



Robustness to bad frames in angular differential imaging: a local weighting approach

O. Flasseur, Loïc Denis, E. Thiébaud, M. Langlois

► To cite this version:

O. Flasseur, Loïc Denis, E. Thiébaud, M. Langlois. Robustness to bad frames in angular differential imaging: a local weighting approach. Astronomy and Astrophysics - A&A, In press, 10.1051/0004-6361/201935859 . ujm-02450679

HAL Id: ujm-02450679

<https://ujm.hal.science/ujm-02450679>

Submitted on 19 Apr 2021

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Robustness to bad frames in angular differential imaging: a local weighting approach

Olivier Flasseur¹, Loïc Denis¹, Éric Thiébaud², and Maud Langlois²

¹ Université de Lyon, UJM-Saint-Etienne, CNRS, Institut d'Optique Graduate School, Laboratoire Hubert Curien UMR 5516, 42023 Saint-Etienne, France
e-mail: surname.name@univ-st-etienne.fr

² Université de Lyon, Université Lyon1, ENS de Lyon, CNRS, Centre de Recherche Astrophysique de Lyon UMR 5574, 69230 Saint-Genis-Laval, France
e-mail: surname.name@univ-lyon1.fr

Received 9 May 2019 / Accepted 7 November 2019

ABSTRACT

Context. The detection of exoplanets by direct imaging is very challenging. It requires an extreme adaptive-optics (AO) system and a coronagraph as well as suitable observing strategies. In angular differential imaging, the signal-to-noise ratio is improved by combining several observations.

Aims. Due to the evolution of the observation conditions and of the AO correction, the quality of the observations may vary significantly during the observing sequence. It is common practice to reject images of comparatively poor quality. We aim to decipher when this selection should be performed and what its impact on detection performance is.

Methods. Rather than discarding a full image, we study the local fluctuations of the signal at each frame and derive weighting maps for each frame. These fluctuations are modeled locally directly from the data through the spatio-temporal covariance of small image patches. The weights derived from the temporal variances can be used to improve the robustness of the detection step and reduce estimation errors of both the astrometry and photometry. The impact of bad frames can be analyzed by statistically characterizing the detection and estimation performance.

Results. When used together with a modeling of the spatial covariances (PACO algorithm), these weights improve the robustness of the detection method.

Conclusions. The spatio-temporal modeling of the background fluctuations provides a way to exploit all acquired frames. In the case of bad frames, areas with larger fluctuations are discarded by a weighting strategy and do not corrupt the detection map or the astrometric and photometric estimations. Other areas of better quality are preserved and are included to detect and characterize sources.

Key words. techniques: image processing – techniques: high angular resolution – methods: statistical – methods: data analysis

1. Introduction

Direct imaging from the Earth is a method of choice for the detection and characterization of exoplanets (Traub & Oppenheimer 2010). By confronting numerical models of estimated astrometry and photometry (Vigan et al. 2010), direct imaging gives access to quantitative physical properties on the detected sources such as age, mass, and effective temperature (Chabrier et al. 2000; Allard et al. 2003, 2007). The most efficient observation strategy is based on angular differential imaging (ADI, see Marois et al. 2006) which consists of tracking the observation target over time (the telescope pupil remaining stable while the whole field of view rotates). A temporal sequence of images is thus acquired in which the companions describe a predictable motion in time while the speckle background remains quasi-static. This observation method is used routinely on the main instruments dedicated to the direct detection of exoplanets such as VLT/SPHERE (Beuzit et al. 2019; Chauvin et al. 2017), GEMINI/GPI (Macintosh et al. 2014; Greenbaum et al. 2018), Keck/VORTEX (Howard et al. 2010; Mawet et al. 2019), Magellan/MagAO (Morzinski et al. 2014; Wagner et al. 2018), and Subaru/SCEXAO (Jovanovic et al. 2015; Currie et al. 2017). The recorded ADI images are then combined by a data-processing algorithm to disentangle the signature of the exoplanets from the speckles (Lafrenière et al. 2007; Marois et al.

2013, 2014; Soummer et al. 2012; Smith et al. 2009; Cantalloube 2016; Gonzalez et al. 2016, 2018; Flasseur et al. 2018a). The statistical modeling of the data plays a central role in these processing algorithms to control the probability of false alarms and prevent spurious detections (Mawet et al. 2014).

The main constraint of this method is due to the very high contrast between the target star and the exoplanets ($\geq 10^5$ in infrared). In addition, the detection performance and the achievable contrast are strongly dependent on both the total parallactic rotation of the field of view and the quality of the observations often impacted by several artifacts. These artifacts can be spatially localized (e.g., in case of defective pixels) or can impact a larger fraction of the field of view in the form of large fluctuations when a decentering of the coronagraph or a sudden degradation of the adaptive optics (AO) correction occurs (e.g., a low-wind effect, Sauvage et al. 2015; Milli et al. 2018). The latter two effects are especially problematic since the exoplanets are most likely to be located close to the center of the field of view which is the area most affected by the fluctuations.

To the best of our knowledge, there is no particular strategy implemented in the state-of-the-art post-processing algorithms to deal with the evolution of the local quality of the acquired frames. Images presenting large fluctuations compared




Fig. 1. Central region of SPHERE-IRDIS images of HIP 72192: *panel a*: measured intensity for five selected frames showing strong temporal fluctuations; *panel b*: map of the estimated temporal scaling parameters $\hat{\sigma}_{k,\ell}$ for the matching frames: $\ell = 1, 3, 8, 13$ and 14 .

to the others are simply flagged as “bad frames” and discarded from the ADI stack, even if some areas of these images contain useful information. In this paper, we present an extension of our recently proposed detection algorithm accounting for the patch covariances of the background (PACO algorithm, see Flasseur et al. 2018a). The proposed improvement is based on the modeling of the temporal variations of the amplitude of the background fluctuations jointly to the spatial covariances. To improve the robustness of the method, we spatially weight each temporal frame in a data-driven fashion based on its relative degree of fluctuation.

Section 2 describes our local modeling of the spatio-temporal fluctuations and Sect. 3 details the extension of the PACO algorithm. In Sect. 4, we use VLT/SPHERE-IRDIS datasets to illustrate the performance of the proposed method in terms of detection capability, achievable contrast, and astrometry accuracy in comparison to two state-of-the-art algorithms routinely used on ADI datasets, TLOCI (Marois et al. 2013, 2014) and KLIP (Soummer et al. 2012), both implemented in the SpeCal software (Galicher et al. 2018) of the SPHERE data center (Delorme et al. 2017). We show that by preventing the suppression of bad frames, the automatic local weighting of the images improves the achievable contrast on the whole field of view.

2. Local modeling of spatio-temporal fluctuations

Even in the absence of sources in the field of view other than the star masked out by the coronagraph, there are fluctuations in images from an ADI sequence. There are several factors that induce these fluctuations, in particular: (i) evolution of the point spread function (PSF) due to changes in the observing conditions and alterations of the quality of the adaptive optics correction; (ii) photon and thermal noise on the camera as well as the erratic response of some uncorrected bad pixels; (iii) partial decentering of the coronagraph; (iv) and evolving noncommon-path aberrations (uncorrected by the AO).

In Fig. 1a, we illustrate the temporal evolution of the intensity measured with the SPHERE-IRDIS instrument. We selected some frames from an observation of HIP 72192 (out of $T = 96$ temporal frames acquired under the 2015-06-11 – 095.C-0389 ESO program with a total apparent rotation of the field of view of 17.3°) that displayed particularly strong fluctuations due to difficult observing conditions. Since these fluctuations are mostly located around the coronagraph, we show a zoom on the central region of the images corresponding to a field of $2.8 \times 2.8 \text{ arcsec}^2$.

The fluctuations are spatially structured and their magnitude varies from one image to the other in the ADI sequence.

In Flasseur et al. (2018a), we modeled the statistical distribution of the patch $\mathbf{r}_{\theta_k, t_\ell}$: the small window made of a few tens of pixels centered at angular location θ_k in the ℓ -th frame. This patch with K pixels was described as the possible superimposition of the signal of point-source (if present) and a background patch $\mathbf{f}_{\theta_k, t_\ell}$, considered as a random realization of a multivariate Gaussian $\mathcal{N}(\mathbf{m}_k, \mathbf{C}_k)$, with spatially variant mean $\mathbf{m}_k \in \mathbb{R}^K$ and covariance $\mathbf{C}_k \in \mathbb{R}^{K \times K}$. At any given location θ_k , the background fluctuations captured through the covariance matrix \mathbf{C}_k were considered stationary (i.e., not evolving with time).

To account for the presence of “bad” frames with stronger background fluctuations than the other frames, we consider in this paper that the background patch $\mathbf{f}_{\theta_k, t_\ell}$ is a random realization of a Gaussian scale mixture (GSM; see e.g. Wainwright & Simoncelli 2000), also known as a compound-Gaussian model (see e.g. Conte et al. 1995):

$$\mathbf{f}_{\theta_k, t_\ell} = \mathbf{m}_k + \sigma_{k,\ell} \mathbf{u}_{k,\ell}, \quad (1)$$

where $\sigma_{k,\ell} > 0$ is a scalar (random or deterministic) variable that acts as a scaling parameter and the vectors $\mathbf{u}_{k,\ell} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_k)$ are independent centered Gaussian random vectors.

Gaussian scale mixtures cover a wide range of distributions; in particular heavy-tailed distributions like gamma distribution, Student’s distribution, or the generalized Laplacian. They can account for the presence of outliers in the data, such as large deviations. They have been introduced in image processing to model the distribution of wavelet coefficients (see Wainwright & Simoncelli 2000; Portilla et al. 2003) and are widely used to model heterogeneous clutter in synthetic aperture radar; see for example Conte et al. (1995), Gini & Greco (2002), Wang et al. (2006), Pascal et al. (2008).

Rather than directly handling the GSM distribution, we estimate the scale parameter $\sigma_{k,\ell}$ for the patch centered on θ_k that is extracted from frame t_ℓ , and consider that $\mathbf{f}_{\theta_k, t_\ell}$ is distributed according to $\mathcal{N}(\mathbf{m}_k, \sigma_{k,\ell}^2 \mathbf{C}_k)$. This corresponds to considering parameter $\sigma_{k,\ell}$ as a deterministic nuisance parameter; see Pascal et al. (2008).

Figure 1b displays maps of the estimated scaling factors $\hat{\sigma}_{k,\ell}^2$ for the five frames shown in Fig. 1a. In these maps, large values of the scaling factors match the areas in which the intensities are much larger than the average. The variations observed within the field of view indicate that temporal fluctuations are spatially nonstationary: abnormally large fluctuations occur only in some areas and cannot be compensated for by a factor common to the whole image. By closer visual inspection, the scaling factors can be seen to vary over very small distances. This indicates why methods based on linear combinations of template on-axis PSFs (TLOCI) or modes (KLIP) fail to capture such small-scale variations, even when a local fitting based on annuli or angular sectors is performed.

