experiment that verifies that the signal fit with PCA is the source of the detection in 4S.

Orbit Fit. We combine the new relative astrometry data point obtained in this work with the relative astrometry points from the three AF Lep b detection papers (R. J. De Rosa et al. 2023; K. Franson et al. 2023; D. Mesa et al. 2023), Hipparcos (see F. van Leeuwen 2007) and Gaia (T. Prusti et al. 2016; A. Vallenari et al. 2023) acceleration, and archival RV data published by R. P. Butler et al. (2017). To combine the data into a single orbit fit, we use the MCMC-based orbit fitting package *orvara* (T. D. Brandt et al. 2021b). We fit for all orbital parameters and the primary and secondary mass, setting the default priors as defined in Table 4 of T. D. Brandt et al. (2021b). To stay comparable to the results in K. Franson et al. (2023), we also set a highly constrained Gaussian prior on the mass of AF Lep: $M = 1.2 \pm 0.06 M_{\odot}$ (P. Kervella et al. 2022). We run the Monte Carlo chain to 500,000 samples over 100 walkers, discarding the first 20% of each chain as burn-in. The parameter distributions and resulting fit values are reported in Figure 11 and Table 2. We report a dynamical mass estimate of $M = 3.74^{+0.53}_{-0.50} M_J$, confirming its status well within the planetary-mass regime. This estimate offers a significant improvement over and is compatible with the masses reported in K. Franson et al. (2023) and R. J. De Rosa et al. (2023).

As the additional VLT-NACO data point is almost precisely opposite to the more recent observations, it offers an ideal opportunity to constrain the orbit's eccentricity. Indeed, we find AF Lep b to move on a circular orbit ($\epsilon = 0.031^{+0.027}_{-0.020}$).

Table 2
Properties of AF Lep b

$\rho_{2011.800}$ (mas)	$323.24^{+6.71}_{-6.44}$
$\theta_{2011.800}$ (deg)	$258.81^{+0.53}_{-0.59}$
$\Delta L'$ (mag)	$10.03^{+0.13}_{-0.12}$
L' (mag)	$14.96^{+0.14}_{-0.13}$
$M_{L'}$ (mag)	$12.81^{+0.14}_{-0.13}$
Primary mass (M_{\odot})	$1.219^{+0.052}_{-0.054}$
Secondary mass (M_J)	$3.74^{+0.53}_{-0.50}$
Semimajor axis (au)	$9.01^{+0.20}_{-0.19}$
Inclination (deg)	$55.8^{+6.2}_{-7.2}$
Ascending node (deg)	$69.7^{+5.6}_{-5.4}$
Mean longitude at 2010.0 (deg)	170^{+11}_{-13}
Eccentricity	$0.031^{+0.027}_{-0.020}$
Argument of periastron (deg)	-10^{+66}_{-52}

Note. Reported uncertainties are the median with the 16% and 84% quantiles of the MCMC samples after burn-in.

Considering that, due to the low mass of the companion, it likely formed in situ in the protoplanetary disk, our measurement agrees much better with tentative evidence of planet-mass companions occurring preferentially on circular orbits (B. P. Bowler et al. 2020). We note that, due to the relative astrometry point being opposite, we do not improve the fit of the inclination. This circular orbit agrees with recent interferometric observations of the planet with VLTI-GRAVITY (W. O. Balmer et al. 2024).

In addition to the 2011 observation, AF Lep was observed a second time with VLT-NACO in 2015 (prog. ID 096.C-0679(A); R. Launhardt et al. 2020). Our reanalysis of this data set did not yield a detection. In 2015, the projected separation of the planet from the star is much smaller ($\lesssim 2\lambda/D$) than in 2011, resulting in higher-contrast requirements. Therefore, the nondetection in 2015 is consistent with our contrast limits and orbit fit.

7. Future Work

The detection of AF Lep b in archival VLT-NACO data demonstrates the potential of 4S to find previously missed

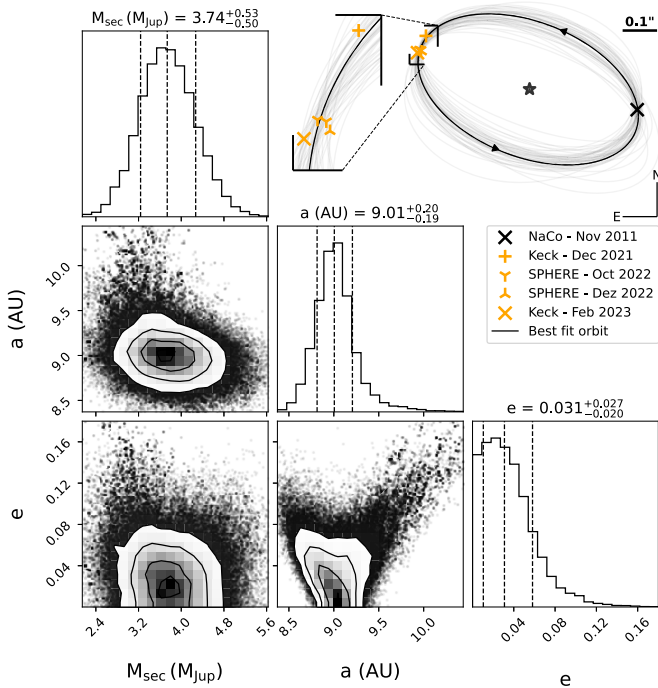


Figure 11. Results from the updated orbit fit of AF Lep b. The corner diagram reports the posterior distribution of an MCMC-based orbit fit using *orvara*. The median with the 16% and 84% quantiles is reported. The top right inset shows all available relative astrometry measurements with the best orbital fit. To illustrate the confidence of the fit, 100 orbits are drawn from the fit posterior (gray). For a summary of all results, see Table 2.

