

ExoDNN: Boosting exoplanet detection with artificial intelligence

Application to *Gaia* Data Release 3

A. Abreu^{1,2,*}, J. Lillo-Box³, A. M. Perez-Garcia⁴, J. Sahlmann⁵, J. H. J. de Bruijne⁶, and C. Cifuentes³

¹ ATG Science & Engineering for the European Space Agency (ESA), ESAC, Spain

² Universidad Complutense de Madrid, Av. Complutense, s/n, Moncloa - Aravaca, 28040 (Madrid), Spain

³ Centro de Astrobiología (CAB, CSIC-INTA), ESAC campus, 28692, Villanueva de la Cañada (Madrid), Spain

⁴ ISDEFE, Beatriz de Bobadilla, 3. 28040 Madrid, Spain

⁵ European Space Agency (ESA), European Space Astronomy Centre (ESAC), Camino Bajo del Castillo s/n, 28692 Villanueva de la Cañada, Madrid, Spain

⁶ European Space Agency (ESA), European Space Research and Technology Centre (ESTEC), Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands

Received 20 May 2025 / Accepted 1 October 2025

ABSTRACT

Context. Transit and radial velocity (RV) techniques are the dominant methods for exoplanet detection, while astrometric exoplanet detections have been very limited thus far. *Gaia* has the potential to radically change this picture, enabling astrometric detections of substellar companions at scale that would allow us to complement the picture of exoplanet architectures given by transit and RV methods.

Aims. Our primary objective in this study is to enhance the current statistics of substellar companions, particularly within regions of the orbital period–mass parameter space that remain poorly constrained by RV and transit detection methods.

Methods. Using supervised learning, we trained a deep neural network (DNN) to recognise the characteristic distribution of the fit quality statistics corresponding to a *Gaia* Data Release 3 (DR3) astrometric solution for a non-single star. We created a deep learning model, ExoDNN, which predicts the probability of a DR3 source to host unresolved companions.

Results. Applying the predictive capability of ExoDNN to a volume-limited sample ($d < 100$ pc) of F, G, K, and M stars from *Gaia* DR3, we have produced a list of 7414 candidate stars hosting companions. The stellar properties of these candidates, such as their mass and metallicity, are similar to those of the *Gaia* DR3 non-single-star sample. We also identified synergies with future observatories, such as PLATO, and we propose a follow-up strategy with the intention of investigating the most promising candidates among those samples.

Key words. astrometry – planets and satellites: detection – binaries: general

1. Introduction

The current census of extrasolar planets contains more than 7000 objects¹, primarily compiled via transit and/or radial velocity (RV) detection methods. These techniques have helped expand our understanding of key physical properties of exoplanets, such as their radius and mass, enabling studies of their bulk density and interior structure (e.g., Seager et al. 2007). However, both the transit and RV methods suffer from known observational biases that favor the detection of close-in orbital companions (e.g., Fischer et al. 2014). To reach longer period orbits, the astrometric detection method can be used (e.g., Binnendijk 1960; Quirrenbach 2011). This involves detecting the reflex motion (wobble) of a star due to an orbiting companion (e.g., exoplanet). The size of this wobble, also known as astrometric orbit signature, scales inversely with distance; therefore, even nearby systems ($d < 100$ pc) yield very small signatures (sub milli-arcsecond). This has resulted in just a handful of brown dwarf or exoplanet detections using astrometry to date (Benedict et al. 2002; Pravdo et al. 2005;

Sahlmann et al. 2013; Curiel et al. 2022). To this end, the ESA *Gaia* telescope (Gaia Collaboration 2016) has provided a unique opportunity for the astrometric detection and characterization of substellar companions to stars. Thanks to its unprecedented precision (end-of-nominal mission estimates for parallax and proper motion are $5 \mu\text{as}$ and $3.5 \mu\text{as yr}^{-1}$, respectively, for $G = 3\text{--}12$ mag de Bruijne et al. 2014), *Gaia* is expected to detect possibly thousands of new planetary and brown dwarf companions in long period orbits down to at least the giant planetary mass regime (Casertano et al. 2008; Sozzetti et al. 2014; Perryman et al. 2014) and up to a few hundred parsec over the whole sky.

In *Gaia* Data Release 3 (Gaia Collaboration 2023c, hereafter, DR3), the *Gaia* non-single-star pipelines (Halbwachs et al. 2023; Holl et al. 2023) have produced $\sim 170\,000$ orbital solutions for binary systems, providing a first glimpse of the great potential of *Gaia* for stellar multiplicity analysis (Gaia Collaboration 2023a). A small subset of these solutions (~ 1800) correspond to sources with companion candidates in the substellar mass regime ($m < 80 M_{\text{Jup}}$) (see Sections 5 and 8.1 in Gaia Collaboration 2023a). However, as these authors remark, *Gaia* DR3 processing limitations do not allow new detections in the planetary mass regime ($m < 20 M_{\text{Jup}}$) to be claimed solely on the basis of astrometry. The next major *Gaia* data release, *Gaia* DR4, is expected to

* Corresponding author: Asier.Abreu@ext.esa.int;
asabre01@ucm.es

¹ Statistics as of early 2025 from: <https://exoplanet.eu/home/>

reach at least an order of magnitude higher astrometric precision than *Gaia* DR3, enabling the expansion of the sensitivity down to the planetary mass regime. Until then, *Gaia* DR3 can provide important insights onto the demographics of the substellar mass range covering $20 M_{\text{Jup}} \lesssim m \lesssim 80 M_{\text{Jup}}$ (see e.g., [Stevenson et al. 2023](#); [Unger et al. 2023](#)). The rich dataset available in DR3, is also well suited for machine learning applications ([Guiglion et al. 2024](#); [van Groeningen et al. 2023](#); [Ranaivomanana et al. 2025](#)) and, since the astrometric time series measurements used to produce the DR3 astrometric solutions are not yet publicly available, machine learning can complement the classical orbital fitting method for identification of substellar companions (e.g., [Sahlmann & Gómez 2025](#)).

In this work, we present a deep learning framework for identifying companion candidates in *Gaia* DR3 astrometry. Although deep learning has been extensively applied to transit detection ([Shallue & Vanderburg 2018](#); [Ansdell et al. 2018](#); [Dattilo et al. 2019](#); [Valizadegan et al. 2022](#); [Dvash et al. 2022](#)), its use in astrometric companion detection remains largely unexplored. Our method is designed to (i) select those DR3 parameters that effectively trace binarity; (ii) model their complex non-linear relationships using a deep neural network, while mitigating any over-fitting; and (iii) ensure model interpretability, so that predictions can be linked to astrophysical indicators. The resulting model, called ExoDNN, identifies non-single star candidates in *Gaia* DR3 with high reliability using astrometric fit-quality statistics and other DR3 parameters available in the main source table ([Hambly et al. 2022](#)).

The paper is structured as follows. Section 2 introduces the ExoDNN algorithm and describes the data and methods used for its training and validation. In Section 3, we describe how we applied ExoDNN to a subset of *Gaia* DR3 stars and present a preliminary list of candidate stars to host companions. Section 4 discusses the results, while Section 5 summarizes the conclusions, highlighting synergies with upcoming exoplanet missions and outlining prospects for follow-up observations.

2. The ExoDNN algorithm

2.1. Detection probability and statistical indicators

The astrometric signature of a binary system is typically defined as an angular measure $\alpha(\text{arcsecond})=a_0/d$, where a_0 is the semimajor axis of the orbit described by the star around the barycentre of the system (in AU) and d is the distance to the observer (in pc). To determine the probability p of detecting this signature, we can establish a detection criterion based on achieving a specific signal-to-noise ratio ($S/N=\alpha/\sigma_m$), where σ_m is the measurement uncertainty (ignoring other possible sources of noise). The S/N will then effectively be a cut-off in probability space, that is, serve as a criterion for acceptable confidence level in the detection (and, thus, $p \propto S/N$). For the case of unknown orbital periods, [Casertano et al. \(2008\)](#) used this criterion and derived a threshold of $S/N > 3$, for a successful astrometric orbit detection at 95% confidence level. Using a similar criterion, [Sahlmann et al. \(2015\)](#) took into account the number of measurements (N_m) to show that a $\sqrt{N_m} \times S/N \geq 20$ is necessary for detection when using astrometric data only.

However, the astrometric signature (α) is not directly observable by *Gaia*. Instead, *Gaia* measures the one-dimensional (1D) position (abscissa) of a star in the along-scan direction, which we hereafter refer to as η . *Gaia* data reduction pipeline, the Astrometric Global Iterative Solution ([Lindegren et al. 2012, 2021](#), AGIS, hereafter), uses a single star model to compute the

Table 1. AGIS astrometric fit quality statistics.

AGIS astrometric fit statistic	Description
ruwe	Re-normalized unit weight error
astrometric_gof_al (gof_{AL})	Goodness-of-fit statistic
astrometric_excess_noise (ϵ)	Excess noise of the source
astrometric_chi2_al (χ^2)	Astrometric goodness-of-fit
astrometric_sigma5d_max (σ_{5d})	Longest semimajor axis of the 5D error ellipsoid

Notes. A detailed description of these statistics is available in [Lindegren et al. \(2021\)](#) Section 5.1.

expected along-scan position of a source in the sky at time t (η^{exp}) and then minimizes the astrometric residuals (observed minus expected) ($\eta^{\text{obs}} - \eta^{\text{exp}}$) to generate an astrometric solution. This astrometric solution provides the estimated (best-fit) astrometric parameters of the source, along with a set of fit quality statistics (see Table 1). In a binary system, the excess motion induced by the unseen companion will affect the observed position of the star η^{obs} , causing the astrometric residuals to be larger than in the single star case and resulting in increased values of fit χ^2 . This has been used by several authors in the past to identify solutions that deviate from the single star model used by AGIS. As an example [Belokurov et al. \(2020\)](#) and [Penoyre et al. \(2022\)](#) used the re-normalized unit weight error (ruwe), which is proportional to the χ^2 of the astrometric fit, to search for astrometric binaries in the *Gaia* catalogue. [Stassun & Torres \(2021\)](#) found a strong correlation between this statistic and the photocentric orbit size of eclipsing binaries. The astrometric excess noise (ϵ) has also been used by [Gandhi et al. \(2022\)](#) in search for X-ray binaries.

However, it should be noted that variability ([Belokurov et al. 2020](#)), stellar crowding ([Luna et al. 2023](#)), or the presence of a proto-planetary stellar disk ([Fitton et al. 2022](#)) has been shown to contribute to high values of ruwe without a binary cause. Besides, both statistics are known to be affected by calibration errors that make them sensitive to the magnitude, color of the source (etc.); although these effects have been mitigated by construction in ruwe, they still affect ϵ in DR3 ([Lindegren et al. 2021](#), Section 5.3). Therefore, these caveats must be taken into account when using AGIS statistics as probes for the presence of unseen companions.

2.2. Deep learning

Despite the limitations mentioned earlier, it is clear that binarity has an effect on the fit quality statistics provided by AGIS, so these can be used to predict the probability, p , that a source from *Gaia* DR3 is non-single, given an array (\bar{x}) of values for those statistics. Now, since we lack an explicit form of the relation between the astrometric signature and the fit quality statistics, finding an analytical form of the probability $p(\bar{x})$, can be challenging. Therefore, we decided to take a machine learning approach and use a special type of artificial neural network, a deep neural network (DNN), to model $p(\bar{x})$.

In general, artificial neural networks can approximate any continuous function to an arbitrary degree of accuracy ([Hornik et al. 1989](#)), and learn such an approximation from data, using provided examples in a setting called supervised learning. If we use a dataset, $\mathcal{D} = \{(\bar{x}_1, y_1), \dots, (\bar{x}_k, y_k)\}$, composed by pairs of examples (\bar{x}_i, y_i) , where \bar{x}_i is array of AGIS fit statistics with

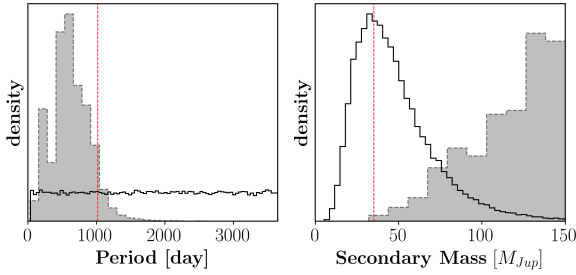


Fig. 1. Left: simulated orbital periods (black) compared to the orbital periods from DR3 NSS Orbital solutions (shade), with the DR3 time baseline marked in vertical. Right: simulated secondary masses (black), compared to the catalogue of DR3 binary masses (shade), with the peak of the distribution of simulated secondary masses in vertical.

known values for a source and y_i only has two possible values ($y_i=0$ representing a single star, and $y_i=1$ representing a system with a companion detectable by *Gaia*), we can teach the network the underlying mapping between \bar{x} and y . The probability density of the random variable y can be described by a continuous Bernoulli distribution (Loaiza-Ganem & Cunningham 2019), expressed as

$$p(y|x; \theta) = \theta^x(1 - \theta)^{1-x}, \quad (1)$$

where $\theta \in [0,1]$ is an unknown parameter that maximizes the probability that $y=y_i$ given \bar{x}_i , and $\bar{x}_i \in [0, 1]^N$. The DNN defined in this manner will effectively learn the parameter of this distribution and produce the output we are looking for, namely, an estimate of the probability $p(\bar{x})$ over many Bernoulli trials, each with a probability expressed as θ .

2.3. Generation of the training set

To train ExoDNN, we need example astrometric solutions for the positive case (stars with a detectable companion) and for the negative case (stars with no detectable companion). Given that a homogeneous, sizeable sample of bona fide single stars was not available for our study, we resorted to compiling a set of synthetic data.

