



HAL
open science

Exoplanet detection in angular differential imaging: combining a statistics-based learning with a deep-based learning for improved detections

Olivier Flasseur, Théo Bodrito, Julien Mairal, Jean Ponce, Maud Langlois,
Anne-Marie Lagrange

► To cite this version:

Olivier Flasseur, Théo Bodrito, Julien Mairal, Jean Ponce, Maud Langlois, et al.. Exoplanet detection in angular differential imaging: combining a statistics-based learning with a deep-based learning for improved detections. SPIE Astronomical Telescopes + Instrumentation, Jul 2022, Montréal, Canada. pp.139, 10.1117/12.2629849 . hal-03988124

HAL Id: hal-03988124

<https://hal.science/hal-03988124>

Submitted on 14 Feb 2023

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.

Exoplanet detection in angular differential imaging: combining a statistics-based learning with a deep-based learning for improved detections

Olivier Flasseur^{a,b}, Théo Bodrito^b, Julien Mairal^c, Jean Ponce^{b,d}, Maud Langlois^e, and
Anne-Marie Lagrange^{a,f}

^aLaboratoire d'Études Spatiales et d'Instrumentation en Astrophysique, Observatoire de Paris,
Univ. PSL, Sorbonne Univ., Univ. Paris Diderot, France

^bInstitut National de Recherche en Informatique et en Automatique, DI/ENS (ENS-PSL,
CNRS, Inria), France

^cUniv. Grenoble Alpes, Institut National de Recherche en Informatique et en Automatique,
Grenoble INP, LJK, CNRS, France

^dCenter for Data Science, New York Univ., USA

^eCentre de Recherche Astrophysique de Lyon, CNRS, Univ. de Lyon, Univ. Claude Bernard
Lyon 1, ENS de Lyon, France

^fUniv. Grenoble Alpes, Institut de Planétologie et d'Astrophysique de Grenoble, France

ABSTRACT

Direct imaging is an active research topic in astronomy for the detection and the characterization of young sub-stellar objects. The very high contrast between the host star and its companions makes detection particularly challenging. In addition to the use of an extreme adaptive optics system and a coronagraph to strongly attenuate the starlight contamination, dedicated post-processing methods combining several images recorded with the pupil tracking mode of the telescope are needed.

In previous works, we have presented the PACO¹⁻³ algorithm capturing the spatial correlations of the data with a multi-variate Gaussian model whose parameters are estimated in a data-driven fashion at the scale of a patch of a few tens of pixels. PACO is parameter free and delivers reliable detection confidences with an improved sensitivity compared to the standard methods of the field (e.g., cADI,⁴ PCA,^{5,6} TLOCI⁷). However, there is a room for improvement in the detection sensitivity due to the approximate fidelity of the PACO statistical model with respect to the observations.

We propose to combine the statistics-based model of PACO with a deep learning approach in a three-step algorithm. First, the data are centered and whitened locally using the PACO framework to improve the stationarity and the contrast in a preprocessing step. Second, a convolutional neural network is trained in a supervised fashion to detect the signature of synthetic sources in the preprocessed science data. The network is trained from scratch with a custom data augmentation strategy allowing to generate a large training set from a single spatio-temporal dataset. Finally, the trained network is applied to the preprocessed observations and delivers a detection map.

We apply our method on eleven datasets from the VLT/SPHERE-IRDIS instrument and compare our method with PACO and other baselines of the field (cADI, PCA). Our results show that the proposed method performs on-par with or better than these algorithms, with a contrast improvement up to half a magnitude with respect to PACO.

Keywords: High-contrast imaging, differential imaging, exoplanet detection, data reduction, processing algorithm, statistical method, deep learning

Further author information: olivier.flasseur@obspm.fr

1. INTRODUCTION

High-contrast imaging is an observational method used to study the close environment of stars.^{8–10} It is particularly adapted to detect young, massive and hot exoplanets.^{11,12} However, despite promising results from direct imaging, only a few dozen exoplanets have been unveiled with this technique since its emergence in the early 2000s.^{13–18} This is mainly due to the very high contrast between the host star and the exoplanets, typically higher than 10^5 in the infrared, making the detection very challenging.

In this context, cutting-edge ground-based facilities are equipped with an (extreme) adaptive optics system and a coronagraph to attenuate as best as possible the starlight. Currently, the non-blocked residual starlight contamination and its temporal evolution remain the main limitations for the detection of point-like sources. It takes the form of spatially-correlated *speckles* that mimic the expected pattern of an exoplanetary signal. The observations are also impacted by additional sources of noise which form with speckles a spatially and non-stationary *nuisance component* corrupting the signal of the off-axis objects of interest.

In order to help distinguish the signal of the objects of interest from that of the nuisance component, high-contrast observations are performed with dedicated strategies. In this paper, we focus on angular differential imaging⁴ (ADI), which consists in tracking the observed target over time, with the telescope derotator tuned to keep the telescope pupil stable while the field of view rotates. Consequently, in the resulting 3-D datasets (2-D + time), the objects of interest follow an apparent motion along a deterministic circular trajectory centered on the star/coronagraph location while the telescope pupil remains static. In that mode, speckles resulting from residual starlight aberrations are *quasi-static*, i.e. they are strongly correlated across exposures. The images can be combined to cancel out most of the speckles while part of the signal from the off-axis objects is preserved.

The final keystone of high-contrast imaging lies in the performance of the post-processing algorithm used to *reduce* the data by a combination of the individual ADI images. The common idea is to estimate a reference image (so-called *on-axis PSF*) of the nuisance component, so that it can be subtracted from the data to reveal the presence of the objects of interest. To do so, a simple solution consists in subtracting the temporal mean or median of the dataset from each frame of the ADI stack. The residual images are then aligned to the true-North so that the signals of the objects of interest are superimposed and can be combined by temporal stacking. This is the principle of the cADI method designed to process the first ADI direct observations.^{4,19} Since then, several more sophisticated strategies have been proposed, see, e.g., Pueyo *et al.*¹⁰ for a review. In particular, the KLIP/PCA^{5,6} algorithm is currently implemented in most of the reduction pipelines^{6,20,21} and it is considered as a standard to process high-contrast observations. The KLIP/PCA algorithm performs a principal component analysis of the data, and a low-rank estimate of the on-axis PSF is formed by keeping the first principal components of the decomposition. Some other methods are based on an inverse problem framework. Among them, the PACO algorithm^{1–3} builds a more consistent statistical model, self-calibrated on the data, that accounts for the spatial correlations of speckles at the scale of small image patches of a few tens of pixels. Given the success of data-driven approaches in solving various high-level imaging tasks (e.g., object segmentation), machine learning and deep learning approaches have also been investigated by the direct imaging community.^{22–24} In particular, Gonzalez *et al.*²² formalize the detection problem as a binary classification task and they develop a fully supervised deep learning approach by means of massive injections of synthetic sources. The approach considers collections of patches preprocessed by KLIP/PCA for different numbers of principal components as input of a random forest or of a convolutional neural network that decides in favor of the presence or on the absence of a point-like sources in each patch. While demonstrating powerful detection capabilities, this algorithm shown to be prone to a high level of false alarms in some cases.²⁵ Besides, the tuning of hyper-parameters remains a critical point making the operating point difficult to reach.

Intensive testing of PACO, both on public²⁵ and on private data challenges, shows that PACO is one of the algorithms of choice to process high-contrast observations. Thanks to its unique data-driven modeling of the data accounting for non-stationary spatial correlations, PACO is especially well suited to process observations in which the typical spatial extent of speckles lies in a patch of a few tens of pixels, as it is the case for SPHERE and GPI observations. The tests also illustrate that the statistical model embedded in PACO is only approximate in case of spatial correlations spread over a patch of a few tens of pixels (e.g., for background-limited observations and/or in case of unstable observing conditions). This is the motivation for this work: we propose to combine

the statistical model of PACO with a supervised deep learning framework. The statistical model of PACO is used to improve the stationarity and the contrast of the data in a preprocessing step and deep learning is in charge of correcting for the (putative) approximate fidelity of the statistical model of PACO to the reality of the observations.