companions. The archive of VLT-NACO observations comprises more than 800 data sets from more than 450 stars (considering only L' -band data taken in pupil tracking mode). It is possible that a reanalysis of these data sets will reveal additional companions. Future work should consider a reanalysis not only for VLT-NACO data but also for other instruments. In this way, a better understanding of the contrast performance of 4S for different observing modes and instruments can be obtained. The use of saliency maps together with 4S may help us to better understand the optics and speckle pattern. If some difficult-to-interpret patterns are systematically visible in the saliency maps of an instrument, they may indicate overlooked effects in the instrument optics. This could open the door to new developments on the hardware side.

The primary drawback of 4S is its computational cost. Processing one data set takes about 10–30 minutes on a single NVIDIA A100/H100 GPU. While this computation time is unproblematic for detection purposes, it quickly becomes infeasible for scientific analysis. Extracting astrometry and photometry, or calculating detection limits, requires reprocessing the same data set hundreds of times. In fact, the most computationally expensive step in 4S is the optimization of the model matrix in step 2 (see Figure 6). We recompute this matrix from scratch for all our fake planet experiments. This might not be necessary. Future work should explore strategies to compute the matrix once and update it as the scientific analysis progresses. Fast GPU-based MCMC implementations compatible with our code are available, for example, in *pyro* (E. Bingham et al. 2019). The L-BFGS optimizer used in this work requires a lot of GPU memory, especially for large science frames. This sets the maximum resolution that can still be processed with one NVIDIA A100/H100 GPU to about

150×150 pixels, depending on the temporal resolution of the data set. We recommend cropping the science frames to benefit from the better contrast performance of 4S at close separations. Unfortunately, this memory requirement currently prevents us from processing the data sets of the exoplanet imaging data challenge (F. Cantalloube et al. 2015). Future work should focus on making our implementation more memory efficient, for example, by using gradient accumulation during optimization.

Further development of the method. The noise model of 4S is a simple linear model. Especially at separations close to the star and in data taken with a coronagraph, nonlinear artifacts can occur. For this purpose, the linear model of 4S in Equation (22) could be replaced by a neural network. Neural networks have already been proposed for speckle subtraction in HCI (see T. N. Wolf et al. 2023). In Section 5, we observed that a substantial part of the information used by 4S to explain the speckle noise are symmetries across the entire frame. Therefore, neural networks should not be used locally, as in T. N. Wolf et al. (2023), but globally. A possible direction for this could be masked autoencoders (K. He et al. 2022).

A number of new postprocessing methods based on PCA have been proposed in the literature. One example is the RSM map of C.-H. Dahlqvist et al. (2020), which replaces the last step in PCA, the averaging along time, and allows combining multiple postprocessing algorithms. Another example is the combination of PCA with forward modeling (FMMF; J.-B. Ruffio et al. 2017). The basic steps of 4S are identical to those of PCA, which is why methods such as RSM and FMMF can also be used with 4S. For this reason, 4S should be seen as a building block for new postprocessing methods.

Our analysis of the saliency maps in Section 5 has shown that the noise model learned by 4S hardly changes over years. This is a strong indicator that 4S could also be used with RDI. For RDI, the noise model of 4S needs to be learned only once and can then be applied to many data sets. This procedure also avoids the high computational costs of 4S for ADI. The parameters of 4S could be conditioned on observation conditions, as suggested by T. D. Gebhard et al. (2022) for half-sibling regression. In combination with RDI, we might be able to obtain a universal model of the speckle noise that utilizes all previously collected data.

8. Summary and Conclusions

In this paper, we have gained a deep understanding for the loss of planet signal in PCA-based postprocessing. Using saliency maps, we identified two main reasons for this loss: first, the convergence of the PCA noise model to the telescope PSF; and second, a problematic loss function that misleads the noise estimate to fit the planet signal. Building on these insights, we developed a new postprocessing algorithm that makes a step toward overcoming the limitations of PCA. We named this method 4S. 4S is characterized by three innovations: (1) a linear noise model with a right reason constraint, (2) a new loss function that is invariant to the planet signal, and (3) a domain-knowledge-informed regularization.

We applied 4S to archival data from the VLT-NACO instrument to compare its contrast performance to PCA. On a sample of 11 data sets, we observed a contrast improvement of up to 1.5 mag, especially at close separations to the star. In addition to achieving deeper contrast, 4S produces residual noise that is more Gaussian compared to PCA (see

Appendix C), reducing the risk of false positives. The main reason for this improvement is the ability of 4S to learn a better representation of the speckle noise. A detailed analysis revealed that 4S can recover the speckle pattern that we would expect from the theory. The increase in contrast performance allowed us to detect the giant exoplanet AF Lep b in archival data from 2011, 11 yr before its discovery. Using this new astrometric point, we significantly improve the constraints on the orbital parameters of the planet.

The introduction of 4S marks a leap for the postprocessing of HCI data. Looking ahead, our findings have the potential to boost exoplanet detection on archival data and for future observations. We believe that the findings from our work will provide inspiration for the development of new postprocessing methods beyond 4S.