An analysis of the empirical distribution of intensities in a patch, in the absence of a source, is carried out in Fig. 2. Three cases are compared: (i) patches extracted at a location θ_k close to the coronagraph, selected to highlight temporal fluctuations (first two rows of the figure); (ii) patches extracted at a location θ_k farther from the coronagraph, showing only moderate temporal fluctuations (following two rows); and (iii) patches extracted all over the field of view (last row). The temporal collection of patches $\{\mathbf{r}_{\theta_k, t_\ell}\}_{\ell=1:T}$ is considered in each case. The empirical distribution of the intensities of the centered patches $\{\mathbf{r}_{\theta_k, t_\ell} - \hat{\mathbf{m}}_k\}_{\ell=1:T}$ is displayed in the left column. Based on a local estimate $\hat{\mathbf{C}}_k$

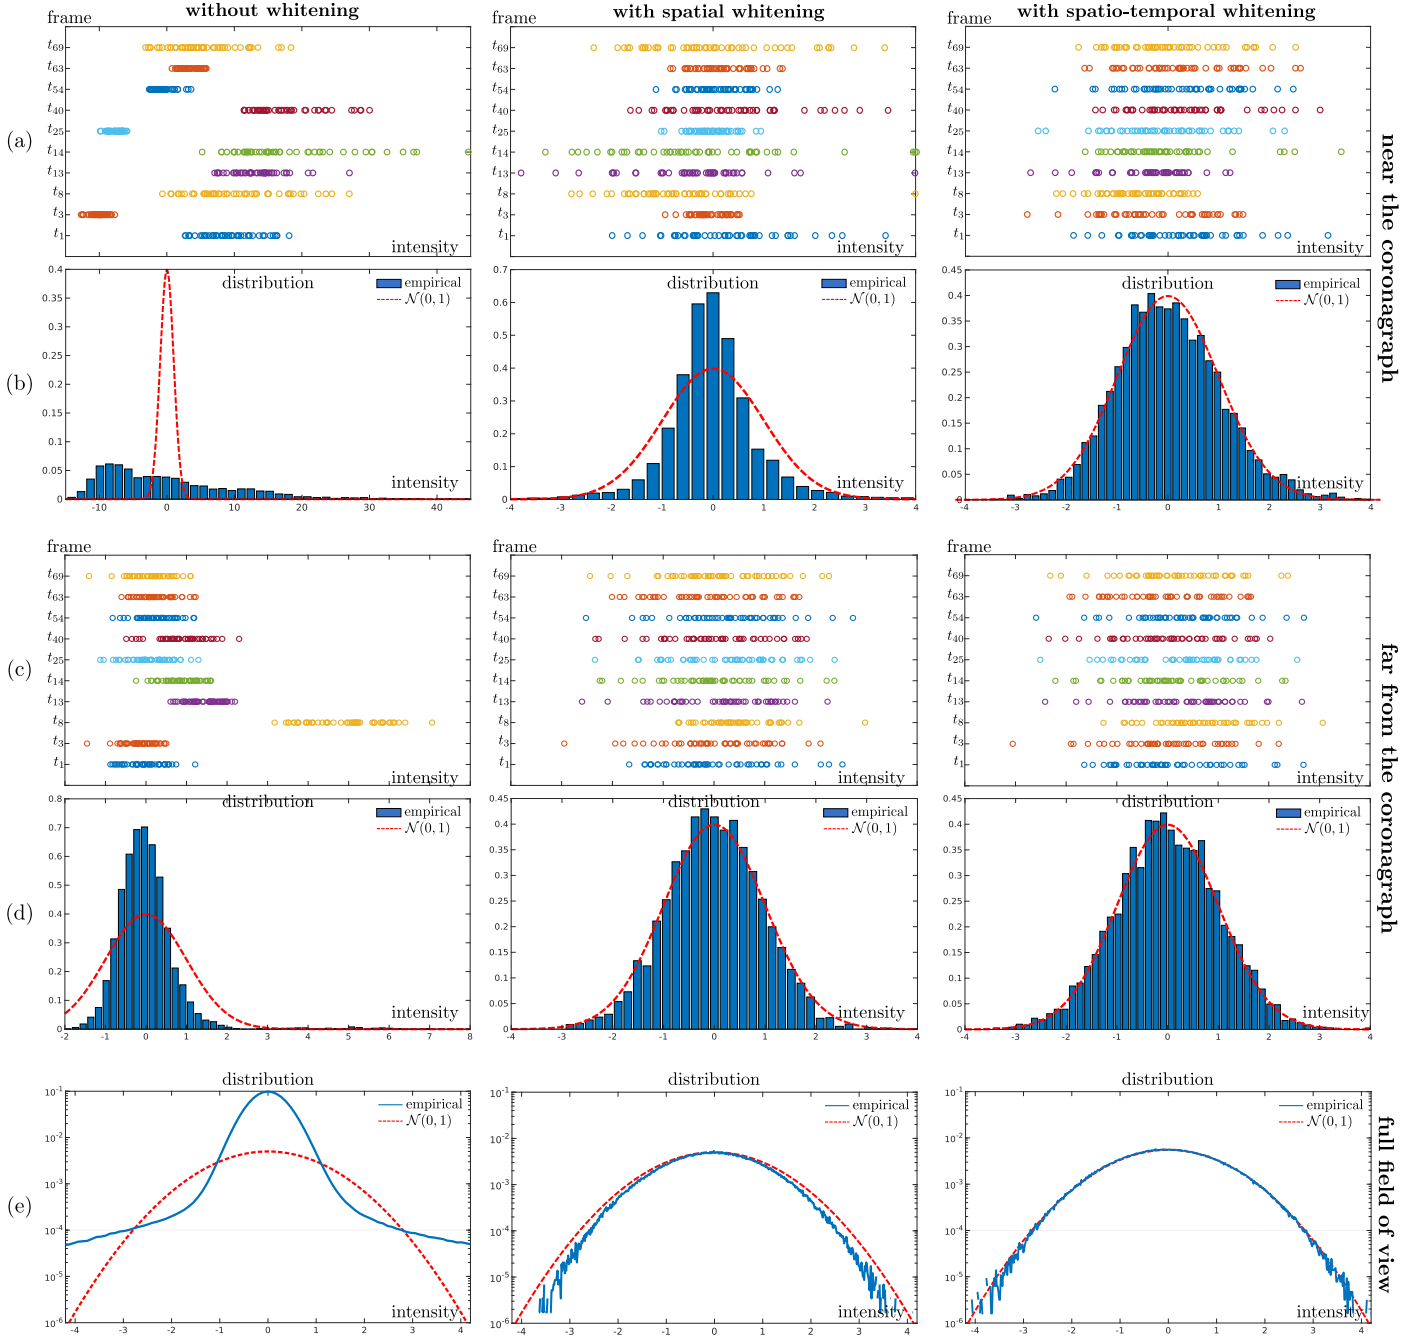


Fig. 2. Temporal evolution and empirical distribution of the intensities within a patch: (a)–(b) in patches extracted at a location θ_k close to the coronagraph; (c)–(d) in patches extracted at a location θ_k farther from the coronagraph; (e) in all the patches from the field of view. Rows *a* and *c*: values within the patches for ten specific frames: $t_1, t_3, t_8, t_{13}, t_{14}, t_{25}, t_{40}, t_{54}, t_{63}$ and t_{69} . Rows *b*, *d*, and *e*: empirical histograms computed over all frames t_1 to t_{69} and, in dashed line, a standard Gaussian. The *first column* corresponds to centered patches $\{\mathbf{r}_{\theta_k, t_\ell} - \bar{\mathbf{m}}_k\}_{\ell=1:T}$. The *second column* corresponds to centered patches that have been both whitened for the spatial correlations and equalized with the temporal scaling factors.

of the spatial covariance matrix, the distribution of spatially whitened and centered patches is then displayed on the central column. The Cholesky factor¹ $\hat{\mathbf{L}}_k$ (i.e., the lower-triangular matrix such that $\hat{\mathbf{L}}_k \hat{\mathbf{L}}_k^T = \hat{\mathbf{C}}_k^{-1}$) is used to obtain the collection of spatially whitened patches $\{\hat{\mathbf{L}}_k^T(\mathbf{r}_{\theta_k, t_\ell} - \bar{\mathbf{m}}_k)\}_{\ell=1:T}$. The empirical distribution of that collection is plotted. The right column

displays the distribution of patch intensities when both a spatial whitening and a time-specific scaling is performed. The empirical distribution of the collection $\{\frac{1}{\hat{\sigma}_{k,d}} \hat{\mathbf{L}}_k^T(\mathbf{r}_{\theta_k, t_\ell} - \bar{\mathbf{m}}_k)\}_{\ell=1:T}$ is plotted in this latter case.

Simply removing the temporal mean does not lead to independent and identically distributed normal residuals, as can be observed in the left column of Fig. 2. The variance of the residuals is much larger near the coronagraph than farther away, a behavior that is generally addressed by introducing a radial scaling of the residuals (Mawet et al. 2014). Locally modeling the

¹ We note that in practice it is more efficient first to compute a Cholesky factorization and then invert the Cholesky factor using a dedicated routine for triangular matrix inversion, if available.

spatial correlations provides a good approximation of the distribution of the patches in most locations θ_k , but fails in regions with large temporal heterogeneity: the distribution in that case is then better described by a GSM; see the central column of Fig. 2b. The presence of frames with larger fluctuations leads to an over-estimation of the variances in $\widehat{\mathbf{C}}_k$, which in turn leads to whitened samples with a variance that is less than one (Fig. 2e). When both a local modeling of the spatial covariance and a time-specific local scaling factor are considered, the empirical distributions are well modeled: after whitening, they closely match a standard Gaussian in all cases (near the coronagraph, farther from the coronagraph, and on average over the whole field of view). Our modeling and the local patch-whitening transform that derives from the GSM model can be seen as an extension of the radial-scaling strategy that is refined in several ways: (i) locality (vs. identical processing of an annulus), (ii) modeling of spatial correlations (vs. uncorrelated noise assumption), and (iii) modeling of time fluctuations (vs. constant correction).

3. Adaptation of the PACO algorithm

3.1. Estimation of the statistics of the background

Estimating, at each spatial location, the mean background \mathbf{m} , the spatial covariance \mathbf{C} , and the scale parameters σ_1 to σ_T is very close to the problem of covariance structure estimation in radar under a compound-Gaussian clutter model (Pascal et al. 2008). There are two differences: (i) in radar the signal and covariance matrices are complex-valued; and (ii) the mean value is zero.

We derive in Appendix A the expression of the maximum likelihood estimates for the scale parameters, the local mean and the local covariance. The expression of each estimator depends on the others and so a fixed point is sought to obtain the joint estimator of the local background statistics. Based on the alternating application of formulas (A.2)–(A.4), an iterative algorithm is obtained. The algorithm is detailed in boxes 1 and 2.

Up to the centering, on accounting for the nonzero mean, the obtained fixed-point algorithm matches the estimator derived in Conte et al. (2002) in the context of clutter with a deterministic texture, and also the approximate maximum likelihood estimator derived in Gini & Greco (2002) under a stochastic model of texture. The convergence of the fixed-point algorithm has been established in Pascal et al. (2008) in the real-valued case as well as in the complex-valued case.

Since our estimates are based on small sample sizes, we include a shrinkage procedure similar to that proposed in Chen et al. (2010) and Flasseur et al. (2018a). To account for the weighting of the samples, we compute an effective number of frames $\widetilde{T} = (\sum_{\ell} 1/\widehat{\sigma}_{\ell}^2)^2 / (\sum_{\ell} 1/\widehat{\sigma}_{\ell}^4)$ that corresponds to the variance reduction reached when performing weighted means with the weights $1/\widehat{\sigma}_{\ell}^2$. This effective number of frames is used to compute the shrinkage factor in Algorithm 3. This shrinkage operation can either be applied at each iteration (during step 4 of Algorithm 1) or after the alternating updates have converged. We tested both approaches and found that they converged to solutions leading to similar detection performances. Applying the shrinkage at each iteration ensures that matrix $\widehat{\mathbf{C}}$ never becomes singular or ill-conditioned.

