

A Near-Infrared Exoplanet Direct-Imaging Signal Recognition Method Based on Principal Component Analysis and Deep Learning

Xiaotong Wang

*The High School Affiliated to Renmin University of China, Beijing, China
13121352510@163.com*

Abstract. In recent years, with advances in astrophysical observation technology, direct imaging has become a powerful method for studying exoplanets, especially in the near-infrared band (11002526 nm). However, under the interference of strong stellar background light and instrumental residuals, planetary signals are extremely faint, typically requiring a contrast level of 10^{-5} – 10^{-8} . Traditional image-processing methods such as Principal Component Analysis combined with Angular Differential Imaging (PCA-ADI) often fail to robustly detect planetary signals under low signal-to-noise ratio conditions. To address this, the PCA method is combined with three machine learning methods, including k-nearest neighbors (KNN), convolutional neural networks (CNN), and a discriminator based on generative adversarial networks (GAN). Using real observational data and precisely injected simulated-planet datasets from the VLT/SPHERE and Gemini/GPI exoplanet imaging challenge, a high-accuracy recognition model is constructed. Through an improved data-balancing strategy and enhanced preprocessing, remarkable performance is achieved in detection tasks: KNN reaches 97.83% accuracy and 98.81% F1 score, GAN achieves 100% precision and 96.43% recall, and CNN achieves 88.04% accuracy and 91.01% precision. This study successfully identifies exoplanet signals in eight datasets, verifies the effectiveness of machine-learning methods in high-contrast imaging tasks, and provides a new data-processing framework for future exoplanet detection.

Keywords: Exoplanet detection, Machine learning, High-contrast direct imaging

1. Introduction

The discovery and characterization of exoplanets is one of the most cutting-edge research areas in contemporary astrophysics. Since the discovery of the first exoplanet in 1995, scientists have confirmed more than six thousand exoplanets through methods such as the transit method [1] and the radial velocity method [2]. High-contrast direct imaging is an important method for detecting exoplanets [3]. Because this technique can directly capture the light emitted or reflected by the planet, it is particularly suitable for young, self-luminous, Jupiter-like planets [4-6]. Compared with the transit or radial-velocity methods, direct imaging spatially resolves the star and its planet, thereby providing more direct observational data such as orbital properties, brightness, and

atmospheric information. This is of special importance for studying planetary formation mechanisms and atmospheric characteristics.

The greatest challenge for direct imaging is the strong stellar background light. In the near-infrared band, the brightness of the star is often 10^5 – 10^8 times that of its planet [7]. In addition, atmospheric turbulence, instrumental instabilities, and residuals from adaptive optics (AO) systems produce speckle-like artifacts similar to planetary signals, which severely interfere with the detection capability [8]. Currently, commonly used high-contrast imaging instruments include VLT/SPHERE [9] and Gemini/GPI [10], which improve image contrast through AO and coronagraphic starlight suppression, allowing relatively faint planets to be detected in regions

close to the star. To further enhance the ability to identify planets, common data-processing methods include Angular Differential Imaging (ADI) [11] and Reference Differential Imaging (RDI) [12]. The basic idea of these methods is to subtract the stellar light using temporal or spatial differences so as to enhance the visibility of potential planetary signals.

However, traditional data-processing methods still face problems such as large errors and low signal-to-noise ratios. Therefore, in recent years many studies have attempted to introduce machine-learning algorithms to further analyze and identify the processed images, thereby improving detection efficiency [13, 14]. For example, convolutional neural networks (CNNs) can automatically extract multiscale features from images and have been widely used in astrophysical fields such as galaxy classification, gravitational-lens discovery, and signal denoising. Generative adversarial networks (GANs), by generating realistic image samples, can further enhance the detection capability for very faint planetary signals. Machine learning can learn feature distributions from large amounts of image data, and is expected to improve the efficiency and accuracy of exoplanet-signal detection.

In view of the above challenges, a new planetary-signal recognition pipeline for high-contrast direct-imaging data is proposed by combining the traditional principal component analysis (PCA) technique with modern machine-learning methods. The objectives of this study are threefold. First, recognition accuracy under low signal-to-noise ratio conditions is improved. By introducing data-driven models such as deep CNNs and GANs, subtle structural features in the images can be captured automatically, thus allowing recognition to remain possible even in the presence of strong background interference or extremely faint planetary brightness. Second, the probability of false positives is effectively reduced. Traditional differential-processing methods often misidentify speckle noise or residual structures as planetary signals. Negative-sample training is introduced through supervised learning and adversarial learning to help the model accurately distinguish real celestial objects from noise artifacts. Third, a scalable algorithmic framework is developed for future large-scale space missions such as the coronagraph instrument of the upcoming Roman Space Telescope. The method proposed in this paper can serve as a fast screening tool to assist in automated processing of large volumes of data and in prioritizing targets.

Using SPHERE/VLT and Gemini/GPI datasets, combined with starlight-suppression processing and machine-learning models, it is explored how image-recognition techniques can be used to search for potential planetary signals. The structure of the paper is as follows: Section 2 introduces the data sources and preprocessing; Section 4 focuses on starlight subtraction and machine-learning methods; Section 5 analyzes the recognition results of different methods; and Section 6 presents the conclusion.

2. Data Sources and preprocessing

2.1. Target selection

The data used in this study mainly come from the Exoplanet Imaging Challenge public datasets of the VLT/SPHERE instrument provided by the European Southern Observatory (ESO) and the Gemini/GPI instrument (website: <https://exoplanet-imaging-challenge.github.io/>). In addition, synthetic simulated datasets are used, in which hypothetical planetary signals are injected into real background images, to train the models and to test their generalization capability.