Acknowledgments

We thank the anonymous referee for a critical and constructive review of the original manuscript, which helped us improve the clarity and 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 have been carried out within the framework of the National Centre of Competence in Research PlanetS supported by the Swiss National Science Foundation (SNSF). J.H. acknowledges the financial support from the Swiss National Science Foundation (SNSF) under project grant No. 200020_200399. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 819155).

Software: This work has made use of a number of open-source Python packages, including `astropy` (The Astropy Collaboration et al. 2013, 2018, 2022), `matplotlib` (J. D. Hunter 2007), `numpy` (C. R. Harris et al. 2020), `pandas` (W. McKinney 2010), `photutils` (L. Bradley et al. 2024), `scikit-image` (S. van der Walt et al. 2014), `scikit-learn` (F. Pedregosa et al. 2011), `scipy` (P. Virtanen et al. 2020), `seaborn` (M. Waskom 2021), `torch` (A. Paszke et al. 2019), and `tqdm` (C. O. da Costa-Luis 2019).

Author Contributions

M.J.B. carried out the main analyses, developed the 4S algorithm, and wrote the manuscript. He programmed the publicly available Python package `fours` and wrote the documentation page. T.D.G. helped with the analysis and the development of 4S. His previous work on half-sibling regression algorithm has set the basis for this paper. F.A.D. carried out the orbit fits of AF Lep b and helped write the corresponding section. S.P.Q. and B.S. provided access to the compute resources needed for this project. All authors discussed the algorithm, analysis, and results. All authors commented on the manuscript.

Appendix A

Linear PSF Subtraction Compared

The right reason mask of 4S introduced in Section 4.1 has similarities with the masks proposed by other postprocessing algorithms. Already the PSF subtraction based on the LOCI algorithm used PSF-sized masks to protect the signal of the planet (D. Lafrenière et al. 2007; C. Marois et al. 2010;




Figure 12. Illustration of the noise estimation implemented by the linear postprocessing algorithms LOCI, PCA, and 4S. The noise estimate $\hat{x}_{t,l}$ at time t and spatial position l is a weighted sum of the values marked in orange, weighted by the parameters in dark red. The dotted circle marks the position of the mask. PCA can be interpreted in two ways, as explained in the text.

R. Soummer et al. 2011). Within this framework, two types of masks can be distinguished: (1) temporal masks, for example, based on field rotation; and (2) spatial PSF-sized masks, applied to the data. Both masks share the same idea: if the planetary signal is excluded from the fit, only the speckle noise can be captured. The temporal exclusion ensures that the planet is not part of the reference frames, while the spatial mask ensures that it is excluded from the frame we want to fit. The same masks are commonly used with PCA (R. Soummer et al. 2012; O. Absil et al. 2013; C. A. Gomez Gonzalez et al. 2017). While the shape of these masks is similar to the right reason mask, the way they work is different. Unlike previous work, the right reason mask of 4S is applied to the model *parameters* during the optimization and not to the *data*.

LOCI, PCA, and 4S are all linear PSF subtraction algorithms. To better understand their differences and masking strategies, let us consider a single science frame \mathbf{x}_t at time t . Using LOCI, PCA, or 4S, we estimate the speckle noise for this particular science frame \hat{x}_t . For simplicity, let us now consider a single pixel position l in this noise estimate $\hat{x}_{t,l}$. All three PSF subtraction algorithms have a different approach to calculating $\hat{x}_{t,l}$ (see Figure 12).

LOCI. This uses a linear combination of reference frames to fit and subtract the speckle noise (D. Lafrenière et al. 2007). In the case of ADI, these reference frames are other science frames from the same data set but excluding temporally adjacent frames (see, e.g., Figure 2 in C. Marois et al. 2010). The noise estimate $\hat{x}_{t,l}$ for LOCI is a weighted sum of the reference frames at the position l :

$$\hat{x}_{t,l}^{\text{LOCI}} = \underbrace{(X_{:,l})^T}_{1 \times N} \cdot \underbrace{\mathbf{c}}_{N \times 1}, \quad (\text{A1})$$

where $\mathbf{c} \in \mathbb{R}^N$ are the weights of the sum. Depending on the temporal exclusion criterion, some values in \mathbf{c} can be zero. The weights \mathbf{c} are fitted while masking a PSF-sized region around the position l . The region on which the fit is based is called the *O*-zone, while the region to which the model is applied to is called the *S*-zone. The *S*-zone can be a single pixel (R. Soummer et al. 2011).

PCA first interpretation. Instead of a linear combination of reference frames, PCA estimates the noise using a linear combination of the principal components \mathbf{U} . Since the number of components is smaller than the number of reference frames used by LOCI ($K \ll N$), PCA can be seen as a regularized version of LOCI. The noise estimate $\hat{x}_{t,l}$ for PCA is a weighted sum of the principal components at the position l :

$$\hat{x}_{t,l}^{\text{PCA}} = \underbrace{\mathbf{U}_{l,:}}_{1 \times K} \cdot \underbrace{\mathbf{U}^T \mathbf{x}_t}_{K \times 1}. \quad (\text{A2})$$

The weights of the sum are given by the projection of the science frame onto the component matrix $\mathbf{U}^T \mathbf{x}_t$. Similar to LOCI, it is possible to exclude a PSF-sized region in this projection by masking the data.