Figure 3 illustrates the fast convergence of the alternating scheme. Each graphic corresponds to a different location θ_k in the field of view, depicted by a red dot. From left to right, they correspond respectively to a region with good temporal stationarity, medium temporal stationarity, and low temporal stationarity

Algorithm 1: Local background statistics estimation

Input: $\{\mathbf{r}_1, \dots, \mathbf{r}_T\}$ (stack of T spatial patches)
(each patch has K pixels)

Output: $\widehat{\mathbf{m}}$ (mean patch)

Output: $\widehat{\mathbf{C}}$ ($K \times K$ spatial covariance)

Output: $\{\widehat{\sigma}_1, \dots, \widehat{\sigma}_T\}$ (scale parameters)

$\widehat{\mathbf{C}} \leftarrow \mathbf{I}$ (initialize at identity matrix)

for $\ell \leftarrow 1$ **to** T **do**

$\widehat{w}_{\ell} \leftarrow 1/T$ (initialize with uniform weights)

$\widehat{\mathbf{m}} \leftarrow \sum_{\ell=1}^T \widehat{w}_{\ell} \cdot \mathbf{r}_{\ell}$ (sample mean)

do

Step 1: Estimate scale parameters

$\{\widehat{\sigma}_1, \dots, \widehat{\sigma}_T\} \leftarrow \text{Alg2}(\{\mathbf{r}_1, \dots, \mathbf{r}_T\}, \widehat{\mathbf{m}}, \widehat{\mathbf{C}})$

Step 2: Update temporal weights

for $\ell \leftarrow 1$ **to** T **do**

$\widehat{w}_{\ell}^{(\text{old})} \leftarrow \widehat{w}_{\ell}$

$\widehat{w}_{\ell} \leftarrow \frac{1/\widehat{\sigma}_{\ell}^2}{\sum_{\ell'=1}^T 1/\widehat{\sigma}_{\ell'}^2}$

Step 3: Update the mean patch

$\widehat{\mathbf{m}} \leftarrow \sum_{\ell=1}^T \widehat{w}_{\ell} \cdot \mathbf{r}_{\ell}$ (weighted mean)

Step 4: Update the covariance

$\widehat{\mathbf{S}} \leftarrow \frac{1}{T} \sum_{\ell=1}^T \frac{1}{\widehat{\sigma}_{\ell}^2} (\mathbf{r}_{\ell} - \widehat{\mathbf{m}})(\mathbf{r}_{\ell} - \widehat{\mathbf{m}})^t$ (sample cov.)

$\widetilde{T} \leftarrow 1 / \sum_{\ell=1}^T \widehat{w}_{\ell}^2$ (effective number of frames)

$\widehat{\mathbf{C}} \leftarrow \text{Alg3}(\widehat{\mathbf{S}}, \widetilde{T})$ (shrinkage covariance estimator)

while $\max_{\ell} |\widehat{w}_{\ell} - \widehat{w}_{\ell}^{(\text{old})}| \geq \epsilon$;

Algorithm 2: Estimate scale parameters

Input: $\{\mathbf{r}_1, \dots, \mathbf{r}_T\}$ (stack of T spatial patches)

Input: $\widehat{\mathbf{m}}$ (estimated mean patch)

Input: $\widehat{\mathbf{C}}$ ($K \times K$ estimated spatial covariance)

Output: $\{\widehat{\sigma}_1, \dots, \widehat{\sigma}_T\}$ (scale parameters)

$\widehat{\mathbf{L}} \leftarrow \text{chol}(\widehat{\mathbf{C}}^{-1})$ (Cholesky factorization)

for $\ell \leftarrow 1$ **to** T **do**

$\check{\mathbf{r}}_{\ell} \leftarrow \widehat{\mathbf{L}}^t (\mathbf{r}_{\ell} - \widehat{\mathbf{m}})$ (whitened patch)

$\widehat{\mu} \leftarrow \frac{1}{K} \sum_{i=1}^K [\check{\mathbf{r}}_{\ell}]_i$ (mean value)

$\widehat{\sigma}_{\ell}^2 \leftarrow \frac{1}{K} \sum_{i=1}^K ([\check{\mathbf{r}}_{\ell}]_i - \widehat{\mu})^2$ (scale parameter)

(close to the coronagraph). The convergence of Algorithm 1 is assessed based on 1000 random initializations of the weights \widehat{w}_{ℓ} (following a uniform distribution on $[0,1]$). The vector of weights at convergence when starting from a constant vector of weights ($\forall \ell, \widehat{w}_{\ell} = 1/T$) is used as a reference. At each iteration, the Euclidean distance to the reference vector of weights is reported, with normalization by the largest distance of all the random draws. Very fast convergence is observed to a unique solution, in all random trials. This empirical evidence indicates that the algorithm also converges when the shrinkage step is included at each iteration.

3.2. Robust computation of a detection map

The PACO algorithm, introduced in Flasseur et al. (2018a), produces a detection map based on a hypothesis test evaluated at each pixel of the field of view. The ADI sequence of images is processed at the scale of image patches. The size of the patches




Fig. 3. Illustration of the convergence of Algorithm 1: 1000 random initializations of the weights $\{\widehat{w}_\ell\}_{\ell=1:T}$ converge to the same solution as that obtained by initializing with constant weights. Convergence at three locations θ_k depicted on the field of view is shown, covering three different cases: small dispersion of the value of scaling parameters (*left*), medium dispersion (*center*), and large dispersion close to the coronagraph (*right*).

Algorithm 3: Shrinkage covariance estimator

Input: $\widehat{\mathbf{S}}$ ($K \times K$ sample covariance matrix)

Input: \widehat{T} (effective number of frames)

Output: $\widehat{\mathbf{C}}$ (covariance matrix estimate)

$$\widehat{\rho} \leftarrow \frac{\text{tr}(\widehat{\mathbf{S}}^2) + \text{tr}^2(\widehat{\mathbf{S}}) - 2 \sum_{i=1}^K [\widehat{\mathbf{S}}]_{ii}^2}{(\widehat{T}+1) \left(\text{tr}(\widehat{\mathbf{S}}^2) - \sum_{i=1}^K [\widehat{\mathbf{S}}]_{ii}^2 \right)} \quad (\text{shrinkage factor})$$

for $i \leftarrow 1$ **to** K **do**

$[\widehat{\mathbf{C}}]_{i,i} \leftarrow [\widehat{\mathbf{S}}]_{i,i}$ (copy diagonal)

for $j \leftarrow i+1$ **to** K **do**

$[\widehat{\mathbf{C}}]_{i,j} \leftarrow (1 - \widehat{\rho})[\widehat{\mathbf{S}}]_{i,j}$ (shrink off-diagonal)

$[\widehat{\mathbf{C}}]_{j,i} \leftarrow (1 - \widehat{\rho})[\widehat{\mathbf{S}}]_{j,i}$ (shrink off-diagonal)

is selected to capture the core of the off-axis PSF, for example a disk of $K = 49$ pixels for SPHERE-IRDIS in K1-K2 observing mode. If a point source object is located at angular position ϕ_0 in some reference frame, then its angular location ϕ_ℓ in the image at time t_ℓ can be deduced from the telescope pointing information by accounting for the rotation of the whole field of view during ADI acquisitions. Source detection is based on the analysis of the collection $\{\mathbf{r}_{[\phi_\ell], t_\ell}\}_{\ell=1:T}$ of the T patches formed by tracking the source through its apparent motion: each patch is extracted at the expected location, rounded to the closest pixel $[\phi_\ell]$, of the point source with reference coordinate ϕ_0 .

During the detection phase, we assume that the observations are dominated by the background signal due to stellar leakages and by the noise so that background statistics can be computed directly from the data (off-axis point sources are considered negligible at that step). The local statistics of the background are then computed with Algorithm 1, for each location $[\phi_\ell]$, based on the collections $\{\mathbf{r}_{[\phi_\ell], t}\}_{t=1:T}$ of observed patches all centered at that same location $[\phi_\ell]$. Once the mean $\widehat{\mathbf{m}}_{[\phi_\ell]}$, spatial covariance $\widehat{\mathbf{C}}_{[\phi_\ell]}$ and scaling factor² $\widehat{\sigma}_{[\phi_\ell], \ell}$ have been estimated for all locations $[\phi_\ell]$ corresponding to the trajectory of a hypothetical point source with reference location ϕ_0 , the likelihood of two hypotheses can be compared:

$$\begin{cases} \mathcal{H}_0 : \{\mathbf{r}_{[\phi_\ell], \ell}\}_{\ell=1:T} &= \{\mathbf{f}_{[\phi_\ell], \ell}\}_{\ell=1:T} \\ (\text{background only}) \\ \mathcal{H}_1 : \{\mathbf{r}_{[\phi_\ell], \ell}\}_{\ell=1:T} &= \alpha \{\mathbf{h}_{[\phi_\ell], \ell}(\phi_\ell)\}_{\ell=1:T} + \{\mathbf{f}_{[\phi_\ell], \ell}\}_{\ell=1:T}, \\ (\text{background + source}), \end{cases} \quad (2)$$

² Only the scaling factor at time t_ℓ is necessary for the detection.

where α is the flux of the point source and $\mathbf{h}_{[\phi_\ell], \ell}(\phi_\ell)$ is a patch of the off-axis PSF extracted around the integer location $[\phi_\ell]$, for a point source located at the subpixel location ϕ_ℓ .

Under the assumption that each of the T background patches of the collection $\{\mathbf{f}_{[\phi_\ell], \ell}\}_{\ell=1:T}$ is an independent realization, each distributed according to its local GSM model, the co-log-likelihood of the data under hypothesis \mathcal{H}_0 is:

$$\begin{aligned} -\log p(\{\mathbf{r}_{[\phi_\ell], \ell}\}_{\ell=1:T} | \mathcal{H}_0) &= \frac{TK}{2} \log 2\pi + \sum_{\ell=1}^T \frac{1}{2} \log \det(\widehat{\sigma}_{[\phi_\ell], \ell}^2 \widehat{\mathbf{C}}_{[\phi_\ell]}) \\ &\quad + \sum_{\ell=1}^T \frac{1}{2} (\mathbf{r}_{[\phi_\ell], \ell} - \widehat{\mathbf{m}}_{[\phi_\ell]})^t \\ &\quad \times (\widehat{\sigma}_{[\phi_\ell], \ell}^2 \widehat{\mathbf{C}}_{[\phi_\ell]})^{-1} (\mathbf{r}_{[\phi_\ell], \ell} - \widehat{\mathbf{m}}_{[\phi_\ell]}). \end{aligned} \quad (3)$$

The co-log-likelihood of the data under hypothesis \mathcal{H}_1 cannot be directly evaluated because the flux α of the point source is not known beforehand. It can however be estimated in the maximum likelihood sense; see Eq. (14) in Flasseur et al. (2018a):

$$\widehat{\alpha} = \arg \min_{\alpha} -\log p(\{\mathbf{r}_{[\phi_\ell], \ell}\}_{\ell=1:T} | \mathcal{H}_1, \alpha) = \frac{\sum_{\ell=1}^T b_\ell}{\sum_{\ell=1}^T a_\ell}, \quad (4)$$

with

$$a_\ell = \frac{1}{\widehat{\sigma}_{[\phi_\ell], \ell}^2} \mathbf{h}_{[\phi_\ell], \ell}(\phi_\ell)^t \cdot \widehat{\mathbf{C}}_{[\phi_\ell]}^{-1} \cdot \mathbf{h}_{[\phi_\ell], \ell}(\phi_\ell), \quad (5)$$

and

$$b_\ell = \frac{1}{\widehat{\sigma}_{[\phi_\ell], \ell}^2} \mathbf{h}_{[\phi_\ell], \ell}(\phi_\ell)^t \cdot \widehat{\mathbf{C}}_{[\phi_\ell]}^{-1} \cdot (\mathbf{r}_{[\phi_\ell], t_\ell} - \widehat{\mathbf{m}}_{[\phi_\ell]}). \quad (6)$$