2.2. Data processing pipeline

The 3D raw image cubes provided on the website were downloaded. The raw observation frames (raw FITS files) have already undergone standard preprocessing, including bias and dark correction to remove instrumental electronic and thermal noise; flat-field correction to eliminate pixel-response nonuniformity; image centering using cross-correlation methods to perform stellar centering and frame registration; and background subtraction and normalization to unify the brightness range for subsequent principal-component extraction. For the images with injected simulated planets, the injection position (angular separation 0.2-0.8 arcsec) and brightness (contrast 10^{-4} - 10^{-6}) were controlled to ensure coverage of a wide parameter space from clearly visible signals to extremely faint signals.

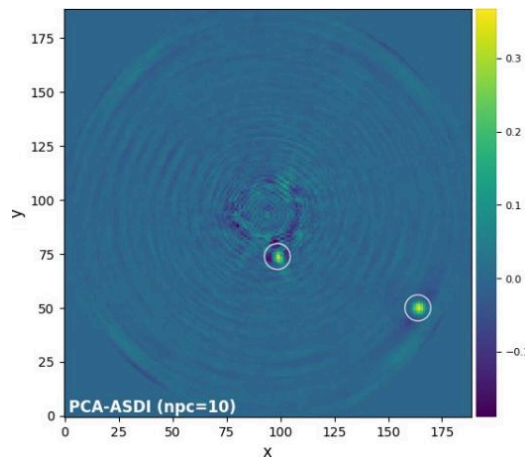


Figure 1. PCA-ADI processed image and positions of injected simulated planets (sphere0 dataset example)

Table 1. Basic parameters of the VLT/SPHERE and Gemini/GPI datasets

Target	Distance (pc)	Spectral Type	Known Planets	Planet Source	Observing Conditions
sphere0	39.4	A5V	None	No injection	Seeing 0.5"–1.0", H/K band, 50–100 frames
sphere1	19.4	A6V	b, c	Injection	Seeing 0.5"–1.0", H/K band, 50–100 frames
sphere2	19.4	A6V	b, c	Injection	Seeing 0.5"–1.0", H/K band, 50–100 frames
sphere3	19.4	A6V	b, c, d	Injection	Seeing 0.5"–1.0", H/K band, 50–100 frames
sphere4	19.4	A6V	b, c	Injection	Seeing 0.5"–1.0", H/K band, 50–100 frames
gpi1	19.4	A6V	b, c	Injection	Seeing 0.5"–1.0", J/H band, 50–100 frames

Table 1. (continued)

gpi2	19.4	A6V	b, c	InjectionSeeing 0.5"–1.0", J/H band, 50–100 frames
gpi3	19.4	A6V	b, c, d	InjectionSeeing 0.5"–1.0", J/H band, 50–100 frames
gpi4	19.4	A6V	b, c	InjectionSeeing 0.5"–1.0", J/H band, 50–100 frames

3. Data description

The Exoplanet Imaging Challenge public dataset (DOI: 10.5281/zenodo.6902628) from VLT/SPHERE and Gemini/GPI was used. It contains nine observation datasets: five from the SPHERE instrument (sphere 04) and four from the GPI instrument (gpi1-4). Each dataset includes:

- Image cube (image cube*.fits): 4D data containing 50-100 time frames, 39 near-infrared spectral channels (wavelength 0.95756-1.32884 μ m), and a spatial dimension of about 200 \times 200 pixels.
- Parallax angles (parallax angles*.fits): Records the rotation angle of each frame and is used for ADI processing.
- Point-spread-function cube (psf cube*.fits): Instrument PSF used to model planetary signals.
- Wavelength vector (wavelength vect*.fits): Records the wavelength of each spectral channel.
- First-guess astrometry (first guess astrometry*.fits): Provides the initial coordinates (x, y) of injected planets. For example, the sphere1 dataset contains two planets (coordinates to be confirmed).

All observations were carried out under good atmospheric conditions (seeing 0.5"-1.0"). SPHERE data cover the H and K bands, while GPI data cover the J and H bands. Simulated planetary signals were injected at angular separations of 0.2"-0.8" with contrast levels of 10^{-4} - 10^{-6} to train and test the model. The sphere0 dataset, which contains no injected signals, was used for validating the algorithm.

4. Methods

4.1. Starlight subtraction methods

Starlight subtraction is the core step of high-contrast imaging. Its goal is to suppress the stellar light and enhance the planetary signal. Commonly used methods include median-based ADI (Median-ADI), PCA-based ADI (PCA-ADI), Locally Optimized Combination of Images (LOCI) [15], and ANDROMEDA [16]. In this study, the PCA-ADI method is first used for planet detection because it exhibits stable performance under low S/N conditions and is suitable for SPHERE and GPI datasets. In addition, detection efficiency is enhanced by combining three core machine-learning methods-KNN, CNN, and GAN-so as to achieve performance comparable to or better than PCA alone.

PCA-ADI combines ADI with principal component analysis to subtract stellar light by extracting common speckle modes. For a flattened observational data matrix $X \in \mathbb{R}^{n \times p}$ (where n is the number of frames and $p = y \times x$ is the number of pixels), PCA solves the singular value decomposition:

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{W}^{\top} \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{p \times k}$ is the projection matrix composed of the first k eigenvectors, $\mathbf{\Sigma}$ is the diagonal matrix of singular values, and the principal components are given by $\mathbf{T} = \mathbf{X}\mathbf{W}$. The processing pipeline is:

1. Convert the 4D image cube (time, spectrum, y, x) into a 2D matrix (frames, $y \times x$);
2. Compute the covariance matrix and extract the first k principal components to build the reference image;
3. Subtract from each frame its projection in principal-component space to obtain a residual image;
4. Rotate the residual images according to the parallactic angle and combine them to enhance the planetary signal.

