

Combining multi-spectral data with statistical and deep-learning models for improved exoplanet detection in direct imaging at high contrast

Olivier Flasseur¹, Théo Bodrito², Julien Mairal³, Jean Ponce^{2,4}, Maud Langlois¹, and Anne-Marie Lagrange^{5,6}

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

²Département d'Informatique de l'École Normale Supérieure (ENS-PSL, CNRS, Inria), France

³Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France

⁴Courant Institute of Mathematical Sciences, Center for Data Science, New York Univ., USA

⁵Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique, Obs. de Paris, Univ. PSL, Sorbonne Univ., Univ. Diderot, France

⁶Univ. Grenoble Alpes, Institut de Planétologie et d'Astrophysique de Grenoble, France

Abstract—Exoplanet detection by direct imaging is a difficult task: the faint signals from the objects of interest are buried under a spatially structured nuisance component induced by the host star. The exoplanet signals can only be identified when combining several observations with dedicated detection algorithms. In contrast to most of existing methods, we propose to learn a model of the spatial, temporal and spectral characteristics of the nuisance, directly from the observations. In a pre-processing step, a statistical model of their correlations is built locally, and the data are centered and whitened to improve both their stationarity and signal-to-noise ratio (SNR). A convolutional neural network (CNN) is then trained in a supervised fashion to detect the residual signature of synthetic sources in the pre-processed images. Our method leads to a better trade-off between precision and recall than standard approaches in the field. It also outperforms a state-of-the-art algorithm based solely on a statistical framework. Besides, the exploitation of the spectral diversity improves the performance compared to a similar model built solely from spatio-temporal data.

Index Terms—detection, supervised deep learning, matched filter, multi-variate data, correlated data

I. INTRODUCTION

Alongside indirect methods for detecting exoplanets (e.g., based on the monitoring of transits or of the radial velocities), direct imaging [1] from Earth is a method of choice to detect and characterize the physical properties (e.g., orbit, atmospheric composition) of exoplanets. This characterization is carried out by comparing the estimated astrometry and spectral energy distribution (SED) with physical models. However, the low number of exoplanets detected and characterized with this technique (a dozen over the last twelve years) testifies to the difficulty of this long-term goal. This is explained by the very large difference in *contrast*, typically greater than five orders

The work of OF, ML, and AML was supported in part by the European Research Council (ERC) under the *European Union's Horizon 2020 research and innovation program* (COBREX; n° 885593). The work of TB and JP was supported in part by the Inria/NYU collaboration, the Louis Vuitton/ENS chair on artificial intelligence and the French Agence Nationale de la Recherche (ANR) as part of the *Investissements d'avenir program* (PRAIRIE 3IA Institute; n° 19-P3IA-0001). The work of JM was supported in part by the ERC (SOLARIS; n° 714381) and by the ANR (3IA MIAI, n° 19-P3IA-0003).




Fig. 1. Illustration of a dataset from the SPHERE-IFS instrument. Left: images at different wavelengths. Right: images at different times (at $\lambda_{25} = 1.4 \mu\text{m}$). Red circles represent the locations of three known exoplanets whose signatures are too faint to be detected without additional processing. Bottom: spatio-spectral and spatio-temporal slice cuts along the white dashed line.

of magnitude in the infrared, between the exoplanets and their host star. Beyond extreme adaptive optics compensating wavefront distortions and the coronagraph partly masking out the star, which are implemented in cutting-edge observational facilities, the detection of faint sources requires a dedicated processing of the data.

We recall in Section II the key principles of high-contrast imaging (HCI) and of the state-of-the-art processing methods. We describe in Section III the proposed method, which is an extension to multi-spectral data of our very recent work dedicated to the processing of spatio-temporal observations [2]. We compare in Section IV the detection performance of the proposed method to state-of-the-art algorithms of the field on data from the SPHERE instrument currently operating at the Very Large Telescope (VLT, Chile).