PCA second interpretation. If we rearrange Equation (A2), we obtain the alternative interpretation of PCA:

$$\hat{x}_{t,l}^{\text{PCA}} = \underbrace{\mathbf{x}_t^T}_{1 \times D} \cdot \underbrace{(\mathbf{U}\mathbf{U}^T)_{:,l}}_{D \times 1}. \quad (\text{A3})$$

The noise estimate $\hat{x}_{t,l}$ is not only a weighted sum of the principal components but also a weighted sum of the pixels in the science frame \mathbf{x}_t . The weights of this sum are given by the l th column vector of the projection matrix $(\mathbf{U}\mathbf{U}^T)_{:,l}$. The saliency maps discussed in Section 3 are based on this interpretation of PCA. A mask can be applied directly to the data \mathbf{x}_t .

4S. Similar to the second interpretation of PCA, the noise estimate $\hat{x}_{t,l}$ in 4S is a weighted sum of the pixels in science frame \mathbf{x}_t :

$$\hat{x}_{t,l}^{\text{4S}} = \underbrace{\mathbf{x}_t^T}_{1 \times D} \cdot \underbrace{\mathbf{B}_{:,l}}_{D \times 1}. \quad (\text{A4})$$

The weights of the sum are given by the l th column of the matrix $\mathbf{B}_{:,l}$. This means that, unlike LOCI, the noise estimate of 4S is a weighted sum not of the N reference frames but of the D pixels in the science frame \mathbf{x}_t .

A.1. The Masks of PCA and 4S

The principal components are not orthogonal to the signal of the planet. This means that a linear combination of the components can fit and subtract the signal of the planet. Similar to the masks used in LOCI, we can apply a mask in the

Saliency Maps at $2.5 \lambda/D$




Figure 13. Comparison of the saliency maps of the masked version of PCA and 4S. The number of components for PCA was selected according to the best contrast performance at $2.5\lambda/D$, which is reached for $K = 90$. The white plus sign represents the position for which the saliency map is calculated. The red star marks the center of the frame.

projection step of PCA in Equation (A2). For comparison with 4S, we use the same PSF-sized masks \mathbf{m}_l as for the right reason mask:

$$\begin{aligned} \hat{x}_{t,l} &= \mathbf{U}_{l,:} \cdot \mathbf{U}^T (\mathbf{x}_t \circ \mathbf{m}_l) \\ &= (\mathbf{x}_t \circ \mathbf{m}_l)^T (\mathbf{U}\mathbf{U}^T)_{:,l} \\ &= \mathbf{x}_t^T ((\mathbf{U}\mathbf{U}^T)_{:,l} \circ \mathbf{m}_l). \end{aligned}$$

This means that instead of masking the data we can mask the l th column vector of the projection matrix $(\mathbf{U}\mathbf{U}^T)_{:,l}$ and obtain the same result. The projection matrix of PCA ($\mathbf{U}\mathbf{U}^T$) and the model matrix of 4S (\mathbf{B}) are both square matrices with shape $D \times D$. In both cases the mask \mathbf{m}_l is applied to the diagonal of the matrix.

In Figure 13 we compare the saliency map of the masked version of PCA with the saliency map of 4S. As shown in the figure, the information used by PCA is very different from the information used by 4S. While the point symmetries of the speckle noise are clearly identified by 4S, PCA uses information from the whole image.

If the masks are applied along the same dimension, what is the cause for this difference?

1. The linear model of 4S is more general: the column vectors of the projection matrix in PCA $(\mathbf{U}\mathbf{U}^T)_{:,l}$ are linearly dependent. The rank of the matrix is set by the number of components:

$$\text{rank}(\mathbf{U}\mathbf{U}^T) = K. \quad (\text{A5})$$

Further, the projection matrix is symmetric, $(\mathbf{U}\mathbf{U}^T)^T = \mathbf{U}\mathbf{U}^T$, and satisfies

$$\mathbf{U}\mathbf{U}^T \cdot \mathbf{U}\mathbf{U}^T = \mathbf{U}\mathbf{U}^T. \quad (\text{A6})$$

In contrast, the column vectors $\mathbf{B}_{:,l}$ are generally not dependent, and the rank of the matrix is unconstrained. Further, \mathbf{B} does not have to be symmetric, $\mathbf{B}^T \neq \mathbf{B}$, and is not a projection matrix, $\mathbf{B}^2 \neq \mathbf{B}$. The linear model of 4S has fewer constraints and can therefore be seen as a generalization of the PCA framework.

2. For PCA, the component matrix \mathbf{U} is the result of the eigenvalue decomposition (or SVD). The mask can be applied during the projection step. For large numbers of principal components, the information used by PCA

converges to a de facto identity (see Section 3.2). Even if we mask this identity in UU^T , there is no incentive for the model to explore other areas of the data. The right reason mask in 4S is embedded in \mathbf{B} . The parameters of 4S are optimized in an end-to-end fashion, taking into account the right reason mask. Since the mask is part of the optimization, the model is forced to search for additional information about the speckle noise.

3. The loss function of PCA is a classical least-squares minimization (see Section 3.3). This loss function is replaced in 4S, making it invariant to the signal of the planet (see Section 4.2). Temporal exclusion criteria as introduced for LOCI (C. Marois et al. 2010) are not needed.