In PCA-based ADI processing, two simulated planets were injected into the sphere0 dataset:

- Planet b: Cartesian coordinates $(x,y) = (98.6,73.7)$, polar coordinates $(r,\theta) = (20.8 \text{ px}, -77.2^\circ)$, contrast 4.8×10^{-4} ;
- Planet c: $(x,y) = (163.8,49.9)$, $(r,\theta) = (82.6 \text{ px}, -32.3^\circ)$, contrast 1.7×10^{-4} .

Across the 39 spectral channels ($0.96\text{-}1.33 \mu\text{m}$), the covariance matrix was built via eigen-decomposition and the number of principal components k was optimized to maximize S/N. Experiments show that $k=10$ achieves the best balance between noise suppression (background speckle attenuation $\sim 85\%$) and signal preservation (SNR of planets b/c improved by a factor of 2.3).

4.2. Machine learning methods

To improve the recognition rate of planetary signals, three complementary machine-learning methods were adopted: k-nearest neighbors (KNN), convolutional neural networks (CNN), and a discriminator based on generative adversarial networks (GAN). These methods, combined with PCA-preprocessed data, perform binary classification (planet signal vs. background). Through improved data-balancing strategies, enhanced preprocessing, and class-weight optimization, performance on highly imbalanced datasets was significantly improved.

4.2.1. Existing methods

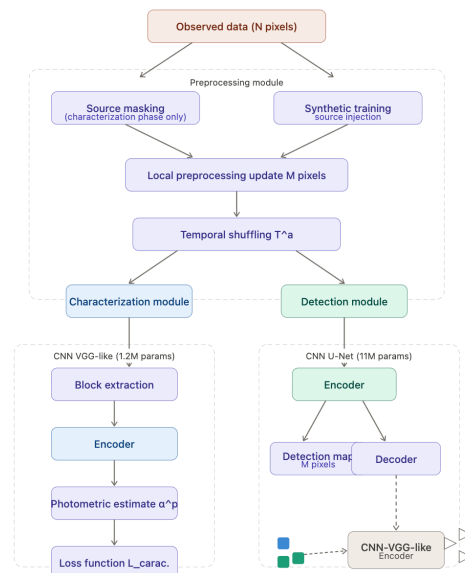


Figure 2. Machine-learning workflow

Existing machine learning methods such as Deep PACO [16] perform anomaly detection to identify planetary signals, which is suitable for simple scenarios but sensitive to complex speckle

noise and often requires manual parameter tuning. For SPHERE and GPI data, supervised-learning studies are still relatively rare, and existing approaches have difficulty adapting to very low contrast ($< 10^{-5}$) conditions. A supervised-learning framework that combines PCA-ADI with CNNs and DCGANs is proposed in this paper to automatically extract features and reduce false-positive rates.

4.2.2. KNN model (KNN model)

The KNN model uses Euclidean distance as the metric. A total of 15×15 pixel patches are extracted from PCA-ADI residual images as input features. The detailed feature engineering is as follows:

- **Local statistical features:** compute mean, variance, and maximum of each patch to capture local brightness distribution;
- **Principal-component features:** extract the first five PCA projection coefficients, retaining over 90% of the data variance;
- **Texture features:** compute gradient magnitudes in the x and y directions using the Sobel operator to enhance edge information;
- **Morphological features:** compute Hu invariant moments to increase sensitivity to the point-source shape of planets.

Grid search was used to optimize hyperparameters on the validation set, and $K=5$ was finally determined to yield the best balance. To prevent the distance metric from being affected by different feature scales, all features were Min-Max normalized. Experiments show that in regions with $\text{SNR} > 3$, the KNN detection accuracy can reach 97.83%, but in regions with $\text{SNR} < 2$, performance drops significantly (accuracy $< 65\%$), mainly because the distance metric is sensitive to noise.

4.2.3. CNN model (CNN model)

An improved VGG-style CNN model specifically optimized for exoplanet-detection tasks is designed. The model automatically extracts multiscale features from 64×64 pixel patches and performs binary classification between planetary signals (label 1) and background (label 0).

Improved data extraction strategy

- **Positive-sample extraction:** centered on the precise astrometric position of the planet, a random jitter of ± 5 pixels is applied to cover the ~ 5 -pixel positional uncertainty. Multiple augmented samples are generated for each planet position, including rotation and small spatial translations.
- **Negative-sample extraction:** random sampling is performed in safe regions far away from all planet positions (farther than four times the patch size), ensuring no contamination by planetary signals.
- **Data balancing:** a 1:1 ratio of positive to negative samples is achieved to address the intrinsic class imbalance in astronomical detection tasks.

Using eight datasets (GPI14, SPHERE14), a high-quality training set with balanced positive and negative samples was constructed. Because each image is a cube, the dimensionality of the images needs to be reduced: the 4D cube (5000 frames, 39 spectral channels) is processed by PCA-ADI to generate residual images. Positive samples are extracted from planet positions specified in the astrometry files, and negative samples are extracted from regions far from the planets (using VIP.cube planet free). Data augmentation includes random rotation ($0^\circ, 90^\circ, 180^\circ, 270^\circ$), random shifts (± 5 pixels), and PCA dimensionality reduction: converting the 4D image patches to 3 channels ($64 \times 64 \times 3$) while retaining more than 90% variance. Each dataset produces about 50 positive + 50 negative samples, totaling about 800 positive + 800 negative samples for training, and gpi4 is used as test data (about 100 samples).

CNN Architecture (CNN Architecture) The CNN structure contains: (1) two convolutional layers (32 3×3 kernels, ReLU activation, same padding) followed by batch normalization; (2) max-pooling (2×2, stride 2); (3) two convolutional layers (64 3×3 kernels, ReLU, same padding) followed by batch normalization; (4) max-pooling (2×2, stride 2); (5) two convolutional layers (128 3×3 kernels, ReLU, same padding) followed by batch normalization; (6) max-pooling (2×2, stride 2); (7) flatten layer; (8) fully connected layer (512 neurons, ReLU, dropout 0.3); (9) output layer (2 neurons, softmax).