II. PRINCIPLE OF DIRECT IMAGING AT HIGH CONTRAST

A classical observation technique in HCI is angular differential imaging (ADI) that consists in recording a 3-D dataset (2-D + time) over a series of short exposures during a few hours of observation. During the acquisition, all the off-axis point-like sources follow an apparent circular motion (typically,

about a few tens of degrees) around the optical axis due to the Earth's rotation. Since the pupil of the telescope is kept fixed, the other structures of the background remain temporally quasi-static. Due to the diffraction and to residual optical aberrations, stellar leakages take the form of spatially correlated speckles. Speckles and other additive sources of noise (thermal background, detector readout, photon noise) form a *nuisance component* that strongly contaminates the images formed downstream of the coronagraph. Besides, there is a significant variability in the amplitude and in the structure of the spatial correlations of the nuisance, which is stronger near the star than farther away. Spectral differential imaging (SDI) is a complementary observation technique allowing to capture images in several spectral channels simultaneously with an integral field spectrograph (IFS). With SDI, the speckle pattern scales approximately linearly with the wavelength due to the diffraction. The temporal and spectral diversity induced by ADI and SDI can be combined advantageously. This hybrid technique, called angular and spectral differential imaging (ASDI), produces 4-D datasets (2-D + time + spectral). In this paper, we consider ASDI. As an example, the SPHERE-IFS instrument produces series of $T \simeq 100$ temporal exposures with $N = 290 \times 290$ pixels each in $L = 39$ spectral bands (between $\lambda_1 = 0.9 \mu\text{m}$ and $\lambda_{39} = 1.6 \mu\text{m}$), see Fig. 1.

To unmix the exoplanet signals from the nuisance, state-of-the-art methods combine the different temporal and spectral images [3]. A simple solution consists in subtracting the temporal mean to each frame of a given spectral channel to attenuate speckles. The residual images are then rotated and scaled to compensate for the (predictable) motion of the off-axis sources, so that their signals are spatially co-localized. The resulting stack is finally combined, e.g., by taking the median, as in the cADI method. Other methods resort to a principal component analysis (PCA) to build an empirical model of the nuisance. These two baseline algorithms [3] are used routinely in HCI. Approaches based on detection theory are more efficient because they model explicitly the attenuation of the sought objects induced by image subtraction. In particular, the PACO algorithm [4] and its extension PACO ASDI [5] to multi-spectral data build a statistical model of the nuisance by capturing their spatial correlations and fluctuations at the scale of small patches of a few tens of pixels. However, it remains room for improvement due to the approximate fidelity of this statistical model to the observations. In this context, we proposed very recently deep PACO [2], a data-driven algorithm for ADI observations combining the statistical model of PACO with supervised deep learning. deep PACO performs better than existing methods, including approaches based on deep learning that remain limited by a high false alarm rate. In this paper, we present deep PACO ASDI, extending deep PACO to multi-spectral data.

III. PROPOSED ALGORITHM

A. Direct Model of the Observed Intensity

An ASDI dataset \mathbf{r} in $\mathbb{R}^{N \times T \times L}$ is formed by N -pixels images, recorded at different times t in $\llbracket 1; T \rrbracket$ and in different

spectral channels ℓ in $\llbracket 1; L \rrbracket$. The contribution \mathbf{r}_ℓ in $\mathbb{R}^{N \times T}$ is:

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

where \mathbf{f}_ℓ in $\mathbb{R}^{N \times T}$ and \mathbf{h}_ℓ in $\mathbb{R}^{N \times T}$ are respectively the contribution of the nuisance and of any point-like source taking the form of the off-axis point-spread function (PSF), at spectral channel ℓ . The contribution of a source $p \in \llbracket 1; P \rrbracket$ is centered at location $\mathcal{F}_{t,\ell}(\phi_p)$ in the t -th image of the ℓ -th channel, where ϕ_p is its initial location on an image at a time t_{ref} and spectral channel λ_{ref} of reference. $\mathcal{F}_{t,\ell}$ is a deterministic geometrical transform (a composition of a rotation and of a radial translation with respect to the star) to account for the apparent rotation induced by ADI and to the translation induced by spatial scaling of factor $\lambda_{\text{ref}}/\lambda_\ell$ applied beforehand to spatially co-align the nuisance's structures, see Fig. 1.