We generated 100 000 examples with roughly equivalent proportion (~50%) of single stars and detectable binary systems, to avoid biasing the neural network towards preferential detection of stars from an over-represented group. These systems were placed randomly in the sky at uniform distances in the range ($1 \leq d \leq 100$ pc). Primary and secondary masses were extracted from the log-normally distributed ranges ($0.5 M_\odot \leq M_1 \leq 1.5 M_\odot$) and ($10 M_{\text{Jup}} \leq M_2 \leq 150 M_{\text{Jup}}$), respectively; the latter roughly covering a mass range from giant planets to low-mass stars (Kroupa 2001; Spiegel et al. 2011; Chabrier et al. 2014). We note that our secondary masses (see Figure 1) are generally lower than the ~1800 candidates from the DR3 NSS catalogue of binary masses (Gaia Collaboration 2023a, Section 5). This is intentional, as our goal is to simulate astrometric signatures induced by very low-mass companions. The orbital periods were drawn from the uniform range ($0.1 \text{ yr} \leq P \leq 10 \text{ yr}$), which is the maximum baseline of *Gaia* observations for the full mission. Once total mass and period of the system were defined, these determined the orbit size (a) through Kepler's third law. The remaining orbital parameters for the Keplerian companion were the eccentricity (e), orbital inclination (i), argument of the periastron (ω), longitude of the ascending node (Ω), and time at the periastron (t_p).

These parameters were randomly drawn from the uniform distributions: $0 \leq e \leq 1$, $0 \leq i \leq 90^\circ$, $0 \leq \omega \leq 360^\circ$, $0 \leq \Omega \leq 180^\circ$, and $0 \leq t_p \leq P$.

We then used the Astromet library from Penoyre et al. (2022) to generate astrometric time series for both single and binary simulated system. Similarly to AGIS, Astromet uses a five-parameter ($5p$) astrometric model to describe the motion of a single source. This model depends on right ascension ($\alpha^* = \alpha \cos \delta$), declination (δ), the proper motions (μ_{α^*} , μ_δ), and the parallax (ϖ) of the star, expressed as

$$\begin{aligned} \eta(t) = & [\alpha^* - \alpha_0^* + \mu_{\alpha^*} (t - t_0)] \sin \phi_t \\ & + [\delta - \delta_0 + \mu_\delta (t - t_0)] \cos \phi_t \\ & + \varpi f_\varpi, \end{aligned} \quad (2)$$

where (α_0^*, δ_0) correspond to the on-sky position at time t_0 , ϕ_t is the position angle of the scan at time t , and f_ϖ is the source along-scan parallax factor ($\partial\eta/\partial\varpi$). ϕ_t and f_ϖ , were obtained from the *Gaia* scanning law².

With the along-scan positions defined, we were able to simulate the 1D observations (η_j^{obs}) that *Gaia* would perform on source j over a time range equivalent to the *Gaia* DR3 time baseline of 34 months. Each simulated observation can be expressed as $\eta_j^{\text{obs}}(t) = \eta_j(t) + \sigma_m$, where $\eta_j(t)$ is given by Eq. (2), and σ_m is the along-scan single measurement uncertainty. The real *Gaia* observation uncertainty displays certain dependencies on time, the across-scan rate, and magnitude (see Lindegren et al. 2021 Appendix A); however, for simplicity, we chose a Gaussian measurement error with zero mean and $50 \mu\text{as}$ standard deviation, the latter corresponding to the median formal DR3 measurement uncertainty at $G=12$ (again Lindegren et al. 2021 Appendix A). We then fit the single star model from Eq. (2) to our simulated observations and generated a five-parameter astrometric solution for each star in our simulation (see example in Figure 2). The least-squares fit implemented by Astromet (see Penoyre et al. 2020 Appendix D), closely follows AGIS implementation, resulting in astrometric solutions that contain the exact same set of fit quality statistics as provided by AGIS. The choice of fitter is therefore very convenient, because a model trained on those astrometric solutions will be applicable to the five-parameter astrometric solutions available in DR3 main source table.

At this stage, however, we only had ($5p$) best fit astrometric solutions ($\hat{\alpha}$, $\hat{\delta}$, $\hat{\mu}_\alpha$, $\hat{\mu}_\delta$, $\hat{\varpi}$), along with their uncertainties and the corresponding fit quality statistics for each solution. To avoid biasing the model towards any positional, motion or distance preference, we dropped the five astrometric solution parameters, and only kept their uncertainties and the fit quality statistics. Also, to have more representative data of the actual DR3 parameter space, we decided to add 16 photometric parameters and 4 spectroscopic DR3 parameters, by conditionally sampling distributions of those parameters derived from real DR3 sources. Details and caveats on the derivation of photometric and spectroscopic information are provided in Appendix A and omitted here for brevity.

This process resulted in a labelled dataset composed of 100 000 examples (\bar{x}_i, y_i) , where \bar{x}_i was an array of 31 DR3 parameters, and y_i was a value (label) denoting either a binary system with a *Gaia* detectable companion ($y_i=1$) or a single star ($y_i=0$). The final list of parameters used to train our model is provided in Table 2.

² *Gaia* scanning law for the full mission can be obtained at: <https://www.cosmos.esa.int/web/gaia/full-mission-scanning-law>

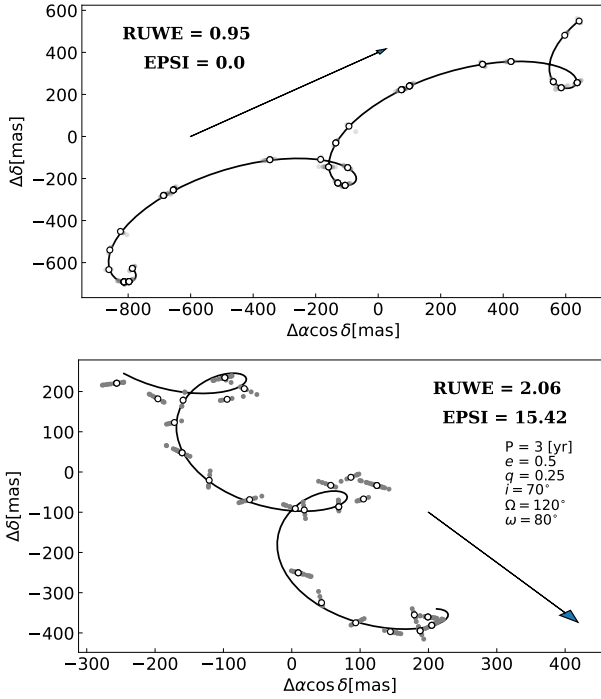


Fig. 2. Top: Barycentre motion of a simulated single star corresponding to Eq. (2) with best fit astrometric parameters. The observed position of the source is marked with solid white circles, and the 1D along-scan observations with gray circles. Bottom: same, but for a simulated binary system, where the observed position of the source is perturbed by a companion (exaggerated for illustration purposes).

2.4. Preprocessing

We then scaled robustly the training dataset using the Scikit-Learn package (Pedregosa et al. 2011). To do this, we centred each model parameter by subtracting its median and then dividing by its interquartile range. Robust scaling is a commonly used scaling technique (e.g., Géron 2017) which takes into account the possible presence of outliers and avoids that the different scales naturally present in the data bias the model towards a particular parameter. The training dataset was then split using a 70-15-15% scheme. 70% of the whole data containing positive and negative examples were used to train our network (training set), 15% (validation set) were used to tune the network hyperparameters and the remaining 15% (the test or hold-out set) were left aside to check how well the model generalizes, this is, how well it predicts over unseen data. Shuffling of the whole data was also performed before the train-val-split to ensure adequate mixture of positive and negative examples across the train, test and validation data splits.

2.5. Model training

2.5.1. Model architecture

ExoDNN is composed of an input layer followed by three densely connected hidden layers, one dropout layer and the final output layer. The input layer contains 31 neurons to map our 31 DR3 parameters. The hidden layers have 64 neurons each. The neurons on these layers use a rectified linear unit (ReLU) gate, as it speeds the training in relative comparison to other non-linear gates. The role of the dropout layer is to reduce over-fitting (Srivastava et al. 2014), a situation in which a model produces a

Table 2. *Gaia* DR3 parameters used to train the model.

Parameter type	Parameter name
Astrometry	astrometric_n_obs_al
	ra_error
	dec_error
	parallax_error
	pmra_error
	pmdec_error
Fit statistic	ruwe
	astrometric_excess_noise
Photometry	astrometric_chi2_al
	astrometric_gof_al
	astrometric_5d_sigma
	phot_g_mean_flux
	phot_g_mean_flux_error
	phot_g_mean_flux_over_error
	phot_g_mean_mag
	phot_bp_mean_flux
	phot_bp_mean_flux_error
	phot_bp_mean_flux_over_error
	phot_bp_mean_mag
	phot_rp_mean_flux
	phot_rp_mean_flux_error
	phot_rp_mean_flux_over_error
phot_rp_mean_mag	
Spectroscopy	phot_bp_rp_excess_factor
	bp_rp
	bp_g
	g_rp
	radial_velocity
	radial_velocity_error
	grvs_mag
grvs_mag_error	

Notes. A detailed description of each of them can be found at: https://gea.esac.esa.int/archive/documentation/GDR3/Gaia_archive/chap_datamodel/sec_dm_main_source_catalogue/ssc_dm_gaia_source.html.

perfect prediction over the data it was trained against, but fails to generate good predictions when exposed to unseen new data. Finally, to ensure that a valid probability is produced by the network, the output layer contains a single neuron with a sigmoid activation function which always produces a predicted probability value \hat{p} in the range $[0, 1]$. We explored network architectures with two to five hidden layers. Configurations with two layers did not achieve the required precision, whereas five layers increased training time without measurable performance gains. The number of neurons per layer was varied between 32 and 256, guided by standard heuristics that balance complexity and computational efficiency. A configuration with 64 neurons per layer yielded the most stable convergence and was therefore adopted. A detailed description of many of the concepts presented in this section such as architecture design, network training, activation gates, back propagation, over-fitting and performance metrics, can be found in Goodfellow et al. (2016).

Table 3 summarizes the final architecture of ExoDNN, along with the training hyper-parameters (optimizer, learning rate, mini-batch size, dropout rate and stopping criterion). To implement ExoDNN we used the Tensorflow package (TensorFlow 2015).

Table 3. Neural network hyper-parameters.

Hyper-parameter		Value
Architecture	No. of hidden layers	3
	Neurons per hidd. layer	64
	Hidd. Neuron gate	ReLU
	Hidd. layer type	Dense
	Output layer gate	Sigmoid
Training	Data split	70-15-15%
	Seed	42
	Shuffle	True
	Learning rate	10^{-4}
	Batch size	64
	Optimizer	Adam
	Loss function	Negative log-loss
	Dropout rate	0.2
	EarlyStopping(monitor)	val data loss
EarlyStopping(patience)	10	
Overfitting		

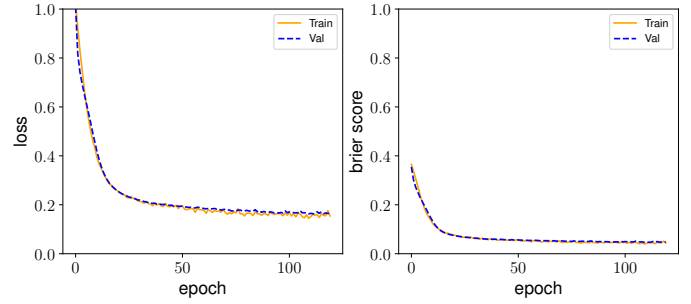
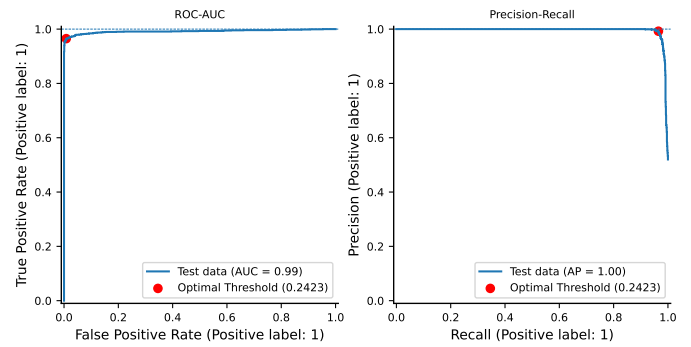
2.5.2. Model optimization

A neural network learns by iteratively adjusting its internal parameters (weights and biases) through forward and back propagation of the training data. During forward propagation, the network makes predictions (using small batches of data) of the probability \hat{p} that a training example \bar{x}_i in the batch belongs to the positive class ($y_i=1$). Then, these predictions are compared to the true class using a loss function that quantifies how far off the predictions are from the actual labels y_i . Our loss function is the average negative log-likelihood of the training set, expressed as

$$l(\theta) = -\frac{1}{N} \sum_i y_i \times \log(\theta) + (1 - y_i) \times \log(1 - \theta), \quad (3)$$

where N is the number of training samples and y_i is the actual class of the i -th sample. The variable θ is the probability computed by the network that the training sample i belongs to the positive class. We note that minimizing the negative log-likelihood is equivalent to maximizing the likelihood of the training data under our parametrized model of the random variable y . During back-propagation, the weights are re-adjusted to lower the loss on the training set. The algorithm responsible for re-adjusting the weights is called optimizer and we chose the Adam optimizer (Kingma & Ba 2015) due to its computational efficiency and fast convergence. The model performs multiple forward and backward passes on iterations called epochs. During an epoch a full pass through the entire training set is performed, and this process is repeated until a stopping criterion is met. In our case, we used a stopping criterion called “early stopping” that halts the training when it detects no improvement on the loss computed over the validation set after a successive number of epochs. This choice was also aimed at avoiding a network over-fitting.

The training process explained earlier is designed to teach ExoDNN to produce probability estimates that are as close as possible to the true underlying class probabilities. Therefore, to evaluate the model performance, we selected a metric called the Brier-score, which is equivalent to the mean-squared error but for a probabilistic classification. This metric has values higher or


Fig. 3. Left: evolution of the training and validation set losses over the different epochs. Right: evolution of the Brier-score during the training.