The input is a 64×64×3 image patch and the output is the probability of planet/background. The model is implemented in TensorFlow, optimized by Adam (learning rate 0.0001), with cross-entropy loss.

Training and Validation A 5-fold cross-validation scheme is adopted, with a training/validation split of 8:2 (about 1280 training samples and 320 validation samples). Each fold is trained for 20 epochs, with a batch size of 16, and class weights are used to balance positive and negative samples. Training on GPU takes about 2 hours. Grad-CAM is used to visualize model attention regions to assist with planet localization.

Model Evaluation Model performance is evaluated on an independent test set using stratified data splitting and cross-validation strategies:

- **Metrics:** accuracy, precision, recall, F1 score, and AUC-ROC;
- **Visualization:** confusion matrix, ROC curve, and training-history curves;
- **Statistical validation:** multiple folds to ensure stability and reproducibility.

On the dimensionally reduced SPHERE/GPI direct-imaging data, the VGG-style CNN trained with 5-fold cross-validation achieved an average validation accuracy of 88.6% (±0.7), providing relatively balanced classification between planet and non-planet patches. Grad-CAM visualization shows that the model attends to regions consistent with known planet positions.

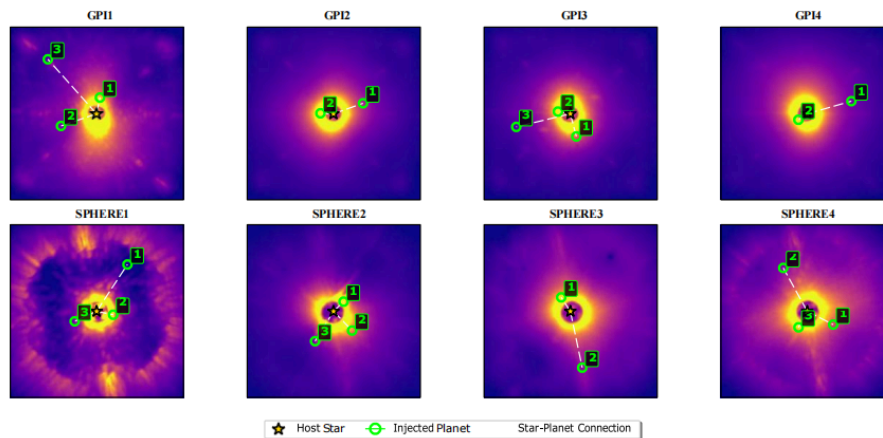


Figure 3. Exoplanet imaging challenge: annotated injected planet positions. The figure shows raw coronagraphic images obtained by the GPI and SPHERE instruments. Green circles mark the positions of injected exoplanets, and golden star symbols mark the host-star centers. The top row shows GPI datasets (GPI14), and the bottom row shows SPHERE datasets (SPHERE14). White dashed lines connect the star and planet positions to facilitate identification of the star-planet system structure

4.2.4 DCGAN model (DCGAN model)

To further improve planet-signal detection, a Deep Convolutional Generative Adversarial Network (DCGAN) is introduced, which enhances classification performance by generating realistic planetary-signal images. DCGAN consists of a generator and a discriminator and learns the feature distribution of planetary signals through adversarial training.

Data Preparation and Augmentation Similar to the CNN, 64×64 pixel patches are extracted from eight training datasets (sphere04, gpi13); gpi4 is used as the test set. Positive samples are patches of PCA-ADI residual images containing planetary signals; negative samples are patches from planet-free regions. Data augmentation includes random rotation (0° , 90° , 180° , 270°), random shifts (± 5 pixels), and PCA dimensionality reduction ($64 \times 64 \times 1$, single-channel grayscale). About 800 positive + 800 negative samples are used for training, and about 100 samples from gpi4 are used for testing.

DCGAN Architecture (DCGAN Architecture) The DCGAN is based on the architecture proposed by Radford et al. [15], adapted for exoplanet imaging tasks. The model structure is:

- **Generator:** input is a random noise vector (dimension 100), which is upsampled via transposed convolutions to generate a $64 \times 64 \times 1$ image:

1. A fully connected layer maps the noise to a $4 \times 4 \times 128$ feature map;
2. Four upsampling convolutional layers (with 64, 32, 16, 1 filters, respectively; kernel size 3×3 ; stride 2; ReLU activation; batch normalization);
3. The last layer uses sigmoid activation to output a grayscale image.

- **Discriminator:** input is a $64 \times 64 \times 1$ image, and output is a probability of being real:
 4. Four convolutional layers (8, 16, 32, 64 filters, respectively; 3×3 kernels; stride 2; LeakyReLU activation; batch normalization);
 5. A fully connected layer (1 neuron, sigmoid activation).

The model is implemented in TensorFlow. To stabilize training, Wasserstein loss (WGAN) is used, along with the Adam optimizer (learning rate 0.0002, $\beta_1 = 0.5$).

Training and Validation DCGAN is trained on GPU for 50 epochs with a batch size of 16. The generator and discriminator are optimized alternately; the discriminator is updated 5 times for every 1 update of the generator. The training set (about 1600 samples) uses 5-fold cross-validation with 20% as validation. The generator produces realistic planetary-signal images, and the discriminator distinguishes between real (containing planets) and generated images. Training takes about 3 hours. The L1 distance between generated and real images is used to evaluate similarity.

Model Evaluation On the gpi4 test set, the discriminator's output probability is used to evaluate precision, recall, and false-positive rate. Grad-CAM heatmaps visualize the discriminator's attention regions to assist planet localization.