B. Step 1: Pre-processing by Statistical Modeling

Based on our previous work [5], we model the random fluctuations of the nuisance component \mathbf{f} by a statistical model whose parameters are self-calibrated on the observations. The model is local, i.e., it is built at the scale of patches of a few tens of pixels to account for the high non-stationarity of the nuisance component. We consider non-overlapping square patches of K pixels, and we note \mathbb{P} the set of locations paving the whole field of view. We model the distribution of a patch $\mathbf{f}_{n,t,\ell}$ centered at pixel n , at time t , and channel ℓ by a scaled multi-variate Gaussian: $\mathcal{N}(\mathbf{m}_{n,\ell}, \sigma_{n,t,\ell}^2 \mathbf{C}_n)$. The sample estimators $\{\widehat{\mathbf{m}}_{n,\ell}, \widehat{\sigma}_{n,t,\ell}^2, \widehat{\mathbf{S}}_n\}$ of $\{\mathbf{m}_{n,\ell}, \sigma_{n,t,\ell}^2, \mathbf{C}_n\}$ are obtained from the TL patches \mathbf{r}_n in $\mathbb{R}^{K \times T \times L}$ in the maximum-likelihood sense with a fixed-point iterative scheme:

$$\begin{cases} \widehat{\mathbf{m}}_{n,\ell} = \left(\sum_t \widehat{\sigma}_{n,t,\ell}^{-2} \right)^{-1} \sum_t \widehat{\sigma}_{n,t,\ell}^{-2} \mathbf{r}_{n,t,\ell}, \\ \widehat{\mathbf{S}}_n = \frac{1}{TL} \sum_{t,\ell} \widehat{\sigma}_{n,t,\ell}^{-2} (\mathbf{r}_{n,t,\ell} - \widehat{\mathbf{m}}_{n,\ell})(\mathbf{r}_{n,t,\ell} - \widehat{\mathbf{m}}_{n,\ell})^\top, \\ \widehat{\sigma}_{n,t,\ell}^2 = K^{-1} (\mathbf{r}_{n,t,\ell} - \widehat{\mathbf{m}}_{n,\ell})^\top \widehat{\mathbf{S}}_n^{-1} (\mathbf{r}_{n,t,\ell} - \widehat{\mathbf{m}}_{n,\ell}), \end{cases} \quad (2)$$

where the scaling factors $\{\widehat{\sigma}_{n,t,\ell}^2\}_{t=1:T, \ell=1:L}$ are the empirical variances of the (whitened) residuals, that allow to neutralize patches with outliers. The total number of samples involved in the computation of $\widehat{\mathbf{S}}_n$ being lower than the number K of pixels in a patch, the sample covariance $\widehat{\mathbf{S}}_n$ is very noisy and can be rank deficient. As in our previous work [5], we regularize it by *shrinkage* [6], i.e. the resulting estimator is a convex combination between the low bias/high variance estimator $\widehat{\mathbf{S}}_n$ and a high bias/low variance estimator $\widehat{\mathbf{F}}_n$:

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

where $\widehat{\mathbf{F}}_n$ is a diagonal matrix with only the sample variances, and the hyper-parameter $\widehat{\rho}_n$ setting a bias-variance trade-off. The latter is estimated [2] for each patch location n through:

$$\widehat{\rho}_n = \frac{\text{tr}(\widehat{\mathbf{S}}_n^2) + \text{tr}^2(\widehat{\mathbf{S}}_n) - 2 \sum_{k=1}^K [\widehat{\mathbf{S}}_n]_{kk}^2}{(Q + 1) \left(\text{tr}(\widehat{\mathbf{S}}_n^2) - \sum_{k=1}^K [\widehat{\mathbf{S}}_n]_{kk}^2 \right)}, \quad (4)$$

with $Q = \left(\sum_{t,\ell} \widehat{\sigma}_{n,t,\ell}^{-2} \right)^2 / \left(\sum_{t,\ell} \widehat{\sigma}_{n,t,\ell}^{-4} \right)$ the equivalent number of patches involve in the computation of $\widehat{\mathbf{C}}_n$ in the presence

of the scaling factors $\{\hat{\sigma}_{n,t,\ell}^2\}_{t=1:T,\ell=1:L}$. Given the statistics of the nuisance, the pre-processed images $\tilde{\mathbf{r}}$ in $\mathbb{R}^{N \times T \times L}$ (i.e., after centering and whitening) are obtained by:

$$\tilde{\mathbf{r}}_{n,t,\ell} = \mathbf{W}_{n,t,\ell} \mathbf{r}_{n,t,\ell} = \hat{\sigma}_{n,t,\ell}^{-1} \hat{\mathbf{L}}_n^\top (\mathbf{r}_{n,t,\ell} - \hat{\mathbf{m}}_{n,\ell}), \quad (5)$$

with $\hat{\mathbf{L}}_n$ the Cholesky's factorization of $\hat{\mathbf{C}}_n^{-1}$ (i.e., $\hat{\mathbf{L}}_n \hat{\mathbf{L}}_n^\top = \hat{\mathbf{C}}_n^{-1}$). Most of the spatial structures of \mathbf{r} are removed in $\tilde{\mathbf{r}}$.