Section 2 presents the main ingredients of the proposed algorithm*. Section 3 evaluates its performance on several high-contrast observations from the VLT/SPHERE-IRDIS instrument.^{26–28} Finally, Sect. 4 presents our conclusions and gives future research prospects.

2. ALGORITHM DESCRIPTION

2.1 Preprocessing by statistical learning of the non-stationary patch covariances

2.1.1 Statistical model of the nuisance component

An ADI dataset $\mathbf{r} \in \mathbb{R}^{N \times T}$ is formed by N -pixels images recorded at different times $t \in \llbracket 1; T \rrbracket$. The direct model for the observed intensity is:

$$\mathbf{r} = \mathbf{f} + \sum_{p=1}^P \alpha_p \mathbf{h}(\phi_p), \quad (1)$$

where $\mathbf{f} \in \mathbb{R}^{N \times T}$ is the nuisance component, and $\mathbf{h}(\phi_p) \in \mathbb{R}^{N \times T}$ stands for the contribution of a point-like source $p \in \llbracket 1; P \rrbracket$ with a contrast α_p that is assumed constant during the few hours of the total observations. The contribution of a source p takes the form of the off-axis PSF centered at location $\mathcal{F}_t(\phi_p)$ in the t -th image where ϕ_p is its initial location on an image at a reference time t_{ref} (e.g., $t_{\text{ref}} = t_1$) and \mathcal{F}_t is a geometrical transform (typically in ADI, a circular translation with respect to the star located at the center of the images) modeling the apparent motion of the field of view between time t_{ref} and time t . The function \mathcal{F}_t is completely deterministic since it depends solely on the measured parallactic angles. Given that very few sources are expected in the field of view, we assume that the measured intensity is the superimposition of the nuisance component and at most one unresolved point-like source p at each pixel location n , i.e., multiple sources do not overlap.

In previous works on the PACO algorithm,^{1–3} we have proposed to describe the random fluctuations of the nuisance component \mathbf{f} by a statistical model whose parameters are learnt in a data-driven fashion. We recall hereafter the main ingredients of this statistical model.

Given the high (spatial) non-stationarity of the nuisance component, the model is built locally at a scale of a patch of a few tens of pixels. It models the distribution of T patches $\mathbf{f}_n = \{\mathbf{E}_{n,t} \mathbf{f}\}_{t=1:T} \in \mathbb{R}^{K \times T}$ extracted around pixel n ($\mathbf{E}_{n,t}$ denotes the K -pixel patch extraction operator at location n and time t) with a multi-variate Gaussian $\mathcal{N}(\mathbf{m}_n, \mathbf{C}_n)$. The covariance matrix \mathbf{C}_n is non-diagonal, i.e., it accounts for the local correlations of \mathbf{f} . The sample estimators $\{\widehat{\mathbf{m}}_n; \widehat{\mathbf{S}}_n\}$ of the local mean and covariances coming from the maximum likelihood are the following:

$$\begin{cases} \widehat{\mathbf{m}}_n = \frac{1}{T} \sum_{t=1}^T \mathbf{E}_{n,t} \mathbf{r} \in \mathbb{R}^K, \\ \widehat{\mathbf{S}}_n = \frac{1}{T} \sum_{t=1}^T (\mathbf{E}_{n,t} \mathbf{r} - \widehat{\mathbf{m}}_n)(\mathbf{E}_{n,t} \mathbf{r} - \widehat{\mathbf{m}}_n)^\top \in \mathbb{R}^{K \times K}. \end{cases} \quad (2)$$

Since the number T of temporal frames is typically lower than the number K of pixels in a patch, the sample covariance $\widehat{\mathbf{S}}_n$ is very noisy and rank deficient. A form of regularization must be enforced to stabilize the estimate and allow the inversion of the covariance matrix involved in the data whitening step (see Sect. 2.1.2). We use a *shrinkage* estimator^{29,30} formed by a convex combination between the low bias/high variance estimator $\widehat{\mathbf{S}}_n$ and a high bias/low variance estimator $\widehat{\mathbf{F}}_n$:

$$\widehat{\mathbf{C}}_n = (1 - \widehat{\rho}_n) \widehat{\mathbf{S}}_n + \widehat{\rho}_n \widehat{\mathbf{F}}_n, \quad (3)$$

*The proposed method will be described in more details in a journal paper currently in preparation.

where $\widehat{\mathbf{F}}_n$ is a diagonal matrix encoding the sample variances:

$$[\widehat{\mathbf{F}}_n]_{kk'} = \begin{cases} [\widehat{\mathbf{S}}_n]_{kk'} & \text{if } k = k' \\ 0 & \text{if } k \neq k'. \end{cases} \quad (4)$$

The hyper-parameter $\widehat{\rho}_n$ plays a key role since it governs a bias-variance trade-off. In our previous works,^{1,31} we have derived its closed-form expression, which is an extension of the results of Chen *et al.*³⁰ in the case of a non-constant valued shrinkage matrix $\widehat{\mathbf{F}}_n$:

$$\widehat{\rho}_n = \frac{\text{tr}(\widehat{\mathbf{S}}_n^2) + \text{tr}^2(\widehat{\mathbf{S}}_n) - 2 \sum_{k=1}^K [\widehat{\mathbf{S}}_n]_{kk}^2}{(T+1) \left(\text{tr}(\widehat{\mathbf{S}}_n^2) - \sum_{i=1}^K [\widehat{\mathbf{S}}_n]_{kk}^2 \right)}. \quad (5)$$

2.1.2 Centering and local whitening of the observations

We consider a set of locations \mathbb{P} where the statistics of the nuisance component are evaluated. The cardinal of \mathbb{P} depends solely on the patch shape and the patch stride used to cover the whole field of view. In this work, we consider non-overlapping square patches of K pixels (i.e., $\text{card}(\mathbb{P}) = \lfloor N/K^2 \rfloor$). The preprocessed images $\tilde{\mathbf{r}} \in \mathbb{R}^{N \times T}$ after centering and whitening are obtained as:

$$\tilde{\mathbf{r}}_n = \mathbf{W}_n \mathbf{r}_n = \widehat{\mathbb{L}}_n^\top (\mathbf{r}_n - \widehat{\mathbf{m}}_n), \forall n \in \mathbb{P}, \quad (6)$$

where \mathbf{W}_n is an operator performing centering and whitening of the collection of patches $\mathbf{r}_n \in \mathbb{R}^{K \times T}$ at location n , such as $\widehat{\mathbb{L}}_n$ is the Cholesky's factorization of $\widehat{\mathbf{C}}_n^{-1}$ (i.e., $\widehat{\mathbb{L}}_n \widehat{\mathbb{L}}_n^\top = \widehat{\mathbf{C}}_n^{-1}$).

2.2 Semantic segmentation by supervised deep learning

We formalize the detection problem as what is sometimes called in the computer vision community a *supervised semantic segmentation* task: starting from a temporal series of preprocessed images including synthetic sources, the goal is to infer a detection map $\widehat{\mathbf{y}} \in [0; 1]^M$, where each pixel value represents a score between 0 and 1 such that a high (resp. low) score values the presence (resp., the absence) of a source centered at that location. Interpreting this score as a true probability of presence of a source requires a control of the uncertainties with dedicated methods that is left for future work. For this reason, in the following, we refer to this score as a *pseudo-probability*.

Section 2.2.1 details the training set construction process, Sect. 2.2.2 describes the selected model architecture, and Sect. 2.2.3 discusses the metrics we consider to evaluate the performance of the proposed method.

2.2.1 Construction of training samples

In high-contrast imaging, obtaining real ground-truth data is a twofold challenge. First, the overall number of positive samples is limited as relatively few point-like sources have been confirmed to date. Second, negative samples are hard to define since some undiscovered sources might be lying in the observed data. To overcome these constraints, we adopted a synthetic training strategy: the training set consists of S pairs $\{\tilde{\mathbf{r}}^{[s]}; \widehat{\mathbf{y}}^{[s]}\}_{s=1:S}$ of samples resulting from the massive injection of synthetic point-like sources. Furthermore, the nuisance component varies drastically from one observation to the other, as it is highly dependent on the observing conditions, the physical properties of the star, and the instrument settings. As a consequence, we followed an observation-dependent approach, and trained a different model on each observation.