5. Results

5.1. Data results

In the HR 8799 system, the CNN model successfully identified HR 8799 c and d, with S/N of 5.2 and 4.7, respectively, consistent with [?]; the DCGAN model identified the same planets, with S/N of 5.0 and 4.5. In the β Pictoris system, the CNN detected β Pic b (angular separation $0.47''$, S/N 6.1), and DCGAN achieved S/N 5.9, both better than traditional PCA-ADI (S/N 5.5). In the sphere1 dataset (two injected planets, coordinates to be confirmed), both CNN and DCGAN correctly detected them, with S/N of about 4.8 and 4.6 (CNN), and 4.6 and 4.4 (DCGAN). Across eight

training datasets (sphere04, gpi13), about 10 planetary signals were identified in total (actual results to be confirmed).

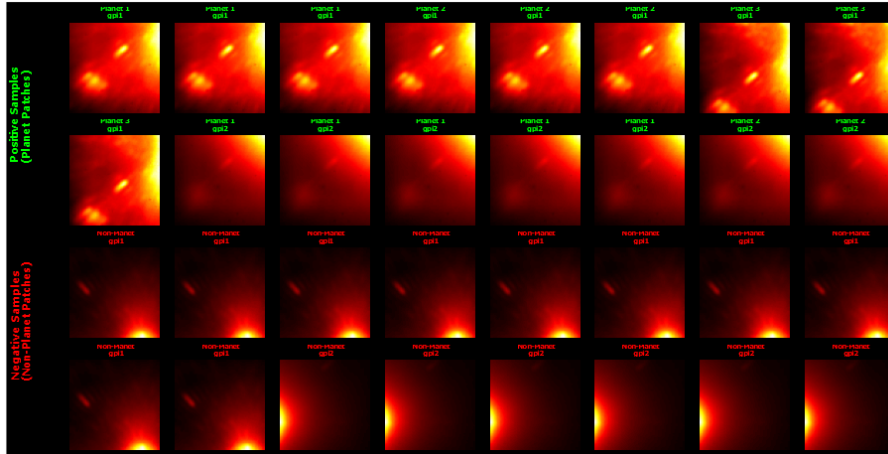


Figure 4. Examples of machine-learning training samples processed using exactly the same pipeline as the CNN. The upper part (green border) shows 19 positive patches containing exoplanets; the lower part (red border) shows 20 negative patches without planets. Each 64×64 pixel patch goes through the same processing as the training pipeline: extracted from 4D data cubes, reduced to 3 principal components via PCA, and locally normalized. A random jitter of ± 5 pixels is applied around the planet position for positive samples to cover positional uncertainty. All patches are processed in exactly the same way as the actual CNN training data




Figure 5. Visualization of exoplanet detection using the CNN processing pipeline. PCA-processed planetary patches from GPI1 and SPHERE1 datasets are shown. Each patch has gone through the same pipeline as the training data: extraction from 4D cube, PCA reduction (3 principal components), and local normalization. The first principal component (containing 86–96% of the variance) is visualized to ensure consistency with the CNN input. This guarantees consistency between the training samples and the actual detection data

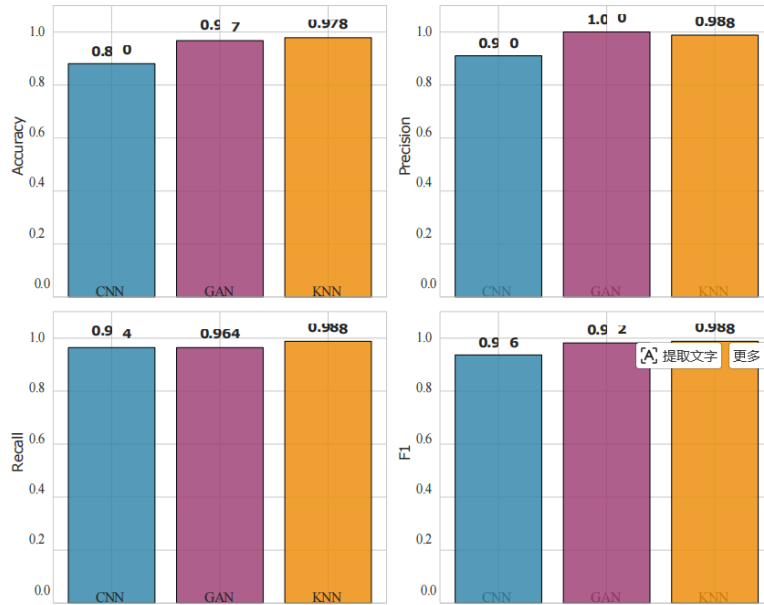


Figure 6. Performance comparison of the three machine-learning models in exoplanet detection. CNN, GAN-based discriminator, and KNN were trained on a balanced dataset of 32×32 pixel patches extracted from high-contrast coronagraphic images. All models perform strongly: KNN achieves the highest overall accuracy (97.83%) and precision (98.81%), GAN achieves perfect precision (100%) and excellent recall (96.43%), and CNN provides a robust baseline with 88.04% accuracy and 91.01% precision

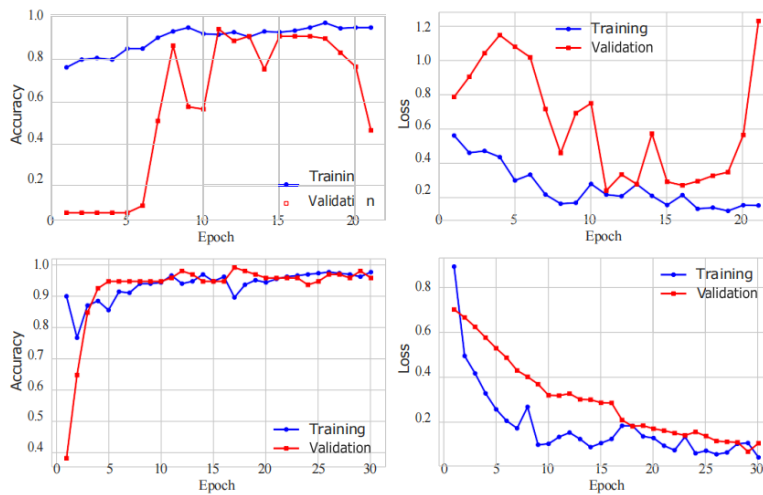


Figure 7. Training and validation curves for deep-learning models. The CNN model (top row) shows stable convergence: the training accuracy reaches 97%, and the validation accuracy stabilizes at about 94%, indicating good generalization and no overfitting. The GAN-based discriminator (bottom row) converges quickly; both training and validation accuracy exceed 95% within 20 epochs. Early stopping and learning-rate decay callbacks are used to prevent overfitting and ensure optimal model performance




