

# K-Stacker: Keplerian image recombination for the direct detection of exoplanets

M. Nowak<sup>1,2</sup>, H. Le Coroller<sup>2</sup>, L. Arnold<sup>3</sup>, K. Dohlen<sup>2</sup>, D. Estevez<sup>5,2</sup>, T. Fusco<sup>4</sup>, J.-F. Sauvage<sup>4</sup>, and A. Vigan<sup>2</sup>

<sup>1</sup> LESIA, Observatoire de Paris, Université PSL, CNRS, Sorbonne Université, Univ. Paris Diderot, Sorbonne Paris Cité, 5 place Jules Janssen, 92195 Meudon, France

e-mail: [mathias.nowak@obspm.fr](mailto:mathias.nowak@obspm.fr)

<sup>2</sup> Aix-Marseille Université, CNRS, CNES, LAM, Marseille, France

e-mail: [herve.lecoroller@lam.fr](mailto:herve.lecoroller@lam.fr)

<sup>3</sup> Aix-Marseille Université, CNRS, OHP (Observatoire de Haute Provence), Institut Pythéas UMS 3470, 04870 Saint-Michel-l'Observatoire, France

<sup>4</sup> ONERA – The French Aerospace Lab BP72 – 29 avenue de la Division Leclerc, 92322 Chatillon Cedex, France

<sup>5</sup> Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, LAPP, 74000 Annecy, France

Received 15 August 2016 / Accepted 30 March 2018

## ABSTRACT

**Context.** Angular differential imaging (ADI) takes advantage of the field rotation naturally induced by altitude-azimuth mounts to reduce static speckle noise. Used with facilities like SPHERE at the VLT, this technique allows one to achieve contrast ratios of  $10^{-6}$ . The ADI method, however, intrinsically limits the useful exposure time on a given target (to about 1–2 h per night). Detecting fainter exoplanets requires the combination of multiple observations acquired on different nights, potentially spread over several weeks or months, but the unknown orbital motion of the planet makes it particularly difficult to properly combine all observations. In the near future, with the upcoming generation of Extremely Large Telescopes (ELTs) with increased resolution, the orbital motion may even become a problem on a single night.

**Aims.** We present a proof of concept for a new algorithm which can be used to detect exoplanets in high-contrast images. The algorithm properly combines multiple observations acquired during different nights, taking into account the orbital motion of the planet.

**Methods.** We simulate SPHERE/IRDIS time series of observations in which we blindly inject planets on random orbits, at random levels of signal-to-noise ratio (S/N), below the detection limit (down to  $S/N \approx 1.5$ ). We then use an optimization algorithm to “guess” the orbital parameters, and take into account the orbital motion to properly recombine the different images and eventually detect the planets.

**Results.** We show that an optimization algorithm can indeed be used to find undetected planets in temporal sequences of images, even if they are spread over orbital time scales. As expected, the typical gain in S/N is  $\sqrt{n}$ ,  $n$  being the number of observations combined. We find that the K-Stacker algorithm is able de-orbit and combine the images to reach a level of performance similar to what could be expected if the planet was not moving. We find recovery rates of  $\approx 50\%$  at  $S/N = 5$ . We also find that the algorithm is able to determine the position of the planet in individual frames at one pixel precision, even despite the fact that the planet itself is below the detection limit in each frame.

**Conclusions.** Our simulations show that K-Stacker can be used to detect planets at very low S/N level, down to  $\approx 2$  in individual frames, for series of ten images. This could be used to increase the contrast limit of current exoplanet imaging instruments and to discover fainter bodies. We also suggest that the ability of K-Stacker to determine the position of the planet in every image of the time series could be used as part of a new observing strategy in which long exposures would be broken into shorter ones spread over months. This could make it possible to determine the orbital parameters of a planet without multiple high-S/N ( $>5$ ) detections.

**Key words.** techniques: image processing – planets and satellites: detection

## 1. Introduction

The direct detection of exoplanetary bodies is an extremely challenging task: for a planet orbiting at 10 AU around a star at a few tens of parsec, the typical angular separation is about 0.1 to 1 arcsec, and typical contrast ratios in the near-infrared (NIR) are expected to be of the order of  $\approx 10^{-5}$  for a young Jupiter-like planet (Marley et al. 2007) to  $10^{-10}$  for an Earth-like planet. The recent development of eXtreme Adaptive Optics (XAO) systems, which are now able to deliver high Strehl ratios by correcting high-order wavefront errors, as well as of coronagraphic imaging systems has led to a new generation of high-contrast imagers (GPI, Macintosh et al., 2014; SPHERE, Beuzit et al., 2008).

However, even with these state-of-the-art facilities, observations are still crippled by speckle noise, originating in atmospheric turbulence and instrumental defects. Since atmospheric speckles have very short decorrelation times, they can be averaged by increasing the exposure time. The situation is a little different when dealing with the so-called pseudo-static (instrumental) speckles, which have much longer correlation times. Even in the most favorable cases, without any specific strategy, static speckles typically limit the useful exposure time to a mere handful of seconds (Macintosh et al. 2005; Soummer et al. 2007; Hinkley et al. 2007).

Numerous image-processing techniques have been proposed to mitigate this static speckle noise, and some are currently

being applied, for example, to SPHERE and GPI data with great success. But subtracting static speckles without affecting any potential exoplanet light is not easy. Some methods (Racine et al. 1999; Sparks & Ford 2002; Thatte et al. 2007) use the fact that, contrary to a planet, the distance of any given speckle from the center of the star scales as  $\lambda/D$ , where  $\lambda$  is the wavelength of observation. This allows one to disentangle the speckles from the planet, and subtract them. On a similar line of thought, one can also use the difference of polarization between the unpolarized starlight and the polarized planet signal to distinguish them, as is done with the Differential Polarimetry method (Canovas et al. 2011). Finally, requiring absolutely no spectral or polarimetric capability whatsoever, the angular differential imaging (ADI; Marois et al. 2006) uses the rotation of the field of view (FoV) naturally induced by alt-az mounts to disentangle pseudo-static speckles from any moving planet.

The ADI and SDI methods are currently implemented in the SPHERE data-reduction pipeline, allowing one to routinely achieve contrast ratios of  $10^{-5}$  and up to  $10^{-6}$  in excellent conditions (Zurlo et al. 2014, 2016; Vigan et al. 2015). However, when using this ADI method, the total time for an observation is limited by the FoV rotation to  $\approx 1$  h, a limitation which was already pointed out by Marois et al. (2006).

An easy way to improve the detection limit of any astronomical instrument is usually to use longer exposure times. But when using a technique like ADI, this can rapidly turn to a very intricate problem: if exposures are limited to about an hour per night, gathering, for example, 10–20 h of exposure requires observing on 10–20 different nights. Due to weather and observatory constraints, this can rapidly lead to extended time sequences, where the different images are possibly taken over different observing runs, with intervals of up to several months. In such cases, the orbital motion of the planet, which up until now has almost always been ignored when combining images, must be taken into account. This is especially true when using a large telescope, probing the innermost part of close-by stellar systems. Furthermore, even data acquired during succeeding nights can be affected by the orbital motion of the planet (see Males et al. 2013 for a detailed discussion of this particular problem).

Males et al. (2015) also worked on the problem of recombining images acquired on longer timescales, in which the planet was moving on a significant part of its orbit. However, they use the assumption that the planet, albeit very faint, could be seen in each individual frame. In this regime, deorbiting can help improve the signal-to-noise ratio (S/N) of the planet, but cannot help with detecting new planets, which could not be seen in a single frame. Males et al. themselves stated that “an important area of investigation will be the performance of [ODI] when there is no prior information with which to determine the orbits.”

This is precisely the problem we intend to tackle with the K-Stacker algorithm.

In a previous paper (Le Coroller et al. 2015), we proposed a new technique to detect exoplanets, called K-Stacker, which allows one to combine multiple observations made over different nights to increase the contrast limit of direct imaging instruments. In the present work we describe in detail the K-Stacker algorithm and we provide a statistical analysis intended as a “proof of concept” of this method.

In Sect. 2, we describe the simulated images on which we tested our K-Stacker algorithm. Section 3 focuses on the problem of recombining multiple observations in which the planet remains undetected, and gives a description of our algorithm, based on the combination of a brute-force and a gradient-descent

method to determine the orbital parameters. In Sect. 4, we present the results of a blind test of this algorithm performed on 50 independent and random simulated observations. We discuss the performances of the algorithm and the required computer resources in Sect. 5. Our conclusions are given in Sect. 6.