This setup implies the design of a custom data-augmentation strategy (i) to prevent overfitting of the model that is trained from a unique temporal series of images, and (ii) to account for our lack of knowledge about real sources –*unknown at training time but that we aim to detect at test time*–. To circumvent these issues, we apply a random permutation of the T images forming the observations \mathbf{r} for each new training sample $s \in \llbracket 1; S \rrbracket$. This operation allows (i) to create artificially different datasets and (ii) to break the temporal consistency of (known

and unknown) real sources. Then, synthetic sources are injected inside the temporally permuted data. At this intermediate stage, each training sample $\bar{\mathbf{r}}^{[s]}$ is obtained by:

$$\bar{\mathbf{r}}^{[s]} := \mathbf{P}^{[s]} \mathbf{r} + \sum_{p=1}^{P^{[s]}} \alpha_p^{[s]} \mathbf{h}(\phi_p^{[s]}), \quad (7)$$

where \mathbf{P} is an operator performing the random temporal permutation and $\mathbf{h}(\phi_p) \in \mathbb{R}^{N \times T}$ represents the spatio-temporal contribution of a synthetic source centered at location ϕ_p on a reference image at time t_{ref} , see Sect. 1[†]. The number of sources P , their contrasts $\{\alpha_p\}_{p=1:P}$ and their initial locations $\{\phi_p\}_{p=1:P}$ are free parameters. In practice, the number P of injected sources in each training sample is drawn uniformly in $\llbracket 1; 10 \rrbracket$. This setting represents a realistic scenario since we expect a few faint point-like sources in the field of view. The initial locations $\{\phi_p\}_{p=1:P}$ of the injected sources are drawn uniformly per angular separation (i.e., the number of sources is in average constant per angular annuli). We train our model on sources which are challenging to detect with other methods: the contrast $\{\alpha_p\}_{p=1:P}$ of the injected sources is drawn uniformly in $\left[3\hat{\sigma}_{\phi_p}^{\text{PACO}}; 12\hat{\sigma}_{\phi_p}^{\text{PACO}} \right]$ where $\hat{\sigma}_{\phi_p}^{\text{PACO}}$ is the 1-sigma contrast reached by PACO at location ϕ_p . This setting covers both sources that are detectable above the standard 5σ detection confidence and sources whose detection confidence remains below the 5σ detection limit reached by PACO.

As preprocessing is expensive and becomes the bottleneck during online data generation, we adopt a local update strategy to reduce its computational cost. Prior to the injection of synthetic sources, the whole dataset is preprocessed, i.e. centered and spatially whitened. We denote $\tilde{\mathbf{r}}$ the precomputed cube. After each injection, the set $\mathbb{S}^{[s]}$ of locations impacted by the signal of the $P^{[s]}$ sources is determined. Outside $\mathbb{S}^{[s]}$, the preprocessed images are obtained from the temporal permutation of $\tilde{\mathbf{r}}$. Inside $\mathbb{S}^{[s]}$, the statistics of the nuisance component are updated given the contamination of the $P^{[s]}$ injected sources, and the preprocessed images are updated with these refined statistics. At this intermediate stage, each training sample $\bar{\mathbf{r}}^{[s]}$ is obtained by:

$$\bar{\mathbf{r}}_n^{[s]} := \begin{cases} \mathbf{W}_n(\bar{\mathbf{r}}_n^{[s]}), & \text{for } n \in \mathbb{S}^{[s]} \cap \mathbb{P}, \\ \tilde{\mathbf{r}}_n, & \text{for } n \in \mathbb{P} - \mathbb{S}^{[s]} \cap \mathbb{P}. \end{cases} \quad (8)$$

Finally, the intermediate images of each training sample are derotated with the opposite of the parallactic angles so that signal of the synthetic sources are spatially co-aligned along the temporal axis:

$$\bar{\mathbf{r}}^{[s]} := \mathbf{D} \bar{\mathbf{r}}^{[s]}, \quad (9)$$

where \mathbf{D} is a derotation operator. This derotation step is mandatory to perform a semantic segmentation with the convolutional neural network we consider (see Sect. 2.2.2) given the limited spatial extent of its receptive field.

The binary ground-truth segmentation map $\mathbf{y}^{[s]}$ is obtained by setting to 1 circular areas of K -pixels centered at the locations $\{\phi_p\}_{p=1:P^{[s]}}$ of the $P^{[s]}$ injected sources. The radius of the circles is set to the full width at half maximum of the off-axis PSF, which corresponds to the expected spatial extent of an exoplanetary signature in the data.

2.2.2 Model and architecture

We chose a U-Net³² with a ResNet18³³ as encoder backbone ($\simeq 11$ millions of free parameters), which is an architecture widely used for image segmentation. Its residual connexions preserve of the input's spatial information along the cascade of convolution and downsampling operations thanks to a direct mapping of the output of each layer of the compression arm into the corresponding layer of the decompression arm. We use the architecture implemented in the SMP package.³⁴ The encoder and decoder parts are formed by four blocks, each one being composed by a series of convolution layers, batch normalization layers, rectified linear unit (ReLU) activations, and max pooling layers. The final layer of the network has a sigmoid activation function to produce a detection map $\hat{\mathbf{y}} \in [0; 1]^M$. The network weights are trained from scratch with the generated data samples.

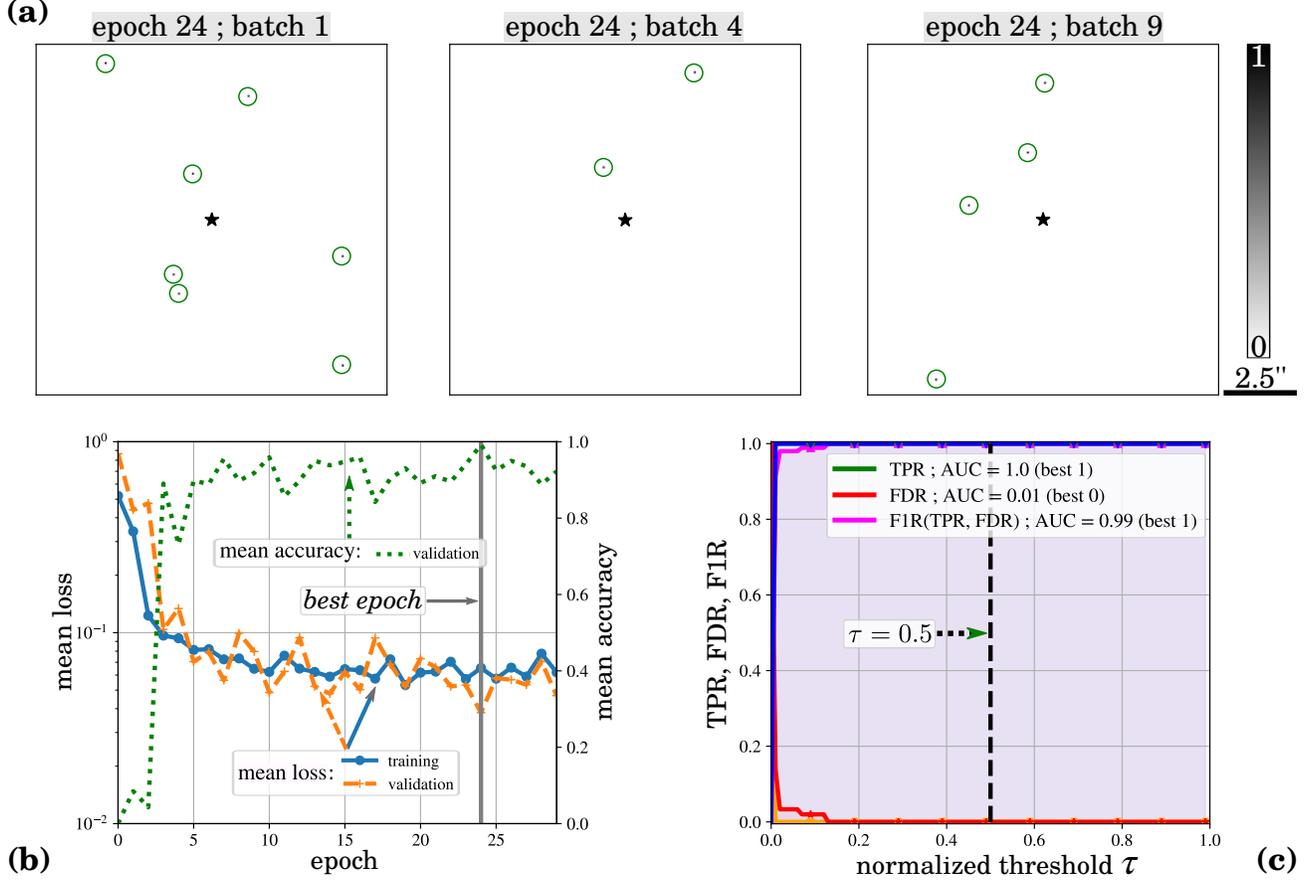


Figure 1. Training and validation results on the HIP 88399 dataset (observing date: 2015-05-10). (a) Examples of detection maps obtained at validation time for the best epoch (number 24). (b) Evolution of the loss function at training and validation time as well as the evolution of the F1R accuracy metric at validation time. (c) ROCs for the best epoch (number 24).