C. Step 2: Exoplanet Detection by Supervised Deep Learning

1) *Training basis*: Starting from a temporo-spectral series of pre-processed images, we aim to infer a detection map $\hat{\mathbf{y}}$ in $[0; 1]^M$, where each pixel-value represents the *pseudo-probability* that a source is centered at that location at time t_{ref} . Obtaining real ground truth data is challenging in HCI: a few real sources have been detected and undiscovered sources might still be present in the data. We thus opt for a synthetic training strategy: the training set consists of S pairs $\{\check{\mathbf{r}}^{[s]}; \mathbf{y}^{[s]}\}_{s=1:S}$ of samples/ground truths, resulting from the massive injection of point-like sources. Besides, we train a different model for each ASDI dataset due to the high variability of the nuisance component from one observation to the other. This data-dependence prevents generalizing the learned detector to a new set of pre-processed data. These peculiarities imply the design of a data-augmentation strategy to prevent over-fitting and to deal with the absence of knowledge about real sources. We thus apply a random permutation of the T images of each spectral channel for each training sample s . This allows (i) to create artificially different datasets, and (ii) to break the temporal consistency of real sources. Synthetic sources are then injected inside the temporally permuted data following the direct model (1) to form the intermediate datasets $\{\check{\mathbf{r}} \in \mathbb{R}^{N \times T \times L}\}_{s=1:S}$ such that:

$$\check{\mathbf{r}}_\ell^{[s]} = \mathbf{P}^{[s]} \mathbf{f}_\ell + \sum_{p=1}^{P^{[s]}} \alpha_{p,\ell}^{[s]} \mathbf{h}_\ell(\phi_p^{[s]}), \quad (6)$$

where operator \mathbf{P} performs the random temporal permutation. For each synthetic source p , we select randomly a synthetic SED α_p in \mathbb{R}^L from a custom library of 10,000 sub-stellar spectra generated with ExoREM [7], a physics-based model accounting for chemistry and for the presence of clouds in the exoplanet's atmosphere. The number P of injected sources is drawn uniformly in $\llbracket 1; 10 \rrbracket$, since we expect few sources in the field of view. The initial source locations $\{\phi_p\}_{p=1:P}$ are drawn uniformly per angular separation, and their mean fluxes (integrated over the SED) are 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 detection limit reached by PACO at location ϕ_p . We thus cover both sources that are detectable and unseen at the classical 5σ detection confidence using PACO. As pre-processing is computationally demanding, we opt for a local update strategy. Prior to the injection of synthetic sources, the dataset \mathbf{r} is pre-processed to form $\tilde{\mathbf{r}}$, see Section III-B. After injections, the set $\mathbb{S}^{[s]}$ of locations impacted by the signal of the $P^{[s]}$ sources is determined. Outside $\mathbb{S}^{[s]}$, the pre-processed images are obtained from the temporal permutation of $\tilde{\mathbf{r}}$. Inside $\mathbb{S}^{[s]}$, the statistics

of the nuisance and the pre-processed images are updated given the contamination of the $P^{[s]}$ injected sources to form $\{\check{\mathbf{r}} \in \mathbb{R}^{N \times T \times L}\}_{s=1:S}$ such that:

$$\check{\mathbf{r}}_{n,t,\ell}^{[s]} = \begin{cases} \mathbf{W}_{n,t,\ell} \tilde{\mathbf{r}}_{n,t,\ell}^{[s]}, & \text{for } n \in \mathbb{S}^{[s]} \cap \mathbb{P}, \\ \mathbf{P}^{[s]} \tilde{\mathbf{r}}_{n,t,\ell}, & \text{for } n \in \mathbb{P} - \mathbb{S}^{[s]} \cap \mathbb{P}. \end{cases} \quad (7)$$

Finally, the apparent motion (rotation and translation with respect to the star) of the sought objects are compensated to co-align spatially the signals of the injected sources along the temporal and spectral axis. It results $\{\check{\mathbf{r}} \in \mathbb{R}^{N \times T \times L}\}_{s=1:S}$.