## 2. Simulated data

In order to develop and test the K-Stacker algorithm, we developed a simple simulator to generate images similar to the one produced by the SPHERE/IRDIS instrument (Dohlen et al. 2008) in its dual-band imaging mode (Vigan et al. 2010), in terms of general characteristics (size of the FoV, width of the PSF, etc.). Whereas we believe that K-Stacker could benefit from an ADI/SDI pre-processing step, which removes most of the pseudo-static speckles, we do not simulate ADI or SDI reduced data.

Even if a number of studies (Hinkley et al. 2007; Soummer et al. 2007; Martinez et al. 2012, 2013) have been done on the temporal evolution of instrumental speckles, the behavior of these quasi-static speckles is hard to simulate because they are related to many different factors (thermal conditions, mechanical stress, instrumental configuration, etc.).

In the present work, we focus on the case of pure atmospheric speckles behind an XAO, which are completely uncorrelated from one image to the other. We also assume that the coronagraph is perfect, and that the coherent part of the light is fully removed. If we do not take into account variations in amplitude of the incoming wavefront, a perfect coherent light suppression can be defined analytically using the following equation (Fusco et al. 2006):

$$I(x, y) = \left| F \left[ \left( e^{i\Delta\phi(a,b)} - e^{\frac{1}{2}i\sigma_{\Delta\phi}^2} \right) P(a, b) \right] (x, y) \right|^2, \quad (1)$$

where  $I(x, y)$  represents the intensity in the focal plane,  $P(a, b)$  the aperture function,  $\Delta\phi(a, b)$  the phase error of the wavefront,  $F$  the Fourier transform, and  $\sigma_{\Delta\phi}$  the standard deviation of the phase error in the incoming wavefront.

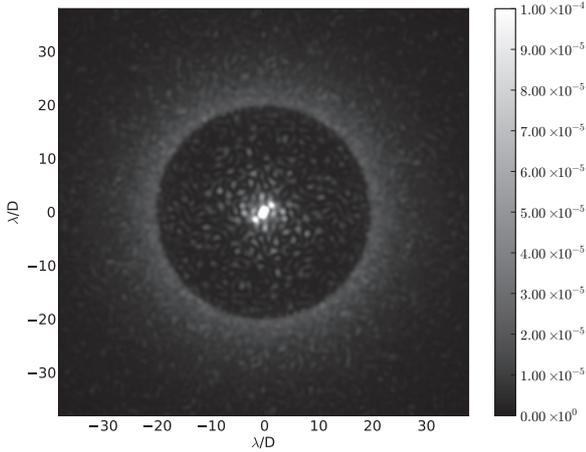
In the above equation, the term  $\Delta\Phi$  is what is ultimately responsible for the speckle noise in the final images, which is known to originate from atmospheric phase variations and/or instrumental defects.

The atmospheric phase masks were simulated using an IDL (Interactive Data Language) code developed by Fusco et al. (2006), which generates phase errors downstream of the adaptive optics system.

The images generated are  $512 \times 512$  pixels, and only represent a part of the bigger SPHERE detector. Our code generates monochromatic images, with a 12.25 mas/pixel spatial sampling corresponding to SPHERE/IRDIS.

In Fig. 1, we show one of our typical exposures, made with our simulated SPHERE/IRDIS instrument, under a sky with 0.8'' seeing. This image was obtained by averaging 100 exposures made with 100 uncorrelated atmospheric masks. The image is normalized to the central peak intensity of the non-coronagraphic PSF. As it was unnecessary for our work, we did not try to calibrate our simulations to match the correct SPHERE/IRDIS photometry, and thus did not include photon noise.

Finally, false planets are injected by directly adding a non-coronagraphic PSF in the images. This means that neither the coronagraph nor the atmospheric and/or instrumental phase errors have any impact on the exoplanet signal.



**Fig. 1.** Simulation of an AO-corrected image at  $1.6\mu\text{m}$  for SPHERE/IRDIS, under a  $0.8''$  seeing sky. The AO corrected area is clearly visible, and extends to about  $20\lambda/D$  ( $1.6''$ ). The image is scaled to the peak value of the non-coronagraphic star PSF.

### 3. Recombining the images: the K-Stacker algorithm

#### 3.1. K-Stacker principle: the recombination as an optimization problem

We consider a set of  $n$  images  $I_1, \dots, I_n$  of a given star, acquired at times  $t_1, \dots, t_n$ , and suppose that an exoplanet is orbiting around this star. In the  $k$ th image, the planet is at position  $(x_k, y_k)$ , but remains undetected. Thanks to the law of orbital mechanics, these positions  $x_k, y_k$  can be related to seven parameters: the six orbital elements, that is, eccentricity  $e$ , semi major-axis  $a$ , epoch at perihelion  $t_0$ , longitude of the ascending node  $\Omega$ , inclination  $i$ , argument at periapsis  $\theta_0$ , and the mass of the central star  $M_*$ . The distance of the star  $d_*$  is also required to project the orbit on the charge coupled device (CCD).

K-Stacker is based on the idea that when trying to recombine the images to detect a hidden orbiting planet, a strong spot should emerge only when the images are recombined along the correct orbit. Otherwise, the speckles should simply average to a certain value, within statistical fluctuations.

To recombine the images along a given orbit, we first compute the position of the planet predicted by the laws of orbital mechanics at times  $t_1$  in image  $I_1$ , at time  $t_2$  in image  $I_2$ , and so on. Then, we rotate/translate each image  $I_k$  for  $k > 0$  to align all these positions, and add all the images. This means that when we recombine the images, we actually know where we expect to see a strong feature if the orbital parameters are correct, which suggests the use of a S/N at this particular position as a measure of the “quality of the recombination”.

For a given orbit, the position of a planet at a time  $t$  is defined by  $x = (t, a, e, t_0, \Omega, i, \theta_0, M_*)$  and assuming that the noise in the different images is fully uncorrelated, the S/N can be computed in the following way:

- First, for each image  $I_k$ , we compute the position  $(x_k, y_k)$  by solving the Kepler equation with a Newton algorithm, projecting the orbital position on the detector plane. This gives a set of one position per image, which can be expressed in polar coordinates:  $(r_1, \theta_1), \dots, (r_n, \theta_n)$ .
- Then, for each  $k$  in  $\{1, \dots, n\}$ , we compute the flux at position  $(r_k, \theta_k)$  in the image  $I_k$ , using a circular integration box of radius equal to the full width at half-maximum (FWHM) of

**Table 1.** Main characteristics of the grid used by our brute-force algorithm, computed for a star of mass  $M_\odot$  at 10 pc.

Param.	Unit	Min val.	Max val.	Step size	Points
$a$	AU	0.9	7.5	0.7	10
$e$	–	0	0.8	0.08	10
$t_0$	yr	0	20	0.2	100
$\Omega$	rad	$-\pi$	$+\pi$	0.18	35
$i$	rad	$-\pi$	0	0.5	7
$\theta_0$	rad	$-\pi$	$+\pi$	0.18	35

the non-coronagraphic PSF of the instrument. To this value, we subtract the estimate of the background flux in image  $I_k$  at radius  $r_k$ . This gives the signal in image  $k$ ,  $s_k(x)$  at the position  $(r_k, \theta_k)$ . The total signal is the sum of all of these values in each image  $S(x) = \sum_k s_k(x)$ .

- Finally, for each  $k$  in  $\{1, \dots, n\}$ , we compute the noise level  $\sigma_k$  at radius  $r_k$  in image  $I_k$ , and add these values quadratically to obtain the global noise value:  $N(x) = \sqrt{\sum_k \sigma_k^2}$ . In this case  $S/N(x) = S(x)/N(x)$ .

This S/N, seen as a function of the six orbital parameters (and of the star mass  $M$  and distance  $d$ , if unknown), can be optimized using a modified brute-force algorithm to detect planets hidden in the individual frames  $I_k$ .