2.2.3 Loss function and accuracy metrics

Our choice for the loss function used for optimizing the network weights at training time is driven by three criteria: (i) handling with the strong class imbalance (the number of background pixels being much larger than the number of pixels from the sources), (ii) being computationally efficient, (iii) matching the astrophysical goals (i.e., having a measure close to a detection accuracy score). We compare losses classically used for semantic segmentation, such as the binary cross-entropy (BCE), ℓ_1 and ℓ_2 norms, mean square error and hinge loss. We have also compared losses based on a similarity measure such as the Dice score³⁵ and hybrid losses combining at least two individual loss measurements (e.g., BCE with Dice score). Our experiments have consistently shown better performances with Dice-based scores. We selected the Dice2 loss (the 2 means for *two* classes), first introduced for biomedical imaging segmentation with very unbalanced classes.^{36,37} Given a set of ground-truth and predicted detection maps $\{\mathbf{y}^{[s]}; \hat{\mathbf{y}}^{[s]}\}$, the Dice2 score is defined by:

$$\mathcal{L}(\mathbf{y}^{[s]}, \hat{\mathbf{y}}^{[s]}) = 1 - \underbrace{\frac{\sum_{m=1}^M \mathbf{y}_m^{[s]} \hat{\mathbf{y}}_m^{[s]} + \epsilon}{\sum_{m=1}^M \mathbf{y}_m^{[s]} + \hat{\mathbf{y}}_m^{[s]} + \epsilon}}_{\text{source error}} - \underbrace{\frac{\sum_{m=1}^M (1 - \mathbf{y}_m^{[s]})(1 - \hat{\mathbf{y}}_m^{[s]} + \epsilon)}{\sum_{m=1}^M 2 - \mathbf{y}_m^{[s]} - \hat{\mathbf{y}}_m^{[s]} + \epsilon}}_{\text{background error}}, \quad (10)$$

[†]When there is no ambiguity, we omit the superscript s on the different quantities.

where ϵ is a minimum-value smoothing and stability parameter added to avoid division by zero. Targeted loss property (i) is satisfied since errors in the source and background areas are penalized equally regardless the relative occurrence of these two classes in $\mathbf{y}^{[s]}$. Property (ii) is also satisfied and we illustrate numerically in the following paragraphs that property (iii) is also reached.

At validation time, we evaluate the capacities of the model to detect point-like sources while avoiding false alarms. In other words, we aim to obtain a model obeying a precision-recall trade-off. For a predicted detection map $\hat{\mathbf{y}}^{[s]}$ in $[0; 1]^M$ thresholded at τ in $[0; 1]$, we count the number of true positives (TP, i.e. true detections), false positives (FP, i.e. false alarms) and false negatives (FN, i.e. missed detections). Following standard practice in direct imaging,^{1,22,25} detections are treated as blobs of one resolution element radius which corresponds to the expected spatial extent of an exoplanetary signature in the data. From TP, FP, and FN, we derive the true positive rate (TPR, i.e., the recall), the false discovery rate (FDR, i.e., the precision), and the F1R score, which is the harmonic mean between TPR and FDR and used as a measure of the precision-recall trade-off:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \in [0; 1] ; \quad \text{FDR} = \frac{\text{FP}}{\text{FP} + \text{TP}} \in [0; 1] ; \quad \text{F1R} = \frac{2}{\frac{1}{\text{TPR}} + \frac{1}{\text{FDR}}} = \frac{2\text{TP}}{2\text{TP} + \text{FN} + \text{FP}} \in [0; 1]. \quad (11)$$

From TPR, FDR, and F1R, receiver operating curves (ROCs³⁸) are built. ROCs are obtained by evaluating the figures of merit defined in Eq. (11) as a function of the detection threshold τ . Finally, the area under the curve (AUC) for the F1R score is computed as an aggregate score of the model performance (best when close to 1).

Figure 1(a) shows some examples of detection maps obtained at validation time for the best validation *epoch*[‡]. These maps illustrate qualitatively the capability of our model to detect synthetic sources while avoiding false alarms. Figure 1(b) shows the evolution of the empirical risk (see Eq. (10)) at training and validation time as well as the evolution of the F1R accuracy metric (see Eq. (11)) at validation time. The loss function does not exhibit significant discrepancy between training and validation steps and the convergence is reached in a few epochs. Besides, the accuracy score is high and well anti-correlated with the loss. This latter observation illustrates that the loss function is a satisfactory estimate of the overall accuracy metric. Finally, Fig. 1(c) gives an illustration of ROCs obtained for the best epoch (gray vertical line in Fig. 1(b)).

3. RESULTS ON ADI SEQUENCES FROM VLT/SPHERE-IRDIS

3.1 Datasets description and reduction strategies

For our comparative analysis, we have selected 11 datasets[§] from the VLT/SPHERE-IRDIS instrument obtained in K1 or in H2 spectral band under various observing conditions. The raw observations were pre-reduced with the data reduction handling pipeline³⁹ of the SPHERE instrument, which performs thermal background subtraction, flat-field correction, anamorphism correction, compensation for spectral transmission, flux normalization, bad pixels identification and interpolation, frame centering, true-North alignment, and frame selection. These operations are complemented by custom routines implemented in the SPHERE data center⁴⁰ in particular, to improve bad pixels correction. Finally, the SPHERE data center combines the pre-reduced observations and delivers the calibrated ADI datasets we consider in this work.

We compare the performance of the proposed method with the cADI, PCA, and PACO algorithms (see Sect. 1). For cADI, we have re-implemented the original method⁴ based on a full-frame estimation of the off-axis PSF and of the signal-to-noise (S/N) map, i.e., without angular-specific processing. We have also used the refined implementation of cADI available in the VIP package,²⁰ which includes a protection angle strategy accounting for a minimal field rotation between successive images when building the off-axis PSF in order to limit the self-subtraction effect. After computation of the off-axis PSF, a signal-to-noise (S/N) map is derived by accounting

[‡]To avoid overfitting, each realization s is unique with no repetition for the different epochs. The notion of *epoch* is used only as a way to evaluate regularly the performances of the model with the validation procedure.

[§]The 11 datasets are obtained on the following stars (and observing dates): HIP 72192 (2015-06-11), HIP 65426 (2017-02-09, 2018-05-13), HIP 88399 (2015-05-10, 2016-04-16, 2018-04-11), HD 131399 (2015-06-12, 2016-05-07) and HD 95086 (2015-05-05, 2018-01-05, 2021-03-11). The diversity in the experienced observing conditions is representative of the SPHERE observations.