2) *Loss function and metrics*: The design of the training loss is driven by three criteria: (i) dealing with the strong class-imbalance, (ii) being computationally efficient, and (iii) representing a measure close to a detection accuracy score. We compared classical losses for pixel-wise classification, and we selected the Dice2 similarity score that was introduced first for biomedical image segmentation [8]. This loss satisfies the targeted criteria, in particular it penalizes equally errors in the source and background areas:

$$\mathcal{L}^{[s]} = 1 - \frac{\sum_m \mathbf{y}_m^{[s]} \hat{\mathbf{y}}_m^{[s]} + \epsilon}{\underbrace{\sum_m \mathbf{y}_m^{[s]} + \hat{\mathbf{y}}_m^{[s]} + \epsilon}_{\text{source error}}} - \frac{\sum_m (1 - \mathbf{y}_m^{[s]}) (1 - \hat{\mathbf{y}}_m^{[s]} + \epsilon)}{\underbrace{\sum_m 2 - \mathbf{y}_m^{[s]} - \hat{\mathbf{y}}_m^{[s]} + \epsilon}_{\text{background error}}}, \quad (8)$$

where $\{\mathbf{y}^{[s]}; \hat{\mathbf{y}}^{[s]}\}$ is a set of ground truth and predicted detection maps, and ϵ is a small stability parameter whose setting has very limited impact on the detection performance. At validation time, 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 derive the true positive rate (TPR) and the false discovery rate (FDR). From TPR and FDR, receiver operating curves (ROCs) are built by varying the detection threshold τ . The area under ROCs (AUC) is used as an overall performance score. We observed experimentally that the loss (8) is well anti-correlated to this accuracy metric.

3) *Architecture and implementation*: We select a U-Net with a ResNet18 as encoder backbone ($\simeq 11$ millions of weights); an architecture widely used for image segmentation. Its implementation is based on the SMP package¹. The encoder and decoder parts have four blocks: each one containing a series of convolution, batch normalization, activation (ReLU), and max pooling layers. The network ends with a sigmoid activation to produce a detection map $\hat{\mathbf{y}}$ in $[0; 1]^M$. At each training step s , the full stack $\check{\mathbf{r}}^{[s]}$ of temporo-spectral pre-processed images with synthetic sources is fed as input of the CNN. Pairs of samples $\{\check{\mathbf{r}}^{[s]}; \mathbf{y}^{[s]}\}$ are generated on the fly, and each realization s is unique to prevent over-fitting. The notion of *epoch* is used only as a way to evaluate regularly the performance and to schedule the learning rate. The number of samples S per epoch is respectively set to 1,000 and to 100 at training and evaluation time. The batch size is equal to 1. The network is trained from scratch with AMSGrad. The parameters of the optimizer and of the scheduler have been

¹https://github.com/qubvel/segmentation_models.pytorch

TABLE I
MEAN AUC FOR ROCs GIVING THE TPR AS A FUNCTION OF THE FDR
FOR REAL AND MASSIVELY INJECTED SYNTHETIC SOURCES FOR
SPHERE-IFS DATASETS. BEST RESULTS ARE IN BOLD FONT.

source type	sep. (")	cADI	PCA	PACO ASDI	proposed
real	[0.0; 1.0]	0.13	0.65	0.80	0.96
synthetic	[0.0; 0.5]	0.24	0.36	0.68	0.76
synthetic	[0.5; 1.0]	0.24	0.52	0.79	0.85

fine tuned on a few datasets (weight decay: 10^{-5} , learning rate: 10^{-3} with 10% decrease every 10 epochs).