#### 3.2. Modified brute-force optimization

##### 3.2.1. Brute force stage

We use a brute-force algorithm to search for the maximum  $S/N$ . We keep a simple shape for the search grid (six-dimensional rectangle, with linear sampling). The algorithm may therefore lose some time exploring unrealistic possibilities. We explore the entire range of possible values, from  $-\pi$  to  $+\pi$  for  $\Omega$  and  $\theta_0$ , and from 0 to  $\pi$  for  $i$ . For  $e$ , we explore values ranging from 0 to 0.8. For  $t_0$ , we also explore the entire possible range, that is, from 0 to  $T$ ,  $T$  being the orbital period. Since the value of the orbital period depends on  $a$ , we take the largest interval: from 0 to  $2\pi \times \sqrt{\frac{a_{\text{max}}^3}{GM}}$ . For  $a$ , we chose to explore values ranging from  $0.09 \times d$  to  $0.75 \times d$ , which correspond to the A.O. corrected area of SPHERE/IRDIS,  $d$  being the distance of the star in parsecs.

The smaller the step size of the grid, the higher the chances are of finding the true maximum, but the longer the computation time is. We estimated the best sampling for each parameter empirically, by looking at the typical width of the global maximum in different configurations.

In Table 1, we summarize the different characteristics of our grid, built for analyzing observations of a star of mass  $1 M_\odot$ , located at 10 pc. Whereas the step sizes and number of points to use are only rough estimates, this table shows that the total number of points to be explored is of the order of  $10^8$ . The simulations have shown that such a grid allows to find a solution in a reasonable computation time (see Sect. 5).

##### 3.2.2. Gradient-descent re-optimization stage

The weak point of the brute-force algorithm is that it may miss the global maximum, likely to be between the points of the finite grid sampling. To circumvent this particular issue, we add a gradient-descent optimization stage to the process, and re-optimize some of the best solutions found by the brute-force

algorithm. To decide how many of the best grid values should be re-optimized, we ran the brute-force algorithm on a set of 25 images generated using our SPHERE/IRDIS simulator (see Sect. 2), in typical sky conditions (seeing of  $0.8''$ ), in which no planet was introduced. This gave us a typical distribution of the noise values sampled by the grid, in which we found that the highest  $S/N$  value was 5.2, and that only 100  $S/N$  values are higher than 4.5 and 1000 higher than 4.2. This means that if we re-optimize the best value found on the grid, a planet will be detected only if the corresponding  $S/N$  grid-maximum reaches a value higher than 5.4. However, if we re-optimize the 100 (resp. 1000) best values, then the planet will be detected if the corresponding grid value is greater than 5.1 (resp. 4.7). Compared to the total  $10^8$  points of the grid, re-optimizing 100 or 1000 values takes little time, and allows us to recover low- $S/N$  planets.

In conclusion of this section, we have built an algorithm that works in three steps:

- First, a brute-force algorithm is used to determine the value of the  $S/N$  function in each point of the grid.
- Then the  $p = 100$  highest values found are re-optimized by a gradient-descent algorithm or similar. We use the Broyden-Fletcher-Goldfarb-Shanno method (Fletcher 1987).
- Finally, we search for planets in all images corresponding to the re-optimized  $S/N$  values above 4.7.

## 4. Results

### 4.1. Blind test

To test the K-Stacker algorithm, we performed a blind experiment. We used the SPHERE/IRDIS simulator presented in Sect. 2 to build 50 sets of ten independent simulated observations, with a seeing value of  $0.8''$ . In each set, a planet was then injected on a random orbit, at a random  $S/N$  level. The orbit was selected by drawing the six parameters randomly using the distributions given in Table 2 allowing the algorithm to redraw a new set of parameters when the first one resulted in an orbit going outside of the AO-corrected area in at least one image. The mass and the distance of the star are expected to be known with a good accuracy, as is usually the case for the bright stars observed in coronagraphy. For each simulation, the planet was not injected at “constant  $S/N$ ”, but rather at “constant flux”, to simulate what is expected from a constant exposure time series of observations. For each series of ten observations, the noise level  $\sigma_1$  at the position of the planet in the first image is computed, and the planet is injected at a flux value of  $F = \text{snr}_1 \times \sigma_1$  on top of the background in each image, where  $\text{snr}_1$  is randomly selected among four possible values: 0,  $5/\sqrt{10}$ ,  $7.5/\sqrt{10}$ , or  $10/\sqrt{10}$ .

Because the local value of the background behind the planet varies in each observation, as does the noise level if the planet is moving with respect to the central star, the exact  $S/N$  differs from one image to the other, and the total expected  $S/N$  level of the simulation can only loosely be expected to be  $S/N \approx \sqrt{10} \times \text{snr}_1$ .

In all these simulations, the star has a mass of  $1 M_\odot$ , is located at a distance of 10 pc, and has a magnitude of 8 in the  $R$  band (AO sensing band). The ten images of each set correspond to observations made at different times, selected arbitrarily to represent a plausible sequence of K-Stacker observations. The times are given in Table 3.

The 50 sets of observations have been prepared using a dedicated computer program, which drew all the random parameters and  $S/N$  values, and stored them in a file. All the simulations

**Table 2.** Parameters used to inject the planet in the 50 simulations of our blind experiment.

Parameter	Range	Distribution
$M_*$	$1 M_\odot$	fixed value
$d_*$	10 pc	fixed value
$a$	[0.2 AU, 7.5 AU]	uniform
$e$	[0, 0.5]	uniform
$t_0$	[−20 yr, 0 yr]	uniform
$\Omega$	[−180 deg, 180 deg]	uniform
$i$	[0, 180 deg]	uniform
$\theta_0$	[−180 deg, 180 deg]	uniform
$\sqrt{10} \times S/N$	{0, 5.0, 7.5, 10.0}	uniform

**Table 3.** Dates used for each of the ten observations constituting every set of the blind test.

Observation number	Date
1	January 1, 2017
2	February 6, 2017
3	March 15, 2017
4	April 20, 2017
5	February 6, 2018
6	March 15, 2018
7	February 6, 2019
8	March 15, 2019
9	February 6, 2020
10	March 14, 2020

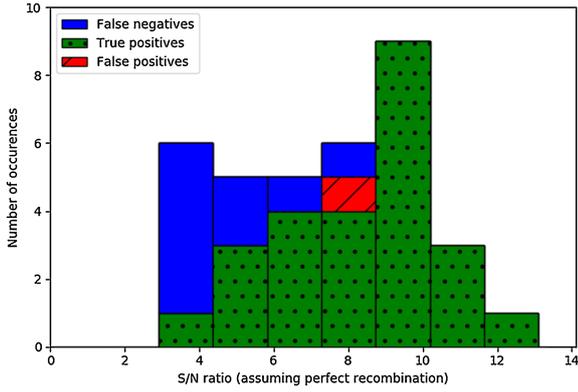
were then processed by the K-Stacker algorithm. For each simulation, K-Stacker produced a set of 100 best optimized recombinations, with their associated  $S/N$  values (before and after the re-optimization step) and orbital parameters. The observer has absolutely no idea of the  $S/N$  and orbital parameters of the planets injected in the data.

The final images produced by K-Stacker can be checked by the observer, as usually done with high-contrast images. The observer can look at the final recombined images and, based on the total recombined  $S/N$  and the shape of the detection found by K-Stacker, assign a flag to each set of observations, validating or not the potential planet candidate.

The file containing the parameters and  $S/N$  levels was only retrieved at the very end of the whole process in an effort to ensure that the observer never knew in advance what to expect from each simulation, and reduce overall bias.

### 4.2. Results obtained

Among the 50 simulations prepared and analyzed in our blind experiment, 15 were assigned by our algorithm to a group of  $S/N \approx 0$ , 12 to  $S/N \approx 5$ , 10 to  $S/N \approx 7.5$ , and 13 to  $S/N \approx 10$ . All of the 15 simulations in which no planet was actually injected were correctly identified as containing no planet candidate by the observer. In the remaining 35 simulations, 25 planets were correctly identified as planet candidates by the observer, 9 planets were missed, and one false positive was found. In Fig. 2, we show the distribution of the planets missed/found as a function of the total real  $S/N$  of the simulation computed afterwards by combining the images using the true set of orbital parameters.



**Fig. 2.** Distribution of the planet candidates found and missed among the 49 simulations as a function of the S/N given a perfect recombination of the images. The true negatives, all grouped at  $S/N = 0$  are not shown.