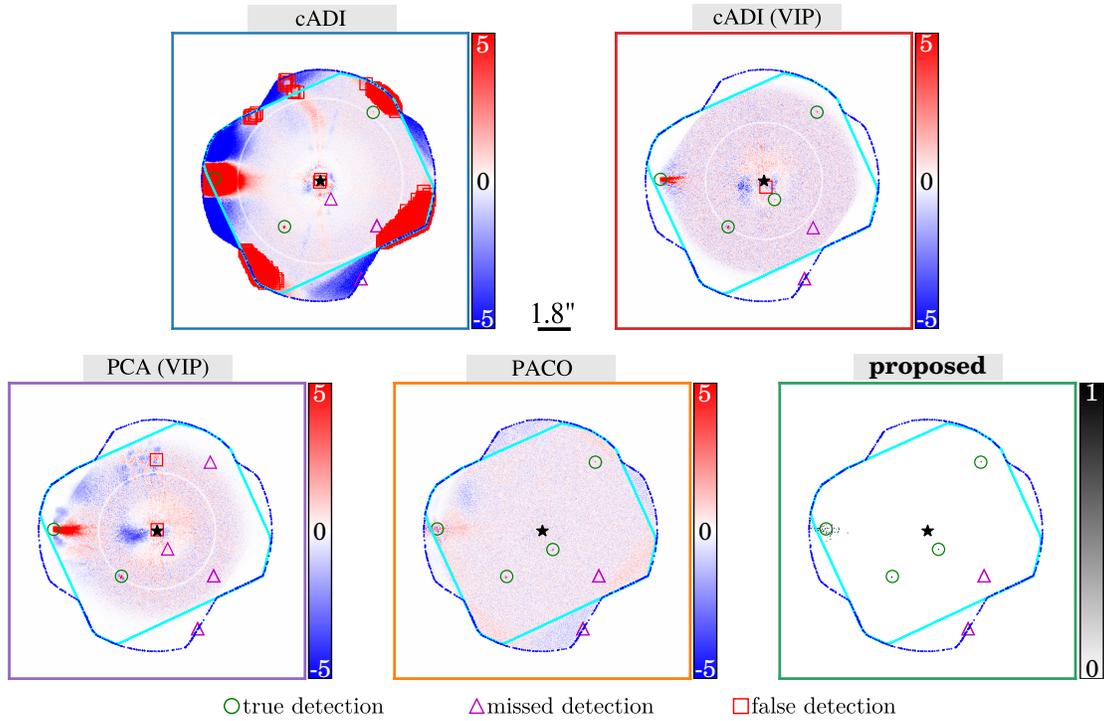


Figure 2. Detection maps obtained with the selected algorithms (see Sect. 3.1). Sources are classified as true, missed and false detections. The detection threshold is set to $\tau = 5$ for cADI, cADI (VIP), PCA(VIP) and PACO. It is set to $\tau = 0.5$ for the proposed algorithm. Dataset HIP 88399 (2015-05-10).

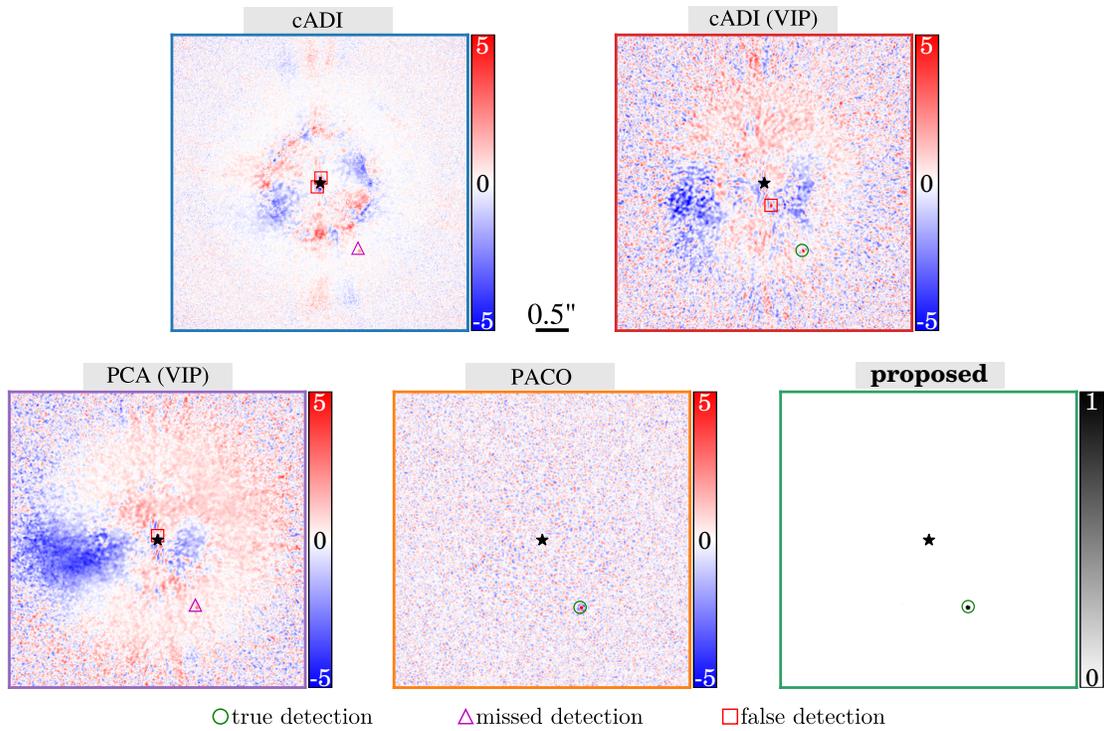


Figure 3. Same caption than Fig. 2. Zoom near the host star. Dataset HIP 88399 (2015-05-10).

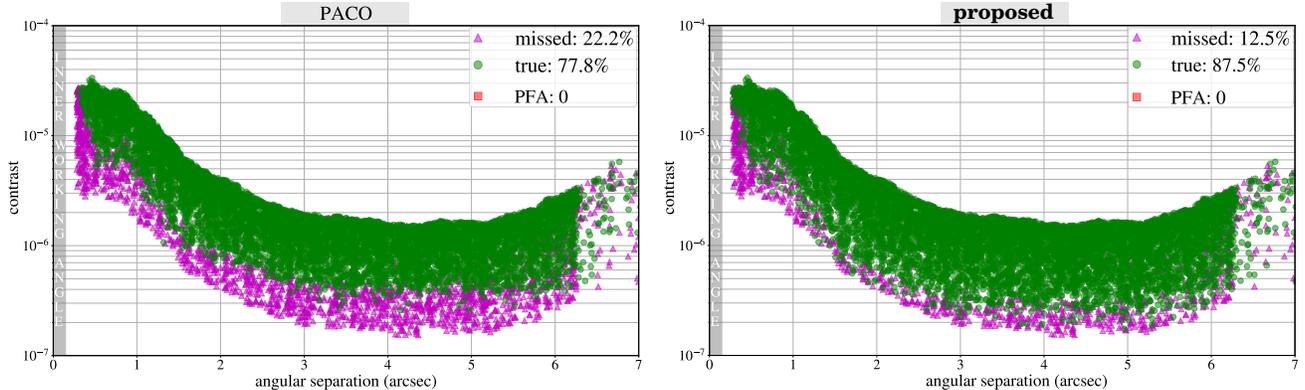


Figure 4. Detection results for 10,000 synthetic sources in a diagram plotting contrast *versus* angular separation. Each synthetic source is classified as missed, true or false detection. Dataset HIP 88399 (2015-05-10).

for an annular-based estimation of the noise in the residual images. We also applied the VIP implementation of the PCA-based algorithm combined with the same protection angle strategy and the annular-based computation of the S/N. For PCA reductions, the number of modes has been optimized by maximizing the S/N of synthetic sources with similar ranges of contrast than the ones we consider for our comparisons. The other parameters of the VIP implementation of cADI and PCA are less critical and are fixed at pre-set values.²⁰ For PACO, we performed the data reduction with our fully unsupervised processing pipeline.²

3.2 Qualitative and quantitative results

We present detailed results for one (HIP 88399, 2015-05-10) typical dataset selected among the 11 observations we consider in this work. Figures 2 and 3 give detection maps produced with the five tested algorithms. The detection threshold is set to $\tau = 5$ for the algorithms producing a S/N map (i.e., cADI, PCA, PACO), and to $\tau = 0.5$ for the proposed method producing a pseudo-probability map. Due to the binary semantic segmentation task we consider for the training step of the proposed method (see Sect. 2.2), its detection map is almost binary (i.e., each pixel value is close either to 0 or 1) so that the setting of the threshold τ is quite flexible. PACO and the proposed method lead to the best qualitative results since they are the only algorithms able to detect four of the six known real sources without any false alarm.