The co-log-likelihood of the data under hypothesis \mathcal{H}_1 , for a source flux $\widehat{\alpha}$, is then:

$$\begin{aligned} -\log p(\{\mathbf{r}_{[\phi_\ell], \ell}\}_{\ell=1:T} | \mathcal{H}_1, \widehat{\alpha}) &= \frac{TK}{2} \log 2\pi + \sum_{\ell=1}^T \frac{1}{2} \log \det(\widehat{\sigma}_{[\phi_\ell], \ell}^2 \widehat{\mathbf{C}}_{[\phi_\ell]}) \\ &\quad + \sum_{\ell=1}^T \frac{1}{2} \mathbf{u}_{[\phi_\ell], \ell}^t (\widehat{\sigma}_{[\phi_\ell], \ell}^2 \widehat{\mathbf{C}}_{[\phi_\ell]})^{-1} \mathbf{u}_{[\phi_\ell], \ell}, \end{aligned} \quad (7)$$

with $\mathbf{u}_{[\phi_\ell], \ell} = \mathbf{r}_{[\phi_\ell], \ell} - \widehat{\alpha} \mathbf{h}_{[\phi_\ell], \ell}(\phi_\ell) - \widehat{\mathbf{m}}_{[\phi_\ell]}$.

After some simplifications, the generalized likelihood ratio test (GLRT) is obtained; see Kay (1998) and Flasseur et al. (2018a):

$$2 \log \frac{p(\{\mathbf{r}_{[\phi_\ell], \ell}\}_{\ell=1:T} | \mathcal{H}_1, \widehat{\alpha})}{p(\{\mathbf{r}_{[\phi_\ell], \ell}\}_{\ell=1:T} | \mathcal{H}_0)} = \frac{(\sum_{\ell=1}^T b_\ell)^2}{\sum_{\ell=1}^T a_\ell} \frac{\mathcal{H}_1}{\mathcal{H}_0} \gtrless \eta. \quad (8)$$

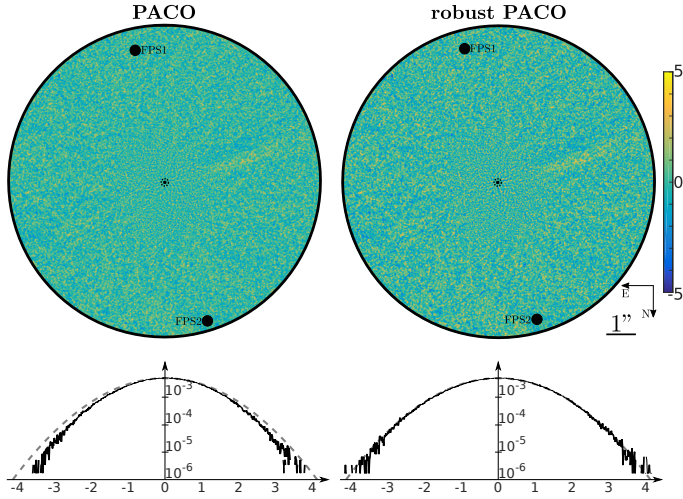


Fig. 4. Detection map computed on a SPHERE-IRDIS ADI dataset of HIP 72192. The two faint point sources in the field of view, FPS1 and FPS2, are marked by black dots. There are no other detectable sources in the field of view. *Left column:* standard PACO algorithm, *right column:* proposed extension to improve robustness. *Bottom row:* solid line showing the empirical distribution of the S/N values of the detection maps, excluding areas FPS1 and FPS2. The dashed line corresponds to a standard normal distribution.

Only a positive flux α makes sense for a source. The estimate $\hat{\alpha}$ should then be replaced by the estimate $\hat{\alpha}^+ = \max(\hat{\alpha}, 0)$ obtained under a positivity constraint. The GLRT can then be recast, for $\eta \geq 0$, into the form:

$$S/N \text{ test : } \frac{\sum_{\ell=1}^T b_{\ell}}{\sqrt{\sum_{\ell=1}^T a_{\ell}}} = \frac{\hat{\alpha}}{\hat{\sigma}_{\alpha}} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \tau, \quad (9)$$

with $\tau = \sqrt{\eta}$, which corresponds to comparing the signal-to-noise ratio (S/N) of the flux estimate to a threshold; see Mugnier et al. (2009) and Flasseur et al. (2018a). The S/N test defined in (9) has the attractive property of being an affine transformation of the observations. Under our GSM model, and thanks to the normalization by the scaling factors $\hat{\sigma}_{[\phi_{\ell}], \ell}$, the ratio $\hat{\alpha}/\hat{\sigma}_{\alpha}$ is distributed according to a standard normal distribution. This simplifies the setting of a detection threshold τ : $\tau = 5$ leads to a probability of false alarms equal to 2.87×10^{-7} , that is, the probability of a Gaussian random variable being greater than five standard deviations.

Figure 4 illustrates that the ratio $\hat{\alpha}/\hat{\sigma}_{\alpha}$ is indeed distributed like a standard normal distribution. Two detection maps are shown for the star HIP 72192. The left part of the figure corresponds to the results of PACO algorithm, as presented in Flasseur et al. (2018a). The corresponding detection map is stationary. When representing the empirical distribution of S/N values in the field of view (after excluding the two areas corresponding to two sources), a good match with a Gaussian distribution is obtained, albeit with a standard deviation slightly below a value of 1. The right-hand part of the figure gives the detection map produced by the extension of PACO introduced in this paper. While both maps are quite similar, it can be noted that time-dependent scaling factors lead to a spatially more stationary map and an improved match to the standard Gaussian distribution. The impact of producing maps with a variance of less than one in the absence of a source with the standard PACO algorithm is to be overly conservative: some detections may be missed at a given false alarm rate while they would be correctly detected




Fig. 5. Contrast curves at 5σ for TLOCI, KLIP, PACO, and the proposed robust extension of PACO, for HIP 72192 (SPHERE-IRDIS, $\lambda = 2.110 \mu\text{m}$). The five triangles and five squares correspond to Monte-Carlo injections (see text).

with the robust version of PACO. We illustrate such behavior in Sect. 4.

The detection map is produced by computing the S/N for each reference location ϕ_0 on a grid. The computational complexity can be significantly reduced by pre-computing terms that can then be shared to evaluate the S/N at several locations. Such an approach, named “fast PACO” in Flasseur et al. (2018a), can also be adapted to our GSM model, provided that the background statistics are computed using the algorithm outlined in Sect. 3.1 and that the formulas are updated to include the scaling factors $\hat{\sigma}_{k, \ell}$ according to Eqs. (5), (6), and (9). Compared to the original fast PACO algorithm, these changes lead to an increase of the computation time by an order of magnitude, corresponding to the typical number of iterations to reach a fixed point in the step that estimates the local background statistics (about ten iterations).

3.3. Robust estimation of photometry and astrometry

In Appendix B, we prove that the GSM model considered in this paper leads to robust estimates of the background statistics and of the flux of a point source in the sense that the estimates remain bounded when a frame takes arbitrarily large values. This contrasts with other data-processing methods in ADI that generally lead to incorrect flux estimates if bad frames with large values are present in the dataset³.

Similarly to the approach in Flasseur et al. (2018a), when characterizing a source found above the detection threshold in the detection map, the source flux is to be re-estimated jointly with the background statistics in order to prevent any self-subtraction that would bias the estimation. In practice, this joint estimation is performed by alternating (i) a re-estimation of the local background statistics using the residuals $\{r_{[\phi_{\ell}], \ell} - \hat{\alpha}h_{[\phi_{\ell}]}(\phi_{\ell})\}_{\ell=1:T}$, where $\hat{\alpha}$ is the current flux estimate, and (ii) a re-estimation of the source flux with the updated background statistics by applying Eq. (4).

³ TLOCI computes medians rather than averages at some steps to improve the robustness. However, such a strategy is generally not sufficient to overcome the presence of large fluctuations on several time frames since some steps include a linear processing of the data.