Appendix B

Fake Planet Experiments for AF Lep b

In Section 6 we estimate the astrometry and photometry of AF Lep b using PCA. Ideally, we would like to use 4S to perform this analysis. Due to the high computational cost, this is currently not possible, a limitation we plan to overcome in future work using the ideas presented in Section 7. To verify that the astrometry and photometry extracted by PCA are consistent with the detection in 4S, we insert a negative fake planet with $\Delta L' = 1003$ mag at the position obtained with PCA (see Table 2). Afterward, we run 4S again on the data without the planet. As shown in the top right panel of Figure 14, the signal of AF Lep b is removed and the residual of 4S is flat. As an additional test, we insert fake planets with $\Delta L' = 1003$ mag at the same separation as AF Lep b but at 40

different position angles. For each fake planet, we rerun 4S and compute the S/N of the artificial companion. On average, the fake planets achieve $S/N = 5.7 \pm 1.5$, which is consistent with the real companion $S/N = 6.8$. Examples of the fake planet residuals and the S/N values computed for all 40 fake planets are shown in the middle and bottom rows of Figure 14.

The S/N varies considerably for different position angles, an effect already reported for VLT-NACO data in M. J. Bonse et al. (2023). Based on this experiment, we conclude that the photometry and position extracted using PCA are consistent with the detection in 4S.

Appendix C

Residual Noise Distribution

A common assumption for the quantification of detections and nondetections in HCI is that the residual noise after postprocessing with PCA is Gaussian distributed (see, e.g., D. Mawet et al. 2014). This assumption is generally justified by the central limit theorem (we average many frames along time) and the whitening effect of PCA. Nevertheless, the residual noise after postprocessing often deviates from Gaussian noise (B. Pairet et al. 2019; M. J. Bonse et al. 2023). This is problematic because the risk of false positives (we claim a detection that is just noise) depends on the type of noise. More specifically, the probability of observing a value of $S/N = 5$ that is just noise is higher if the noise is heavy tailed, that is, the frequency of bright values such as speckles is higher than for Gaussian noise. M. J. Bonse et al. (2023) found that if the residual noise follows a Laplacian distribution, the final contrast could be overestimated by up to 1 mag if the test of D. Mawet et al. (2014) is used. In practice, the true noise distribution is usually unknown, making it difficult to compare limits across data sets. This is also true if we want to compare different postprocessing techniques. We study the residuals of PCA and 4S for data sets #1–#9 (see Table 1) to get a better understanding of their noise distributions. Data sets #10 and #11 are excluded owing to strong saturation at close separations. We consider two regions: (1) within $(2.5\text{--}4.5)\lambda/D$ distance to the star, and (2) within $(5.0\text{--}7.0)\lambda/D$ distance to the star. For each data set and region we plot Q–Q plots and calculate the coefficient of determination R^2 (see, e.g., B. Pairet et al. 2019, for a detailed explanation). A comparison of the residual noise for data sets #5 and #9 is shown in Figure 15. The Q–Q plots shown in the figure compare the pixel noise distribution with Gaussian noise. If the residual noise were perfectly Gaussian, we would expect all values to lie on the gray diagonal line. The better the observed noise can be explained by Gaussian noise, the closer the coefficient of determination R^2 is to 1. For PCA, the large values are above the diagonal line, indicating that the noise is heavy tailed. In the case of 4S, the residual noise is much closer to Gaussian, although it is still not perfectly Gaussian. The quantification based on the R^2 values confirms this observation. A summary of the results for all data sets is shown in Figure 16. In general, the residual noise at close separations tends to be more affected by speckle and is therefore often more non-Gaussian. For most data sets, especially those for which PCA deviates the most from Gaussian noise, 4S improves the Gaussianity of the residuals. This means that 4S not only improves the contrast limits, as shown in Section 5, but also reduces the risk of false positives.




Figure 14. Verification of the astrometry and photometry estimated with the help of PCA. The top two panels show the residuals of 4S for AF Lep b (left) and after insertion of a negative fake planet (right) based on the results given in Table 2. In the middle row, fake planet residuals are shown that are inserted into the data at the same separation and contrast as AF Lep b, but at different position angles. The bottom panel summarizes the S/N values of the 40 fake planet experiments described in the text.