IV. RESULTS

A. Datasets Description and Tested Algorithms

The proposed algorithm is compared to the spectral version of baseline methods (cADI and PCA) as implemented in VIP [9], an open source package for HCI. Details about the setting of these methods can be found in [2]. We also used the unsupervised PACO ASDI [5] algorithm to ground the benefits of the deep learning stage embedded in the proposed algorithm. To evaluate the benefits of a joint spectral processing, we also compare the proposed method against PACO and deep PACO, that do not take benefits of the spectral diversity. The detection threshold is set to $\tau = 5$ for cADI, PCA, PACO (ASDI), which yields a SNR map as detection scores. We set $\tau = 0.5$ for deep PACO and the proposed algorithm, that produce a pseudo-probability map. We selected three datasets obtained with the SPHERE-IFS instrument ($L = 39$) by the observation of the following stars: β Pictoris (observed in 2018-09-17, *dataset 1*), β Pictoris (2018-12-15, *dataset 2*), and HR 8799 (2015-07-04, *dataset 3*). We also selected the eleven datasets from the SPHERE-IRDIS instrument we considered in our recent work on deep PACO [2]. SPHERE-IRDIS produces ASDI datasets but in dual band only (i.e., $L = 2$). In [2], we processed the first spectral channel solely. In this paper, we revisit these datasets with the proposed method and we compare our results with deep PACO. We exemplify detailed results for three of these eleven datasets: HIP 88399 (2015-05-10, *dataset 4*), HD 95086 (2021-03-11, *dataset 5*), and HD 95086 (2015-05-05, *dataset 6*).

B. Qualitative and Quantitative Performance

Figure 2 compares detection maps produced by the four multi-spectral algorithms we consider on the three datasets of SPHERE-IFS. The proposed algorithm, PCA and PACO ASDI are able to detect the five known exoplanets without any false alarm at the prescribed detection threshold. First line of Table I complements these results with AUC under ROCs built by varying the detection threshold. Even when the number of real sources is limited, these results emphasize that the proposed approach leads to a better trade-off between precision and recall than the other methods. We also resort to massive injection of simulated synthetic sources (considered in small batches) with various fluxes and positions to ground in more detail the performance of the proposed method. Second and third lines of Table I report the resulting AUC under ROCs




Fig. 2. Detection maps obtained with the proposed approach on considered SPHERE-IFS datasets comparatively to three state-of-the-art methods in the field. Sources are classified as true, missed and false detections based on the prescribed detection threshold. Due to the apparent rotation of the field induced by ASDI, the detection can be performed on a larger field of view (dashed blue line) than the spatial extent of the sensor (light blue line).




Fig. 3. Detection maps obtained with the proposed algorithm on dataset 3. Left: injected training sources have a flat SED. Right: SED of injected sources is based on atmospheric models. The detection threshold is set to $\tau = 0.5$, and sources are classified as true, missed and false detections.

for two regimes of angular separations (i.e., distance to the star and associated noise regimes). The proposed approach outperforms the other methods in both cases.

Figure 3 completes this study by comparing detection maps produced by the proposed algorithm using synthetic training




Fig. 4. ROCs showing the TPR as a function of the FDR averaged over eleven SPHERE-IRDIS datasets. Straight lines are for a joint multi-spectral analysis (two channels), while dashed line stands for a monochromatic analysis.

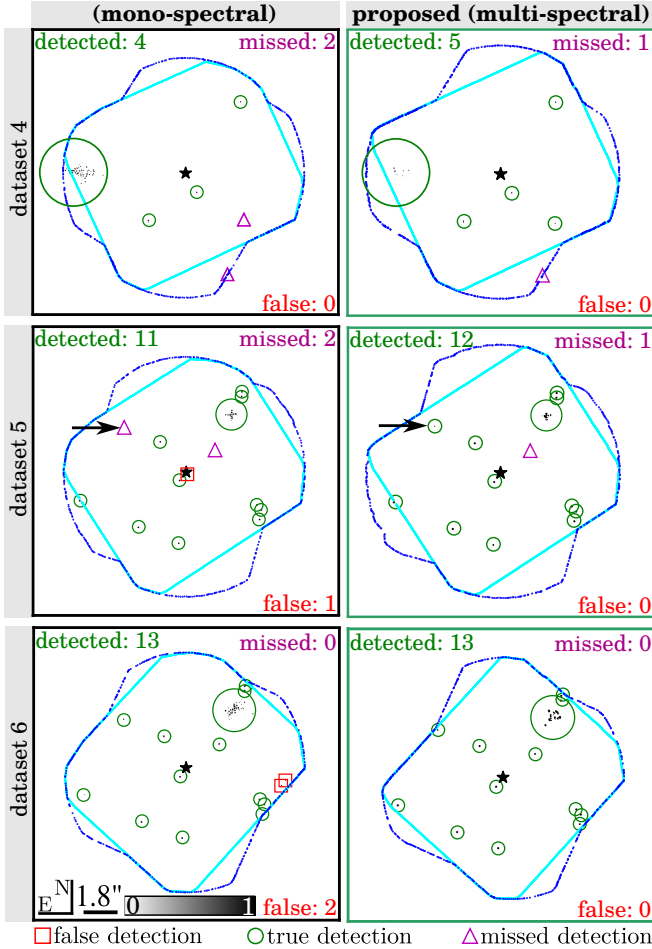


Fig. 5. Comparison of SPHERE-IRDIS detection maps obtained by the proposed algorithm (right) and its monochromatic version [2].