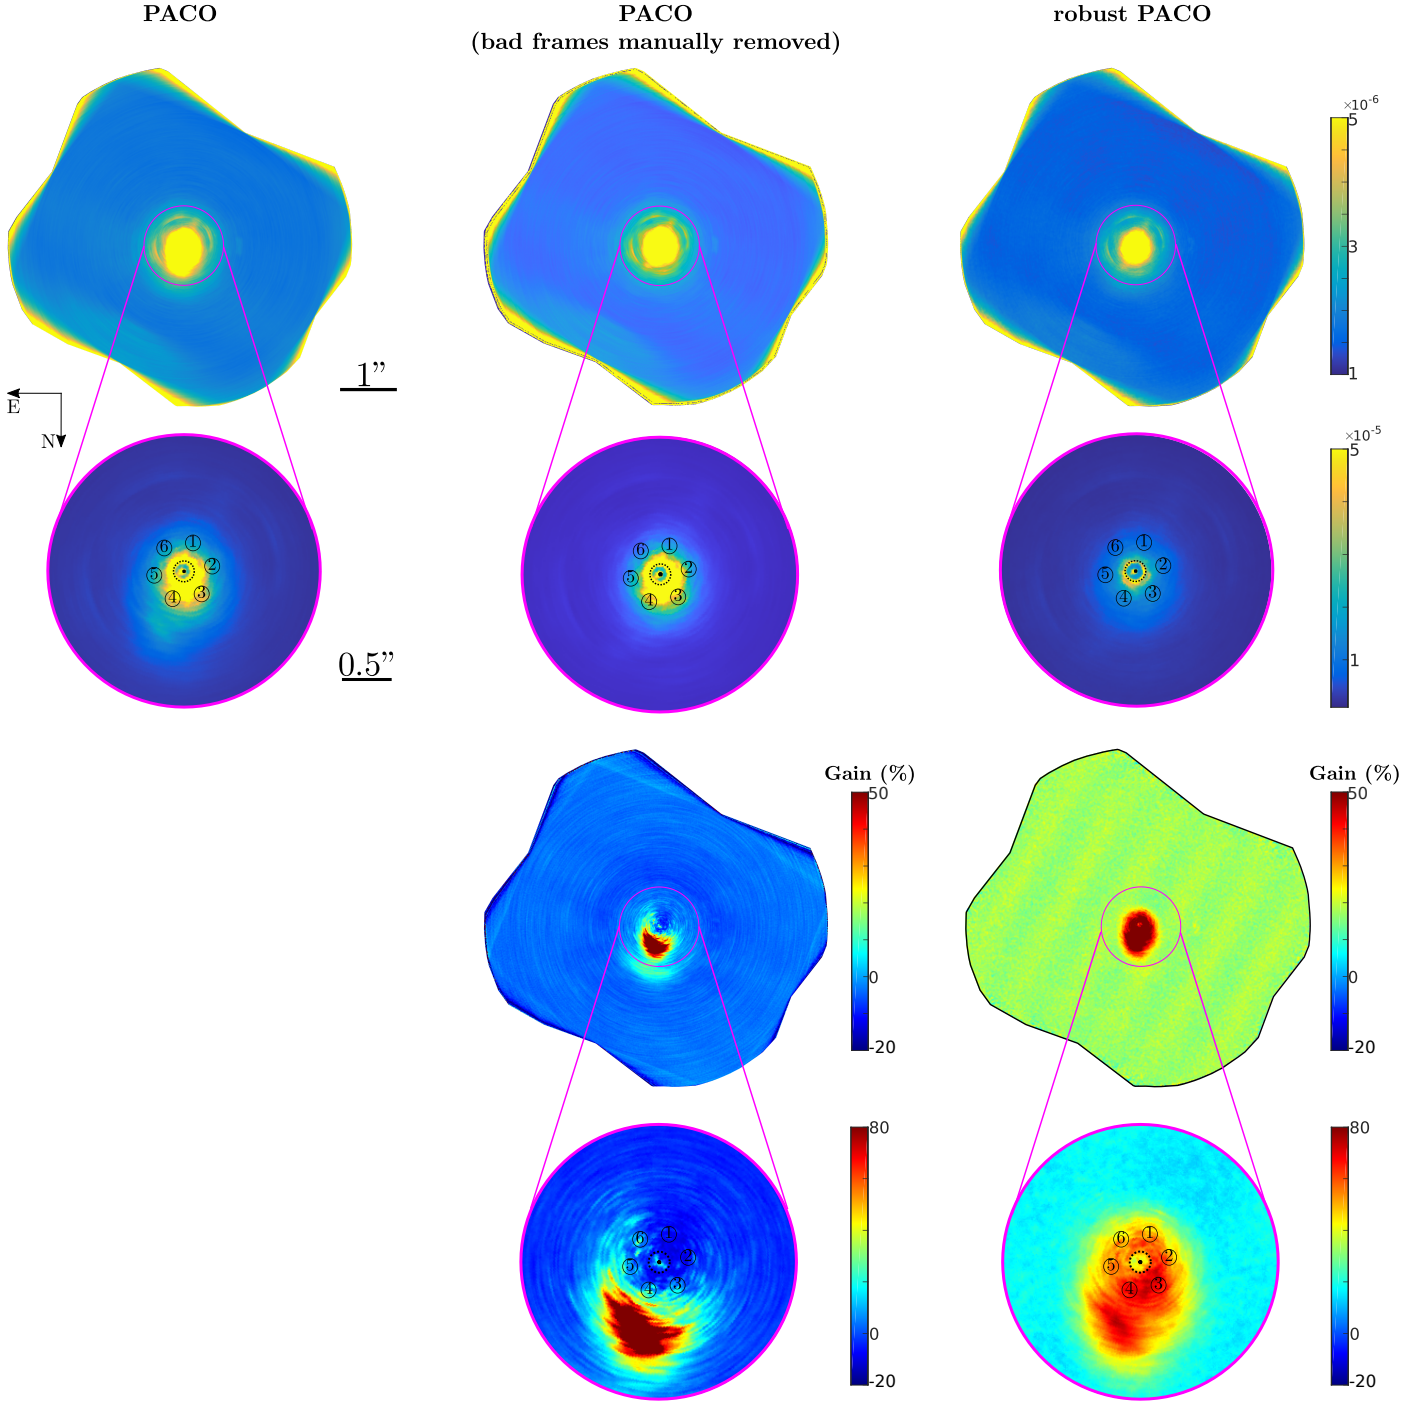


Fig. 6. Comparison of contrast maps for HIP 72192 (SPHERE-IRDIS). *Top row:* largest contrast between the host star and a point source that would theoretically lead to a probability of detection greater than or equal to 50% at the detection threshold $\tau = 5$. Three methods are compared: the standard PACO algorithm applied on the whole ADI stack (*left*), PACO applied to the subset of the ADI stack obtained by removing the eight frames with the largest fluctuations (*center*), and the proposed robust extension of PACO applied on the whole stack (*right*). The bottom row gives the gain in contrast with respect to PACO algorithm.

The accurate astrometric estimation of the source is obtained by evaluating the S/N, with the unbiased flux estimate, over a refined subpixel grid around the detected location.

Once a source has been characterized both in terms of astrometry and photometry, its contribution to the data can be subtracted and the detection map updated accordingly, similarly to source extraction with the CLEAN algorithm; see a more detailed description in [Flasseur et al. \(2018b\)](#).

4. Characterization of detection, and astrometric and photometric performances

The maximum achievable contrast between the host star and an off-axis point source is important information to characterize the overall performance of the instrument including the data processing. Reliable evaluations of the detection limit (expressed in contrast) typically require the injection of a fake source at

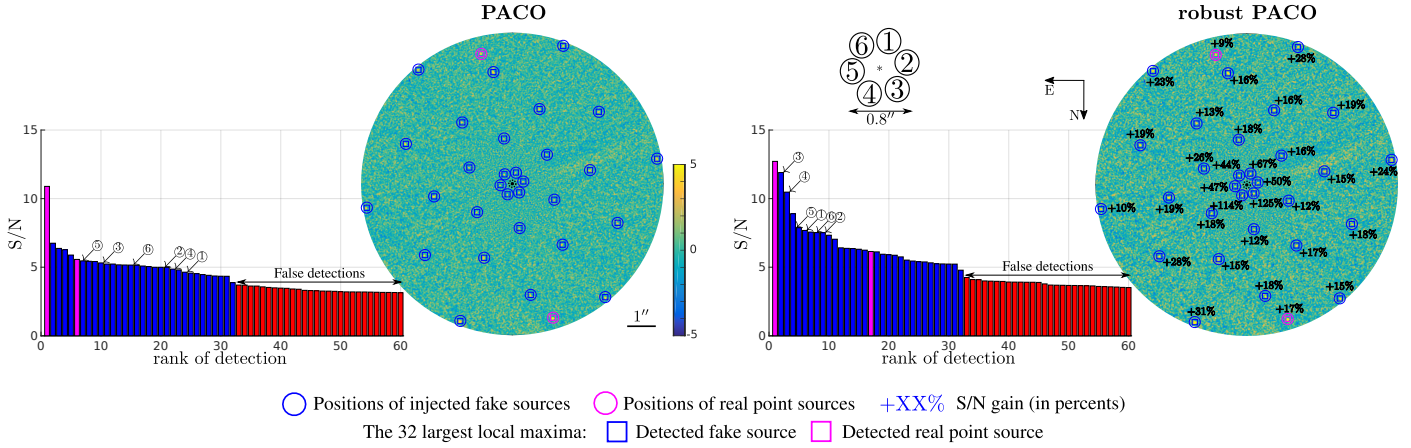


Fig. 7. Comparison of PACO (*left*) and the robust extension of PACO introduced in this paper (*right*) on the SPHERE-IRDIS ADI dataset with 30 fake point sources injected in addition to the two known faint point sources (same level of contrast as presented in Flasseur et al. 2018a). The detection maps are given together with the S/N values of the 60 first detections. The six fake sources closest to the host star are denoted by digits ① to ⑥. The two real sources are indicated in pink. The increase in S/N brought by the robust approach is given as a percentage on the detection map.

different locations of the field of view and for different values of contrast. This Monte-Carlo simulation approach is computationally very costly, and therefore other methods to evaluate contrast are followed. These other methods are based on a statistical modeling and lead to predictions that are valid only to the extent that the statistical model itself is valid. It was shown in Flasseur et al. (2018b) that contrast curves produced by reference algorithms like TLOCI and KLIP were over-optimistic due to a coarse statistical model of the residues. Contrast curves obtained with PACO are closer to the performance achieved by Monte-Carlo, however we pointed out in Flasseur et al. (2018a) that they only provided a lower bound, which was not reached in practice when the statistics of the background were learnt, in the detection step, on patches that contained both the background and the source(s).

Figure 5 presents contrast curves computed on a SPHERE-IRDIS ADI dataset (at $\lambda = 2.110\mu\text{m}$) of HIP 72192. Four contrast curves are reported, corresponding to TLOCI, KLIP, PACO, and our robust extension of PACO. Contrast curves of TLOCI and KLIP should be analyzed with caution because they were computed based on a threshold at 5σ that does not correspond to a probability of false alarm of 2.87×10^{-7} since the residuals are not Gaussian distributed. Due to the nonstationarity of TLOCI and KLIP S/N maps, a spatially variant threshold should be used to reach a constant false-alarm rate and compute reliable contrast values. For PACO and robust PACO, we give both the theoretical contrast curve and the results of Monte-Carlo simulations computed for five different angular separations (triangles for PACO, squares for robust PACO). The theoretical contrast curves, corresponding to fluxes α such that $P(\hat{\alpha}/\hat{\sigma}_\alpha \geq 5) = 50\%$, are obtained by computing $5\hat{\sigma}_\alpha$, that is, $5/\sqrt{\sum_\ell a_\ell}$, at each point of the field of view, with intensities expressed relative to the intensity of the host star. PACO and robust PACO achieve better contrast than TLOCI and KLIP. The discrepancy between the theoretical contrast and the Monte-Carlo simulations is reduced with robust PACO compared to PACO. Moreover, the contrast is clearly improved at small separations.

Figure 6 gives the 2D contrast maps for three detection methods, on the same ADI dataset as that studied in Fig. 5. From left to right are displayed contrast maps of the standard PACO algorithm, PACO algorithm after manually removing the eight frames

that display the largest fluctuations of the series, and the proposed robust extension of PACO. The gains in contrast obtained with respect to standard PACO are given on the bottom row of the figure. Manually removing the bad frames has two effects: (i) an improvement of the contrast in the area impacted by the large fluctuations of the bad frames, and (ii) a slight degradation of the contrast in the rest of the field of view due to the reduction of the size of the ADI dataset. By locally adapting the scaling factors, robust PACO improves the contrast everywhere. Even in regions that are not affected by strong temporal heterogeneity, it is beneficial to account for fluctuations of the scaling factors. As expected, the gain is the largest in the area affected by the strongest fluctuations. There, the gain in contrast reaches 80%.

Standard PACO and its robust extension are also compared on the same SPHERE-IRDIS dataset as introduced in Flasseur et al. (2018a). In this dataset, 30 fake point sources were injected in addition to the two known faint point sources, at contrasts such that they could all be detected without a false alarm (see Table 3 of Flasseur et al. (2018a) for the corresponding levels of contrast). Figure 7 gives the detection maps produced by PACO and robust PACO. We refer the reader to Fig. 10 in Flasseur et al. (2018a) for a comparison with other detection methods. By including time-specific scaling factors, robust PACO improves the S/N of all sources. The largest increase occurs close to the host star where the S/N is improved by more than 100% for two point sources.