## 5. Discussion

### 5.1. Success rate and required computer resources

Figure 2 indicates that very high success rates (100%) can be expected of K-Stacker when the number of images acquired, and/or the planet signal strength, is high enough so that the total recombined S/N can reach at least a value of 9. When the S/N is lower (down to 6), the algorithm can still detect the planet in most cases ( $\approx 80\%$  in total). When the S/N lies around 5, a 50% recovery rate seems to be achievable using the K-Stacker algorithm. This is very encouraging, as it is comparable to what can be expected using more conventional techniques.

These success rates have been obtained using the computer cluster at Laboratoire d’Astrophysique de Marseille (LAM). We used a total of 50 cores of the cluster for an estimated total computing power of about 150 GFlops, and each of the simulations (10 images) took about 10–15 h for the grid search, plus about 1–2 h for the gradient-descent reoptimization. These numbers scale linearly with the number of frames processed in an observation.

### 5.2. The problem of false positives

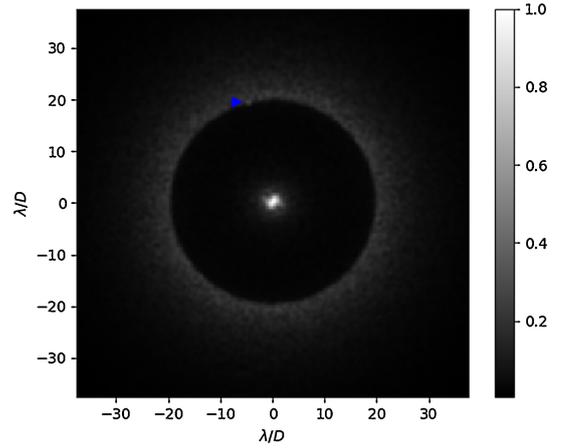
The success rate of a planet search algorithm, which defines its ability to effectively find planets and avoid “false negatives”, is an important characteristic, but its ability to avoid “false positives” (FP), where a speckle is falsely flagged as a planet, is also fundamental.

In a standard single-frame high-contrast observation, the probability of getting a false detection can be written as the product of the typical number of speckles in the images and the probability of any one speckle to be luminous enough to be mistaken for a planet:

$$P_{\text{FP}} = N_{\text{speckles}} \times P(S/N_{\text{speckle}} > S/N_t), \quad (2)$$

where  $S/N_t$  is the threshold S/N, corresponding to the minimum value of S/N a feature in the image has to reach in order to be considered as a potential planet.

Considering only the AO-corrected area, an instrument like SPHERE/IRDIS typically has a number of speckles  $N_{\text{speckles}} \approx 1000$ . Assuming that speckle noise due to atmospheric turbulence follows a normal distribution of variance 1 (in unit of S/N), the broadly used value of  $S/N_t = 5$  leads to a probability of false



**Fig. 3.** One of the individual frames of simulation 48, which lead to the only false positive of the series. The planet that has not been detected is on the edge of the AO-corrected area (indicated by the blue arrow).

detection of  $\approx 3 \times 10^{-4}$ , or about one false detection every 3500 observations.

At first glance, in the case of K-Stacker, the situation seems much more problematic since each of the  $10^8$  orbits tried by the algorithm can potentially lead to a false detection. This particular problem was already pointed out by Males et al. (2013). A first estimate of the false positive probability can be done assuming that the S/N value of each orbit tried by K-Stacker is independent of all the other values. In this case, the probability of getting a false detection on any K-Stacker run can be written as:

$$P_{\text{FP}} = 1 - [1 - P(S/N_{\text{orbit}} > S/N_t)]^{N_{\text{orbits}}}. \quad (3)$$

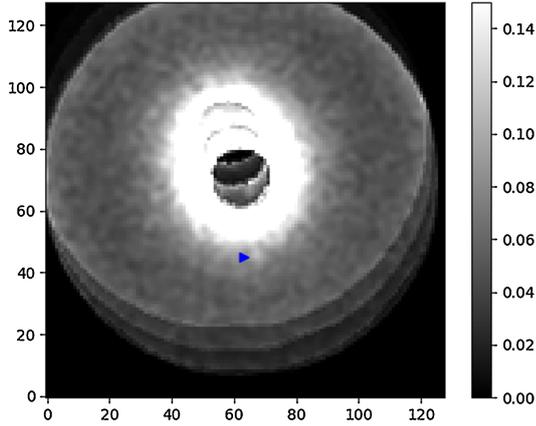
If we also assume that each orbit leads to an S/N value distributed according to a normal law of mean 0 (counting the background subtraction) and variance 1 (see Sect. 3), then the threshold  $S/N_t = 5$  leads to  $P(S/N_{\text{orbit}} > S/N_t) = 2.9 \times 10^{-7}$ , and the probability of getting a false positive is almost exactly 1.

However, in all the 50 simulations of our blind test, only one gave a false positive result. In this simulation, a planet was injected at a high S/N value (about 2–3 in each individual frame), but very close to the edge of the AO-corrected area (see Fig. 3).

The planet itself was not detected, but another feature was falsely flagged as a planet candidate. However, the observer also noted that the S/N value found by K-Stacker was low (4.80 before the reoptimization, 4.98 after), and that the shape of the recombined spot had an apparent lack of central symmetry (Fig. 4). The result was flagged as a planet candidate, but with a comment saying that it should be taken with caution. There is no doubt that if this case was a real one, the observer would have requested further observations before claiming a planet detection.

At this point, it must also be pointed out that with recent instruments like SPHERE, which provide images at several wavelengths (IFS), color-magnitude diagrams are usually used to help discriminate between potential planet candidates and brown dwarves, background features, or speckles. In our monochromatic blind test, this check cannot be done, and the rate of false detections is necessarily higher than in classical coronagraphy, which makes our result pessimistic.

Overall, the K-Stacker algorithm seems to be much more resilient to false positives than what could be expected from Eq. (3). We believe that two reasons can explain this.



**Fig. 4.** Recombined K-Stacker image where the false planet candidate can be seen (indicated by the arrow). The observer noted the asymmetry of the planet spot.

Firstly, the previous reasoning assumes that all the orbits tested by K-Stacker are independent. This leads to a very high “number of trials”, and thus to a high false positive probability. It is unlikely, though, that all the  $10^8$  orbits tested could really be independent. Different sets of orbital parameters can correspond to very similar orbits, especially when the total time spanned by the observations is small compared to the orbital period. This effectively reduces the number of independent orbits that should be taken into account in Eq. (3). We did not try to thoroughly test this hypothesis, but interestingly enough, this idea, based only on the empirical results of our K-Stacker runs and on our experience using it, agrees with the conclusion reached by Males et al. (2013), using a more sophisticated theoretical reasoning.

Secondly, we also noticed that when the optimization algorithm ends on a “noise maximum” which could lead to a false positive, the resulting image does not show a clear PSF-shaped spot, as is the case when the algorithm finds the planet maximum (see Appendix A for some examples). In Table 4, we give, for each of the 15 simulations in which no planet was injected, the highest S/N value of the 100 best solutions found during the grid search as well as the minimum and maximum S/N values found by the gradient-descent algorithm. This table shows that if a simple S/N threshold were to be used to flag planet candidates, the number of false positives would be much higher. This result emphasizes the possibility of using a “recombined spot shape” criterion to disentangle false positives from true planets, and the importance of the observer’s judgement for the reliability of this technique in its current implementation.

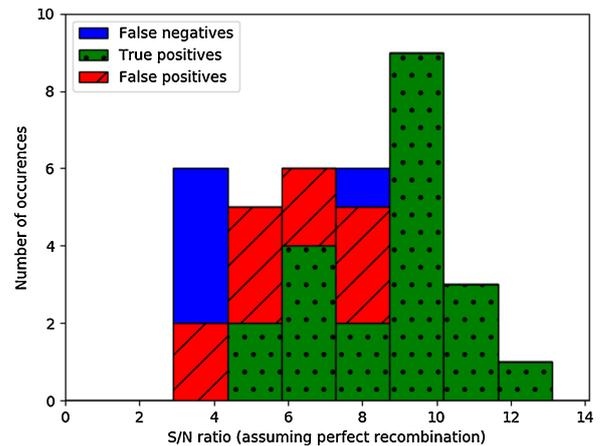
### 5.3. Can we get rid of the astronomer?