Fig. 4. Left: classification threshold computation using the area under curve. Right: classification threshold computation using the precision-recall curve. The optimal threshold that achieves the best compromise between false positive and true positive rates is marked as a red dot. This is our chosen $p_0=0.242$.

equal than zero, and is defined as

$$Brier = \frac{1}{N} \sum_{i=1}^N (\hat{p}_i - y_i)^2, \quad (4)$$

where N is the number of examples, \hat{p}_i is the probability computed by our model for example, \bar{x}_i , and y_i is the actual value of the random variable y for that example. Figure 3 shows the evolution of the log-loss (left) and brier-score (right) during the training. Since ExoDNN outputs probabilities, to evaluate its performance over a test set with “hard” labels (positive or negative class label), we first needed to establish a classification threshold. This threshold determines when a given predicted probability \hat{p}_i belongs to the positive class (binary star). The optimal threshold was computed by making predictions over our hold-out (test) set while varying the probability threshold and finding which of those thresholds maximized a chosen binary classification metric. The selected metric was the $f1$ -score metric because it represents a trade-off between precision and recall. In this context it is important that we correctly classify stars that host detectable companions as such (precision), and that we capture as many of those as possible (recall). The $f1$ -score has values between zero and one, and is defined as

$$f1 = \frac{2 * Precision * Recall}{Precision + Recall}. \quad (5)$$

ExoDNN retained a very low false positive rate as the classification threshold was varied, as shown in Figure 4 (left), and also maintained high precision and recall at the same time, as shown

Table 4. Model evaluation results.

Model	$f1$ -score	Brier score
ExoDNN	0.901	0.034
Logistic Regression	0.817	0.125
Random Forest	0.812	0.159

Notes. Binary classification scores for ExoDNN and alternative classifiers over the hold-out dataset.

in Figure 4 (right). The optimal balance between the true positive and the false positive rates over the test data was found for a probability threshold value of $p_0=0.242$. This means that a given example is considered to belong to the positive class, this is, star possibly hosts *Gaia* detectable companion(s), when the predicted probability on output, \hat{p} , is strictly higher than p_0 , $\hat{p}(\bar{x}_i) > p_0$.

2.5.3. Model evaluation

Using the 15% hold-out set from the train–validation–test split, we evaluated the model on unseen data. According to our performance metrics (Section 2.5.2), a well-performing model should exhibit a low Brier score (approaching zero) and a high $f1$ -score (approaching one). For benchmarking, we compared our model against commonly used binary classifiers: the Logistic Regression (Cox 1958) and the Random Forest (Breiman 2001), a couple of classifiers that are commonly used for binary classification. A good overview of these models is available in Hastie et al. (2001, Chapters 4,11,15) and details on their parameter settings are provided in Appendix B. Table 4 shows a summary of the performance metrics obtained by ExoDNN and the alternative classifiers. We found that ExoDNN outperformed the alternative binary classifiers, showing both a low Brier-score and a high $f1$ -score, supporting its adoption over standard classifiers.

2.6. Model inference and interpretation

We describe above how ExoDNN was built and trained, but we might also consider how it comes up with a predicted probability value for a given input. Our model has learned a non-linear approximation f (using weights and biases) of the true conditional probability, $P(y=1|\bar{x}_i)$, where \bar{x}_i is an array of 31 DR3 parameters for a given star. When asked to predict over the input \bar{x}_i , it will produce an estimate \hat{p} by computing $\hat{p}(\bar{x}_i) = \sigma(f(\bar{x}_i))$, where σ is the sigmoid function,

$$\hat{p}(\bar{x}_i) = \frac{1}{1 + e^{-f(\bar{x}_i)}}. \quad (6)$$

Now, while the model will use all 31 DR3 parameters to compute \hat{p} , not all of them have the same weight towards that computation. To better understand the relative importance of those 31 DR3 parameters in the prediction, we used SHapley Additive exPlanations (SHAP) (Lundberg & Lee 2017), a commonly used tool to interpret machine learning models. Figure 5 (top) shows the relative importance of the most influential parameters towards the model prediction. The x -axis corresponds to the mean shap value. This is a quantitative measure of how much each feature changes the prediction compared to the average prediction, and gives an idea of its relative importance against other features. We can see how *ruwe* and *astrometric_excess_noise* (*epsi*) have the highest relative

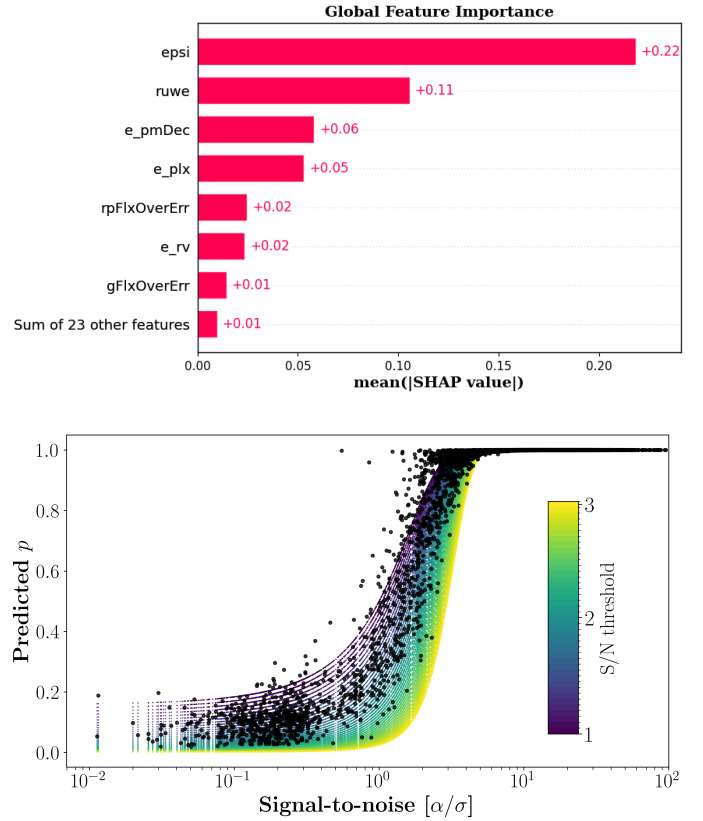


Fig. 5. Top: parameter impact on the model output as estimated by the neural network model ranked by decreasing impact. High parameter values are color-coded in red, while low values are displayed in blue. Bottom: predicted probability for each binary star example on the test dataset (black dots), together with a detection probability proxy computed as $p = 1 - \Phi(S/N, \delta)$, where Φ is the CDF of the normal distribution, $S/N = \alpha/\sigma$, and δ is a variable S/N threshold from 1 to 3.

importance for ExoDNN prediction, which is expected based on our analysis on how binarity affects the astrometric fit statistics (Section 2.1). Then we have a group of parameters that trace the formal uncertainties in the astrometric solution, the proper motion (declination) and parallax errors (*pmDecErr*, *plxErr*), which are also expected to increase due to deviation from a purely linear motion model in the binary case. The RV error (*radVelErr*) traces the periodic velocity wobble induced by an orbiting companion in the host star, and therefore we would also expect this parameter to have a significant impact on the model decision. Finally, we have a group of photometric uncertainty related parameters (*rpFlxOverErr*, *gFlxOverErr*). Sullivan et al. (2025) show how unresolved binarity has a direct impact on flux contamination over the *Gaia* detection window, resulting in larger photometric residuals. The remaining 23 model parameters were still used, but we note that they have less of an impact in the computation of \hat{p} .

To complement this analysis, we display in Figure 5 (bottom) the computed probability \hat{p} for each of the simulated binaries in the test data sample, along a detection probability proxy computed for different S/N levels. We can see how \hat{p} scales with the astrometric S/N. This is also reassuring with respect to ExoDNN’s behavior, indicating that the model has learned a proxy for the detection probability based on the astrometric S/N ratio noted in Section 2.1.

Table 5. Model testing results.

Data source	Total	Exp.	Det.	Acc.
NSS ORBT	384	384	289	~75%
NSS ASB1	1332	1332	1250	~94%
CARMENES	48	48	34	~71%
Control	5000	0	58	~99%

Notes. Expected (Exp.) column refers to the number of objects from the given sample that meet the chosen detectability criterion. Detected (Det.) column refers to the number of stars in the corresponding sample for which the model has computed a probability $p(y=1) > p_0$. Accuracy (Acc.) reports model accuracy (detected/expected) on the given dataset.

2.7. Model testing

The hold-out dataset used to check ExoDNN generalization (see Section 2.5), is fully synthetic. However, such data cannot capture the intrinsic DR3 calibration features mentioned in Section 2.1. Therefore, we used unresolved binary systems detected by *Gaia* in DR3 (true positives) to test the behavior of ExoDNN when exposed to unseen and real data. These are: (1) the 384 DR3 NSS OrbitalTargetedSearch* and OrbitalAlternative* solutions (NSS ORBT) from [Gaia Collaboration \(2023a, Table 1\)](#), within 100pc; (2) the 1332 DR3 NSS AstroSpectroscopic solutions (combined single lined spectroscopic and astrometric binary) (NSS ASB1) from [Gaia Collaboration \(2023a\) Table 1](#), within 100pc; (3) the 48 M-dwarfs from the CARMENES input catalogue ([Cifuentes et al. 2025](#)) that have have a corresponding DR3 NSS solution. For all these systems, we had access to an existing orbital solution either from literature or from the *Gaia* DR3 NSS solutions, with an estimate of the size of the semimajor axis of the photocentric orbit (a_0). We then adopted the astrometric S/N detection criterion from [Sahlmann et al. \(2015\)](#), namely: $\sqrt{N_{\text{obs}}} \alpha / \sigma_{\eta} > 20$, where $\alpha = a_0/d$, N_{obs} corresponds to the astrometric_n_good_observations_al DR3 parameter of the given source, and $\sigma_{\eta} = 150 \mu\text{s}$ is the median along-scan measurement uncertainty estimated for DR3 sources at $G=12$ (see Figure 3 in [Holl et al. 2023](#)).

This computation is merely a consistency check, since by construction all the selected real sources induce enough astrometric signal to be detectable by *Gaia*. Therefore, the number of expected detections on the given test sample is the same as the total number of sources for that sample. Finally, to measure ExoDNN baseline false positive rate, we generated a further test sample composed of 5000 new synthetic single stars using the approach described in Section 2.3. A perfect model should predict close to zero probability of hosting a companion for these true negative sources. The false positive rate, as determined from the control sample, resulted in approximately 1.2%.

Table 5 shows the summary results of this process over the entire set of samples. Our model accuracy (detected/expected) varied depending on the sample, from ~70% for the CARMENES M-dwarf sample, to over 90% for the *Gaia* NSS single lined astro-spectroscopic binaries.

Figure 6 illustrates how the astrometric signature levels induced by the binary systems in our simulation are (intentionally) comparable to the signatures induced unresolved binaries so far detected by *Gaia*, and typically higher than currently confirmed exoplanets. Although the overall prediction results are not as optimal as we would like across all real samples, the performance of a model will always decrease when presented with real

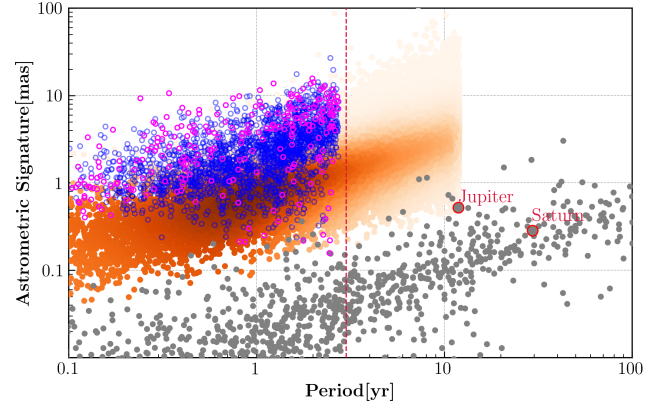


Fig. 6. Astrometric signature vs. orbital period for the simulated binary sources (orange), the 384 DR3 NSS ORBT solutions (magenta), the 1332 DR3 NSS ASB1 solutions (blue) and the currently confirmed exoplanets ([Schneider et al. 2011](#)) (gray). The red dashed vertical line represents the DR3 time baseline of 34 months in comparison to the shorter periods of the *Gaia* DR3 unresolved binary orbital solutions. For the DR3 NSS solutions we used the expression $\alpha = a_0/d$ to compute the astrometric signature. a_0 was computed using the Thiele-Innes elements (A, B, F, G) of the corresponding *Gaia* DR3 NSS solution and the recipe available in [Halbwachs et al. \(2023, Appendix A\)](#), and d was computed as the inverse of the source parallax. For the simulated sources we used $\alpha = M_2/M_1 \times a_1/d$ from [Perryman et al. \(2014\)](#), where M_1 and M_2 are the primary and secondary masses, respectively, and a_1 is the semimajor axis of the orbit of the secondary.

data. In particular, the characteristic DR3 calibration issues mentioned in Section 2.1 are more likely to affect the fainter and redder M-dwarfs from CARMENES sample. We therefore settled for a version of the model that showed a good compromise between the true and false positive rate, and decided to apply specific post-processing to tackle (at least partially) the mentioned calibration issues. We detail this approach in later sections.

3. Application to *Gaia* DR3

3.1. ExoDNN application

With our trained and tested deep learning model, we turned to making predictions over real stars from *Gaia* DR3. We selected a volume limited sample of astrometrically well observed main sequence and giant stars at moderate distances ($d < 100$ pc), from the *Gaia* DR3 Astrophysical Parameters sample (Apsis, [Creevey et al. 2023](#)). This dataset provides estimates for key stellar parameters such as mass, effective temperature and metallicity (among others), which were necessary to characterize the host stars of our potential candidates. To this sample, we applied a set of filtering criteria similar to those used for obtaining the FGKM golden sample of astrophysical parameters ([Gaia Collaboration 2023b, CR23](#) hereafter), but relaxed to reduce the heavy filtering of M stars that takes place in the golden sample. As such, our input for model prediction, fully contains the DR3 FGKM golden sample as a subset, but has a higher fraction of M stars in comparison (see Figure 7). Full details on this process are provided in Appendix C and omitted here for brevity.