Figure 8 compares TLOCI, KLIP, standard PACO, and robust PACO algorithms on an ADI stack of HD 95086 ($T = 52$ temporal frames acquired under the 2015-05-05 – 095.C-0298(A) ESO program with a total apparent rotation of the field of view of 18.2°). The HD 95086 hosts an exoplanet of $5M_J$ in mass (HD 95086 b) discovered (Rameau et al. 2013a) and confirmed (Rameau et al. 2013b) by direct imaging with the SPHERE instrument. In addition, six known background point sources are in the field of view. Figure 8 shows that both the PACO method and its robust extension achieve better detection performance than reference algorithms; in particular, the S/N values are larger for all sources (except for the background star 2 which poses no detection problem in itself: it is so bright that it is visible directly in a single image of the ADI dataset). With PACO and robust PACO, it is possible to detect without any false alarm the seven known sources by thresholding the S/N

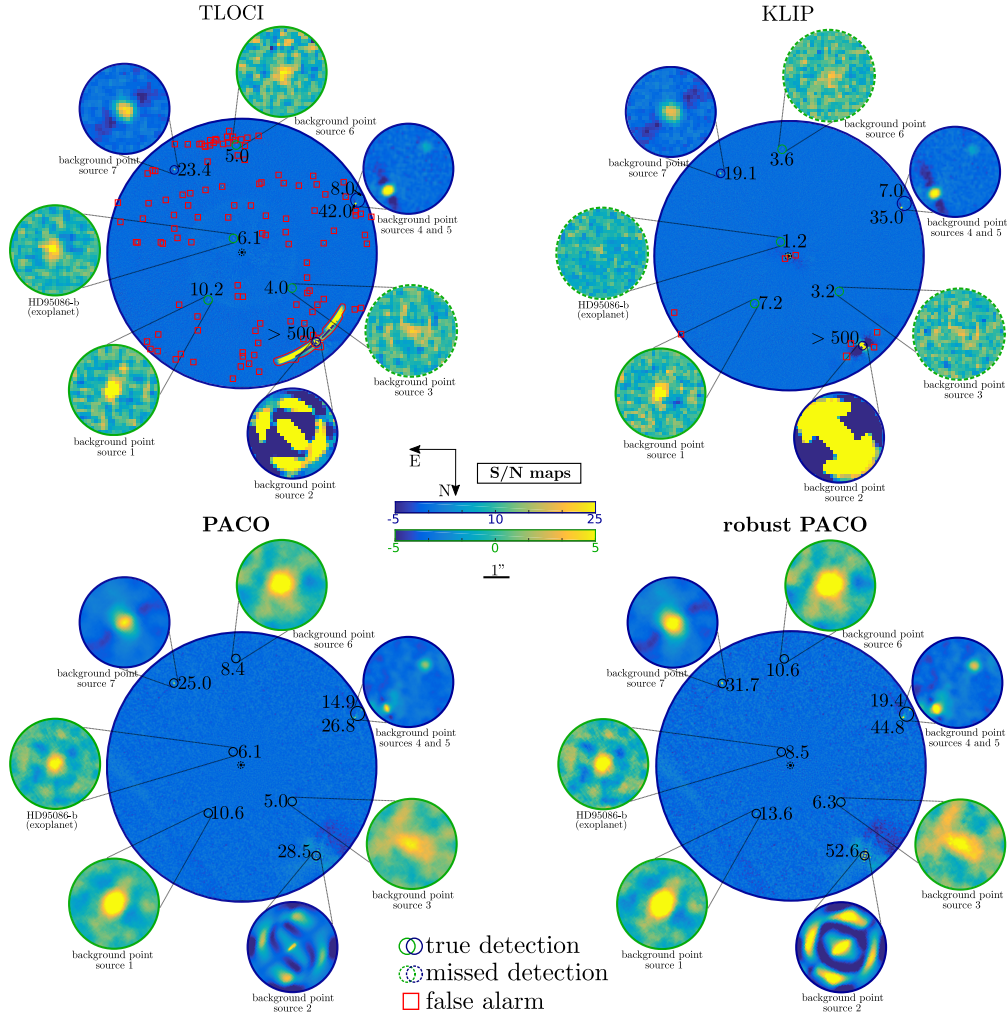


Fig. 8. Maps of S/N from TLOCI, KLIP, PACO, and the robust extension of PACO introduced in this paper on an ADI dataset of HD 95086. Seven circular insets show a zoom of the S/N maps around the known point sources. The S/N maps of these insets are evaluated on a subpixel grid (four nodes per pixel) for PACO and robust PACO. True detections are marked by straight circles, the missed detections ($S/N < 5$) are marked by dashed circles, and false alarms are identified by red squares.

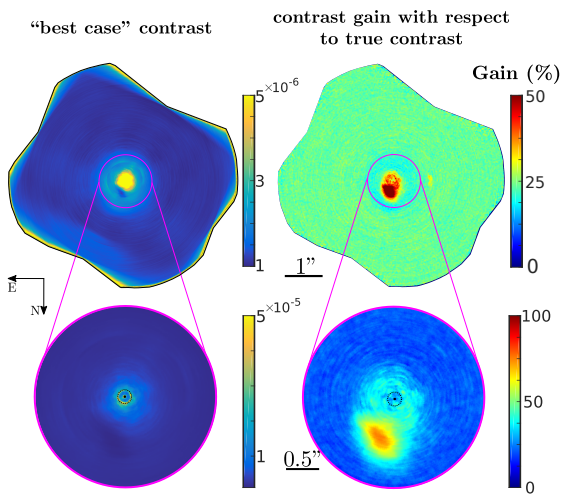


Fig. 9. *Left:* contrast that would be obtained for HIP 72192 under ideal observation conditions (i.e., as good as that of the best image, during the entire ADI acquisition). *Right:* gain in contrast compared to the actual contrast, computed by accounting for temporal fluctuations of the scale of the residuals.

maps at $\tau = 5$. In comparison, S/N maps from TLOCI and KLIP present several artifacts due to the presence of outliers in the datasets, preventing an automatic analysis of these maps. The robust extension of PACO increases the detection confidence of all sources present in the field of view, especially for the ones with the smallest angular separations, such as the exoplanet HD 95086 b ($S/N = 6.1$ with PACO vs. 8.5 with robust PACO).

The presence of frames of poorer quality in an ADI stack degrades the detection performance with respect to an ADI stack of constant quality. It may be useful to the astronomers who planned the observations to compare the contrast achieved given the observation conditions experienced to the contrast that would have been reached should the conditions have been as good as that of the best frames. Such a difference in contrast can be assessed by computing a “best case” contrast where $\hat{\sigma}_{[\phi_\ell], \ell}^2$ in the computation of a_ℓ in Eq. (5) is replaced by $\min_{\ell'} \hat{\sigma}_{[\phi_\ell], \ell'}^2$ (i.e., the smallest scaling factor at that location $[\phi_\ell]$). Figure 9 shows the “best case” contrast for HIP 72192. Compared to the actual contrast predicted based on the estimated time-specific scaling factors, the “best case” contrast is about 25% better in most of the field of view, and close to 70–80% in the region most impacted by large temporal fluctuations. If the aim of the observations was to characterize a faint point source that falls in that

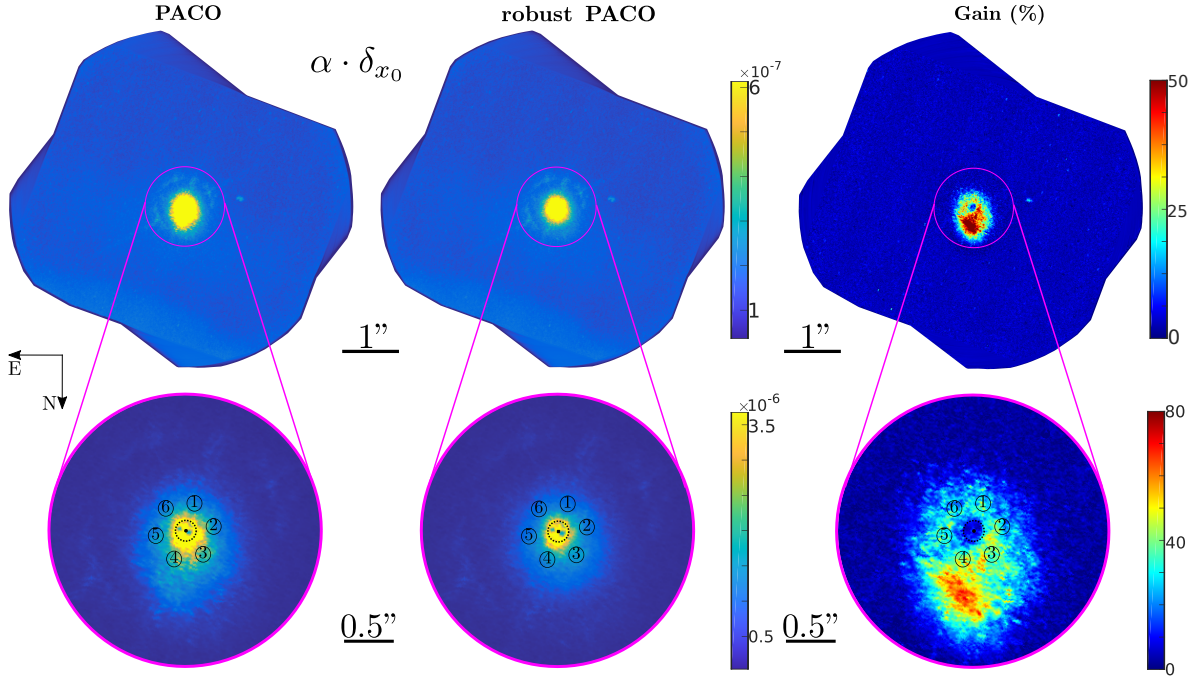


Fig. 10. Theoretical accuracy (minimal standard deviation given by the Cramér-Rao lower bounds) δ_{x_0} on the estimated astrometric location x_0 from PACO and its robust extension presented in this paper around the HIP 72192 star. Here, δ_{x_0} maps are multiplied by the flux α of the exoplanet (expressed in arcsec \times flux). The gain (in percent) brought by the robust approach is given on the right. The six fake sources closest to the host star are denoted by digits ① to ⑥.

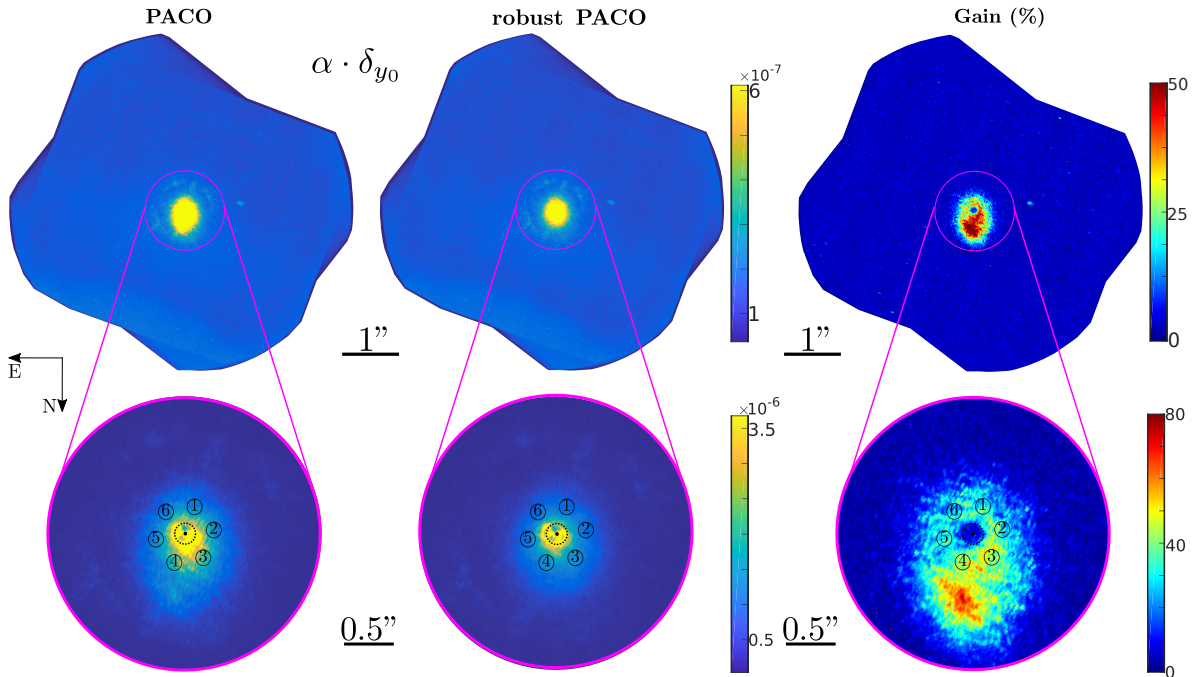


Fig. 11. Theoretical accuracy (minimal standard deviation given by the Cramér-Rao lower bounds) δ_{y_0} on the estimated astrometric location y_0 from PACO and its robust extension presented in this paper around the HIP 72192 star. Here, δ_{y_0} maps are multiplied by the flux α of the exoplanet (expressed in arcsec \times flux). The gain (in percent) brought by the robust approach is given on the right. The six fake sources closest to the host star are denoted by digits ① to ⑥.

area, it might be worth considering re-observing in the hope of experiencing better conditions and getting closer to the best-case contrast.