In its current implementation, the K-Stacker algorithm produces 100 images for each time series analyzed, and requires the intervention of a skillful observer to check these recombined images for the presence of a planet. This step, while also used in classical high-contrast imaging, is tedious, and may introduce uncontrolled bias in the data-reduction process. For the time being, however, and as already discussed in Sect. 5.2, this step is absolutely necessary to avoid large numbers of false positives.

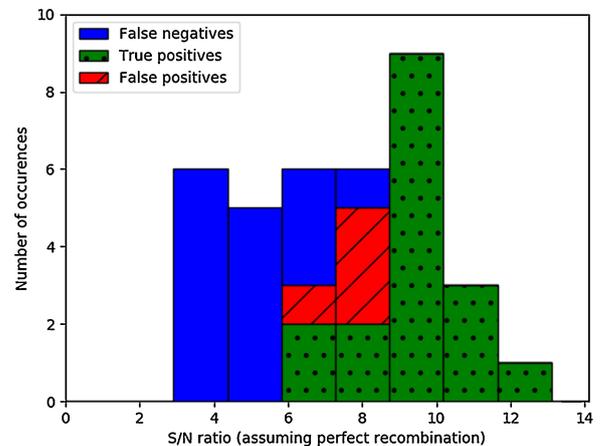
In Fig. 5 (resp. Fig. 6), we show what the results of our blind test would have been if we had used only a S/N threshold to flag planet candidates. For each simulation, the best solution found

**Table 4.** Values of the 100th best S/N found by the grid search, and minimum/maximum values found by the reoptimization algorithm for each of the 15 simulations where no planet was injected.

Simulation	Grid search	Gradient descent	
	S/N 100th value	Minimum S/N	Maximum S/N
1	4.27	4.62	5.64
6	4.45	4.60	4.95
8	4.58	4.93	5.74
11	4.76	4.82	5.78
13	4.50	4.61	5.26
15	5.50	5.60	5.96
21	4.80	4.86	5.59
26	4.57	4.59	5.24
27	4.55	4.67	5.67
29	4.45	4.60	5.80
37	4.82	4.89	5.68
38	4.75	4.82	5.69
41	4.86	4.92	5.55
45	4.57	4.66	6.55
49	4.46	4.52	5.70



**Fig. 5.** Resulting distribution of the planet candidates found and missed among the 33 simulations where a planet was effectively injected, when using only a S/N threshold  $>5$  to flag planet candidates.



**Fig. 6.** As in Fig. 5, but for a S/N threshold of 7.

by K-Stacker is simply flagged as a planet candidate if the S/N is above 5 (resp. 7), without any intervention of an observer. As expected, the number of false positives (10, resp. 5 in total) is much higher, especially at low S/N. In our blind test, the observer was able to detect that these “best solutions” found by K-Stacker were actually “super-speckles”, and he eliminated them (see also Fig. A.1). For most of them, he was also able to find the planet candidates.

As discussed in Sect. 5.5, the shape of the recombined spot is important. To try to take that into account, we switched from our simple circular photometric aperture to a slightly more sophisticated Gaussian-weighted aperture. We found no real difference with our simple circular aperture, and this new method did not alleviate the need for the observer to carefully look at each image to assess the presence of the planet. A more complex algorithm, maybe one based on machine-learning trained to disentangle a planet PSF from recombined speckles, will be required in order to get rid of the observer’s judgement. This would be a major improvement for K-Stacker, but is left for future work. Further astrophysical information (spectral, polarimetry) will also be used to remove the false positives.

#### 5.4. Orbital parameters determination

K-Stacker can also provide an estimate of the orbital parameters, as a by-product of the optimization algorithm. The precision with which these parameters are estimated depends on many different factors, such as the actual orbit on which the planet is moving, the total time spanned by the observing sequence, and so on. It is clear, for example, that if the planet does not move significantly during the sequence, one should not expect to have reliable information about the orbital parameters. In such cases, K-Stacker is able to recenter the images and detect planets not reachable with other methods, but several very different sets of orbital parameters can lead to a good recombination.

In Fig. 7, we give, for each of the 25 simulations in which the planet was found, the mean distance between the true position of the planet in each image and the one predicted according to the set of parameters optimized by the K-Stacker algorithm, as a function of the total path travelled, and of the S/N. The two points with a mean error of more than 3 pixels correspond to two simulations in which the planet is very close to the edge of the AO-corrected area in certain images of the sequence, reducing the ability of the algorithm to properly constrain the orbit.

A linear fit on the other points reveals that the mean distance error is slightly increasing with the total path length travelled by the planet ( $3 \times 10^{-3}$  pixel per pixel travelled by the planet).

When the planet travels along a large portion of its orbit in the sequence, we expect the orbit to be better constrained, and the algorithm to find a better fit to it. But at the same time, a fixed error on the orbital elements has a much greater impact on the estimated positions if the planet moves along a large portion of its orbit. The result of Fig. 7 shows that we are still dominated by the second effect, suggesting that the reoptimization stage could still be improved.

In Fig. 8, we show the error on each of the six parameters individually, as a function of the total path length travelled by the planet. On some parameters (e.g.,  $a$ ,  $i$ ,  $t_0$ ), there seems to be a clear advantage of having longer displacements of the planet to get better estimates. On some others (e.g.,  $e$ ,  $\Omega$ ,  $\theta_0$ ), it is not as clear. These parameters, which are not better fitted for a longer travelled path, may be responsible for the overall increasing slope of Fig. 7.

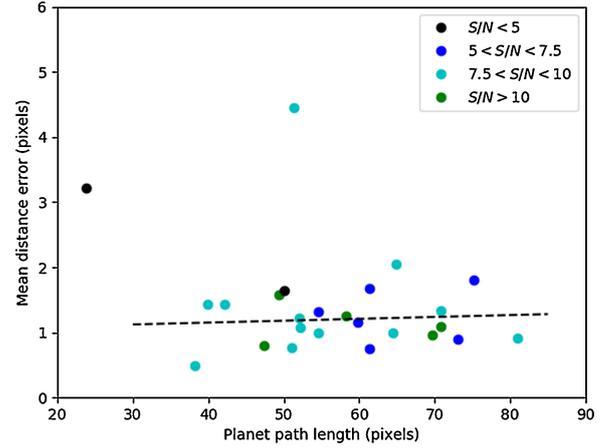


Fig. 7. Mean distance between the true position of the planet and the position found by K-Stacker as a function of the total path length travelled by the planet, for the 25 simulations in which the planet was detected.