This selection resulted in a total of 102 117 stars with estimated spectral types F, G, K, and M within 100pc to which we applied the same pre-processing steps used for the training data, namely the same robust scaling of the model attributes that was described in Section 2.5. The mentioned spectral types

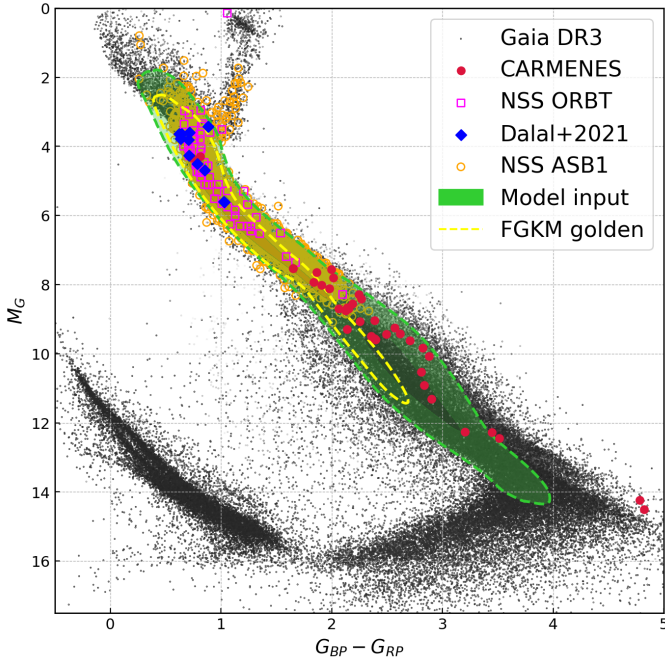


Fig. 7. *Gaia* BP minus RP color vs. absolute magnitude for the model input stars (green region). The FGKM golden sample is shown in yellow. Absolute magnitude is computed as: $m_G = phot_g_mean_mag + 5.0 \cdot \log(parallax) - 10$. Close binary systems where the smaller companion is a brown dwarf or a giant planet are displayed as diamonds (blue) and NSS Orbital* solutions are displayed as squares (magenta). CARMENES sources are displayed as circles (orange). Sources in correspond to the *Gaia* Catalog of Nearby Stars within $d < 100$ pc.

correspond to the `spectraltype_esphs` parameter of the DR3 astrophysical parameters, which is computed using the low-resolution spectra from the BP-RP instrument onboard *Gaia*. We then used ExoDNN to compute the probability of each of these stars to host unseen companions. Our model returned 14 606 stars with a predicted probability $p > p_0 (= 0.242)$, which were considered initial candidates. The breakdown per spectral type (`spectraltype_esphs`) of these initial candidates is: 6017 M-type stars, 4180 K-type star, 2311 G-type stars and 2098 F-type stars.

3.2. Post-processing

We then applied a series of post-processing steps over ExoDNN prediction. These were aimed at obtaining a list of consolidated candidates sources to host unresolved companions. First, we identified candidates that are already known binaries or with confirmed companions by doing a cross-match with a cone search radius of 1.8 arc-seconds against four different catalogues, the 9th Catalog of Spectroscopic Binaries (Pourbaix et al. 2004), the subset of *Gaia* DR3 NSS orbital solutions (Gaia Collaboration 2023a) within 100pc, the Washington Double Stars catalogue (Mason et al. 2001) and the online Encyclopedia of Exoplanetary Systems³. This cross-match identified 4383 stars which are present in either one or several of these catalogues. Some illustrative examples of known binaries that ExoDNN model has correctly identified among the initial list of candidates, are 61 Cyg (WDS J21069+3845AB), one of the closest binary systems, Struve 2398 (WDS J18428+5938AB),

³ Accessed early May 2025 at: <https://exoplanet.eu/home/>

composed of two M-dwarfs, η Cas (WDS J00491+5749AB), or the spectroscopic binaries 70 Oph (SBC9 1022) and HD137763 (SBC9 1638). Since the model is trained on binary systems without any internal hierarchical structure, it cannot distinguish between binaries and triples or quadruples. Therefore, systems such as κ Tuc (WDS J01158-6853) or σ CrB (WDS 16147+3352), where *Gaia* has resolved the components, are also reported as candidates by ExoDNN. Other prominent examples of stars within the prediction range of 100pc with confirmed substellar companions that the model correctly reports as candidates are: ϵ Eri, a young K-dwarf at ~ 3 pc with a confirmed $\sim 1.6 M_{\text{Jup}}$ companion (Benedict et al. 2006), TZ Ari, a close M5.0V star located at ~ 4.8 pc that hosts an RV detected Saturn-type exoplanet ($\sim 0.21 M_{\text{Jup}}$) (Quirrenbach et al. 2022), ν And at ~ 13.5 pc orbited by several exoplanets including a Jupiter-like companion detected through RV (Curiel et al. 2011), and BD+05 5218 (HIP 117179) a G3V-dwarf 88 pc away that hosts an astrometrically detected $\sim 44 M_{\text{Jup}}$ brown dwarf (Stevenson et al. 2023).

Second, we tackled the calibration issues affecting AGIS statistics (Section 2.1), by following Lindegren et al. (2021), Section 5.3 recommendations. We removed sources with `astrometric_excess_noise_sig` < 2 , and with colors ($G_{BP} - G_{RP} \leq 0.4$), or ($G_{BP} - G_{RP} \geq 4.0$). We extended the red end color limit to avoid filtering out known systems which our model correctly considers candidates and would otherwise be filtered out using the recommended constrain of ($G_{BP} - G_{RP} > 3.0$). We also removed sources with (`phot_g_mean_mag` < 6) and (`phot_g_mean_mag` > 16), as `ruwe` can be unreliable beyond those magnitude. This filter removed 1579 sources.

Third and final, we used the `ipd_frac_multi_peak` parameter as an additional discriminator. This statistic is sensitive to the presence of multiple point-spread-functions (PSF) within the same detection window. This situation can occur depending on the scan direction when a partially resolved binary is observed by *Gaia*. We adopted the suggested value for filtering on this statistic from Fabricius et al. (2021), and only retained sources with `ipd_frac_multi_peak` < 2 . Because this parameter is stored as an integer percent, a value of 2 actually means that 0–1% of the transits detected for a given source were multi-peaked. Hence, this filter aims to obtain clean single-peaked sources where the chances of any astrometric perturbation due to a close (partially unresolved) companion of stellar nature is close to zero. This filter removed 1230 sources.

In Figure 8 we illustrate (from left to right) the different cuts in color, magnitude and fraction of multiple PSF found in the detection window, applied on each successive post-processing filter to obtain the final candidates. All together, these three post-processing steps removed 7 192 stars from the original candidate list, resulting in a total of 7414 stars which were considered final candidates. The breakdown per spectral type of the final 7414 candidates is: 3757 M-type stars, 1852 K-type star, 950 G-type stars, and 855 F-type stars.

3.3. List of candidates

We show in Figure 9 a color-magnitude diagram of the *Gaia* Catalog of Nearby Stars (Gaia Collaboration 2021), where we highlight the final 7414 consolidated candidates reported by ExoDNN, along with a few sources to help showcase ExoDNN predictions. These showcase examples are *Gaia*-4b and *Gaia*-5b (Stefánsson et al. 2025), a massive planet and a brown dwarf, respectively, with masses of $m = 11.8 \pm 0.7 M_{\text{Jup}}$ and $m = 20.9 \pm 0.5 M_{\text{Jup}}$, the only substellar companions from DR3 that have

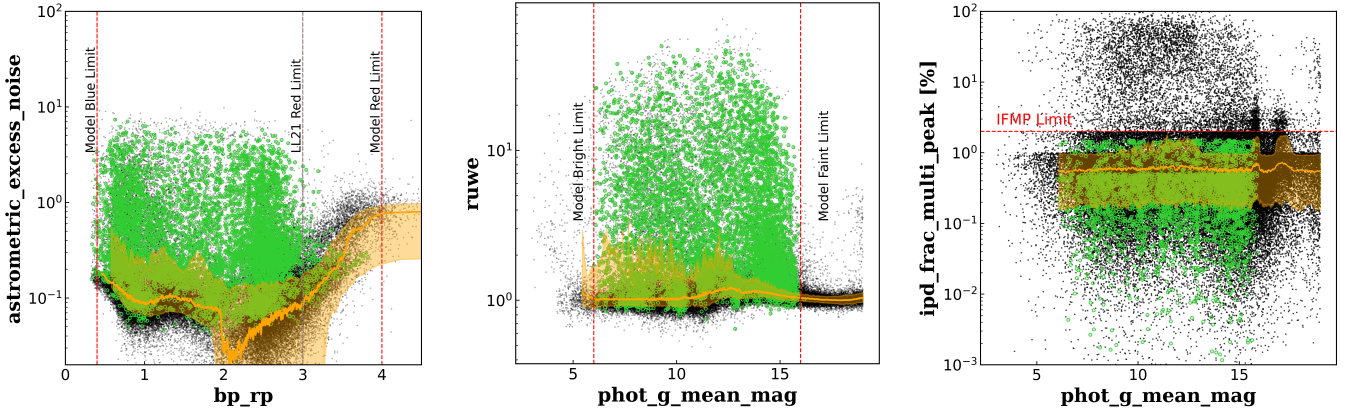


Fig. 8. Left: astrometric_excess_noise vs. *Gaia* color *bp_rp*. Model inputs are displayed in black and candidates in green. The vertical red lines show the applied color cuts. Orange line corresponds to the running median of the excess noise and the shaded orange region corresponds to the 16th-84th percentiles. Middle: equivalent plot but for *ruwe* vs. *Gaia* *G* band magnitude. The vertical red lines show the applied bright and faint *Gaia* magnitude cuts. Right: equivalent plot but for *ipd_frac_multi_peak* vs. *Gaia* *G* magnitude. The horizontal red line corresponds to the selected cut-off limit for this parameter.

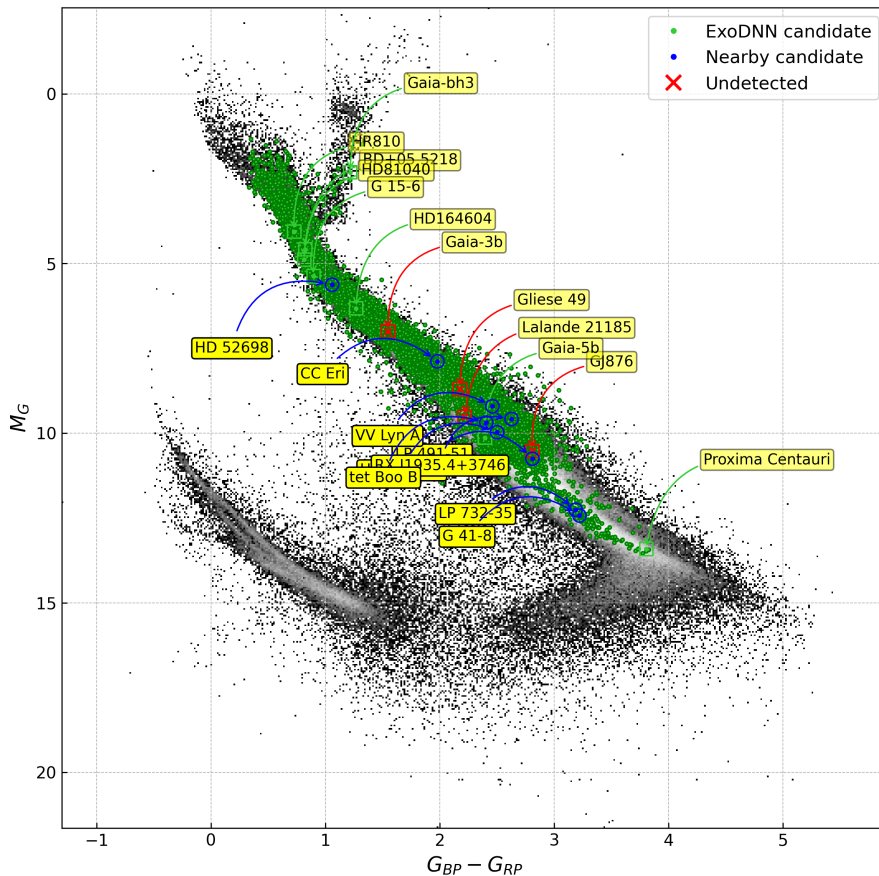


Fig. 9. *Gaia* BP minus RP color vs. absolute magnitude diagram of *Gaia* DR3 Catalog of Nearby Sources. The 7414 final candidates reported by ExoDNN after filtering are displayed in green. Prediction for a few reference objects, Proxima Centauri, LHS1610, BD+05 5128, HD 164604, *Gaia*-3b, *Gaia*-4b and *Gaia*-5b, etc., is also shown. Sources with known companions detected by ExoDNN are marked by the green boxes, while red-crosses refer to reference objects for which model predicted probability $\hat{p} < p_0$, that is, undetected. ExoDNN closest new candidates are highlighted to the left of the main sequence with blue circles.

been externally confirmed so far via RV follow-up observations. Both *Gaia*-4b and *Gaia*-5b are reported as candidates by ExoDNN. G 15-6, where Sahlmann & Gómez (2025) have recently identified a brown dwarf companion with an astrometric mass $m=62.4 \pm 6.2 M_{\text{Jup}}$, is also reported as a candidate. We also show *Gaia* black hole 3 (*Gaia* Collaboration 2024),

whose host star is also detected as a candidate. We omit *Gaia*-1b and *Gaia*-2b since these are close-in transiting planets (Panahi et al. 2022) with very short orbital periods (~ 3 days) that do not produce enough gravitational pull to be astrometrically detectable in DR3 (or even DR4/5). Another interesting example is HD 164604, a K-dwarf at ~ 40 pc orbited by a $\sim 14 M_{\text{Jup}}$ super

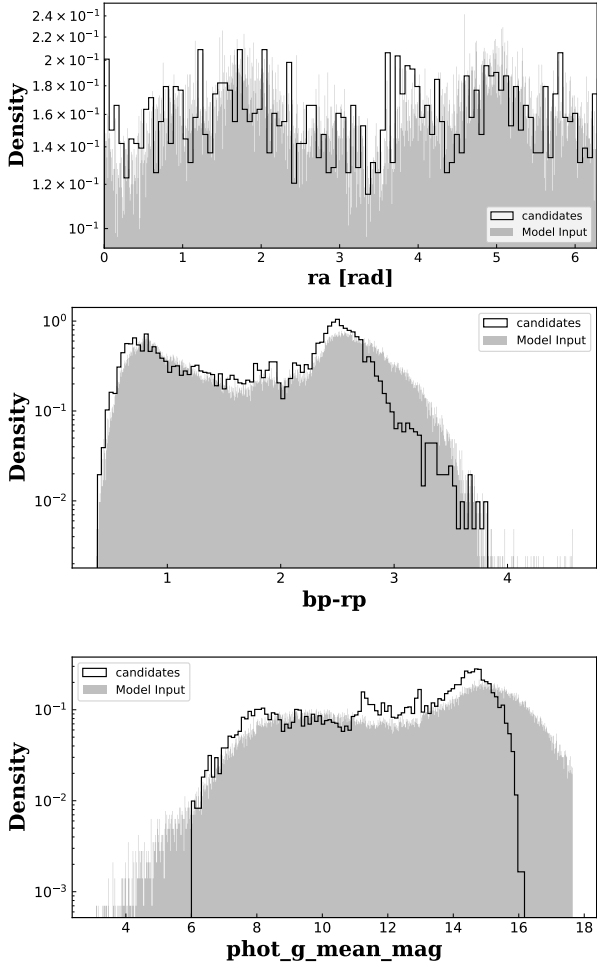


Fig. 10. Top: distribution of the right ascension values for candidates sources (solid black line) and the original input for the model (gray). Middle: equivalent display but for *Gaia* BP minus RP color ($bp-rp$) distribution. Bottom: equivalent display but for the *Gaia* magnitude ($phot_g_mean_mag$) in this case.