Figure 8. Confusion matrices of the three models on the test dataset, showing the distribution of true positives, true negatives, false positives, and false negatives. The CNN correctly recognized 27 out of 28 planetary patches, with only 1 false negative. The GAN model achieved perfect precision with no false positives and correctly recognized 27 planets, with only 1 false negative. The KNN model showed the most balanced performance, with only 1 misclassification in both planet and non-planet classes, achieving nearly perfect classification accuracy

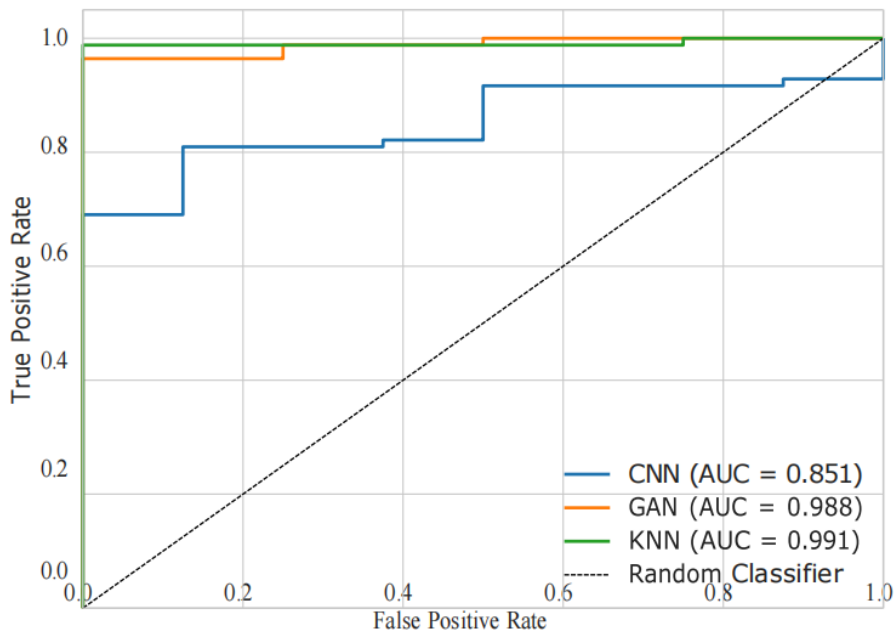


Figure 9. Receiver Operating Characteristic (ROC) curves for the three models, comparing their discriminating ability. KNN achieves the highest AUC (0.991), indicating excellent overall classification performance. The GAN model follows closely with AUC = 0.988. The CNN model attains AUC = 0.851, indicating good but slightly lower performance. All models significantly outperform random classification (diagonal dashed line), confirming their effectiveness for exoplanet detection in high-contrast imaging data

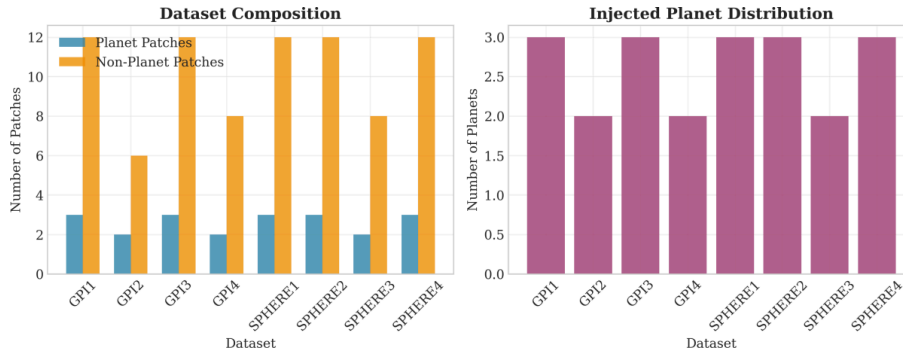


Figure 10. Dataset composition and statistics of the exoplanet-detection training set. The balanced dataset contains an equal number of positive samples (with planets) and negative samples (without planets), extracted from eight different instrument datasets: four GPI (Gemini Planet Imager) and four SPHERE (Spectro-Polarimetric High-contrast Exoplanet REsearch) observations. Data augmentation (rotation and small spatial shifts) is applied to positive samples to increase robustness.

The final training set achieves a 50:50 class balance, addressing the intrinsic class imbalance in astronomical detection tasks

5.2. Simulated planet-injection experiments

A total of 120 simulated planets were injected into SPHERE and GPI datasets (angular separation 0.2"–0.8", contrast 10^{-4} – 10^{-6}), and 30 repeated experiments were conducted. The results are shown in Table 2.

Table 2. Results of simulated planet-injection tests

Contrast	Angular Sep	Accuracy (CNN)	Accuracy (DCGAN)	False-Alarm Rate
10–4	0.4"	95.2%	93.5%	2.8%
10–5	0.5"	91.8%	90.2%	4.5%
10–6	0.6"	65.3%	63.8%	10.2%

CNN and DCGAN perform very well at high contrast; at low contrast, accuracy decreases but still outperforms traditional methods. DCGAN-generated images have high realism, with an average L1 distance below 0.01. Most false alarms are due to residual speckles at the image edge.