The astrometric and photometric accuracies can be characterized over the whole field of view by computing the Cramér-Rao lower bounds of the vector \mathbf{p} of parameters (the 2D angular

location ϕ_0 and the flux α) that characterize a point source. The Cramér-Rao lower bound is a good estimate of the covariance of the maximum likelihood estimator when the number of samples is large enough (Kendall et al. 1948). We follow the approach of Flasseur et al. (2018a): the Fisher information matrix \mathbf{I}^F on the vector \mathbf{p} is given by

$$[\mathbf{F}(\mathbf{p})]_{i,j} = \sum_{\ell=1}^T \frac{1}{\hat{\sigma}_{[\phi_0],\ell}^2} \frac{\partial [\alpha \mathbf{h}_{[\phi_0]}(\phi_0)]^\ell}{\partial p_i} \hat{\mathbf{C}}_{[\phi_0]}^{-1} \frac{\partial [\alpha \mathbf{h}_{[\phi_0]}(\phi_0)]}{\partial p_i}, \quad (10)$$

where i and j vary in value between 1 and 3 (p_1 and p_2 represent the two components of the 2D location ϕ_0 and p_3 corresponds to the flux α). The product $\alpha \mathbf{h}_{\theta_k}$ models the signal of a point source of flux α . As in Flasseur et al. (2018a), we use a continuous model of the off-axis PSF (an isotropic Gaussian can be used) to simplify the computation of the spatial derivatives. The standard deviations δ_i for each of the three parameters are obtained from the diagonal of the inverse of Fisher information matrix:

$$[\delta(\mathbf{p})]_i = \sqrt{[\mathbf{F}(\mathbf{p})^{-1}]_{i,i}}. \quad (11)$$

Figures 10 and 11 give the astrometric Cramér-Rao lower bounds on the whole field of view obtained with PACO and robust PACO. They show that the patch weighting included in the robust approach improves the estimation accuracies on the whole field of view, especially in the presence of high stellar leakages (more than 50% improvement near the coronagraph).

5. Conclusion

Most data-processing techniques for point-source detection and characterization in ADI datasets are not robust: the presence of images with poor AO correction can considerably degrade their performance. It is therefore necessary to detect and eliminate these frames in a preprocessing step. In contrast, the approach described in this paper models the temporal variations of the amplitude of the background fluctuations jointly to the spatial correlations using a Gaussian scale mixture model. Since these fluctuations are accounted for, the estimations are robust to large fluctuations. Rather than discarding full images, we use a local modeling in order to spatially adapt the processing and give a very small statistical weight only to the areas subject to large fluctuations. That way, the proposed robust PACO processing makes the most of the available data.

The robust variant of the PACO algorithm introduced in this paper leads to improved detection performances, in particular at close separations where the stellar residuals dominate. Interestingly, it is also possible to estimate the achievable contrast not only by taking into account the actual image quality of the ADI dataset, but also the contrast that would have been reached if the AO correction had been as good as the best correction in the ADI time series. This information can be highly valuable in order to plan subsequent re-observations.

Coronagraphic observations obtained with integral field spectrographs could also benefit from a modeling of the spatial and temporal fluctuations. We will investigate in a future work how to include a modeling of the spectral correlations for these instruments.

Acknowledgements. We thank A. Boccaletti (LESIA, Paris, France) who provided the transmission of the SPHERE coronagraphs. We also thank the anonymous Referee for his careful reading of the manuscript as well as his insightful comments and suggestions. This work has made use of the SPHERE Data Centre, jointly operated by OSUG/IPAG (Grenoble, France), PYTHEAS/LAM/CESAM (Marseille, France), OCA/Lagrange (Nice, France),

Observatoire de Paris/LESIA (Paris, France), and Observatoire de Lyon/CRAL (Lyon, France).

References

- Allard, F., Allard, N. F., Homeier, D., et al. 2007, *A&A*, **474**, L21
Allard, F., Guillot, T., Ludwig, H. G., et al. 2003, *Symposium-International Astronomical Union* (Cambridge University Press), 211, 325
Beuzit, J.-L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, **631**, A15
Cantalloube, F. 2016, Ph.D. Thesis, Grenoble Alpes
Chabrier, G., Baraffe, I., Allard, F., & Hauschildt, P. 2000, *ApJ*, **542**, 464
Chauvin, G., Desidera, S., Lagrange, A.-M., et al. 2017, *A&A*, **605**, L9
Chen, Y., Wiesel, A., Eldar, Y. C., & Hero, A. O. 2010, *IEEE Trans. Signal Process.*, **58**, 5016
Conte, E., De Maio, A., & Ricci, G. 2002, *IEEE Trans. Signal Process.*, **50**, 1908
Conte, E., Lops, M., & Ricci, G. 1995, *IEEE Aerosp. Electron. Syst. Mag.*, **31**, 617
Currie, T., Guyon, O., Tamura, M., et al. 2017, *ApJ*, **836**, L15
Delorme, P., Meunier, N., Albert, D., et al. 2017, in *SF2A-2017: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, held 4-7 July, 2017 in Paris, eds. C. Reylé, P. Di Matteo, F. Herpin, et al., 347
Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2018a, *A&A*, **618**, A138
Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2018b, in *2018 25th IEEE Int. Conf. Image Process.*, 2735
Galicher, R., Boccaletti, A., Mesa, D., et al. 2018, *A&A*, **615**, A92
Gini, F., & Greco, M. 2002, *Signal Process.*, **82**, 1847
Gonzalez, C., Absil, O., Absil, P.-A., et al. 2016, *A&A*, **589**, A54
Gonzalez, C., Absil, O., & Van Droogenbroeck, M. 2018, *A&A*, **613**, A71
Greenbaum, A. Z., Pueyo, L., Ruffio, J. B., et al. 2018, *GPI spectra of HR8799 c, d, and e in HK bands with KLIP Forward Modeling*
Howard, A. W., Johnson, J. A., Marcy, G. W., et al. 2010, *ApJ*, **721**, 1467
Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, *PASP*, **127**, 890
Kay, S. M. 1998, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Upper Saddle River, NJ, USA: Prentice Hall), 1
Kendall, M. G., Stuart, A., & Ord, J. K. 1948, *The Advanced Theory of Statistics* (JSTOR), 1
Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, E. 2007, *ApJ*, **660**, 770
Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *Proc. Nat. Acad. Sci.*, **111**, 12661
Marois, C., Correia, C., Galicher, R., et al. 2014, in *Adaptive Optics Systems IV*, Int. Soc. Opt. Photonics, 9148, 91480U
Marois, C., Correia, C., Vigan, J.-P., & Currie, T. 2013, *Proc. Int. Astron. Union*, **8**, 48
Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, **641**, 556
Mawet, D., Hirsch, L., Lee, E. J., et al. 2019, *ApJ*, **157**, 33
Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, **792**, 97
Milli, J., Kasper, M., Bourget, P., et al. 2018, in *Adaptive Optics Systems VI*, Int. Soc. Opt. Photonics, 10703, 107032A
Morzinski, K. M., Close, L. M., Males, J. R., et al. 2014, in *Adaptive Optics Systems IV*, Int. Soc. Opt. Photonics, 9148, 914804
Mugnier, L. M., Cornia, A., Sauvage, J.-F., et al. 2009, *J. Opt. Soc. Am. A*, **26**, 1326
Pascal, F., Chitour, Y., Ovarlez, J.-P., Forster, P., & Larzabal, P. 2008, *IEEE Trans. Signal Process.*, **56**, 34
Portilla, J., Strela, V., Wainwright, M. J., & Simoncelli, E. P. 2003, *IEEE Trans. Image Process.*, **12**
Rameau, J., Chauvin, G., Lagrange, A.-M., et al. 2013a, *ApJ*, **772**, L15
Rameau, J., Chauvin, G., Lagrange, A.-M., et al. 2013b, *ApJ*, **779**, L26
Sauvage, J. F., Fusco, T., Guesalaga, A., et al. 2015, *Adaptive Optics for Extremely Large Telescopes 4-Conference Proceedings*, 1
Smith, I., Ferrari, A., & Carillet, M. 2009, *IEEE Trans. Signal Process.*, **57**, 904
Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJ*, **755**, L28
Traub, W. A., & Oppenheimer, B. R. 2010, *Exoplanets* (University of Arizona Press), 111
Vigan, A., Moutou, C., Langlois, M., et al. 2010, *MNRAS*, **407**, 71
Wagner, K., Follete, K. B., Close, L. M., et al. 2018, *ApJ*, **863**, L8
Wainwright, M. J., & Simoncelli, E. P. 2000, *Advances in Neural Information Processing Systems*, 855
Wang, J., Dogandzic, A., & Nehorai, A. 2006, *IEEE Trans. Signal Process.*, **54**, 3884

Appendix A: Estimation of the local statistics of the background

Let $\{\mathbf{r}_{\theta_k, t_\ell}\}_{\ell=1:T}$ be a collection of 2D patches centered at location θ_k . Under our model, the ℓ -th patch of the collection is distributed under \mathcal{H}_0 according to a normal distribution with mean \mathbf{m}_k and spatial covariance $\mathbf{C}_{k,\ell} = \sigma_{k,\ell}^2 \mathbf{C}_k$. In this Appendix, we derive the expression of the maximum likelihood estimators $\widehat{\mathbf{m}}_k$, $\widehat{\sigma}_{k,\ell}^2$ and $\widehat{\mathbf{S}}_k$ of \mathbf{m}_k , $\sigma_{k,\ell}^2$ and \mathbf{C}_k , respectively. In the paper, the notation $\widehat{\mathbf{C}}_k$ is kept to denote the shrinkage estimator obtained by reducing the value of off-diagonal elements of the maximum likelihood estimator $\widehat{\mathbf{S}}_k$.

Under the assumption that each patch is an independent realization of our model of the background, the co-log-likelihood \mathcal{L} of the collection is given by:

$$\begin{aligned} \mathcal{L} &= -\log p(\{\mathbf{r}_{\theta_k, t_\ell}\}_{\ell=1:T} | \mathbf{m}_k, \{\sigma_{k,\ell}^2\}_{\ell=1:T}, \mathbf{C}_k) \\ &\propto \sum_{\ell=1}^T \log \det(\sigma_{k,\ell}^2 \mathbf{C}_k) + \sum_{\ell=1}^T \frac{1}{\sigma_{k,\ell}^2} \bar{\mathbf{r}}_{\theta_k, t_\ell}^t \mathbf{C}_k^{-1} \bar{\mathbf{r}}_{\theta_k, t_\ell} + \text{const.}, \end{aligned} \quad (\text{A.1})$$

with $\bar{\mathbf{r}}_{\theta_k, t_\ell} = \mathbf{r}_{\theta_k, t_\ell} - \mathbf{m}_k$. The maximum likelihood estimators correspond to the location of the minimum of (A.1), where the gradient of the co-log-likelihood is equal to zero.