Jupiter and detected via astrometry (Gaia Collaboration 2023a), that is also reported as a candidate by ExoDNN. We searched for signs of preferential detection of candidates on particular sky positions, distances, *Gaia* colors, or *Gaia* magnitudes. Figure 10 shows the distributions of some of those parameters, which do not reveal any significant bias in the model output candidates. When we analyse the nearby candidates reported by ExoDNN, shown in Table 6, we can see that the majority are M-dwarfs. This is not totally surprising since the astrometric signature of an hypothetical companion will be stronger in an M-dwarf due to the higher mass ratio. This, together with their abundant numbers ($\sim 75\%$ of the stellar population Henry & Jao 2024), makes M stars usual targets for RV exoplanet surveys (e.g., Ribas et al. 2023). We find among these closest candidates an illustrative example (HD 52698) on how the model decision is driven by the most influential parameters $ruwe$ and $astrometric_excess_noise$. This spectroscopic binary has intrinsic high values of both astrometric fit statistics, which drives ExoDNN to make a strong prediction of the probability that this source has companions, 0.99 probability against the average of the other 12 candidates 0.37.

To check whether any biases had been introduced by adopting a deep learning approach in this domain, we compared

ExoDNN predictions with those from Sahlmann & Gómez (2025). In that work, the authors used a semisupervised anomaly detection framework that combines feature importance and ensemble classifiers to identify potential companions to stars. The scientific goal (astrometric detection of substellar companions) is equivalent to ours, but the technique is different, thus making the comparison relevant. We fed ExoDNN with the 22 high-probability candidate sources reported by these authors (see their Table 4). To do this, we preprocessed the required model parameters corresponding to each of those 22 *Gaia* DR3 sources and then computed the probability of each of them to host a substellar companion using ExoDNN. We present the predictions made by ExoDNN over those 22 candidates in Table 7. ExoDNN predicts that 16 out of 22 ($\sim 72\%$) also are candidates to host companions. However, as noted by Sahlmann & Gómez (2025) (see their section 4.2), some of these candidates (e.g., * 54 Cas and BD+75 510), were identified after DR3 release as false positives. Finally, the six sources not considered as candidates by ExoDNN have, in relative comparison to the others, a combination of $ruwe$ and $astrometric_excess_noise$ that results in a lower probability than our established threshold. The rest of the candidates (16) are detected by ExoDNN with computed probabilities to hold companions in the range (0.25 to 0.99) depending on the relative values of the two main astrometric fit quality indicators. We find cases such as HD 207740, HD 40503, GSC 09436-01089 and G15-6, where the relative high values of both the $ruwe$ and $astrometric_excess_noise$ drive ExoDNN prediction to a high predicted probability of holding companions. This result is reassuring, as it represents an independent check of the ExoDNN behavior.

4. Discussion

4.1. Effects of filtering

The significant reduction ($\sim 50\%$) from the original 14 606 candidates to the final 7414 has its origin in the post-processing filters we applied (Section 3.2). Over 30% of the original candidates were identified during the cross-match (first filter) as previously known binary or multiple systems. These candidates are considered true positives, which ExoDNN has correctly identified and thus, they were removed because the focus of our study is on finding genuinely new candidates. The remaining filtering steps remove 20% of the initial candidates but affect differently across the spectral types. While the astrometric excess noise significance filter (second filter) affects the fainter (M-type) stars more heavily, the multi peak discriminator (third filter) affects more F, G and K-type stars. The net effect of these filters in the total numbers per spectral type is illustrated in Figure 11.

Another prominent feature of our final candidates, is the apparent depletion of faint candidates ($M_G > 12$). This is due to the strong filtering induced by the second post-processing step (Section 3.2) for faint magnitudes. The fraction of DR3 sources with significant (> 2) excess noise decreases from 97% at $9 \leq G \leq 12$ to only 20% at $G = 14$ (Lindegren et al. 2021, Table 4), and therefore that filter is particularly efficient in removal of M-type candidates in the magnitude range ($12 \leq G \leq 15$). These sources frequently exhibit excess noise significance values lower than 2, see Figure 11 (s) values. Also, the significant depletion of $\sim 50\%$ F and G-type initial candidates is caused by our third filtering step. The general increased multiplicity fraction of F and G dwarfs, $\sim 60\%$ for F-G stars vs. 25–40% for M-dwarfs (Duchêne & Kraus 2013), will result in relatively higher numbers of the $ipd_frac_multi_peak$ parameters for the spectral

Table 6. ExoDNN closest-candidates prediction.

Name	Alt. name	<i>Gaia</i> designation	Parallax[mas]	Sp Type	ruwe	ϵ [mas]	\hat{p}
G 158-27	GJ 1002	<i>Gaia</i> DR3 2441630500517079808	206.35	M5.5V ⁷	1.34	0.36	0.29
CD Cet	GJ 1057	<i>Gaia</i> DR3 3179036008830848	116.28	M4.5V ⁸	1.45	0.35	0.27
Wolf 1069	GJ 1253	<i>Gaia</i> DR3 2188318517720321664	104.44	dM5.0 ⁶	1.42	0.31	0.33
LP 732-35	GJ 3668	<i>Gaia</i> DR3 3560746351897773952	86.90	M4.5V ¹	1.90	0.36	0.52
CC Eri	GJ 103	<i>Gaia</i> DR3 4946938113149426944	86.61	K7.V C ¹⁰	2.26	0.38	0.33
G 41-8	–	<i>Gaia</i> DR3 605079910398496640	85.80	M6.0V ²	1.12	0.20	0.35
VV Lyn A	GJ 277A	<i>Gaia</i> DR3 898506093171335424	83.38	M4.5V ¹	2.51	0.41	0.37
HD 304043	GJ 442	<i>Gaia</i> DR3 5339892367684811520	78.89	M3.5V ¹	1.01	0.34	0.31
RX J1935.4+3746	–	<i>Gaia</i> DR3 2051984436005124480	69.12	M3.0V ³	2.20	0.42	0.50
θ Boo B	GJ 549 B	<i>Gaia</i> DR3 1604859408265509632	68.81	M2.5V ⁴	1.15	0.35	0.36
G 268-38	GJ 3053	<i>Gaia</i> DR3 2371032916186181760	66.83	M4.5V ¹	1.53	0.28	0.29
HD 52698	GJ 259	<i>Gaia</i> DR3 2920772722738017920	67.80	K1.V B ⁵	13.03	2.16	0.99
LP 491-51	–	<i>Gaia</i> DR3 3968118399284369280	65.84	M3.0V ³	3.17	0.54	0.43

Notes. ruwe and astrometric_excess_noise ϵ values are from *Gaia* DR3 main source table. The column \hat{p} , is the probability computed by ExoDNN that the given source hosts a detectable companion. (1) Reid et al. (1995); (2) Newton et al. (2014); (3) Shkolnik et al. (2009); (4) Alonso-Floriano et al. (2015); (5) Torres et al. (2006); (6) Schweitzer et al. (2019); (7) Shan et al. (2024); (8) Lépine et al. (2013); (9) Keenan & McNeil (1989); (10) Torres et al. (2006).

Table 7. ExoDNN prediction on external candidates.

Name	Alt. name	<i>Gaia</i> designation	Parallax[mas]	ruwe	ϵ [mas]	\hat{p}
HD 207740	HIP 107821	<i>Gaia</i> DR3 1576108450508750208	14.020275	5.338995	1.204246	0.99
HD 40503	HIP 28193	<i>Gaia</i> DR3 5323844651848467968	11.477482	5.047571	0.909602	0.99
GSC 09436-01089	–	<i>Gaia</i> DR3 3921176983720146560	15.452359	3.036620	0.617942	0.95
G15-6	HIP 73800	<i>Gaia</i> DR3 1897143408911208832	20.304506	3.340308	0.644446	0.94
HD 206484	–	<i>Gaia</i> DR3 2171489736355655680	13.936007	2.631042	0.536911	0.79
HD 49264	HIP 32253	<i>Gaia</i> DR3 3913728032959687424	14.752394	2.584024	0.482179	0.75
HD 212620	BD +75 825	<i>Gaia</i> DR3 5148853253106611200	13.800585	2.746541	0.417314	0.75
HD 106888	HIP 59933	<i>Gaia</i> DR3 1156378820136922880	13.872044	3.363915	0.353119	0.67
BD+3292	HIP 2666	<i>Gaia</i> DR3 2540855308890440064	10.091830	2.012354	0.272028	0.51
HD 128717	BD +58 1514	<i>Gaia</i> DR3 4545802186476906880	18.444739	2.028104	0.315474	0.40
HD 104289	BD +60 1368	<i>Gaia</i> DR3 364792020789523584	12.960809	1.581556	0.272007	0.36
TYC 4998-437-1	–	<i>Gaia</i> DR3 3909531609393458688	4.687014	1.720975	0.276329	0.36
LP 769-9	–	<i>Gaia</i> DR3 1878822452815621120	16.052830	1.917616	0.236748	0.34
* 54 Cas	HIP 10031	<i>Gaia</i> DR3 5484481960625470336	8.776449	1.982503	0.246423	0.30
BD-06 2423A	–	<i>Gaia</i> DR3 6330529666839726592	7.563384	1.626470	0.193307	0.30
BD+75510	BD +75 510	<i>Gaia</i> DR3 5773484949857279104	4.225611	1.437869	0.176325	0.25
HD 153376	BD +15 3089	<i>Gaia</i> DR3 2280560705703031552	7.321575	1.419174	0.193269	0.21
HD 206848	BD +23 4384	<i>Gaia</i> DR3 1610837178107032192	13.560875	1.320370	0.214271	0.21
TYC 277-599-1	–	<i>Gaia</i> DR3 2884087104955208064	25.506822	1.408193	0.156793	0.13
HD 78631	CD-51 3504	<i>Gaia</i> DR3 522135261462534528	37.011715	0.876396	0.112947	0.12
HD 99251	BD+08 2498	<i>Gaia</i> DR3 1712614124767394816	28.249344	1.019851	0.083416	0.11
BD+24 4592	HIP 110894	<i>Gaia</i> DR3 3067074530201582336	14.480729	1.580362	0.170151	0.05

Notes. ruwe and astrometric_excess_noise ϵ values are from *Gaia* DR3 main source table. The column \hat{p} , is the probability computed by ExoDNN that the given source hosts a detectable companion.

types F, G and K over M-type sources, see Figure 11 (*f*) values. Therefore the third filtering step removes primarily the brighter and more massive initial candidates. We plan to improve these features in further versions of ExoDNN by applying refined selection criteria for sources around the lower astrometric excess noise significant limit and for those with multiple psfs within the detection window.

4.2. Host characterization

Thanks to the selection of the model input from the DR3 Astrophysical Parameters dataset, we have access to stellar parameters such as the mass, metallicity, and age of the candidate stars. These parameters have been derived by the astrophysical parameter inference system (Apsis, Fouesneau et al. 2023), using the

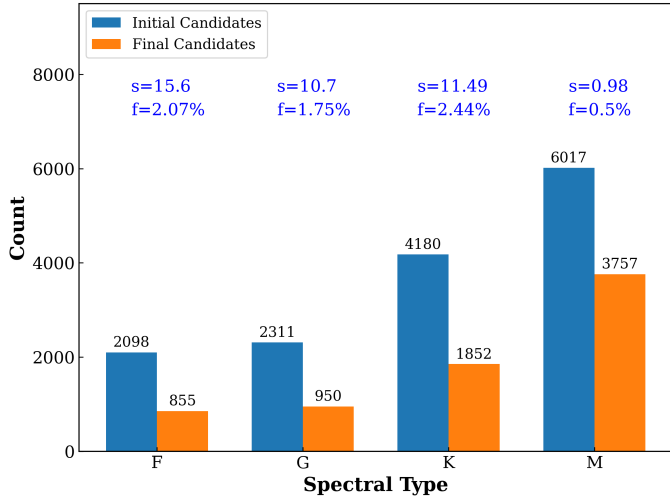


Fig. 11. Overall number of candidates per spectral type (spectraltype_esphs), pre-filtering (blue) and post-filtering (orange). Blue text labels on top of each group of bars correspond respectively to (*s*): median values of the astrometric_excess_noise_sig parameter for the given spectral type and (*f*): median values of the ipd_frac_multi_peak parameter for the given spectral type.

low-resolution spectra provided by *Gaia*'s blue and red photometers (BP/RP), and provide valuable information to characterize our candidate stars population. If we analyse the corresponding distributions of mass, metallicity and age for the candidates reported by ExoDNN, we find a peak mass $M_{\text{cand.}} = 0.7 \pm 0.1 M_{\odot}$ and a peak metallicity of $M/H_{\text{cand.}} = -0.1 \pm 0.01$ dex (see Figure 12). This would suggest that the candidates found by ExoDNN are similar, in terms of stellar properties, to the NSS astro-spectroscopic binaries, but not to the NSS eclipsing binaries (see [Gaia Collaboration 2023a](#)). The model has been trained on a distribution of periods that overlaps with the one from the NSS astro-spectroscopic binaries (typically hundreds of days), which results in a model that is more sensitive to those orbital configurations rather than to the much shorter orbital periods of the eclipsing binaries (typically a few days).

It should be noted that the parameters derived by the Apsis may suffer from calibration problems in the case of unresolved binaries ([Fouesneau et al. 2023](#), Section 3.5.2). This is due to the lack of sufficiently high-quality synthetic models of BP and RP spectra of unresolved binaries and, therefore, our comparison can only be regarded as qualitative.