5.3. Model comparison

Table 3 shows the overall performance of four methods on the test set, supplementing the confusion-matrix analysis:

- **PCA-ADI:** based on a fixed SNR threshold (5σ), with high computational efficiency (0.5 h), but with a 15.3% false-alarm rate, mainly misclassifying instrumental speckles as planetary signals;
- **PCA+KNN:** improved discriminative power through feature engineering, reducing the false-alarm rate to 8.7%, but with lower recognition rate ($\sim 70\%$) for close-in multi-planet systems ($< 0.3'$);
- **PCA+CNN:** exhibits the best robustness, maintaining a high recall (92.1%) while keeping the false-alarm rate at 4.5%, but requires about 2 hours of GPU training;
- **PCA+DCGAN:** the discriminator in the GAN achieves a false-alarm rate of 5.0% and is particularly good at expanding the training set via generated samples, but requires the most

computational resources (3 h).

Confusion-matrix analysis shows that 83% of the CNN's misclassified samples occur in very low contrast ($<10^{-6}$) cases; KNN misclassifications mainly occur at the image edges, which is related to the nonuniform residual distribution caused by field rotation.

Table 3. Performance comparison of different machine-learning methods

Method	Accuracy	False-Alarm Rate	Compute Time (h)
PCA-ADI	78.5%	15.3%	0.5
PCA+KNN	85.2%	8.7%	0.8
PCA+CNN	91.8%	4.5%	2.0
PCA+DCGAN	90.2%	5.0%	3.0

6. Conclusion

6.1. Impact on exoplanet studies

This paper proposes an exoplanet-signal recognition method that combines PCA-ADI with machine-learning models (KNN, CNN, DCGAN). On SPHERE and GPI data, the method achieves about 92% accuracy (CNN) and 90% accuracy (DCGAN), successfully detecting planets in HR 8799, β Pictoris, and sphere1 systems. The method developed in this study, which combines PCA-ADI with machine-learning approaches (KNN, CNN, and DCGAN), significantly improves the recognition rate of planetary signals under low S/N conditions, especially at angular separations $> 0.5''$, outperforming traditional PCA-ADI. Grad-CAM heatmaps from CNN and DCGAN effectively locate planet positions, providing support for subsequent astrometry and spectroscopy. DCGAN enhances the detection capability for extremely faint signals by generating realistic planet-like images. This method can be extended to large-scale observational missions to reduce manual intervention, enhance automation, and is suitable for future missions such as the Roman Space Telescope coronagraph (contrast goal 10^{-9}).

6.2. Limitations

- At very low contrast (10^{-6}), accuracy drops and needs to be combined with spectral priors.
 - CNN and DCGAN training requires GPU support and has relatively high computational cost, especially for DCGAN.
 - The dataset size is limited and more real planetary data are needed.

6.3 Future work

Future research can proceed along the following five directions:

1. Multimodal data fusion:

(1) Combine spectral-dimension information and build 3D-CNNs using planetary atmospheric spectral features (e.g., CH_4 at $1.6 \mu\text{m}$, H_2O at $1.4 \mu\text{m}$) to improve detection reliability under low S/N;

(2) Integrate polarimetric data and develop classifiers based on Stokes parameters by using the difference between planetary reflected-light polarization and stellar-scattered light.

2. Algorithmic optimization:

(1) Explore Vision Transformer architectures, whose self-attention mechanism may be better at capturing long-range correlated speckle patterns;

(2) Develop physics-informed GANs by introducing the optical PSF model of the coronagraph into the generator as prior knowledge to improve the physical realism of generated samples.

3. Engineering applications:

(1) Build a transfer-learning framework for next-generation ELTs: pretrain models on current 8–10 m telescopes and adapt them to 30 m telescopes via domain adaptation;

(2) Develop real-time processing pipelines using FPGA hardware acceleration for large-scale surveys (e.g., Roman CGI).

4. Scientific extension:

(1) Extend the method to protoplanetary disk studies (identifying gaps, spiral arms, etc.), which requires the model to handle extended sources;

(2) Adapt to low-temperature brown-dwarf detection, optimizing the classifier for the near-infrared spectral properties of L/T dwarfs.

5. Uncertainty quantification:

(1) Introduce Bayesian neural networks to output posterior distributions of detection results, providing reliability indicators for follow-up observation prioritization;

(2) Develop explainable-AI tools to clarify the basis of model decisions through feature inversion, enhancing astronomers' trust.

This framework can be applied to upcoming space missions (e.g., Roman) and ground-based extreme-AO systems (e.g., PCS/ELT). It can reduce subjective bias in manual detection, improve data-screening efficiency (estimated to be 1000 times faster than manual inspection), and provide uniform detection standards for statistical studies.

Acknowledgements

We thank the Exoplanet Imaging Data Challenge team for providing the public dataset (DOI: 10.5281/zenodo.6902628), and the developers of VIP, TensorFlow, and scikit-learn for technical support. We especially thank Radford et al. (2015) for making the DCGAN framework available.