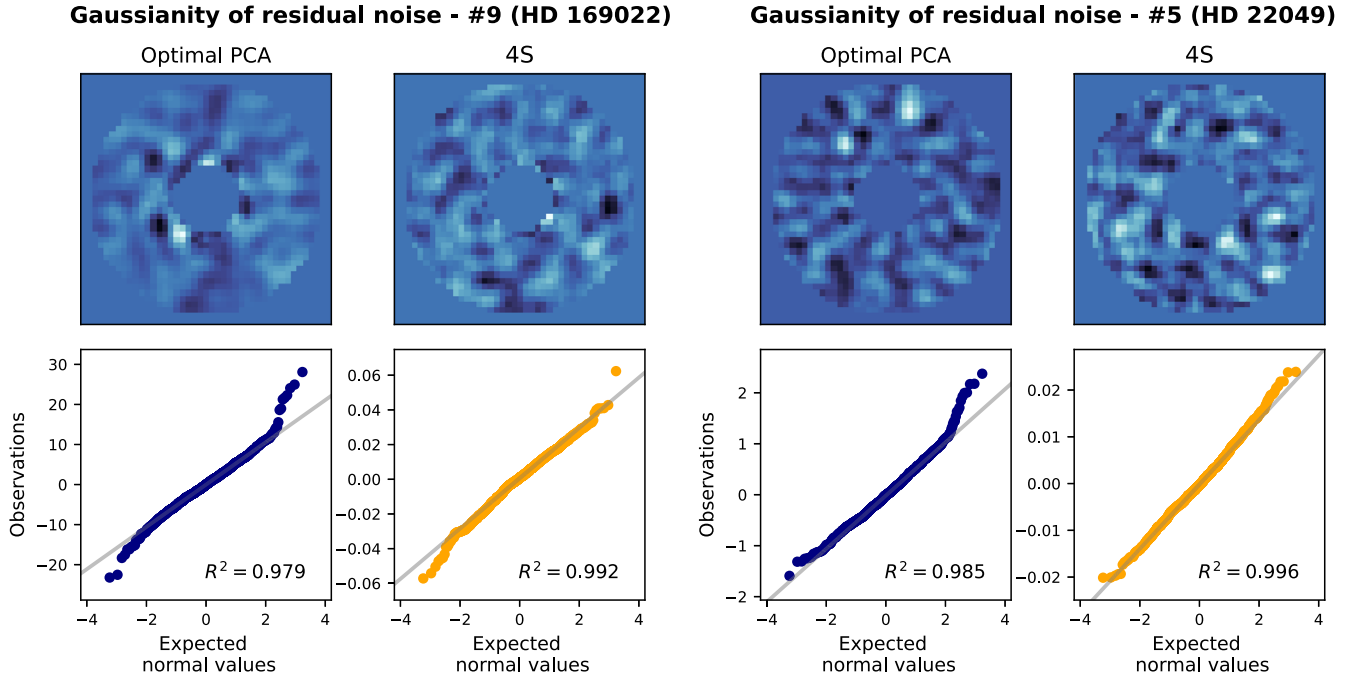


Figure 15. Two examples of the Gaussianity test performed in Figure 16. The top row shows examples of the residuals obtained with PCA and 4S. The pixels outside $(2.5-4.5)\lambda/$ were masked to extract the noise. Q-Q plots are shown in the bottom row. If the noise were perfectly Gaussian, we would expect all values to be on the gray diagonal line and the value of R^2 to be close to 1.




Figure 16. Comparison of the Gaussianity of the noise in the residuals of 4S and PCA for data sets #1–#9. If the coefficient of determination R^2 is closer to 1, the noise is better explained by Gaussian noise. We choose the number of principal components K and the regularization strength λ of 4S that maximize the contrast (optimal PCA).




Figure 17. PCA saliency maps for different numbers of principal components K and separations from the star. The saliency maps show which information from the science frames is used by PCA to estimate the noise at the position marked by the white circle. Especially for large K and separations close to the star, only the information within the white circle is used, resulting in a substantial loss of the planet signal. The red star marks the center of the frame.















Residuals - AF Lep b (2011)

Figure 18. Residual images of the AF Lep data set shown in Figure 10 for several different settings of the algorithm’s hyperparameters (principal components K for PCA and regularization strength λ for 4S). The planet AF Lep b is barely visible in the PCA residuals. For 4S, any choice of parameters yields a clear detection, highlighting the robustness of the method.

ORCID iDs

Markus J. Bonse  <https://orcid.org/0000-0003-2202-1745>
 Timothy D. Gebhard  <https://orcid.org/0000-0001-9310-8579>
 Felix A. Dannert  <https://orcid.org/0000-0002-5476-2663>
 Olivier Absil  <https://orcid.org/0000-0002-4006-6237>
 Faustine Cantalloube  <https://orcid.org/0000-0002-3968-3780>
 Valentin Christiaens  <https://orcid.org/0000-0002-0101-8814>
 Gabriele Cugno  <https://orcid.org/0000-0001-7255-3251>
 Emily O. Garvin  <https://orcid.org/0000-0003-2530-9330>
 Jean Hayoz  <https://orcid.org/0000-0003-3768-5712>
 Markus Kasper  <https://orcid.org/0000-0002-8425-6606>
 Elisabeth Matthews  <https://orcid.org/0000-0003-0593-1560>
 Bernhard Schölkopf  <https://orcid.org/0000-0002-8177-0925>
 Sascha P. Quanz  <https://orcid.org/0000-0003-3829-7412>

