Astronomy & Astrophysics manuscript no. main
July 15, 2025
©ESO 2025

Blended starlight reveals millions of candidate binaries in Gaia spectra.

Millions of Main-Sequence Binary Stars from Gaia BP/RP Spectra

Jiadong Li[1], Hans-Walter Rix[1], Yuan-Sen Ting[2,3,1], Johanna Müller-Horn[1], Kareem El-Badry[4,1], Chao Liu[5,6,7,8], Rhys Seeburger[1], Gregory M. Green[1], and Xiangyu Zhang[1]

[1] Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany
e-mail: jdli@mpia.de
[2] Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA
[3] Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA
[4] Department of Astronomy, California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA
[5] Key Lab of Space Astronomy and Technology, National Astronomical Observatories, Beijing, 100101, China
[6] Institute for Frontiers in Astronomy and Astrophysics of Beijing Normal University, Beijing, 100875, China
[7] University of Chinese Academy of Sciences, Beijing, 100049, China
[8] Zhejiang Lab, Hangzhou, 311121, China

The catalogue identifies candidates; it does not confirm every system individually.

July 15, 2025

arXiv:2507.09622v1 [astro-ph.SR] 13 Jul 2025

ABSTRACT

We present an extensive catalog of likely main-sequence binary stars systems (MSMS), derived from Gaia Data Release 3 BP/RP (XP) spectra through the comparison of single- and binary-star model fits. Leveraging the large sample of low-resolution Gaia XP spectra, we use a neural network to build a forward modeling for spectral luminosities of single-stars, as a function of stellar mass and photometric metallicity. Applying this model to XP spctra, we find that this enables the identification of binaries with mass ratios between 0.5 and 1.0 and flux ratios > 0.1 as "poor" fits to the data, either in spectral shape or flux normalization. From an initial sample of 35 million stars within 1 kpc, we identify 14 million possible binary candidates, and a high-confidence "golden sample" of 1 million binary systems. This large, homogeneous sample of SED-based binaries enables population studies of luminous MSMS binaries, and – in conjunction with kinematic or astrometric probes – permits to identify binaries with dark or dim companions, such as white dwarfs, neutron stars and black hole candidates, improving our understanding of compact object populations.

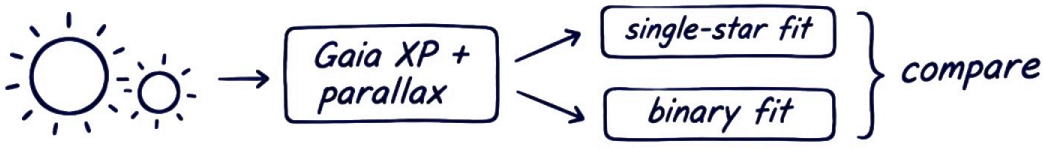
1. Introduction

Binary star systems constitute fundamental components of the stellar population, with half of solar-type stars residing in multiple systems (Raghavan et al. 2010; Duchêne & Kraus 2013; Moe et al. 2019). Binary systems provide an important context and set of constraints for most sub-fields of astrophysics (Rix et al. 2019). The statistical properties of binary populations—including the binary fraction, period distribution, mass ratio distribution, and eccentricity distribution—provide constraints on star formation theories and dynamical evolution models. Moreover, these characteristics vary systematically with stellar properties such as stellar mass, age, and metallicity, creating distinctive patterns across different stellar populations and offering valuable insights into the physical mechanisms governing both the formation and subsequent evolution of stellar systems (Duchêne & Kraus 2013; Moe et al. 2019; El-Badry et al. 2021).

The periods of binary systems span a wide range in timescales (from minutes to Myr). Current stellar surveys such as *Gaia*, LAMOST, SDSS-IV/APOGEE and SDSS-V (Gaia Collaboration et al. 2016; Liu et al. 2020; Majewski et al. 2017; Kollmeier et al. 2017) primarily identify binaries through spectroscopic (El-Badry et al. 2018; Price-Whelan et al. 2020; Zhang et al. 2022) or astrometric signatures of orbital motion (e.g. Holl et al. 2023; Penoyre et al. 2022a,b). Previous studies of unresolved binaries have generally relied on spectroscopic methods that are most effective for systems exhibiting radial velocity variations, such as identifying double-lined spectroscopic binaries (SB2s) through cross-correlation techniques (e.g., Li et al. 2021) or disentangling component spectra from spectra taken at different orbital phases (e.g., Seeburger et al. 2024). These methods are sensitive only to systems with notable velocity separations ($v_{\rm los} \gtrsim 10 - 30$ km s^{-1}) and high mass ratios ($q \geq 0.7$) (Zhang et al. 2022). This approach inherently limits detection to relatively short-period binaries, which represent only about one-third of the total binary population (Duchêne & Kraus 2013).

Similarly, the identification of single-lined spectroscopic binaries (SB1s) through radial velocity variability is also biased toward shorter period systems. Traditional methods are effective within specific period ranges but miss many systems outside their sensitivity windows. Spectroscopic techniques, such as identifying single-lined (SB1) and double-lined (SB2) spectroscopic binaries, are limited to systems with periods of approximately 1 to 1000 days, where radial velocity (RV) variations are detectable (El-Badry 2024). Astrometric methods, on the other hand, are most sensitive to binaries with periods comparable to *Gaia*'s observational baseline (Castro-Ginard et al. 2024). In contrast, Spectral Energy Distribution (SED)-based techniques remain effective across the full range of periods for unresolved binaries with distinct components, as they do not rely on radial velocity or astrometric variability.