Figure 4 shows detection results on a sample of 10,000 synthetic sources in a diagram contrast *versus* angular separation for PACO and the proposed method. Each synthetic source is classified as missed, true or false detection using the detection thresholds defined previously. For PACO, setting the detection threshold at $\tau = 5$ corresponds to a realistic control¹⁻³ of the probability of false alarms (PFA) at 5σ (i.e., $\text{PFA} \simeq 3 \times 10^{-7}$). While the PFA should theoretically be controlled by the other algorithms producing a S/N map (cADI and PCA), we have shown in previous works¹⁻³ that the contrast curves are over-optimistic for these algorithms (i.e., there are significantly more false alarms than expected) due to a miss-modeling of the nuisance component. This claim is also supported by the detection maps given in Figs. 2 and 3. For the proposed method, converting pseudo-probabilities into S/N scores is not feasible given that the pseudo-probabilities are very close either to 0 or 1. For this reason, we check empirically that the targeted false alarm rate at 5σ is satisfied. By thresholding the 11 detection maps of the proposed approach at $\tau = 0.5$, we experienced fewer false alarms than statistically expected at 5σ so that our results can be fairly compared with the PACO results. Figure 4 illustrates the capability of the proposed method to detect fainter sources than PACO.

Figures 5 and 6 give ROCs representing the true positive rate as a function of the false discovery rate (see Sect. 2.2.3 and Eq. (11)) for the HIP 88399 (2015-05-10) dataset. This type of representation gives a comparison of the precision-recall trade-off reached by each method, regardless the detection quantity (S/N or pseudo-probability) they produce. These curves are obtained by counting the number of true positives (TP), and false alarms (FA) for the full range of possible detection thresholds, i.e. $\tau \in [0; 1]$ for the proposed method, and $\tau \in [\min(\hat{\mathbf{y}}); \max(\hat{\mathbf{y}})]$ for cADI, PCA, and PACO. Figure 5 is obtained by considering only the six real known

sources in the field of view for the HIP 88399 (2015-05-10) dataset. Figure 6 is obtained with the 10.000 synthetic sources considered for results presented in Fig. 4. For the synthetic sources, the results are split in four different angular separation ranges: $[0; 2]^\circ$, $[2; 4]^\circ$, $[4; 6]^\circ$, and $[6; 7]^\circ$. Figures 5 and 6 are complemented by Figs. 7 and 8 presenting similar results averaged over the 11 datasets we consider in this study. These results illustrate the benefits of the proposed method in terms of precision-recall trade-off: the AUC under ROC is improved by at least 7% with respect to the comparative algorithms for the four angular separation ranges we consider.

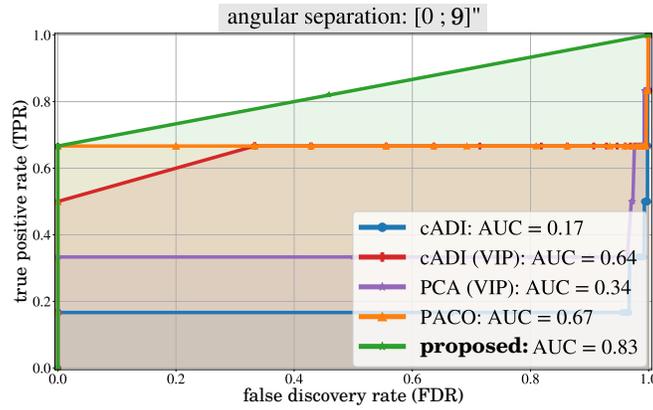


Figure 5. ROCs true positive rate as a function of the false discovery rate for known real sources. Dataset: HIP 88399 (2015-05-10).

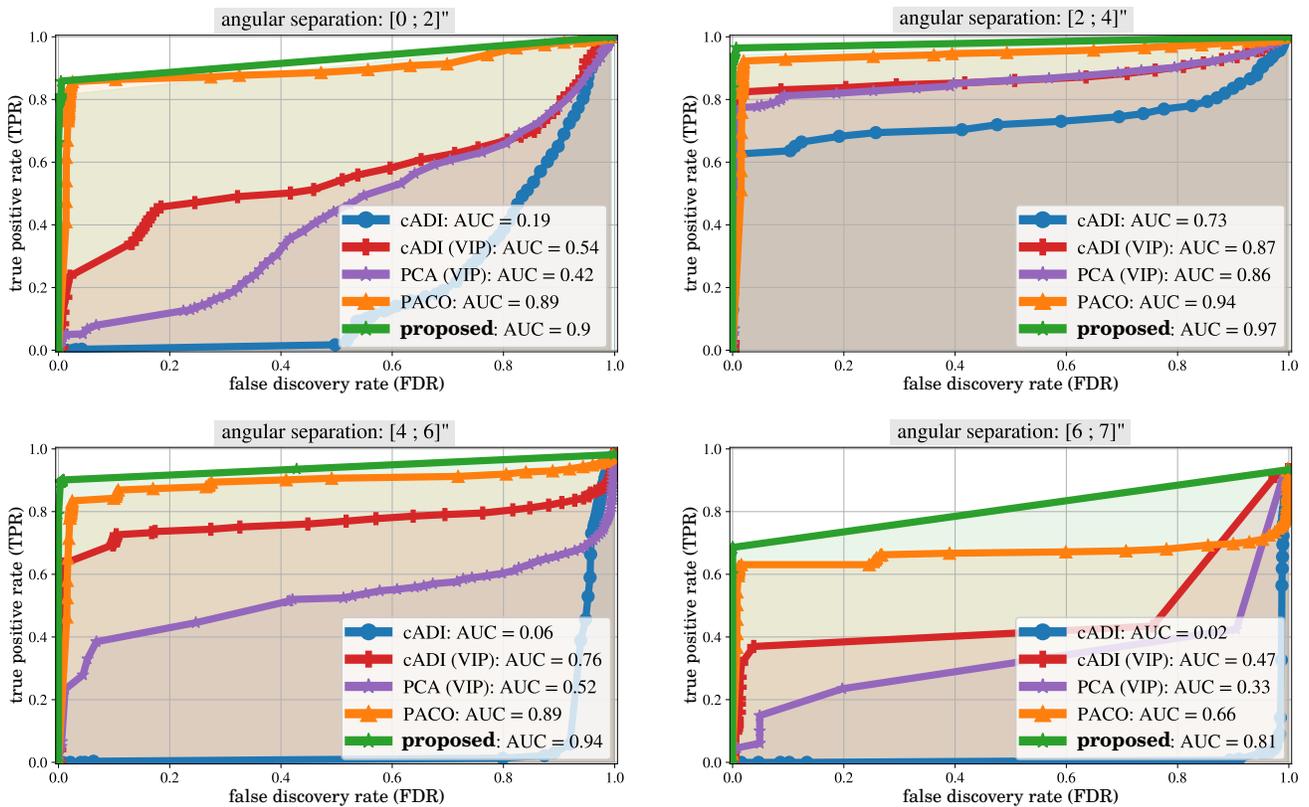


Figure 6. ROCs true positive rate as a function of the false discovery rate for injected synthetic sources. Dataset: HIP 88399 (2015-05-10).

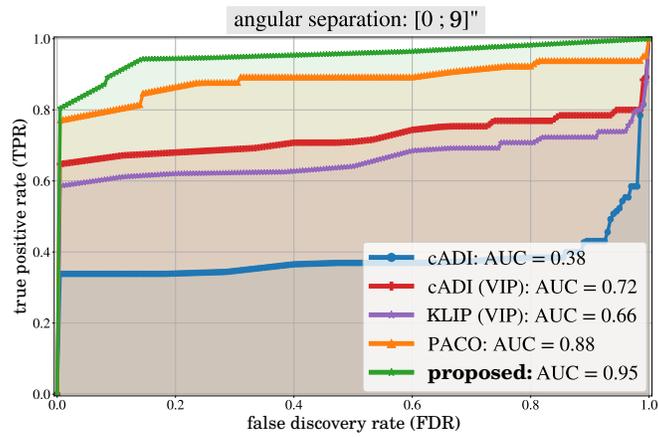


Figure 7. Same caption than Fig. 5. Mean results over 11 datasets.