4.3. Caveats

We remark that in its current version, ExoDNN reports candidate stars to host companions detectable by *Gaia*, but it cannot assess the nature (stellar or substellar) of the potential companions. In general, when the current ExoDNN reports a source as a candidate, we may face different scenarios. One, where the signatures of existing companions in the system are of the order of the DR3 measurement uncertainty. In this case, ExoDNN prediction may be just detecting the signatures of existing companions in the system or alternatively, indicating the presence of a yet unknown companion in a longer period orbit. Another possible scenario occurs when the source has no reported companions yet, but ExoDNN predicts their presence. And yet another scenario may occur, which represents an astrophysical false positive. In

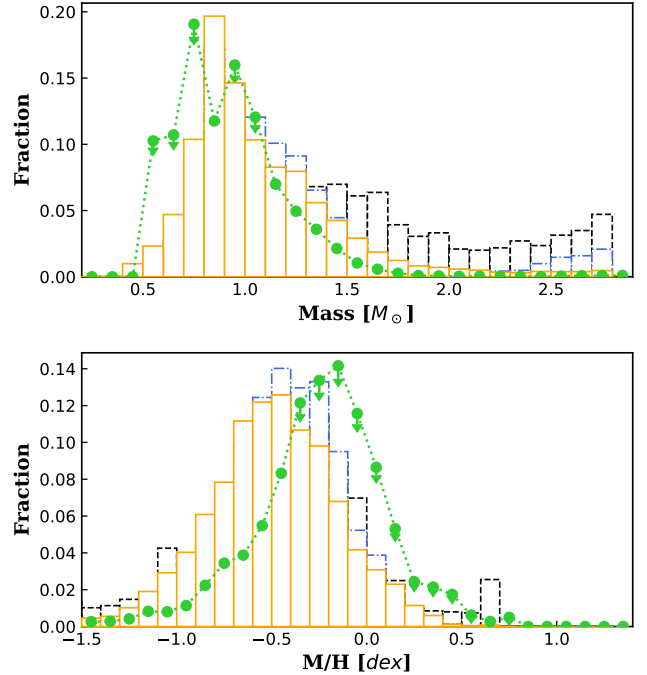


Fig. 12. Top: primary mass distribution of the candidate stars reported by ExoDNN (green). Also added for comparison are the mass distributions corresponding to the NSS Orbital* (orange), AstroSpectroSB1 (blue), and Eclipsing* binary (black) solutions from the DR3 NSS. The ExoDNN predictions show the 6% model false positive rate as downward arrows. Bottom: equivalent distributions but for metallicity of the candidates reported by ExoDNN.

this latter case, the AGIS statistics are inflated but for reasons other than multiplicity, triggering a false positive. Irrespective of the case, supporting evidence such as additional RV measurements, epoch astrometry or imaging data, are necessary to first confirm or rule out the existence of companions, and only then be able to provide estimates on their masses.

5. Conclusions

Using the rich dataset provided in *Gaia* DR3, we created a deep learning model, ExoDNN, that uses the available DR3 parameters to predict the probability of a *Gaia* DR3 star to host one or more companions. To do this, ExoDNN takes advantage of the existing correlation between AGIS fit quality statistics and the deviation from a single source model. We have shown how ExoDNN is able to correctly detect different types of multiplicity, from known binaries on wide orbits to star+brown dwarf configurations. When applied to a sample of stars from the *Gaia* DR3 catalogue within 100 pc, ExoDNN generated a list of 7414 new candidates to host unresolved companions. The number of candidate stars detected is comparable in order of magnitude to predictions from earlier studies (e.g., [Casertano et al. 2008](#); [Perryman et al. 2014](#)). The post-processing steps we applied are conservative, which resulted in the removal of approximately $\sim 50\%$ of the initial candidates found by ExoDNN, but this step was necessary to avoid contamination from sources possibly affected by DR3 AGIS calibration limitations. A false positive rate of $\sim 1.2\%$ is expected from the current version of ExoDNN and should be taken into account when assessing the proposed candidates by performing additional scrutiny on a per-source basis. The main limitation of the current version of ExoDNN is

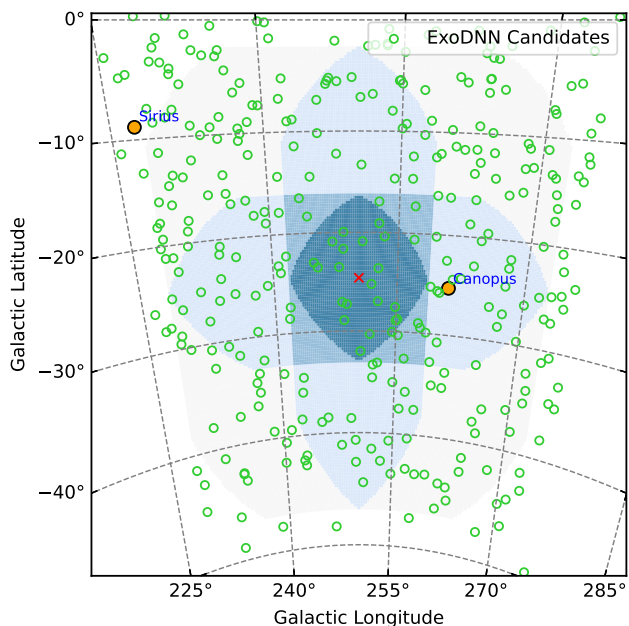


Fig. 13. ExoDNN G-type star candidates overlay in PLATO first field of view (LOPS1).

the inability to determine the nature (stellar or substellar mass) of potential companions without incorporating external data, such as RV measurements or epoch astrometry. In addition, since the ExoDNN predictions are generated from DR3 astrometry, those predictions are limited by the available precision in DR3 measurements, which has been the main driver in applying ExoDNN to volume-limited sample of stars ($d < 100$ pc). Technically speaking, however, nothing prevents applying the model to sources further away (one extreme example is ExoDNN prediction over *Gaia* BH3, located at $d \sim 590$ pc). Despite these limitations and until *Gaia* DR4 is published along with a list of exoplanet candidates, ExoDNN candidates could be a powerful tool for the community for benchmarking different theories of planetary formation given its statistical relevance. In addition, the proposed solar-type (G) candidates could be interesting in the context of future missions such PLATO (Rauer et al. 2022). In Figure 13, we display the G-type candidates found by ExoDNN over the PLATO first field of view, FoV South 1 (240 G-dwarf candidates).

We plan to further enhance the current model performance in several ways. First, we would like to improve the model sensitivity specially over the faint end covering M-dwarfs. Training ExoDNN on additional examples of real M-type stars that are known to host companions could help it learn (at least partially), the intrinsic *Gaia* DR3 calibration artifacts affecting these sources. This, in turn, could eliminate the current need for applying specific excess-noise cuts to M-type candidates. Then, we will also revise the current model architecture to include uncertainty estimations. This will introduce the ability to inform the user on the uncertainty in a given probability prediction and require the evolution of the current neural network architecture towards a deep probabilistic architecture (e.g., Neal 2012; Blundell et al. 2015; Lakshminarayanan et al. 2017; Abdar et al. 2021). Last but not least, to gain the ability to assess the nature of reported candidates, we will need to add external datasets, such as RV values or astrometric time series. This will allow us to constrain the candidate companion masses and increase the

current scope of the model from astrometric detection of substellar companions to a broader one, namely, the characterization of substellar companions.

Some of our candidates are close enough to make them suitable targets for follow-up observations by different means. For candidates at moderate distances, RV follow-up observations with high-resolution spectrographs such as HARPS (Cosentino et al. 2012) or ESPRESSO (Pepe et al. 2021), would allow for the presence of companions to be confirmed or for orbital configuration and masses to be at least partially constrain. We intend to assess the suitability for follow-up observations of the most promising candidates reported in this work and identify the best follow-up strategy.

Data availability

The full version of the 7414 new *Gaia* DR3 candidate stars to host companions is available at the CDS via <https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/704/A150>

Acknowledgements. C.C. acknowledges financial support by the Consejo Superior de Investigaciones Científicas (CSIC) through the internal project 2023AT003 associated to the RYC2021-031640-I. J.L.-B. is funded by the Spanish Ministry of Science and Universities (MICIU/AEI/10.13039/501100011033) and NextGenerationEU/PRTR grants PID2019-107061GB-C61, PID2021-125627OB-C31, CNS2023-144309, and PID2023-150468NB-I00. This research made use of the following software: *astropy*, (a community-developed core Python package for Astronomy, *Astropy Collaboration* 2022), *SciPy* (Virtanen et al. 2020), *matplotlib* (a Python library for publication quality graphics Hunter 2007), *numpy* (Harris et al. 2020) and *TensorFlow* (TensorFlow 2015). This research has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement. We also made use of NASA's Astrophysics Data System (ADS) Bibliographic Services, the SIMBAD database, operated at CDS, and the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program. We also used the data obtained from or tools provided by the portal exoplanet.eu of The Extrasolar Planets Encyclopaedia.

References

- Abdar, M., Pourpanah, F., Hussain, S., et al. 2021, *Information Fusion*, 76, 243
- Alonso-Floriano, F. J., Morales, J. C., Caballero, J. A., et al. 2015, *A&A*, 577, A128
- Ansdell, M., Ioannou, Y., Osborn, H. P., et al. 2018, *ApJ*, 869, L7
- Astropy Collaboration (Price-Whelan, A. M., et al.) 2022, *ApJ*, 935, 167
- Belokurov, V., Penoyre, Z., Oh, S., et al. 2020, *MNRAS*, 496, 1922
- Benedict, G. F., McArthur, B. E., Forveille, T., et al. 2002, *ApJ*, 581, L115
- Benedict, G. F., McArthur, B. E., Gatewood, G., et al. 2006, *AJ*, 132, 2206
- Binnendijk, L. 1960, *Properties of Double Stars: A Survey of Parallaxes and Orbits* (Philadelphia: University of Pennsylvania Press)
- Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. 2015, in *International conference on machine learning*, PMLR, 1613
- Breiman, L. 2001, *Mach. Learn.*, 45, 5
- Bressan, A., Marigo, P., Girardi, L., et al. 2012, *MNRAS*, 427, 127
- Casertano, S., Lattanzi, M. G., Sozzetti, A., et al. 2008, *A&A*, 482, 699
- Chabrier, G., Johansen, A., Janson, M., & Rafikov, R. 2014, in *Protostars and Planets VI*, eds. H. Beuther, R. S. Klessen, C. P. Dullemond, & T. Henning (Tucson: University of Arizona Press), 619
- Cifuentes, C., Caballero, J. A., González-Payo, J., et al. 2025, *A&A*, 693, A228
- Cosentino, R., Lovis, C., Pepe, F., et al. 2012, *SPIE Conf. Ser.*, 8446, 84461V
- Cox, D. R. 1958, *J. R. Stat. Soc.*, 20, 215
- Creevey, O. L., Sordo, R., Pailler, F., et al. 2023, *A&A*, 674, A26
- Curiel, S., Cantó, J., Georgiev, L., Chávez, C. E., & Poveda, A. 2011, *A&A*, 525, A78
- Curiel, S., Ortiz-León, G. N., Mioduszewski, A. J., & Sanchez-Bermudez, J. 2022, *AJ*, 164, 93