#### 5.5. Possible improvements of the algorithm

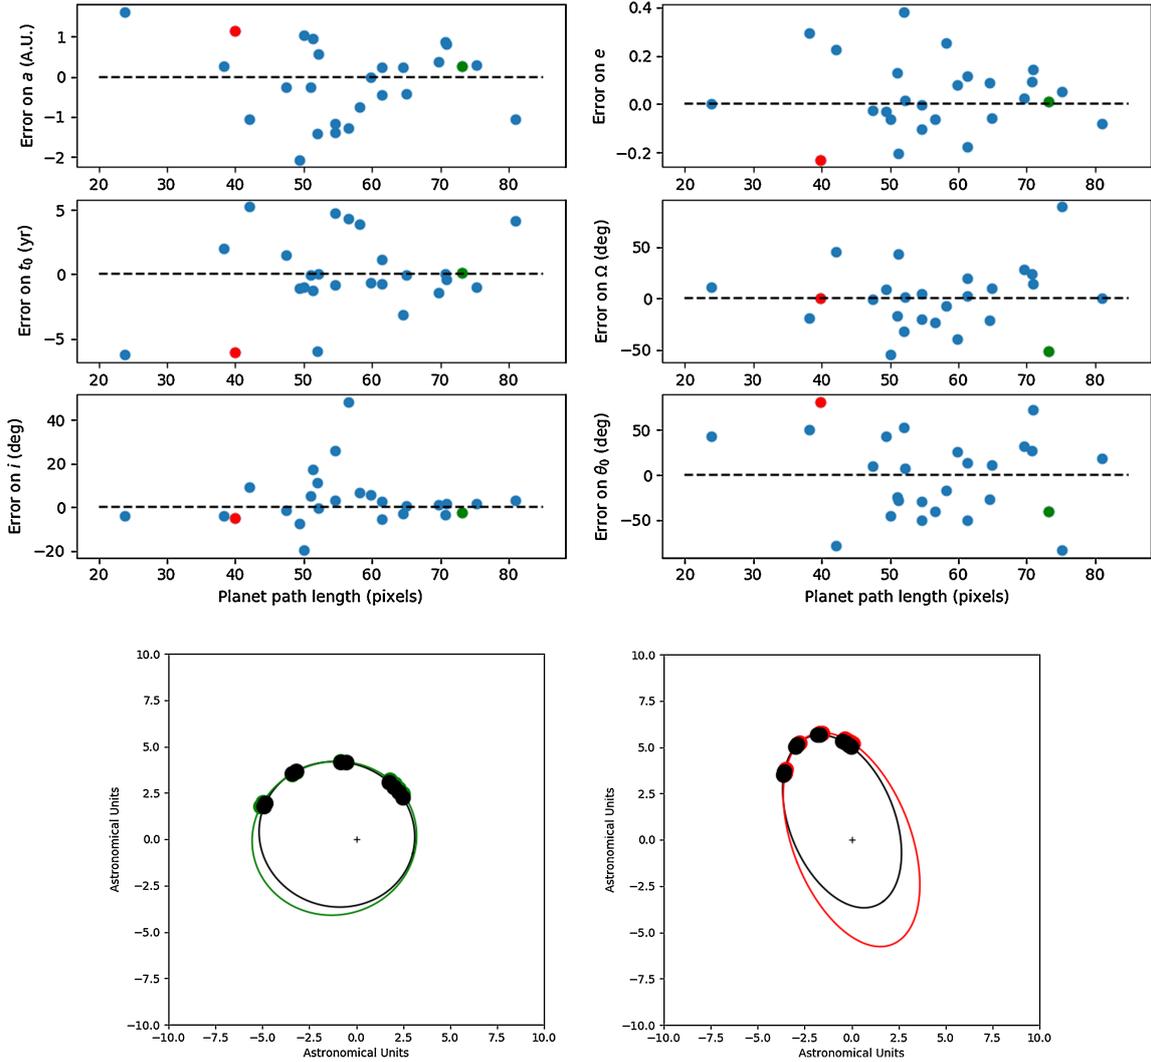
We find that the current version of the K-Stacker algorithm can achieve a 50% detection rate on targets with a total  $S/N \approx 5$ . Whereas this result already shows that it is possible to use K-Stacker to find very faint planets (S/N below 2 in individual frames for a series of ten observations), we strongly believe that there is still room for improvement in the algorithm.

One sure way to improve the results of the K-Stacker algorithm would be to optimize the search grid. For example, a larger step in  $\theta_0$  could be used for smaller values of  $a$ . In our version of the algorithm, we are exploring the same range of values for  $t_0$  no matter what the actual value of  $a$  is. But it is well known that the semi major-axis and the orbital period are directly related, and this could be used to reduce the range explored for smaller values of  $a$ .

In fact, an optimal grid may be constructed for the K-Stacker algorithm by looking for a set of orbits which diverge by at least one FWHM of the instrument PSF in at least one image of the sequence. This could be done numerically, or maybe even analytically. This would ensure that the minimum possible number of orbits are used, and hence would help in reducing the total computing time, and the rate of false positives.

Also, one area which has yet to be studied in detail is the definition of the function to be optimized usually referred to as the “gain function” in optimization problems. As pointed out in Sect. 5.2, when the algorithm is trying to recombine speckles to create a “super-speckle”, the resulting pattern is “blurred” and does not resemble a real PSF. The most efficient gain function might be one that takes into account the S/N but also the shape of the recombined spot. When searching for planets close to the star, small sample statistics should be taken into account to compute the S/N (Mawet et al. 2014).

Along this line of thought, we also tried to use a noise-weighted averaging of the different images (close to what is suggested in Bottom et al., 2017), to take into account the fact that in our series, because of the motion of the planet with respect to the central star, some exposures might be better than others (e.g., when the planet is far away from the star). This modification did not drastically change our results. The same blind test lead to 25 planets recovered out of 34 (instead of the 24 in the version presented here), with no false negative (instead of one



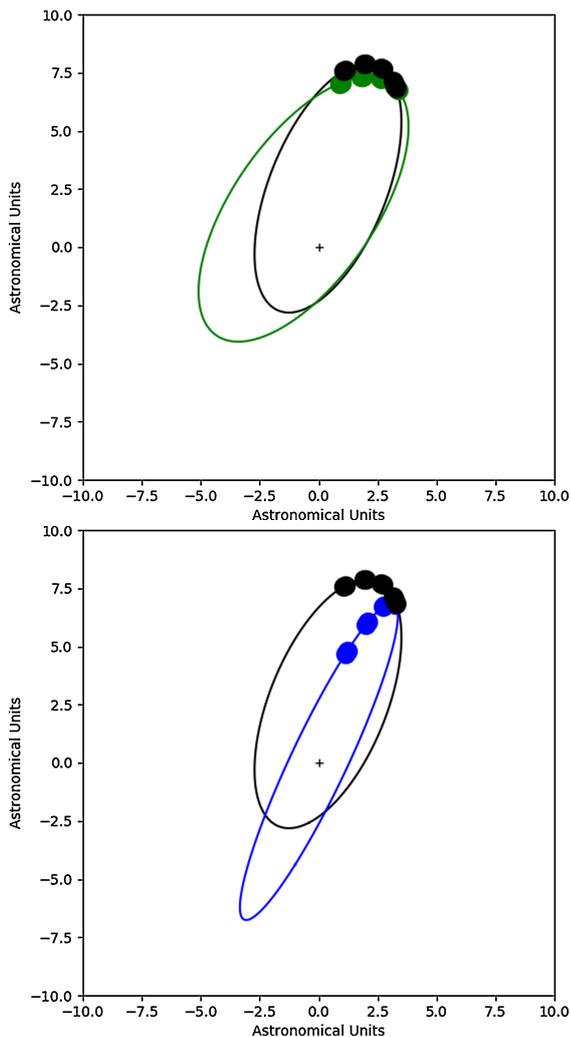
**Fig. 8.** *Upper six panels:* errors made on the orbital parameters as a function of the total path travelled by the planet, for each of the simulations in which a planet was correctly identified. In these panels, two cases, corresponding to a good fit (green dots) and a bad fit (red dots) of the orbit are highlighted. *Two lower panels:* corresponding orbits with the real orbit of the planet in black, and the orbit found by K-Stacker in green (good fit case), and red (bad fit case). For each orbit, the position of the planet in all ten images is also represented.

found here). A more sophisticated algorithm will be necessary to drastically improve the results.

We also noted an unexpected behavior of the algorithm when using the noise-weighted average. This behavior is exemplified in one particular simulation, in which the planet was found by the algorithm, but the orbit was not properly recovered (see Fig. 9). An in-depth investigation revealed that this was due to some sort of interaction between the way the planets are injected in our images, and the noise-weighted averaging. We recall that in our images, the planet is injected at constant flux, rather than constant S/N, to reflect the reality of a constant exposure time. This means that from one image to another, the actual S/N may vary. In this particular case, the planet was injected at a higher S/N in the first few images of the sequence than in the others ( $S/N \approx 1.9$  vs. 1.2). Because of the noise structure in high-contrast images, the noise-averaged combination also has a natural tendency to favor images in which the computed position is far from the central star. The combination of these two factors made for a situation in which a highly eccentric orbit correctly fitting the first part of the real orbit and moving rapidly towards the central star in the last part (as in Fig. 9, bottom panel) yields

a good total S/N. The classical combination algorithm, which gives the same importance to each image of the sequence, is less subject to this effect.