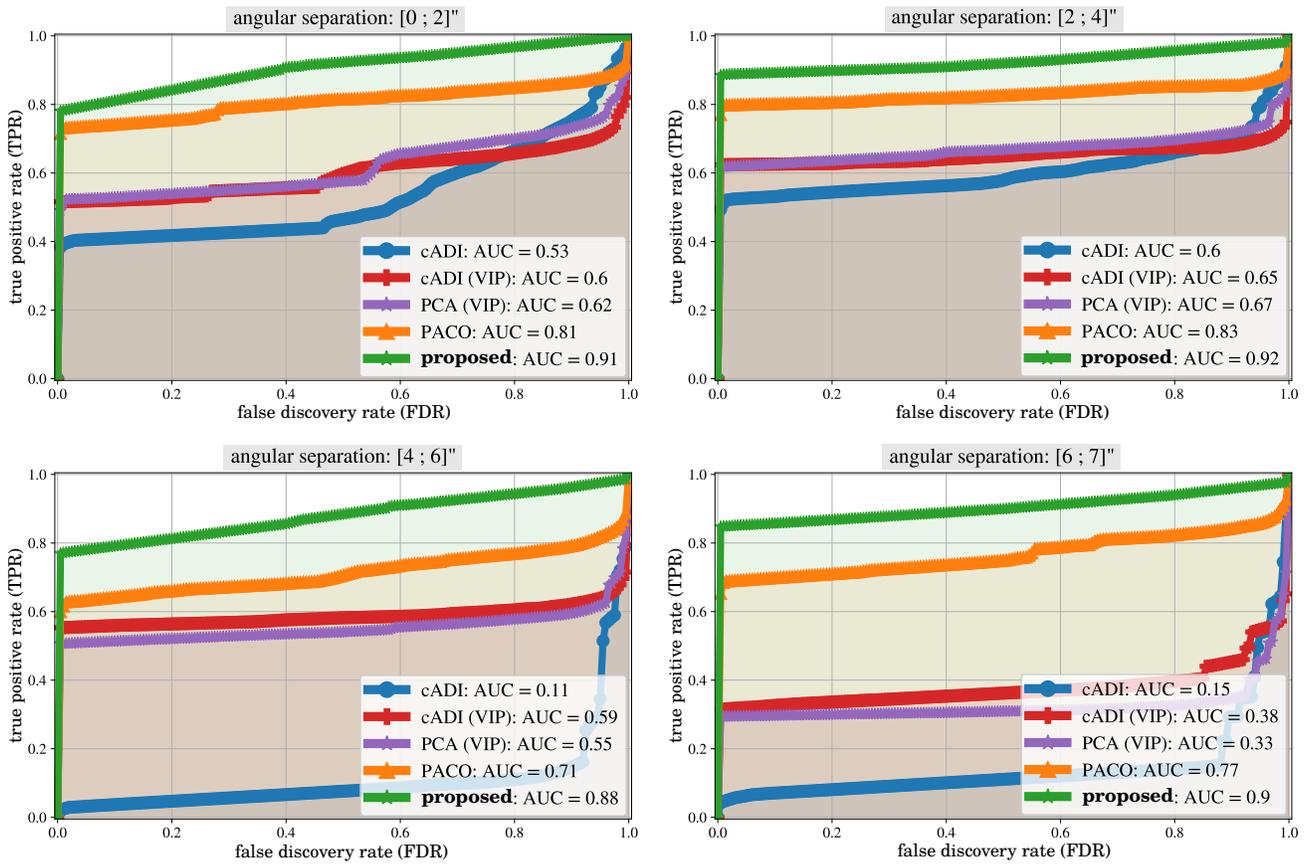


Figure 8. Same caption than Fig. 6. Mean results over 11 datasets.

4. CONCLUSION

We have described the key principles of a new algorithm for detecting point-like sources at high contrast from ADI observations. It combines the statistics-based model of PACO with deep learning in a three step procedure: (i) the data are centered and whitened using the PACO framework, (ii) a CNN is trained to detect synthetic sources from the preprocessed images, and (iii) the detection map is inferred. While the CNN itself works as a black-box approach, the proposed method encompasses prior domain knowledge such as the apparent motion of sources and the expected shape of the exoplanetary signal inside the ADI datasets. More importantly, the proposed approach capitalizes on the statistical model of the nuisance component embedded in PACO to improve the stationarity and the contrast during the preprocessing step. Tested on VLT/SPHERE-IRDIS datasets, the proposed method performs on-par with or better than PACO and other standard algorithms of the field. We are currently working on the extension of the model to multi-spectral data at low resolution (e.g., for the VLT/SPHERE-IFS data). We also plan to work on the control of the uncertainties: such a control is mandatory to interpret the inferred confidence scores as statistically-grounded probabilities.

ACKNOWLEDGMENTS

This project is supported in part by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (COBREX; grant agreement n° 885593). The work of JP and TB was supported in part by the Inria/NYU collaboration, the Louis Vuitton/ENS chair on artificial intelligence and the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR19-P3IA0001 (PRAIRIE 3IA Institute). The work of JM was supported in part by the ERC grant number 714381 (SOLARIS project) and by ANR 3IA MIAI@Grenoble Alpes (ANR-19-P3IA-0003).

This work has made use of the SPHERE Data Center, 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

- [1] Flasseur, O., Denis, L., Thiébaud, É., and Langlois, M., “Exoplanet detection in angular differential imaging by statistical learning of the nonstationary patch covariances - The PACO algorithm,” *Astronomy & Astrophysics* **618**, A138 (2018).
- [2] Flasseur, O., Denis, L., Thiébaud, É., and Langlois, M., “An unsupervised patch-based approach for exoplanet detection by direct imaging,” in [*IEEE International Conference on Image Processing*], 2735–2739 (2018).
- [3] Flasseur, O., Denis, L., Thiébaud, É. M., and Langlois, M., “Exoplanet detection in angular and spectral differential imaging: local learning of background correlations for improved detections,” in [*SPIE Astronomical Telescopes + Instrumentation*], **10703**, 107032R, International Society for Optics and Photonics (2018).
- [4] Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., and Nadeau, D., “Angular differential imaging: A powerful high-contrast imaging technique,” *The Astrophysical Journal* **641**(1), 556 (2006).
- [5] Soummer, R., Pueyo, L., and Larkin, J., “Detection and characterization of exoplanets and disks using projections on Karhunen-Loève eigenimages,” *The Astrophysical Journal Letters* **755**(2), L28 (2012).
- [6] Amara, A. and Quanz, S. P., “PYNPOINT: an image processing package for finding exoplanets,” *Monthly Notices of the Royal Astronomical Society* **427**(2), 948–955 (2012).
- [7] Marois, C., Correia, C., Véran, J.-P., and Currie, T., “TLOCI: A fully loaded speckle killing machine,” **8**(S299), 48–49 (2013).
- [8] Traub, W. A. and Oppenheimer, B. R., “Direct imaging of exoplanets,” *Exoplanets*, 111–156 (2010).
- [9] Bowler, B. P., “Imaging extrasolar giant planets,” *Publications of the Astronomical Society of the Pacific* **128**(968), 102001 (2016).
- [10] Pueyo, L., “Direct imaging as a detection technique for exoplanets,” *Handbook of Exoplanets*, 705–765 (2018).