- Dattilo, A., Vanderburg, A., & Shallue, C. 2019, *AAS Meeting Abstracts*, 233, 140.16
- de Bruijne, J. H. J., Rygl, K. L. J., & Antoja, T. 2014, *EAS Pub. Ser.*, 67-68, 23
- Duchêne, G., & Kraus, A. 2013, *ARA&A*, 51, 269
- Dvash, E., Peleg, Y., Zucker, S., & Giryes, R. 2022, *AJ*, 163, 237
- Eker, Z., Soydugan, F., Soydugan, E., et al. 2015, *AJ*, 149, 131
- Eker, Z., Soydugan, F., Bilir, S., et al. 2020, *MNRAS*, 496, 3887
- Fabricsius, C., Luri, X., Arenou, F., et al. 2021, *A&A*, 649, A5
- Fischer, D. A., Howard, A. W., Laughlin, G. P., et al. 2014, *Exoplanet Detection Techniques* (Tucson: University of Arizona Press)
- Fitton, S., Tofflemire, B. M., & Kraus, A. L. 2022, *Res. Notes Am. Astron. Soc.*, 6, 18
- Fouesneau, M., Frémat, Y., Andrae, R., et al. 2023, *A&A*, 674, A28
- Gaia Collaboration (Prusti, T., et al.) 2016, *A&A*, 595, A1
- Gaia Collaboration (Smart, R. L., et al.) 2021, *A&A*, 649, A6
- Gaia Collaboration (Arenou, F., et al.) 2023a, *A&A*, 674, A34
- Gaia Collaboration (Creevey, O. L., et al.) 2023b, *A&A*, 674, A39
- Gaia Collaboration (Vallenari, A., et al.) 2023c, *A&A*, 674, A1
- Gaia Collaboration (Panuzzo, P., et al.) 2024, *A&A*, 686, L2
- Gandhi, P., Buckley, D. A. H., Charles, P. A., et al. 2022, *MNRAS*, 510, 3885
- Géron, A. 2017, *Hands-on Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems* (Sebastopol, CA: O'Reilly Media)
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, *Deep Learning* (Cambridge, MA: MIT Press)
- Gosset, E., Damerджи, Y., Morel, T., et al. 2025, *A&A*, 693, A124
- Guiglion, G., Nepal, S., Chiappini, C., et al. 2024, *A&A*, 682, A9
- Halbwachs, J.-L., Pourbaix, D., Arenou, F., et al. 2023, *A&A*, 674, A9
- Hambly, N., Andrae, R., De Angeli, F., et al. 2022, Gaia DR3 Documentation, Chapter 20: Data Model Description, Gaia Data Processing and Analysis Consortium (DPAC), European Space Agency
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357
- Hastie, T., Tibshirani, R., & Friedman, J. 2001, *The Elements of Statistical Learning, Springer Series in Statistics* (New York, NY, USA: Springer New York Inc.)
- Henry, T. J., & Jao, W.-C. 2024, *ARA&A*, 62, 593
- Holl, B., Sozzetti, A., Sahlmann, J., et al. 2023, *A&A*, 674, A10
- Hornik, K., Stinchcombe, M., & White, H. 1989, *Neural Netw.*, 2, 359
- Hunter, J. D. 2007, *CSE*, 9, 90
- Keenan, P. C., & McNeil, R. C. 1989, *ApJS*, 71, 245
- Kingma, D. P., & Ba, J. 2015, in ICLR (Poster), eds. Y. Bengio, & Y. LeCun
- Kroupa, P. 2001, *MNRAS*, 322, 231
- Lakshminarayanan, B., Pritzel, A., & Blundell, C. 2017, *Advances in Neural Information Processing Systems*, 30
- Lépine, S., Hilton, E. J., Mann, A. W., et al. 2013, *AJ*, 145, 102
- Lindegren, L., Lammers, U., Hobbs, D., et al. 2012, *A&A*, 538, A78
- Lindegren, L., Hernández, J., Bombrun, A., et al. 2018, *A&A*, 616, A2
- Lindegren, L., Klioner, S. A., Hernández, J., et al. 2021, *A&A*, 649, A2
- Loaiza-Ganem, G., & Cunningham, J. P. 2019, *Advances in Neural Information Processing Systems*, 32
- Luna, A., Marchetti, T., Rejkuba, M., & Minniti, D. 2023, *A&A*, 677, A185
- Lundberg, S. M., & Lee, S.-I. 2017, *Advances in Neural Information Processing Systems*, 30
- Mason, B. D., Wycoff, G. L., Hartkopf, W. I., Douglass, G. G., & Worley, C. E. 2001, *AJ*, 122, 3466
- Morton, T. D. 2015, Astrophysics Source Code Library [[record ascl:1503.010](#)]
- Neal, R. M. 2012, *Bayesian Learning for Neural Networks* (Springer Science & Business Media), 118
- Newton, E. R., Charbonneau, D., Irwin, J., et al. 2014, *AJ*, 147, 20
- Panahi, A., Zucker, S., Clementini, G., et al. 2022, *A&A*, 663, A101
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, 12, 2825
- Penoyre, Z., Belokurov, V., Evans, N. W., Everall, A., & Koposov, S. E. 2020, *MNRAS*, 495, 321
- Penoyre, Z., Belokurov, V., & Evans, N. W. 2022, *MNRAS*, 513, 5270
- Pepe, F., Cristiani, S., Rebolo, R., et al. 2021, *A&A*, 645, A96
- Perryman, M., Hartman, J., Bakos, G. Á., & Lindegren, L. 2014, *ApJ*, 797, 14
- Pourbaix, D., Tokovinin, A. A., Batten, A. H., et al. 2004, *A&A*, 424, 727
- Pravdo, S. H., Shaklan, S. B., & Lloyd, J. 2005, *ApJ*, 630, 528
- Quirrenbach, A. 2011, *Exoplanets*, ed. S. Seager (Tucson, AZ: University of Arizona Press), 526, 157
- Quirrenbach, A., Passegger, V. M., Trifonov, T., et al. 2022, *A&A*, 663, A48
- Ranaivomanana, P., Uzundag, M., Johnston, C., et al. 2025, *A&A*, 693, A268
- Rauer, H., Aerts, C., Deleuil, M., et al. 2022, *European Planetary Science Congress*, 453
- Reid, I. N., Hawley, S. L., & Gizis, J. E. 1995, *AJ*, 110, 1838
- Ribas, I., Reiniers, A., Zechmeister, M., et al. 2023, *A&A*, 670, A139
- Sahlmann, J., & Gómez, P. 2025, *MNRAS*, 537, 1130
- Sahlmann, J., Lazorenko, P. F., Ségransan, D., et al. 2013, *A&A*, 556, A133
- Sahlmann, J., Triard, A. H. M. J., & Martin, D. V. 2015, *MNRAS*, 447, 287
- Schneider, J., Dedieu, C., Le Sidaner, P., Savalle, R., & Zolotukhin, I. 2011, *A&A*, 532, A79
- Schweitzer, A., Passegger, V. M., Cifuentes, C., et al. 2019, *A&A*, 625, A68
- Scott, D. W. 2015, *Multivariate Density Estimation: Theory, Practice, and Visualization* (Hoboken: Wiley)
- Seager, S., Kuchner, M., Hier-Majumder, C. A., & Militzer, B. 2007, *ApJ*, 669, 1279
- Shallue, C. J., & Vanderburg, A. 2018, *AJ*, 155, 94
- Shan, Y., Revilla, D., Skrzypinski, S. L., et al. 2024, *A&A*, 684, A9
- Shkolnik, E., Liu, M. C., & Reid, I. N. 2009, *ApJ*, 699, 649
- Silverman, B. W. 1986, *Density Estimation for Statistics and Data Analysis*, Chapman & Hall/CRC monographs on statistics and applied probability (London: Chapman and Hall)
- Sozzetti, A., Giacobbe, P., Lattanzi, M. G., et al. 2014, *MNRAS*, 437, 497
- Spiegel, D. S., Burrows, A., & Milsom, J. A. 2011, *ApJ*, 727, 57
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. 2014, *JMLR*, 15, 1929
- Stassun, K. G., & Torres, G. 2021, *ApJ*, 907, L33
- Stefánsson, G., Mahadevan, S., Winn, J. N., et al. 2025, *AJ*, 169, 107
- Stevenson, A. T., Haswell, C. A., Barnes, J. R., & Barstow, J. K. 2023, *MNRAS*, 526, 5155
- Sullivan, K., Kraus, A. L., Berger, T. A., & Huber, D. 2025, *AJ*, 169, 29
- TensorFlow, G. 2015, *Google Research*, 10, s15326985ep4001
- Torres, C. A. O., Quast, G. R., da Silva, L., et al. 2006, *A&A*, 460, 695
- Unger, N., Ségransan, D., Barbato, D., et al. 2023, *A&A*, 680, A16
- Valizadegan, H., Martinho, M. J. S., Wilkens, L. S., et al. 2022, *ApJ*, 926, 120
- van Groeninge, M. G. J., Castro-Ginard, A., Brown, A. G. A., Casamiquela, L., & Jordi, C. 2023, *A&A*, 675, A68
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Methods*, 17, 261

Appendix A: Adding photometric and spectroscopic information

To add photometric and spectroscopic information for every simulated source we first built two differentiated samples, one corresponding to the single stars (negative case) and another for the binary stars (positive case). Each sample contained a combination of 50000 main sequence and giant real stars. In the absence of a large enough sample of bona fide single stars to construct our single stars subset, we resorted to the Gaia DR3 FGKM golden sample (Gaia Collaboration 2023b). These are set of about three million sources with very-high-quality photometry, astrometry, and stellar parameters. From the full FGKM sample we selected sources with no appreciable signs of binarity or multiplicity. To do this we required well-observed single-star solutions with a low *ruwe* ($ruwe \leq 1.2$ Lindegren et al. 2018; Fabricius et al. 2021) and at least 11 visibility periods. We also added further filtering on *ipd_frac_multi_peak* to avoid double stars with a large separation (≥ 1.5 arcsecond. Fabricius et al. 2021) and *ipd_gof_harmonic_amplitude* to also reject pairs with smaller separations. Finally, to minimize contamination from binaries we excluded all sources listed in the Gaia DR3 Non-Single Star (NSS) catalogue (Halbwachs et al. 2023; Gosset et al. 2025) (astrometric and spectroscopic binaries) using the *non_single_star* flag. We joined the resulting subset with the astrophysical parameters DR3 table *gaiadr3.astrophysical_parameters* to obtain the mass estimate for each source. This can be translated into the following query in the Gaia archive:

```
SELECT TOP 100000 A.*, B.* FROM
gaiadr3.gold_sample_fgkm_stars AS B
JOIN gaiadr3.gaia_source AS A USING (source_id)
WHERE A.ruwe < 1.2
AND A.visibility_periods_used > 11
AND A.ipd_frac_multi_peak < 2
AND A.ipd_gof_harmonic_amplitude < 0.1
AND A.non_single_star = 0
```

For the binary stars we did the opposite operation, selecting astrometrically well-observed sources from the main Gaia DR3 table with high *ruwe* value. We removed the cuts on initial parameter determination (IPD) parameters and set the *non_single_star* flag to value 1, which indicates that the source is identified as an astrometric binary in the DR3 NSS catalogue. This can be translated into the following query in the Gaia archive:

```
SELECT TOP 100000 A.*, B.* FROM
gaiadr3.gaia_source AS A
JOIN gaiadr3.astrophysical_parameters AS B USING
(source_id)
WHERE A.ruwe > 1.4
AND A.visibility_periods_used > 11
AND A.non_single_star = 1
```

A description of all the parameters used in these queries is available in the online Gaia documentation⁴.

⁴ Gaia DR3 main source table description: https://gea.esac.esa.int/archive/documentation/GDR3/Gaia_archive/chap_datamodel/sec_dm_main_source_catalogue/ssec_dm_gaia_source.html

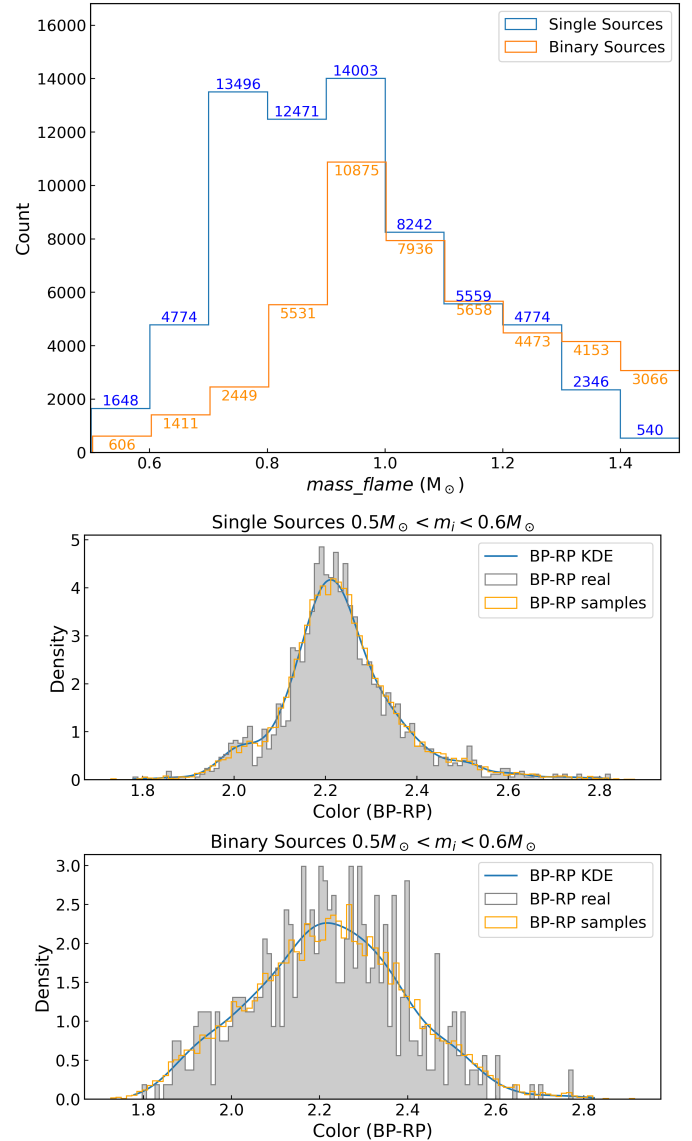


Fig. A.1. TOP: Mass distribution of single (blue) and binary (orange) samples with bins of $0.1M_{\odot}$ on the variable *mass_flame*. MIDDLE: Distribution of color indices for single sources (gray) compared to sampled colors (orange), obtained from a 1D KDE fit corresponding to the low end of the simulation mass range. BOTTOM: Equivalent plot as in the middle but for the binary sources in this case.

For each generated sampled we only retained those stars that have a value of the *mass_flame*⁵ parameter that falls within our simulated primary mass range of ($0.5M_{\odot} \leq M_{flame} \leq 1.5M_{\odot}$). The *mass_flame* provides an estimate of the stellar mass and comes from the FLAME module of the Astrophysical Parameters Inference System (Apsis, Creevey et al. 2023). For single stars this estimated mass refers to the mass of the star, and for binaries it corresponds to the mass of the primary object. Then, to generate synthetic stellar parameters for each star (either single or binary) we used a 2D kernel density estimation (KDE) approach (e.g., Silverman 1986; Scott 2015). This allowed us to conditionally

⁵ A description of this parameter can be found in the Gaia DR3 astrophysical parameters table: https://gea.esac.esa.int/archive/documentation/GDR3/Gaia_archive/chap_datamodel/sec_dm_astrophysical_parameter_tables/ssec_dm_astrophysical_parameters.html

sample values from a 2D KDE fit of the joint probability density function $p(c, X)$ where c is the Gaia $BP - RP$ color of a source and X is a given DR3 parameter. Since our samples contain a combination of dwarf, sub-giants and giants, we considered that a non-parametric fitting approach such as KDE was more adequate to handle the potential multi-modality present in certain stellar parameters, and opted for kernel density estimation versus parametric distribution fitting. The average number of sources per $0.1M_{\odot}$ mass bin in the range ($0.6M_{\odot} \leq M_{flame} \leq 1.4M_{\odot}$) was in the order of thousands, see Figure A.1 (top panel), so this guaranteed that we had sufficient statistics for the KDE sampling. The generation of synthetic parameters was done as follows:

1. We first divided each sample into mass bins of size $0.1M_{\odot}$ using the parameter `mass_flame`.
2. Then, for each star in the given sample and mass bin we extracted the pairs (c_i, X_i) , where c_i is the `bp_rp` color and X_i is the value of the stellar parameter of interest (e.g., magnitude) for the i -th star in the sample.
3. We then constructed a 2D kernel density estimator over the pairs (c_i, X_i) to estimate the joint probability density function $p(c, X)$:

$$\hat{p}(c, X) = \frac{1}{N} \sum_{i=1}^N K((c - c_i, X - X_i); H), \quad (\text{A.1})$$

where $K(x; H)$ is a Gaussian bivariate kernel and H is a so-called bandwidth matrix. The bandwidth matrix H controls the smoothness of the KDE, and is defined as :

$$H = \begin{bmatrix} h_c^2 & h_{cX} \\ h_{cX} & h_X^2 \end{bmatrix}, \quad h_j \propto \sigma_j, \quad j \in \{c, X\}, \quad (\text{A.2})$$

where h_j is proportional to the standard deviation of the corresponding variable σ_j and the off-diagonal elements h_{cX} capture the correlation between color, c , and the variable, X . In practice $K(x; H)$ is automatically computed using `gaussian_kde` from `scipy.stats`.