sources having either a flat SED or realistic SEDs based on physical ExoREM models, see Section III-C1. It emphasizes the importance of using reliable exoplanet SEDs since a flat SED (in contrast unit, i.e., same spectrum than the star) leads to additional detection peaks that we identified as false alarms based on the visual inspection of the detection maps and of the data themselves (i.e., the detection peaks are not static along the spectral axis or correspond to bright speckles).

Finally, we briefly discuss the benefits of joint processing

of multi-spectral data. Figure 4 gives ROCs averaged over the eleven SPHERE-IRDIS datasets (i.e., $L = 2$) containing 59 off-axis real known point-like sources in total. These datasets were processed with the four considered multi-spectral algorithms. Our results are also compared to a monochromatic analysis (on the first spectral channel) with PACO and deep PACO, as done in our recent work [2]. The detection maps from the datasets with the largest differences between mono- and multi-spectral processing are shown in Fig. 5. These results illustrate the benefits of a joint multi-spectral analysis since it allows to avoid some false alarms, especially near the star where the nuisance component and its fluctuations are the strongest, as well as far from the star where border effects occur with mono-spectral processing (the two spectral channels having a slightly different field of view). Besides, the point-like source marked by a black arrow on dataset 5 corresponds to a new faint source (likely a background star seen in the projected field of view) that we have consistently identified on two datasets of the same star (including dataset 6) with deep PACO in [2]. It is now consistently detected on a third (worse) observation (dataset 5) by the proposed approach while it was not with its monochromatic version [2].

V. CONCLUSION

We have proposed in this paper a new algorithm for exoplanet detection by the joint processing of spatio-temporal-spectral series of images. The algorithm includes a statistical modeling capturing most of the correlations of the data in a pre-processing step. A CNN is then trained to detect the residual signatures of injected synthetic sources. Our results show a better trade-off between precision and recall than state-of-the-art methods. The joint exploitation of the spectral diversity improves the performance of our recent work based on a similar framework for spatio-temporal data. The control of the uncertainty and the construction of a model from multiple datasets jointly will be addressed in future work.

REFERENCES

- [1] W.A. Traub and B.R. Oppenheimer, "Direct imaging of exoplanets," *Exoplanets*, 2010.
- [2] O. Flasseur, T. Bodrito, J. Mairal, J. Ponce, M. Langlois, and A.-M. Lagrange, "deep PACO: Combining statistical models with deep learning for exoplanet detection and characterization in direct imaging at high contrast," *minor rev. MNRAS*, <https://arxiv.org/abs/2303.02461>, 2023.
- [3] L. Pueyo, "Direct imaging as a detection technique for exoplanets," *Handbook of Exoplanets*, pp. 705–765, 2018.
- [4] O. Flasseur, L. Denis, É. Thiébaud, and M. Langlois, "Exoplanet detection in angular differential imaging by statistical learning of the nonstationary patch covariances - The PACO algorithm," *A&A*, 2018.
- [5] —, "PACO ASDI: an algorithm for exoplanet detection and characterization in direct imaging with integral field spectrographs," *A&A*, 2020.
- [6] Y. Chen, A. Wiesel, Y.C. Eldar, and A.O. Hero, "Shrinkage algorithms for MMSE covariance estimation," *IEEE TSP*, 2010.
- [7] B. Charnay, B. Bézard, J.-L. Baudino, M. Bonnefoy, A. Boccaletti, and R. Galicher, "A self-consistent cloud model for brown dwarfs and young giant exoplanets: comparison with photometric and spectroscopic observations," *AJ*, 2018.
- [8] L. Wang, C. Wang, Z. Sun, and S. Chen, "An improved Dice loss for pneumothorax segmentation by mining the information of negative areas," *IEEE Access*, 2020.
- [9] C.A Gomez Gonzalez, O. Wertz, O. Absil, V. Christiaens, D. Defrère *et al.*, "VIP: Vortex image processing package for high-contrast direct imaging," *AJ*, 2017.