From the condition $\left. \frac{\partial \mathcal{L}}{\partial \sigma_{k,\ell}^2} \right|_{\sigma_{k,\ell}^2 = \widehat{\sigma}_{k,\ell}^2} = 0$, we get:

$$\frac{K}{\widehat{\sigma}_{k,\ell}^2} - \frac{1}{\widehat{\sigma}_{k,\ell}^4} \bar{\mathbf{r}}_{\theta_k, t_\ell}^t \mathbf{C}_k^{-1} \bar{\mathbf{r}}_{\theta_k, t_\ell} = 0,$$

which leads to

$$\widehat{\sigma}_{k,\ell}^2 = \frac{1}{K} \bar{\mathbf{r}}_{\theta_k, t_\ell}^t \mathbf{C}_k^{-1} \bar{\mathbf{r}}_{\theta_k, t_\ell}. \quad (\text{A.2})$$

The variance at time t_ℓ around pixel k is therefore estimated by computing the sample variance of a spatially whitened version of the ℓ -th patch.

From the condition $\nabla_{\mathbf{m}_k} \mathcal{L} |_{\mathbf{m}_k = \widehat{\mathbf{m}}_k} = \mathbf{0}$, we get:

$$-2 \sum_{\ell=1}^T \frac{1}{\sigma_{k,\ell}^2} \mathbf{C}_k^{-1} (\mathbf{r}_{\theta_k, t_\ell} - \widehat{\mathbf{m}}_k) = \mathbf{0}.$$

Since \mathbf{C}_k^{-1} is not singular, we obtain:

$$\widehat{\mathbf{m}}_k = \frac{\sum_{\ell=1}^T \frac{1}{\sigma_{k,\ell}^2} \mathbf{r}_{\theta_k, t_\ell}}{\sum_{\ell=1}^T \frac{1}{\sigma_{k,\ell}^2}}. \quad (\text{A.3})$$

The mean background in patches around the k -th pixel corresponds to a weighted average where a patch at time t_ℓ with large variance is given a small weight.

From the condition $\nabla_{\mathbf{C}_k} \mathcal{L} |_{\mathbf{C}_k = \widehat{\mathbf{S}}_k} = \mathbf{0}$, we derive:

$$T \cdot \widehat{\mathbf{S}}_k^{-1} - \widehat{\mathbf{S}}_k^{-1} \left(\sum_{\ell=1}^T \frac{1}{\sigma_{k,\ell}^2} \bar{\mathbf{r}}_{\theta_k, t_\ell} \bar{\mathbf{r}}_{\theta_k, t_\ell}^t \right) \widehat{\mathbf{S}}_k^{-1} = \mathbf{0},$$

which gives:

$$\widehat{\mathbf{S}}_k = \frac{1}{T} \sum_{\ell=1}^T \frac{1}{\sigma_{k,\ell}^2} \bar{\mathbf{r}}_{\theta_k, t_\ell} \bar{\mathbf{r}}_{\theta_k, t_\ell}^t. \quad (\text{A.4})$$

The spatial covariance estimator corresponds to the sample covariance upon a proper rescaling by a factor $1/\sigma_{k,\ell}$ of residual patches $\bar{\mathbf{r}}_{\theta_k, t_\ell}$.

Appendix B: Proof of the robustness of the estimators

In this appendix, we show that our estimators are robust in the sense that, should the patch $\mathbf{r}_{\theta_k, t_{\ell^*}}$ be replaced by a scaled version $\mathbf{r}'_{\theta_k, t_{\ell^*}} = \eta \cdot \mathbf{r}_{\theta_k, t_{\ell^*}}$ with an arbitrarily large factor η , at a given frame ℓ^* , the estimators of the background statistics $\widehat{\mathbf{m}}'_k$ and $\widehat{\mathbf{S}}'_k$ as well as the flux estimator $\widehat{\alpha}'$ computed on the collection with the scaled patch remain bounded. In fact, when η grows to infinity, we show that the patch $\mathbf{r}'_{\theta_k, t_{\ell^*}}$ gets completely discarded in the estimation of the flux.

Since the estimators $\widehat{\mathbf{m}}_k$, $\widehat{\mathbf{C}}_k$ and $\widehat{\sigma}_{k,\ell}^2$ are defined implicitly through a fixed-point procedure, directly proving the robustness is difficult. Instead, we show that when η approaches infinity, the following estimates correspond to a fixed point:

$$\begin{cases} \widehat{\sigma}_{k,\ell}^{2'} = \widehat{\sigma}_{k,\ell}^2 & \text{when } \ell \neq \ell^* \\ \widehat{\sigma}_{k,\ell^*}^{2'} \rightarrow \infty \\ \widehat{\mathbf{m}}'_k = \widehat{\mathbf{m}}_k \\ \widehat{\mathbf{S}}'_k = \frac{T-1}{T} \widehat{\mathbf{S}}_k + \frac{K}{T} \frac{\mathbf{r}_{\theta_k, t_{\ell^*}} \mathbf{r}_{\theta_k, t_{\ell^*}}^t}{\mathbf{r}_{\theta_k, t_{\ell^*}}^t (\widehat{\mathbf{S}}'_k)^{-1} \mathbf{r}_{\theta_k, t_{\ell^*}}}, \end{cases} \quad (\text{B.1})$$

where $\widehat{\sigma}_{k,\ell}^2$, $\widehat{\mathbf{m}}_k$ and $\widehat{\mathbf{S}}_k$ correspond to the estimators computed on the collection $\{\mathbf{r}_{\theta_k, t_\ell}\}_{\ell \neq \ell^*}$ of the patches with the ℓ^* -th frame removed.

Starting with the initial values given in (B.1), the application of Eq. (A.2) leaves the estimated variances $\widehat{\sigma}_{k,\ell}^{2'}$ unchanged for all ℓ different from ℓ^* . For ℓ^* , $\widehat{\sigma}_{k,\ell^*}^{2'} \sim \eta^2 \widehat{\sigma}_{k,\ell}^2$ which tends to infinity when η approaches infinity.

From Eq. (A.3), we see that the patch $\mathbf{r}'_{\theta_k, t_{\ell^*}}$ with associated infinite variance $\widehat{\sigma}_{k,\ell^*}^{2'}$ has zero weight: $\frac{1}{\widehat{\sigma}_{k,\ell^*}^{2'}} \mathbf{r}'_{\theta_k, t_{\ell^*}} \sim \frac{1}{\eta^2 \widehat{\sigma}_{k,\ell}^2} \mathbf{r}_{\theta_k, t_{\ell^*}}$, so that $\lim_{\eta \rightarrow \infty} \frac{1}{\widehat{\sigma}_{k,\ell^*}^{2'}} \mathbf{r}'_{\theta_k, t_{\ell^*}} = \mathbf{0}$. The mean patch then corresponds to that obtained on the collection of patches with the ℓ^* -th patch excluded.

The approximations $\widehat{\sigma}_{k,\ell^*}^{2'} \sim \frac{\eta^2}{K} \mathbf{r}_{\theta_k, t_{\ell^*}}^t (\widehat{\mathbf{S}}'_k)^{-1} \mathbf{r}_{\theta_k, t_{\ell^*}}$ and $\bar{\mathbf{r}}'_{\theta_k, t_{\ell^*}} \sim \eta \mathbf{r}_{\theta_k, t_{\ell^*}}$ together with the update of the sample covariance estimate (A.4) lead to $\widehat{\mathbf{S}}'_k = \frac{T-1}{T} \widehat{\mathbf{S}}_k + \frac{K}{T} \frac{\mathbf{r}_{\theta_k, t_{\ell^*}} \mathbf{r}_{\theta_k, t_{\ell^*}}^t}{\mathbf{r}_{\theta_k, t_{\ell^*}}^t (\widehat{\mathbf{S}}'_k)^{-1} \mathbf{r}_{\theta_k, t_{\ell^*}}}$.

In conclusion, the set of parameters defined in (B.1) is a fixed-point. Moreover, neither the mean background patch $\widehat{\mathbf{m}}'_k$ nor the spatial covariance $\widehat{\mathbf{S}}'_k$ depend on η (i.e., they both remain bounded when $\eta \rightarrow \infty$). It remains to show the robustness of the flux estimator $\widehat{\alpha}$ of a source.

If the scaled patch is not in the collection of all patches that contain the PSF of the source, then its influence is limited since we have just shown that the background statistics are robustly estimated. In the case of the superimposition of a scaled background and of the source of interest (as is the case when a poor correction of the AO leads to large stellar leakages at some locations of the field of view), the estimated flux becomes:

$$\widehat{\alpha}(\phi_{t_\ell}) = \frac{\sum_{\ell=1}^T \frac{1}{\widehat{\sigma}_{k,\ell}^{2'}} \mathbf{h}_{[\phi_{t_\ell}]}(\phi_{t_\ell})^t \cdot \widehat{\mathbf{C}}_{[\phi_{t_\ell}]}^{-1} \cdot (\mathbf{r}'_{[\phi_{t_\ell}], t_\ell} - \widehat{\mathbf{m}}_{[\phi_{t_\ell}]}')}{\sum_{\ell=1}^T \frac{1}{\widehat{\sigma}_{k,\ell}^{2'}} \mathbf{h}_{[\phi_{t_\ell}]}(\phi_{t_\ell})^t \cdot \widehat{\mathbf{C}}_{[\phi_{t_\ell}]}^{-1} \cdot \mathbf{h}_{[\phi_{t_\ell}]}(\phi_{t_\ell})}. \quad (\text{B.2})$$

Therefore, when η becomes large, the term $(\mathbf{r}'_{[\phi_{t_\ell^\star}]_{t_\ell^\star}} - \widehat{\mathbf{m}}_{[\phi_{t_\ell^\star}]}) / \hat{\sigma}_{k,\ell^\star}^{2'}$ tends to zero and we obtain:

$$\widehat{\alpha}(\phi_{t_\ell}) \approx \frac{\sum_{\ell \neq \ell^\star} \frac{1}{\hat{\sigma}_{k,\ell}^{2'}} \mathbf{h}_{[\phi_{t_\ell}]}(\phi_{t_\ell})^t \cdot \widehat{\mathbf{C}}_{[\phi_{t_\ell}]}^{-1} \cdot (\mathbf{r}'_{[\phi_{t_\ell}]_{t_\ell}} - \widehat{\mathbf{m}}_{[\phi_{t_\ell}]})}{\sum_{\ell \neq \ell^\star} \frac{1}{\hat{\sigma}_{k,\ell}^{2'}} \mathbf{h}_{[\phi_{t_\ell}]}(\phi_{t_\ell})^t \cdot \widehat{\mathbf{C}}_{[\phi_{t_\ell}]}^{-1} \cdot \mathbf{h}_{[\phi_{t_\ell}]}(\phi_{t_\ell})}, \quad (\text{B.3})$$

which indicates that the scaled frame $\mathbf{r}'_{[\phi_{t_\ell^\star}]_{t_\ell^\star}}$ tends to be completely discarded when the scaling factor η is large. Estimation of the photometry and of the astrometry is therefore robust to the presence of a frame with an arbitrarily large background.