Finally, it must be noted that the K-Stacking method does not necessarily imply the use of a brute-force optimization algorithm. Any type of optimization algorithm could be used to optimize the S/N function: simulated annealing, genetic algorithms, amoeba, and so on. However, we believe that the brute-force method is one of the most appropriate for K-Stacker. Despite not being known for its efficiency, this method has a very interesting property: the sets of orbital parameters on which the S/N has to be calculated are known in advance and never change (these are the points of the search grid). This means that adding a new image to a set of  $n$  observations already processed by K-Stacker does not take much more computation time if the “signal” and “noise” terms computed in each image and for each point of the search grid are systematically stored. In this case, these terms can be combined to generate the S/N values using the equation  $S/N = \sum_k S_k / \sqrt{\sum_k \sigma_k^2}$ . Then to add a  $n + 1$ th image to a set of  $n$  images, one can simply compute the values of  $S_{n+1}$  and  $\sigma_{n+1}$ , and



**Fig. 9.** Illustration of an unexpected behavior of the K-Stacker algorithm when using a noise-weighted averaged combination of images. *Top panel:* best orbit found by the classical combination algorithm, and *bottom panel:* best orbit found when using the noise-weighted averaged combination for the same simulation. In this case, the planet has a  $S/N \approx 1.9$  in the first four images, and  $\approx 1.2$  in the last six. The real orbit of the planet is in black.

recompute the  $S/N$ . The new image can be processed separately and it is not necessary to reprocess the entire set of  $n + 1$  images. Inversely, when using Monte-Carlo methods, in which the optimization path is generated dynamically by the algorithm based on previous values found, adding a new image to a set of already processed observations requires a complete re-processing of the whole set, which can prove extremely time-consuming.

On the other hand, a Monte-Carlo-type algorithm would more thoroughly explore the local search space around the best orbits, helping to determine error bars on the orbital parameters. The best way to proceed might be a combination of both a grid search to reduce the search space, and an MCMC (Monte-Carlo Markov Chain)-based algorithm.

## 6. Conclusion: what can we use K-Stacker for?

In this study, we have shown how the K-Stacking method could be used to combine multiple high-contrast images obtained during different nights to detect exoplanets. We

simulated SPHERE/IRDIS coronagraphic images taken over several months with planets at low  $S/N$  (below the detection limit) in individual frames, and used the K-Stacker algorithm to combine these images and detect the planets. We find that when the total number  $n$  of available images is large enough to get  $\sqrt{n} \times S/N_{\text{ind}} \geq 7$  (where  $S/N_{\text{ind}}$  is the  $S/N$  level in individual frames), the K-Stacker algorithm is able to detect the planet with a very high level of reliability ( $>90\%$ ). For cases where  $5 \leq \sqrt{n} \times S/N_{\text{ind}} \leq 6$ , we find a recovery rate of about 50%.

The detection rate can probably be improved by refining and optimizing the search grid. This will be done in a future work, using real data. This first study was limited to simulated images in order to have complete control over all the important parameters (star magnitude, turbulence conditions, XAO performances, etc.). It has shown that K-Stacker could achieve good recovery rates (i.e., similar to what usual high-contrast imaging techniques are expected to provide) on simulated data, in an acceptable amount of time. Future studies on this subject should be done using real data, possibly with false planets injected.

In its current implementation, K-Stacker still heavily relies on the ability of the observer to identify false positives in the recombined images. Whereas our blind test has shown that this could be done reliably (only one false positive, which was identified as an uncertain planet candidate by the observer, among our 49 simulations), it is clear that this has to be tested using real data.

In its current state, though, the K-Stacker technique could already prove useful to detect very faint planets. The ADI technique intrinsically limits the useful observing time that can be spent each night on any single star (up to about 1–2 h). Adding to this that high-contrast imaging instruments require the best observing conditions, that these instruments are sharing telescope time with others, and that stars are not visible year-round in the night sky, accumulating about 10 h of ADI observation on a single target in a few months (to limit orbital motion) can prove excessively complicated. Using the K-Stacker would drastically improve the chances of doing so, and reduce the complexity of observing schedules.

Finally, we also want to point out that this method could also be used as part of a new scheme of observation, in which exposures would not be made sequentially on one night, but would be spread over multiple nights, to better constrain the orbit. For the same contrast in the final image, the K-Stacker technique would also yield an estimate of the orbital parameters, which cannot be obtained with a single night's observation.

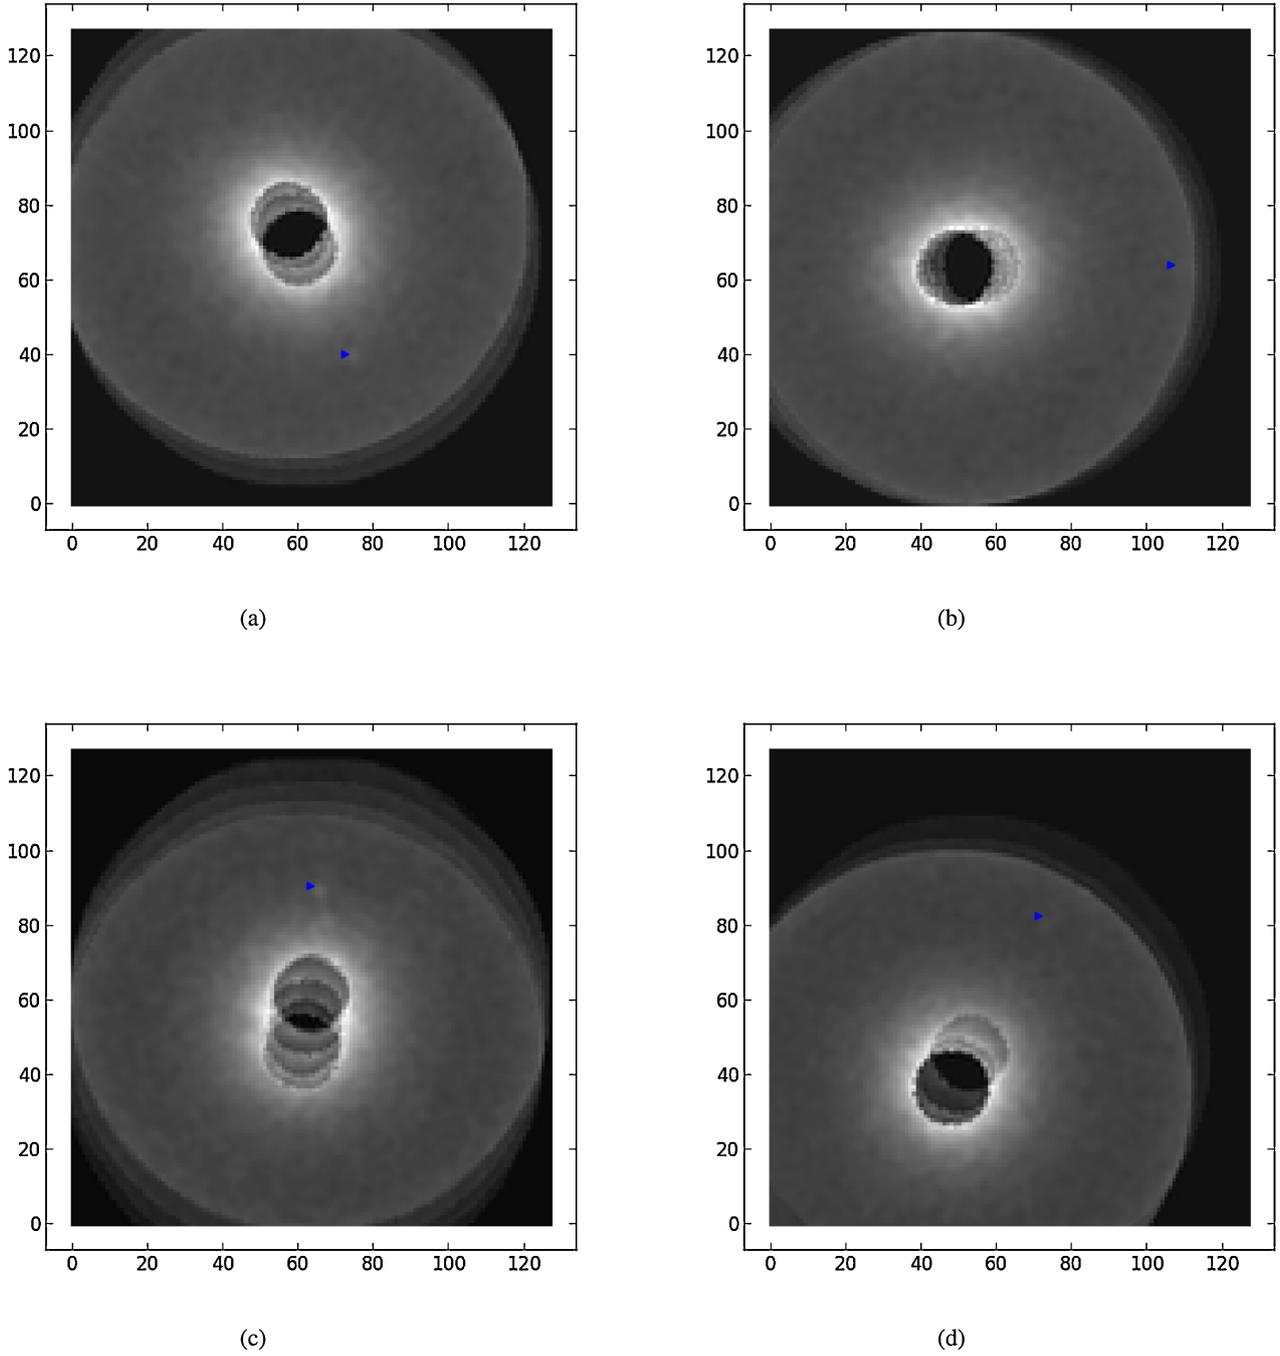
*Acknowledgements.* We thank the anonymous referee for his/her careful reading of the manuscript, and his/her helpful suggestions and comments. This article was significantly improved after his/her review. This research has been funded by the PNP/ASHRA/INSU. This paper has used the cluster facilities of LAM operated by the CeSAM data center.