References

- Absil, O., Milli, J., Mawet, D., et al. 2013, *A&A*, **559**, L12
 Ali, S., Abuhmed, T., El-Sappagh, S., et al. 2023, *Inf. Fusion*, **99**, 101805
 Amara, A., & Quanz, S. P. 2012, *MNRAS*, **427**, 948
 Arcidiacono, C., & Simoncini, V. 2018, *Proc. SPIE*, **10703**, 1070331
 Balmer, W. O., Franson, K., Chomez, A., et al. 2025, *AJ*, **169**, 30
 Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. 2018, *JMLR*, **18**, 1
 Bell, C. P. M., Mamajek, E. E., & Naylor, T. 2015, *MNRAS*, **454**, 593
 Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, **631**, A155
 Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, *JMLR*, **20**, 1
 Bishop, C. M. 2006, *Pattern Recognition and Machine Learning*, Information Science and Statistics (New York: Springer)
 Bloemhof, E. E. 2007, *OExpr*, **15**, 4705
 Bloemhof, E. E., Dekany, R. G., Troy, M., & Oppenheimer, B. R. 2001, *ApJL*, **558**, L71
 Bonse, M. J. 2024a, Intermediate Results for: use the 4S (Signal-Safe Speckle Subtraction), v0.1, Zenodo, doi:[10.5281/zenodo.11457071](https://doi.org/10.5281/zenodo.11457071)
 Bonse, M. J. 2024b, Raw Data for: Use the 4S (Signal-Safe Speckle Subtraction), v0.1, Zenodo, doi:[10.5281/zenodo.11456704](https://doi.org/10.5281/zenodo.11456704)
 Bonse, M. J., Garvin, E. O., Gebhard, T. D., et al. 2023, *AJ*, **166**, 71
 Bonse, M. J., Quanz, S. P., & Amara, A. 2018, arXiv:[1804.05063](https://arxiv.org/abs/1804.05063)
 Bowler, B. P., Blunt, S. C., & Nielsen, E. L. 2020, *AJ*, **159**, 63
 Bradley, L., Sipőcz, B., & Robitaille, T. 2024, *astropy/photutils: v1.12.0*, Zenodo, doi:[10.5281/zenodo.10967176](https://doi.org/10.5281/zenodo.10967176)
 Brandt, G. M., Dupuy, T. J., Li, Y., et al. 2021a, *AJ*, **162**, 301
 Brandt, T. D., Dupuy, T. J., Li, Y., et al. 2021b, *AJ*, **162**, 186
 Butler, R. P., Vogt, S. S., Laughlin, G., et al. 2017, *AJ*, **153**, 208
 Cantalloube, F., Gomez-Gonzalez, C., Absil, O., et al. 2020, *Proc. SPIE*, **11448**, 114485A
 Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. 2015, *A&A*, **582**, A89
 Cantero, C., Absil, O., Dahlqvist, C. H., & Van Droogenbroeck, M. 2023, *A&A*, **680**, A86
 Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017, *A&A*, **605**, L9
 Cheetham, A., Ségransan, D., Peretti, S., et al. 2018, *A&A*, **614**, A16
 Christiaens, V., Gonzalez, C. A. G., Farkas, R., et al. 2023, *JOSS*, **8**, 4774
 Cugno, G., Pearce, T. D., Launhardt, R., et al. 2023, *A&A*, **669**, A145
 Currie, T., Biller, B., Lagrange, A.-M., et al. 2023, in *Protostars and Planets VII* (San Francisco, CA: ASP)
 da Costa-Luis, C. O. 2019, *JOSS*, **4**, 1277
 Daglayan, H., Vary, S., Leplat, V., Gillis, N., & Absil, P.-A. 2023, arXiv:[2304.03619](https://arxiv.org/abs/2304.03619)
 Dahlqvist, C.-H., & Absil, O. 2021, *A&A*, **646**, A49
 Dahlqvist, C.-H., Cantalloube, F., & Absil, O. 2020, *A&A*, **633**, A95
 De Rosa, R. J., Nielsen, E. L., Wahhaj, Z., et al. 2023, *A&A*, **672**, A94
 Flasseur, O., Bodrito, T., Mairal, J., et al. 2024, *MNRAS*, **527**, 1534
 Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2018, *A&A*, **618**, A138
 Franson, K., Bowler, B. P., Zhou, Y., et al. 2023, *ApJL*, **950**, L19
 Fulton, B. J., Rosenthal, L. J., Hirsch, L. A., et al. 2021, *ApJS*, **255**, 14
 Gebhard, T. D., Bonse, M. J., Quanz, S. P., & Schölkopf, B. 2022, *A&A*, **666**, A9
 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
 Gray, R. O., Corbally, C. J., Garrison, R. F., et al. 2006, *AJ*, **132**, 161
 Halko, N., Martinsson, P. G., & Tropp, J. A. 2011, *SIAMR*, **53**, 217
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 He, K., Chen, X., Xie, S., et al. 2022, in *IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)* (New York: IEEE), 15979
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Jaderberg, M., Simonyan, K., Zisserman, A., et al. 2015, arXiv:[1506.02025](https://arxiv.org/abs/1506.02025)
 Janson, M., Squicciarini, V., Delorme, P., et al. 2021, *A&A*, **646**, A164
 Jensen-Clem, R., Mawet, D., Gomez Gonzalez, C. A., et al. 2017, *AJ*, **155**, 19
 Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, **127**, 890
 Kenworthy, M. A., Codona, J. L., Hinz, P. M., et al. 2007, *ApJ*, **660**, 762
 Kervella, P., Arenou, F., & Thévenin, F. 2022, *A&A*, **657**, A7
 Kingma, D. P., & Ba, J. 2015, arXiv:[1412.6980](https://arxiv.org/abs/1412.6980)
 Kuhn, J. R., Potter, D., & Parise, B. 2001, *ApJL*, **553**, L189
 Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, **660**, 770
 Lagrange, A. M., Philipot, F., Rubini, P., et al. 2023, *A&A*, **677**, A71
 Langlois, M., Gratton, R., Lagrange, A.-M., et al. 2021, *A&A*, **651**, A71
 Launhardt, R., Henning, T., Quirrenbach, A., et al. 2020, *A&A*, **635**, A162
 Lewis, B., Fitzgerald, M. P., Dodkins, R. H., Davis, K. K., & Lin, J. 2023, *AJ*, **165**, 59
 Liu, D. C., & Nocedal, J. 1989, *MatPr*, **45**, 503
 Long, J. D., Males, J. R., Haffert, S. Y., et al. 2023, *AJ*, **165**, 216
 Macintosh, B., Graham, J. R., Barman, T., et al. 2015, *Sci*, **350**, 64
 Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *PNAS*, **111**, 12661
 Males, J. R., Close, L. M., Miller, K., et al. 2018, *Proc. SPIE*, **10703**, 1070309
 Males, J. R., Fitzgerald, M. P., Belikov, R., & Guyon, O. 2021, *PASP*, **133**, 104504
 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., Macintosh, B., & Véran, J.-P. 2010, *Proc. SPIE*, **7736**, 77361J
 Mawet, D., Absil, O., Delacroix, C., et al. 2013, *A&A*, **552**, L13
 Mawet, D., Absil, O., Montagnier, G., et al. 2012, *A&A*, **544**, A131
 Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, **792**, 97
 Mawet, D., Serabyn, E., Liewer, K., et al. 2009, *ApJ*, **709**, 53
 McKinney, W. 2010, in *Proc. the 9th Python in Science Conference*, ed. S. van der Walt & J. Millman (Austin, TX: SciPy), 56
 Mesa, D., Gratton, R., Kervella, P., et al. 2023, *A&A*, **672**, A93
 Meshkat, T., Kenworthy, M. A., Quanz, S. P., & Amara, A. 2013, *ApJ*, **780**, 17
 N'Diaye, M., Pueyo, L., & Soummer, R. 2015, *ApJ*, **799**, 225
 Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019, *AJ*, **158**, 13
 Otten, G. 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
 Paszke, A., Gross, S., Massa, F., et al. 2019, arXiv:[1912.01703](https://arxiv.org/abs/1912.01703)
 Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825
 Perrin, M. D., Sivaramakrishnan, A., Makidon, R. B., Oppenheimer, B. R., & Graham, J. R. 2003, *ApJ*, **596**, 702
 Prusti, T., de Bruijne, J., H., J., Brown, A. G. A., et al. 2016, *A&A*, **595**, A1
 Pueyo, L. 2016, *ApJ*, **29**
 Quanz, S. P., Schmid, H. M., Geissler, K., et al. 2011, *ApJ*, **738**, 23
 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, *A&A*, **553**, A60
 Ren, B., Pueyo, L., Zhu, G. B., Debes, J., & Duchêne, G. 2018, *ApJ*, **852**, 104
 Ribak, E. N., & Gladysz, S. 2008, *OExpr*, **16**, 15553
 Ribeiro, M. T., Singh, S., & Guestrin, C. 2016, in *Proc. the 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining* (New York: ACM), 1135
 Ross, A. S., Hughes, M. C., & Doshi-Velez, F. 2017, arXiv:[1703.03717](https://arxiv.org/abs/1703.03717)
 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
 Samland, M., Bouwman, J., Hogg, D. W., et al. 2021, *A&A*, **646**, A24
 Schramowski, P., Stammer, W., Teso, S., et al. 2020, *Nat. Mach. Intell.*, **2**, 476
 Simonyan, K., Vedaldi, A., & Zisserman, A. 2013, arXiv:[1312.6034](https://arxiv.org/abs/1312.6034)
 Smilkov, D., Thorat, N., Kim, B., Viégas, F., & Wattenberg, M. 2017, arXiv:[1706.03825](https://arxiv.org/abs/1706.03825)
 Snik, F., Otten, G., Kenworthy, M., et al. 2012, *Proc. SPIE*, **8450**, 224
 Soummer, R., Hagan, J. B., Pueyo, L., et al. 2011, *ApJ*, **741**, 55
 Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJL*, **755**, L28
 Sparks, W. B., & Ford, H. C. 2002, *ApJ*, **578**, 543