- [11] Schneider, J., Dedieu, C., Le Sidaner, P., Savalle, R., and Zolotukhin, I., “Defining and cataloging exoplanets: the exoplanet.eu database,” *Astronomy & Astrophysics* **532**, A79 (2011).
- [12] Nielsen, E. L., De Rosa, R. J., Macintosh, B., Wang, J. J., Ruffio, J.-B., Chiang, E., Marley, M. S., Saumon, D., Savransky, D., Ammons, S. M., et al., “The Gemini planet imager exoplanet survey: Giant planet and brown dwarf demographics from 10 to 100 au,” *The Astronomical Journal* **158**(1), 13 (2019).
- [13] Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafrenière, D., and Doyon, R., “Direct imaging of multiple planets orbiting the star HR 8799,” *Science* **322**(5906), 1348–1352 (2008).
- [14] Lagrange, A.-M., Gratadour, D., Chauvin, G., Fusco, T., Ehrenreich, D., Mouillet, D., Rousset, G., Rouan, D., Allard, F., Gendron, É., et al., “A probable giant planet imaged in the β Pictoris disk: VLT/NaCo deep L-band imaging,” *Astronomy & Astrophysics* **493**(2), L21–L25 (2009).
- [15] Nielsen, E. L., Liu, M. C., Wahhaj, Z., Biller, B. A., Hayward, T. L., Boss, A., Bowler, B., Kraus, A., Shkolnik, E. L., Tecza, M., et al., “The Gemini NICI planet-finding campaign: Discovery of a multiple system orbiting the young A star HD 1160,” *The Astrophysical Journal* **750**(1), 53 (2012).
- [16] Macintosh, B., Graham, J., Barman, T., De Rosa, R., Konopacky, Q., Marley, M., Marois, C., Nielsen, E., Pueyo, L., Rajan, A., et al., “Discovery and spectroscopy of the young jovian planet 51 Eri b with the Gemini Planet Imager,” *Science* **350**(6256), 64–67 (2015).
- [17] Chauvin, G., Desidera, S., Lagrange, A.-M., Vigan, A., Gratton, R., Langlois, M., Bonnefoy, M., Beuzit, J.-L., Feldt, M., Mouillet, D., et al., “Discovery of a warm, dusty giant planet around HIP 65426,” *Astronomy & Astrophysics* **605**, L9 (2017).
- [18] Keppler, M., Benisty, M., Müller, A., Henning, T., Van Boekel, R., Cantalloube, F., Ginski, C., Van Holstein, R., Maire, A.-L., Pohl, A., et al., “Discovery of a planetary-mass companion within the gap of the transition disk around PDS 70,” *Astronomy & Astrophysics* **617**, A44 (2018).
- [19] Lagrange, A.-M., Bonnefoy, M., Chauvin, G., Apai, D., Ehrenreich, D., Boccaletti, A., Gratadour, D., Rouan, D., Mouillet, D., Lacour, S., et al., “A giant planet imaged in the disk of the young star β Pictoris,” *Science* **329**(5987), 57–59 (2010).
- [20] Gonzalez, C. A. G., Wertz, O., Absil, O., Christiaens, V., Defrère, D., Mawet, D., Milli, J., Absil, P.-A., Van Droogenbroeck, M., Cantalloube, F., et al., “VIP: Vortex Image Processing package for high-contrast direct imaging,” *The Astronomical Journal* **154**(7), 12pp (2017).
- [21] Galicher, R., Boccaletti, A., Mesa, D., Delorme, P., Gratton, R., Langlois, M., Lagrange, A.-M., Maire, A.-L., Le Coroller, H., Chauvin, G., et al., “Astrometric and photometric accuracies in high contrast imaging: The SPHERE speckle calibration tool (SpeCal),” *Astronomy & Astrophysics* **615**, A92 (2018).
- [22] Gonzalez, C., Absil, O., and Van Droogenbroeck, M., “Supervised detection of exoplanets in high-contrast imaging sequences,” *Astronomy & Astrophysics* **613**, A71 (2018).
- [23] Yip, K. H., Nikolaou, N., Coronica, P., Tsiaras, A., Edwards, B., Changeat, Q., Morvan, M., Biller, B., Hinkley, S., Salmond, J., et al., “Pushing the limits of exoplanet discovery via direct imaging with deep learning,” in *[Joint European Conference on Machine Learning and Knowledge Discovery in Databases]*, 322–338, Springer (2019).
- [24] Gebhard, T. D., Bonse, M. J., Quanz, S. P., and Schölkopf, B., “Half-sibling regression meets exoplanet imaging: PSF modeling and subtraction using a flexible, domain knowledge-driven, causal framework,” *arXiv preprint arXiv:2204.03439* (2022).
- [25] Cantalloube, F., Gomez-Gonzalez, C., Absil, O., Cantero, C., Bacher, R., Bonse, M., Bottom, M., Dahlqvist, C.-H., Desgrange, C., Flasseur, O., et al., “Exoplanet imaging data challenge: benchmarking the various image processing methods for exoplanet detection,” in *[Adaptive Optics Systems VII]*, **11448**, 1027–1062, SPIE (2020).
- [26] Beuzit, J.-L., Vigan, A., Mouillet, D., Dohlen, K., Gratton, R., Boccaletti, A., Sauvage, J.-F., Schmid, H. M., Langlois, M., Petit, C., et al., “SPHERE: the exoplanet imager for the Very Large Telescope,” *Astronomy & Astrophysics* **631**, A155 (2019).
- [27] Dohlen, K., Langlois, M., Saisse, M., Hill, L., Origine, A., Jacquet, M., Fabron, C., Blanc, J.-C., Llored, M., Carle, M., et al., “The infra-red dual imaging and spectrograph for SPHERE: design and performance,” in *[SPIE Astronomical Telescopes + Instrumentation]*, 70143L–70143L, International Society for Optics and Photonics (2008).

- [28] Dohlen, K., Saisse, M., Origne, A., Moreaux, G., Fabron, C., Zamkotsian, F., Lanzoni, P., and Lemarquis, F., “Prototyping of differential optics for the SPHERE IRDIS dual imaging planet finder camera,” in [*SPIE Astronomical Telescopes + Instrumentation*], **7018**, 701859, International Society for Optics and Photonics (2008).
- [29] Ledoit, O. and Wolf, M., “A well-conditioned estimator for large-dimensional covariance matrices,” *Journal of Multivariate Analysis* **88**(2), 365–411 (2004).
- [30] Chen, Y., Wiesel, A., Eldar, Y. C., and Hero, A. O., “Shrinkage algorithms for MMSE covariance estimation,” *IEEE Transactions on Signal Processing* **58**(10), 5016–5029 (2010).
- [31] Flasseur, O., Denis, L., Thiébaud, É., Langlois, M., et al., “REXPACO: An algorithm for high contrast reconstruction of the circumstellar environment by angular differential imaging,” *Astronomy & Astrophysics* **651**, A62 (2021).
- [32] Ronneberger, O., Fischer, P., and Brox, T., “U-net: Convolutional networks for biomedical image segmentation,” in [*International Conference on Medical image computing and computer-assisted intervention*], 234–241, Springer (2015).
- [33] He, K., Zhang, X., Ren, S., and Sun, J., “Deep residual learning for image recognition,” in [*Proceedings of the IEEE conference on computer vision and pattern recognition*], 770–778 (2016).
- [34] Yakubovskiy, P., “Segmentation Models Pytorch.” https://github.com/qubvel/segmentation_models.pytorch (2020).
- [35] Milletari, F., Navab, N., and Ahmadi, S.-A., “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” in [*2016 fourth international conference on 3D vision (3DV)*], 565–571, IEEE (2016).
- [36] Sudre, C. H., Li, W., Vercauteren, T., Ourselin, S., and Jorge Cardoso, M., “Generalised Dice overlap as a deep learning loss function for highly unbalanced segmentations,” in [*Deep learning in medical image analysis and multimodal learning for clinical decision support*], 240–248, Springer (2017).
- [37] Wang, L., Wang, C., Sun, Z., and Chen, S., “An improved Dice loss for pneumothorax segmentation by mining the information of negative areas,” *IEEE Access* **8**, 167939–167949 (2020).
- [38] Kay, S. M., [*Fundamentals of statistical signal processing: estimation theory*], Prentice-Hall, Inc. (1993).
- [39] Pavlov, A., Möller-Nilsson, O., Feldt, M., Henning, T., Beuzit, J.-L., and Mouillet, D., “SPHERE data reduction and handling system: overview, project status, and development,” in [*SPIE Astronomical Telescopes + Instrumentation*], **7019**, 701939, International Society for Optics and Photonics (2008).
- [40] Delorme, P., Meunier, N., Albert, D., Lagadec, E., Le Coroller, H., Galicher, R., Mouillet, D., Boccaletti, A., Mesa, D., Meunier, J. C., Beuzit, J. L., Lagrange, A. M., Chauvin, G., Sapone, A., Langlois, M., Maire, A. L., Montargès, M., Gratton, R., Vigan, A., and Surace, C., “The SPHERE Data Center: a reference for high contrast imaging processing,” in [*Annual meeting of the French Society of Astronomy and Astrophysics*], (Dec. 2017).