## References

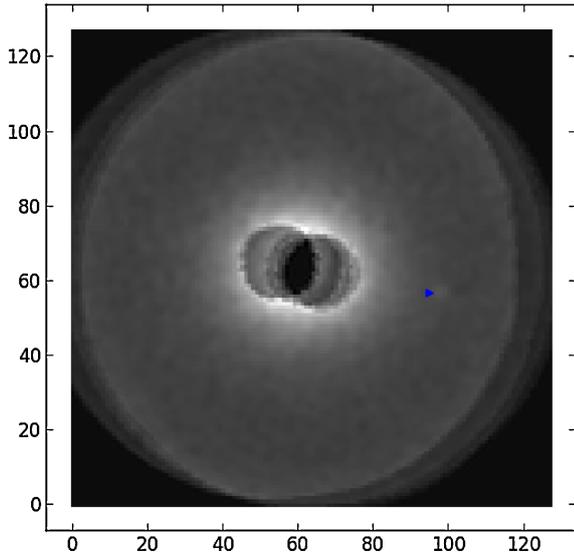
- Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2008, *SPIE Conf. Ser.*, 7014, 18
- Bottom, M., Ruane, G., & Mawet, D. 2017, *Res. Notes Am. Astron. Soc.*, 1, 30
- Canovas, H., Rodenhuis, M., Jeffers, S. V., Min, M., & Keller, C. U. 2011, *A&A*, 531, A102
- Dohlen, K., Langlois, M., Saisse, M., et al. 2008, *SPIE Conf. Ser.*, 7014, 3
- Fletcher, R. 1987, *Practical Methods of Optimization*, 2nd edn. (Chichester, NY: Wiley-Interscience Publication)
- Fusco, T., Rousset, G., Sauvage, J.-F., et al. 2006, *Opt. Express*, 14, 7515
- Hinkley, S., Oppenheimer, B. R., Soummer, R., et al. 2007, *ApJ*, 654, 633
- Le Coroller, H., Nowak, M., Arnold, L., et al. 2015, in *Proceedings of Colloquium "Twenty Years of Giant Exoplanets"* held at Observatoire de Haute

- Provence, France, October 5–9, 2015, eds. I. Boisse, O. Demangeon, F. Bouchy, & L. Arnold, [59](#)
- Macintosh, B., Poyneer, L., Sivaramakrishnan, A., & Marois, C. 2005, in *Astronomical Adaptive Optics Systems and Applications II*, eds. R. K. Tyson & M. Lloyd-Hart, *Proc. SPIE*, [5903](#), [170](#)
- Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *Proc. Natl. Acad. Sci. USA*, [111](#), [12661](#)
- Males, J. R., Skemer, A. J., & Close, L. M. 2013, *ApJ*, [771](#), [10](#)
- Males, J. R., Belikov, R., & Bendek, E. 2015, in *Techniques and Instrumentation for Detection of Exoplanets VII*, *Proc. SPIE*, [9605](#), [960518](#)
- Marley, M. S., Fortney, J. J., Hubickyj, O., Bodenheimer, P., & Lissauer, J. J. 2007, *ApJ*, [655](#), [541](#)
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, [641](#), [556](#)
- Martinez, P., Loose, C., Aller Carpentier, E., & Kasper, M. 2012, *A&A*, [541](#), [A136](#)
- Martinez, P., Kasper, M., Costille, A., et al. 2013, *A&A*, [554](#), [A41](#)
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, [792](#), [97](#)
- Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, *PASP*, [111](#), [587](#)
- Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007, *ApJ*, [669](#), [642](#)
- Sparks, W. B. & Ford, H. C. 2002, *ApJ*, [578](#), [543](#)
- Thatte, N., Abuter, R., Tecza, M., et al. 2007, *MNRAS*, [378](#), [1229](#)
- Vigan, A., Moutou, C., Langlois, M., et al. 2010, *MNRAS*, [407](#), [71](#)
- Vigan, A., Gry, C., Salter, G., et al. 2015, *MNRAS*, [454](#), [129](#)
- Zurlo, A., Vigan, A., Mesa, D., et al. 2014, *A&A*, [572](#), [A85](#)
- Zurlo, A., Vigan, A., Galicher, R., et al. 2016, *A&A*, [587](#), [A57](#)

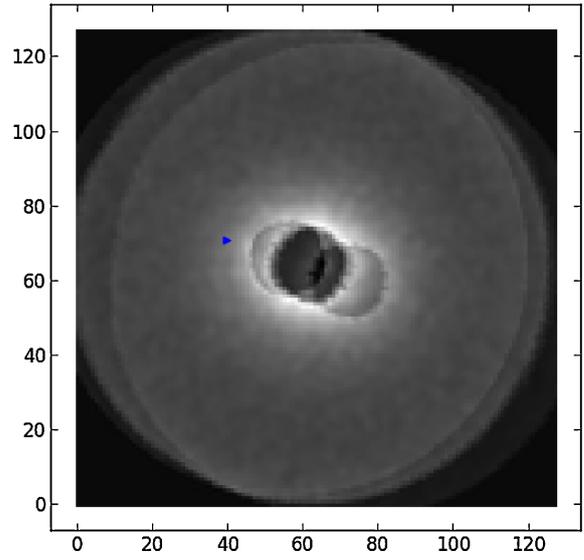
**Appendix A: Example of solutions found by K-Stacker**



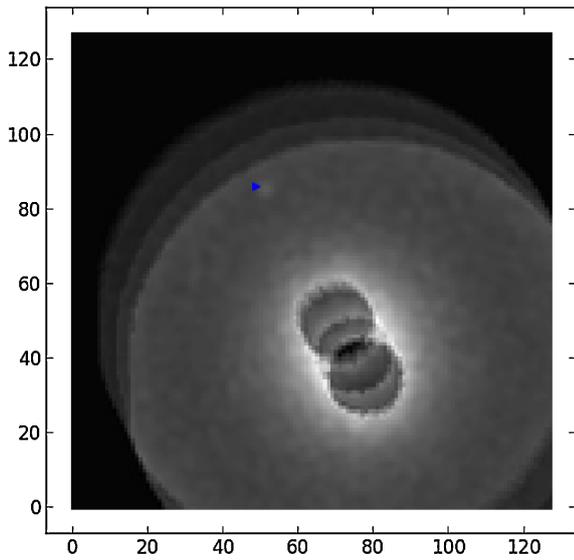
**Fig. A.1.** A set of best recombined images as found by K-Stacker, for different reoptimized  $S/N$  values, and for cases where the algorithm found a true/false planet candidate. In each panel, the blue arrow shows the position of the planet candidate found by K-Stacker. *Panel a* (resp. *panel b*):  $S/N = 5.7$  with a correct (resp. false) planet candidate. *Panel c* (resp. *panel d*):  $S/N = 6.0$  with a correct (resp. false) planet candidate. *Panel e* (resp. *panel f*):  $S/N = 6.5$  with a correct (resp. false) planet candidate. *Panel g*:  $S/N = 8.3$ , with a true planet candidate. We note that the observer has been able to correctly identify a false positive ( $S/N > 6$ ) without planets injected in simulations: *panel b*, *d*, and *e* using the shape of the spot, without any a priori information on the planet injection.



(e)



(f)



(g)

Fig. A.1. continued.