4. Now, for any desired color, c^* , we compute the conditional probability distribution of X as

$$\hat{p}(X | c^*) = \frac{\hat{p}(c^*, X)}{\int \hat{p}(c^*, X) dX}. \quad (\text{A.3})$$

which yields a smooth, continuous estimate of $p(X|c)$ for any value of c .

This procedure generated a library of $M \times N$ joint probability density functions $p(X|c)$ where M is the number of mass bins and N is the number of synthetic parameters. Still, to use these joint PDFs, we needed to generate the missing value of the Gaia $BP-RP$ color parameter for a given star with simulated mass M_{\star} . To do this we used the `isochrone` software package (Morton 2015) and performed a basic 1D interpolation using a main-sequence MIST isochrone (Bressan et al. 2012). We assumed solar metallicity and fixed 4.6GYr age for every star in our simulated samples, even though we now this is only roughly valid for a subset of the stars we used to generate the 2D KDEs library. However, given that: (1) the influence of photometry and spectroscopy parameters in the model decision is of second-order importance compared to the astrometric fit statistics, and (2) preserving correlations between parameters is of higher importance in this scenario than a high intrinsic parameter accuracy, we opted for this overall process of synthetic parameters generation versus one based on full usage of physical stellar models.

Since the primary goal was to generate a plausible set of photometric and spectroscopic parameters, we considered that this approach was sufficient.

Nevertheless, to avoid any possible unphysical parameter combinations in our generated synthetic photometry, we implemented a coherence test that evaluates whether a simulated set of Gaia ($BP - RP$) color, absolute Gaia magnitude (M_G), and mass (M_{\star}) values are consistent with the physical properties of a single main-sequence star. To do this, we estimated the absolute G -band magnitude of a star with mass, M_{\star} , and color, $BP - RP$, using the following expressions,

$$M_G^{\text{pred}}(M_{\star}, BP - RP) = M_{\text{bol}}(M_{\star}) - BC_G(BP - RP), \quad (\text{A.4})$$

where M_{bol} is the bolometric magnitude associated with a simulated stellar mass M_{\star} , and BC_G is the bolometric correction in the Gaia G band. The bolometric magnitude of the simulated star was computed as

$$M_{\text{bol}}(M_{\star}) = M_{\text{bol},\odot} - 2.5 \log_{10} \left(\frac{L(M_{\star})}{L_{\odot}} \right), \quad (\text{A.5})$$

where $M_{\text{bol},\odot} \simeq 4.74$ and $L(M_{\star}) \simeq M_{\star}^{4.0}$ (Eker et al. 2015). The Gaia G -band bolometric correction (BC_G) was determined by interpolating the empirical calibration grid presented in (Eker et al. 2020, Table 6). In this work the authors provide empirically derived bolometric corrections as a function of intrinsic stellar parameters for solar-type main-sequence stars. We then checked the difference between the predicted absolute magnitude for a given star and the synthetically generated one and only retained those synthetic values for which this difference was lower than a threshold: $|M_G^{\text{pred}} - M_G^{\text{synth}}| < \tau = 2\sigma \text{ mag}$, where σ is the standard deviation of the computed residuals. We show in Figure A.2 the distribution of residuals $M_G^{\text{pred}} - M_G^{\text{synth}}$ and a comparison of the synthetic color-magnitude diagram obtained from our 2D KDE sampling procedure for both single and binary samples.

Now, while our selection procedure for single stars is designed to produce a sample of astrometrically well behaved (this is single) stars, this sample could still be contaminated with short-period binaries, near-equal-mass binaries with face-on orbits, or long-period binaries with a longer period than the DR3 time baseline. To assess the possible impact of introducing wrongly labelled binaries as single systems, we consider the $f1$ -score metric definition,

$$f1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{\text{TP}}{\text{TP} + \frac{1}{2}(\text{FP} + \text{FN})}, \quad (\text{A.6})$$

where TP refers to true positives, FP to false positives, and FN to false negatives. Since mislabelling binaries as single stars effectively increases FN, the $f1$ -score is expected to decrease as a consequence. We can analyse how a small perturbation ΔFN will affect the sensitivity of the $f1$ -score. If we hold TP and FP fixed and take the partial derivative with respect to FN:

$$\frac{\partial f1}{\partial \text{FN}} = -\frac{1}{2} \frac{\text{TP}}{(\text{TP} + \frac{1}{2}(\text{FP} + \text{FN}))^2}. \quad (\text{A.7})$$

Hence, for a small change ΔFN ,

$$\begin{aligned} \Delta f1 &\approx -\frac{1}{2} \frac{\text{TP}}{(\text{TP} + \frac{1}{2}(\text{FP} + \text{FN}))^2} \Delta\text{FN}, \\ \frac{\Delta f1}{f1} &\approx -\frac{\Delta\text{FN}}{2(\text{TP} + \frac{1}{2}(\text{FP} + \text{FN}))}. \end{aligned} \quad (\text{A.8})$$

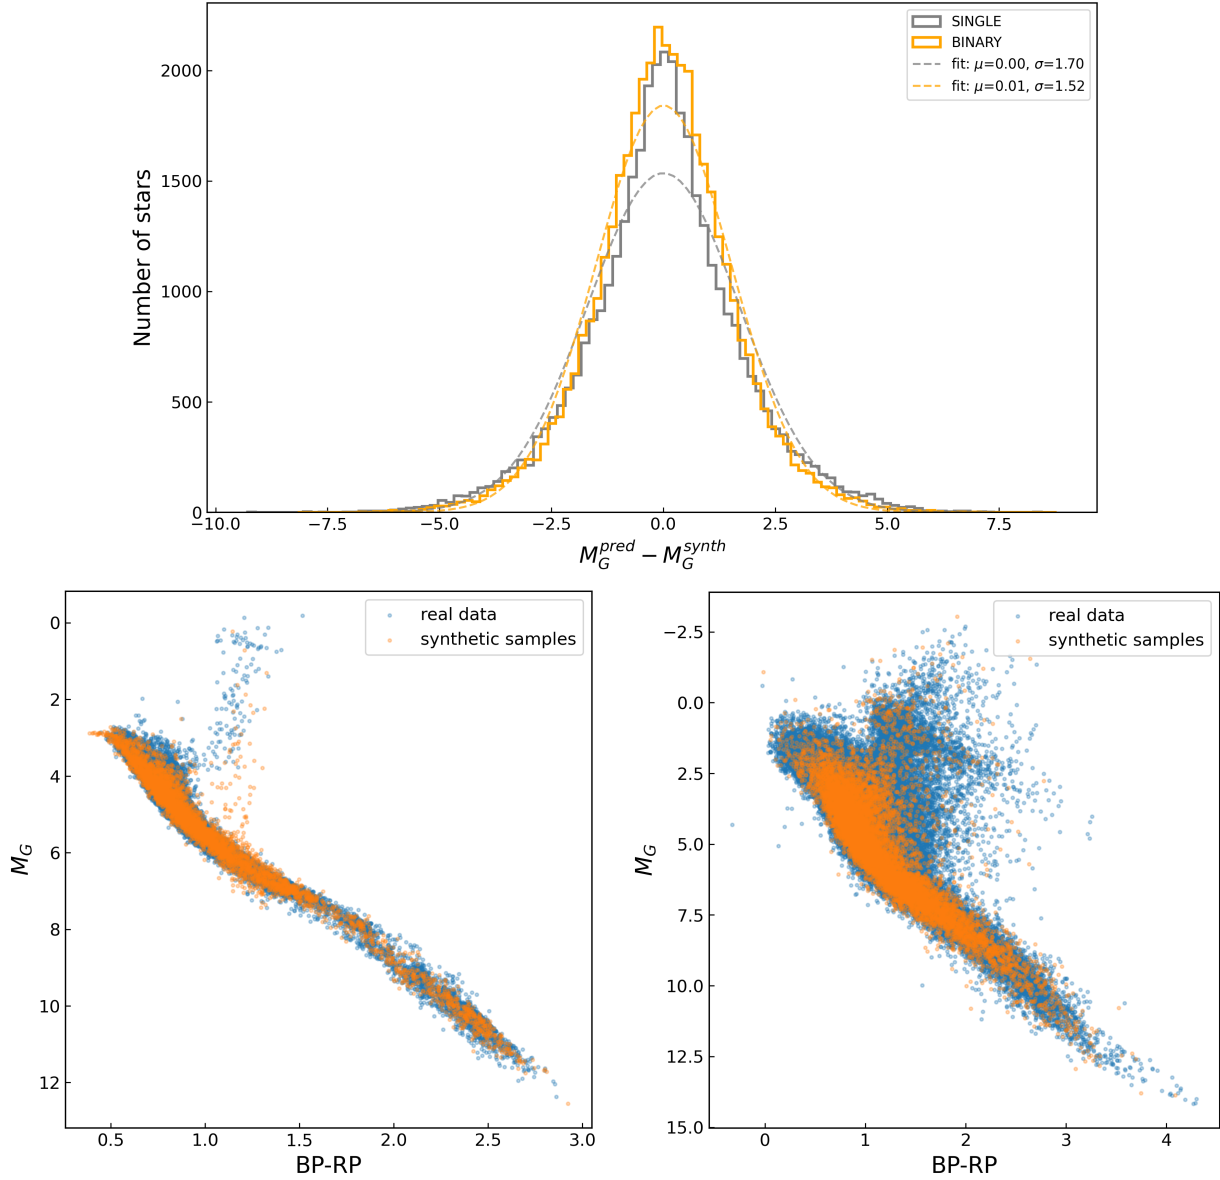


Fig. A.2. TOP: Distribution of the residuals $M_G^{pred} - M_G^{synth}$ for single and binary samples with corresponding Gaussian fits to the residuals. BOTTOM LEFT: Colour-magnitude diagram of the 100000 single (blue) sources and synthetic color-magnitude diagram with conditionally sampled color and magnitude values (orange). BOTTOM RIGHT: equivalent plot for the binaries with real binaries in (blue) and synthetic color-magnitude diagram with conditionally sampled color and magnitude values (orange).

The magnitude of the drop in $f1$ -score depends mainly on the class balance. Using the derived formulae, we can analyse the expected degradation of performance for a balanced, strong model such as our ExoDNN, when a small fraction of FN is added. For example, assuming $TP = 900$, $FP = 100$, $FN = 100$, and using equation A.6 we compute the $f1$ -score as:

$$f1 = \frac{900}{900 + \frac{1}{2}(100 + 100)} = 0.9.$$

The sensitivity in this score is then given by Equation A.7:

$$\frac{\partial f1}{\partial FN} = -\frac{1}{2} \frac{900}{1000^2} = -0.00045.$$

Therefore, if we increase $\Delta FN = 10$ (i.e., +10%), using Equation A.8 we have:

$$\Delta f1 \approx -0.00045 \times 10 = -0.0045, \quad f1' \approx 0.8955,$$

and we expect an absolute drop of 0.0045 in the $f1$ -score in this case. For $\Delta FN = 5$, the drop is halved, $f1' \approx 0.8978$. We can see how for a balanced model the impact of introducing a small fraction ($\leq 10\%$) of false negatives is relatively small.

Appendix B: Additional models

For a relative performance comparison against ExoDNN, the selected binary classifiers were logistic regression and random forest. The logistic regression is similar to our network in the sense that it is a probabilistic classifier whose objective is to minimize the log-loss, but it is fundamentally a linear model which cannot capture complex non-linear relationships unlike our ExoDNN. Random Forest, on the other side, is a non-parametric approach that uses an ensemble of decision trees to learn decision boundaries via recursive feature splits.

We trained both classifiers using the same training data used to train ExoDNN and applied equivalent pre-processing steps, namely, robust scaling of the data before training. Table B.1 and Table B.2 list the different hyper-parameters used to train the alternative classifiers.

Table B.1. Best hyper-parameters for logistic regression

Hyper-parameter	Value
C	0.8
max_iter	100
solver	newton-cg
tol	0.0001

Table B.2. Best hyper-parameters for random forest

Hyper-parameter	Value
bootstrap	true
max_depth	1
max_features	4
min_samples_leaf	1
min_samples_split	2
n_estimators	50

Appendix C: Model input selection

To generate the input for the deep learning model prediction we selected sources from the Gaia main source table (`gaiadr3.gaia_source`) in the Gaia archive via the following query:

```
SELECT A.*, B.* FROM gaiadr3.gaia_source AS A
JOIN gaiadr3.astrophysical_parameters AS B USING
(source_id)
WHERE A.parallax >10
AND A.visibility_periods_used>11
AND A.teff_gspphot > 2500
AND A.teff_gspphot < 7500
AND A.phot_variable_flag NOT LIKE 'VARIABLE'
```

The join with the astrophysical parameters DR3 table (`gaiadr3.astrophysical_parameters`) is required to obtain estimated stellar parameters for each source generated by the astrophysical parameters inference system (Apsis, Creevey et al. 2023). Figure C.1 shows the number of sources retrieved by the query with a breakdown per spectral type.

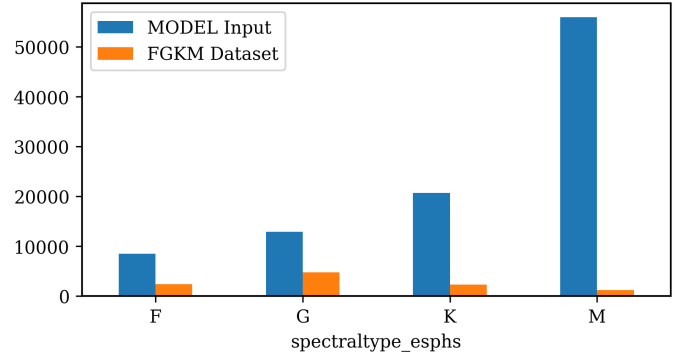


Fig. C.1. Comparison of the number of sources per spectral type between the inputs used to make model prediction and the FGKM golden sample (Creevey et al. 2023)