- Stolker, T., Bonse, M. J., Quanz, S. P., et al. 2019, *A&A*, **621**, A59
- Sundararajan, M., Taly, A., & Yan, Q. 2017, in ICML (Seattle, WA: PMLR), 3319, <https://proceedings.mlr.press/v70/sundararajan17a.html>
- Sutskever, I., Martens, J., Dahl, G., & Hinton, G. 2013, in ICML (Seattle, WA: PMLR), 1139, <https://proceedings.mlr.press/v28/sutskever13.html>
- The Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, **935**, 167
- The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- The Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Thompson, W., & Marois, C. 2021, *AJ*, **161**, 236
- Tobin, T. L., Currie, T., Li, Y., et al. 2024, *AJ*, **167**, 205
- Vallenari, A., Brown, A. G. A., Prusti, T., et al. 2023, *A&A*, **674**, A1
- van der Walt, S., Schnberger, J. L., Nunez-Iglesias, J., et al. 2014, *PeerJ*, **2**, e453
- van Leeuwen, F. 2007, *A&A*, **474**, 653
- Vigan, A., Fontanive, C., Meyer, M., et al. 2021, *A&A*, **651**, A72
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Wahhaj, Z., Cieza, L. A., Mawet, D., et al. 2015, *A&A*, **581**, A24
- Wang, J. J., Ruffio, J.-B., De Rosa, R. J., et al., 2015 pyKLIP: PSF Subtraction for Exoplanets and Disks, Astrophysics Source Code Library, ascl:1506.001
- Waskom, M. 2021, *JOSS*, **6**, 3021
- Wolf, T. N., Jones, B. A., & Bowler, B. P. 2024, *AJ*, **167**, 92
- Zúñiga-Fernández, S., Bayo, A., Elliott, P., et al. 2021, *A&A*, **645**, A30