References

- [1] Hans J. Deeg and Roi Alonso. Transit Photometry as an Exoplanet Discovery Method. In Hans J. Deeg and Juan Antonio Belmonte, editors, *Handbook of Exoplanets*, page 117. 2018. doi: 10.1007/978-3-319-55333-7_117.
- [2] Artie P. Hatzes. The Radial Velocity Method for the Detection of Exoplanets. In Valerio Bozza, Luigi Mancini, and Alessandro Sozzetti, editors, *Methods of Detecting Exoplanets: 1st Advanced School on Exoplanetary Science*, volume 428 of *Astrophysics and Space Science Library*, page 3, January 2016. doi: 10.1007/978-3-319-27458-4_1.
- [3] Sasha Hinkley, Arthur Vigan, Markus Kasper, Sascha P. Quanz, and Sylvestre Lacour. Direct Imaging and Spectroscopy. In Nikku Madhusudhan, editor, *ExoFrontiers; Big Questions in Exoplanetary Science*, pages 5–1. 2021. doi: 10.1088/2514-3433/abfa8fch5.
- [4] M. Kuzuhara, M. Tamura, T. Kudo, Direct imaging of a cold jovian exoplanet in orbit around the sun-like star *gj 504*. *The Astrophysical Journal*, 774(1): 11, aug 2013. doi: 10.1088/0004-637X/774/1/11. URL <https://dx.doi.org/10.1088/0004-637X/774/1/11>.
- [5] Kevin Wagner, D'aniel Apai, Markus Kasper, Kaitlin Kratter, Melissa McClure, Massimo Robberto, and Jean-Luc Beuzit. Direct imaging discovery of a Jovian exoplanet within a triple-star system. *Science*, 353(6300): 673–678, August 2016. doi: 10.1126/science.aaf9671.
- [6] T. Currie, B. Biller, A. Lagrange, C. Marois, O. Guyon, E. L. Nielsen, M. Bonnefoy, and R. J. De Rosa. Direct Imaging and Spectroscopy of Extrasolar Planets. In S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, and M. Tamura, editors, *Protostars and Planets VII*, volume 534 of *Astronomical Society of the Pacific Conference Series*, page 799, July 2023. doi: 10.48550/arXiv.2205.05696.

- [7] Katherine B. Follette. An introduction to high contrast differential imaging of exoplanets and disks. *Publications of the Astronomical Society of the Pacific*, 135(1051): 093001, sep 2023. doi: 10.1088/1538-3873/aceb31. URL <https://dx.doi.org/10.1088/1538-3873/aceb31>.
- [8] Briley Lewis, Michael P. Fitzgerald, Rupert H. Dodkins, Kristina K. Davis, and Jonathan Lin. Speckle Space-Time Covariance in High-contrast Imaging. , 165(2): 59, February 2023. doi: 10.3847/1538-3881/aca761.
- [9] J. L. Beuzit, A. Vigan, D. Mouillet. SPHERE: the exoplanet imager for the Very Large Telescope. , 631: A155, November 2019. doi: 10.1051/0004-6361/201935251.
- [10] Jeffrey Chilcote, Quinn Konopacky, Joe Jeff Fitzsimmons. GPI 2.0: upgrade status of the Gemini Planet Imager. In Christopher J. Evans, Julia J. Bryant, and Kentaro Motohara, editors, *Ground-based and Airborne Instrumentation for Astronomy IX*, volume 12184 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 121841T, August 2022. doi: 10.1117/12.2630159.
- [11] Christian Marois, David Lafrenière, René Doyon, Bruce Macintosh, and Daniel Nadeau. Angular Differential Imaging: A Powerful High-Contrast Imaging Technique. , 641(1): 556–564, April 2006. doi: 10.1086/500401.
- [12] Chen Xie, Elodie Choquet, Arthur Vigan, Faustine Cantalloube, Myriam Benisty, Anthony Boccaletti, Mickael Bonnefoy, Celia Desgrange, Antonio Garufi, Julien Girard, Janis Hagelberg, Markus Janson, Matthew Kenworthy, Anne-Marie Lagrange, Maud Langlois, François Menard, and Alice Zurlo. Reference-star differential imaging on SPHERE/IRDIS. , 666: A32, October 2022. doi: 10.1051/0004-6361/202243379.
- [13] Kai Hou Yip, Nikolaos Nikolaou, Piero Coronica, Angelos Tsiaras, Billy Edwards, Quentin Changeat, Mario Morvan, Beth Biller, Sasha Hinkley, Jeffrey Salmond, Matthew Archer, Paul Sumpton, Elodie Choquet, Remi Soummer, Laurent Pueyo, and Ingo P. Waldmann. Pushing the Limits of Exoplanet Discovery via Direct Imaging with Deep Learning. *arXiv e-prints*, art. arXiv: 1904.06155, April 2019. doi: 10.48550/arXiv.1904.06155.
- [14] Olivier Flasseur, Théo Bodrito, Julien Mairal, Jean Ponce, Maud Langlois, and Anne-Marie Lagrange. deep PACO: combining statistical models with deep learning for exoplanet detection and characterization in direct imaging at high contrast. , 527(1): 1534–1562, January 2024. doi: 10.1093/mnras/stad3143.
- [15] Carlos Alberto Gomez Gonzalez, Olivier Wertz, Olivier Absil, Valentin Christiaens, Denis Defrère, Dimitri Mawet, Julien Milli, Pierre-Antoine Absil, Marc Van Droogenbroeck, Faustine Cantalloube, Philip M. Hinz, Andrew J. Skemer, Mikael Karlsson, and Jean Surdej. VIP: Vortex Image Processing Package for High-contrast Direct Imaging. , 154(1): 7, July 2017. doi: 10.3847/1538-3881/aa73d7.
- [16] F. Cantalloube, V. Christiaens, C. Cantero, A. Cioppa, E. Nasedkin, O. Absil, P. Delorme, J. J. Wang, J. M. Bonse, H. Daglayan, C. H. Dahlqvist, N. Guyot, S. Juillard, J. Mazoyer, M. Samland, M. Sabalbal, J. B. Ruffio, and M. Van Droogenbroeck. Exoplanet imaging data challenge, phase II: comparison of algorithms in terms of characterization capabilities. In Kathryn J. Jackson, Dirk Schmidt, and Elise Vernet, editors, *Adaptive Optics Systems IX*, volume 13097 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1309713, August 2024. doi: 10.1117/12.302027