This methodology bridges detection gaps, complementing short-period sensitivity of eclipsing binary searches (requiring specific orbital alignments) and long-period capabilities of direct imaging (typically exceeding 10,000 days for Gaia). SED-based approaches can detect the estimated 73% of binaries with periods beyond 10 years (El-Badry et al. 2018; Duchêne & Kraus 2013) that remain invisible to conventional RV surveys (Gao et al. 2014; Badenes et al. 2018), while simultaneously providing independent verification for systems within the detection windows of other methods. For binaries with periods too long for dynamical detection yet too tight for spatial resolution by *Gaia*,

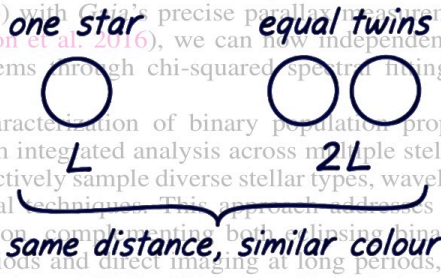


our XP-based methodology may represent a key identification technique,

The SED method identifies binary stars by modeling the composite spectral energy distribution of two stellar components and detecting deviations from single-star models, with enhanced sensitivity when the components exhibit obvious temperature or luminosity contrasts (e.g., Aznar Cuadrado & Jeffery 2001). The SED-based methods complement these approaches by operating independently of orbital dynamics, identifying binaries through their composite spectral signatures regardless of period. These signals are most pronounced when the two components have different spectral features, as is the case with white dwarf–main-sequence binaries (Li et al. 2025).

However, the presence of an unresolved companion changes the observable spectrum in all cases where the two stars are not identical. When two main-sequence stars have similar masses, the normalized spectral features or normalized SED show minimal differences between single and binary systems, creating a potential degeneracy (El-Badry et al. 2018). However, with precise distance measurements (e.g., from *Gaia* parallaxes), this degeneracy can be resolved: a binary system's total luminosity is the sum of its two components, so for stars with similar spectral features, a binary would appear twice as luminous as a single star. When forced to fit such a binary with a single-star model, the excess luminosity is misinterpreted as a hotter (bluer) temperature to compensate for the unphysically high flux at the known distance.

Together, these complementary detection techniques enable stellar surveys to deliver samples of binary stars and stellar companions orders of magnitude larger than previously known, spanning all stages of stellar evolution. Even without detectable velocity offsets, unresolved binary spectra contain distinctive signatures of unseen companions. By combining the quality of *Gaia* DR3 low-resolution BP/RP (XP) spectra (Gaia Collaboration et al. 2023b) with *Gaia*'s precise parallax measurements (Gaia Collaboration et al. 2016), we can now independently identify binary systems through chi-squared spectral fitting (Li et al. 2025).



The characterization of binary population properties will emerge from integrated analysis across multiple stellar surveys, which collectively sample diverse stellar types, wavelengths, and observational techniques. This approach addresses gaps in binary detection, complementing both eclipsing-binary searches at short periods and direct imaging at long periods, while capturing the estimated 70% of binaries with periods exceeding 10 years that traditional RV surveys miss (Duchêne & Kraus 2013; El-Badry et al. 2018) for binaries with periods too long for dynamical detection yet too close for spatial resolution.

In this paper, we focus on main-sequence main-sequence binary (MSMS) stars from *Gaia* XP spectra, presenting an inventory of unresolved binary systems in which both components are main-sequence stars. We employed the high quality of Gaia DR3 low-resolution XP spectra (Gaia Collaboration et al. 2023b), combined with precise parallax measurements from *Gaia*, (Gaia Collaboration et al. 2016), we conduct a homogeneous search for binary systems throughout the Galaxy, identifying and characterizing these systems with completeness and consistency.

Our approach is based on the methodology introduced by Li et al. (2025). We utilized *Gaia* XP spectra, which provide low-resolution spectrophotometry for over 220 million sources in the Galaxy. Although these spectra have a lower resolution ($R \sim 20 - 100$) than traditional spectroscopic surveys, they offer sufficient signal-to-noise and self-consistent flux calibration (De Angeli et al. 2023). By modeling the composite spectral energy distributions of binary components, we can identify binaries regardless of their velocity separation, including long-period systems with negligible velocity offsets between components, provided they have favorable mass ratios ($0.5 \lesssim q \lesssim 1.0$).

In this work, we present a MSMS catalog derived from *Gaia* XP spectra, containing over 14 million binary candidates. The remainder of the paper is structured as follows: In Section 2, we describe the *Gaia* DR3 data and our sample selection criteria. Our methodology, described in Section 3, builds upon our forward model for stellar spectra mapping intrinsic parameters: stellar mass (M), photometric metallicity ($[M/H]_{phot}$), and extinction (E) to predicted spectra. In section 4, we present the validation results and the method to classify the single and binary stars using parameter-dependent thresholds. Section 5 details the statistical properties of our sample, including our binary candidate selection process and the creation of a high-confidence "golden sample" of approximately 1 million binary systems. We discuss the implications of our findings and summarize our conclusions in Section 6. This catalog represents a valuable resource for investigating the formation and evolution of binary systems throughout the Galaxy. The size and diversity of our sample enables robust statistical analyses of binary system properties across different Galactic environments, offering new insights into the processes governing binary star formation and evolution.

Parallax supplies the distance needed to compare absolute brightness.

2. Data

In this section, we first describe the properties and processing of the *Gaia* BP/RP spectra, followed by a detailed outline of the sample selection criteria and pre-processing steps.

2.1. XP: The Gaia BP/RP Low-Resolution Spectra

Gaia, DR3 provides low-resolution spectrophotometric data (XP spectra) for about 220 million sources, obtained with the Blue Photometer (BP, 330–680 nm) and Red Photometer (RP, 640–1050 nm) instruments onboard the spacecraft (Gaia Collaboration et al. 2023b). These spectra cover the optical to near-infrared range with a wavelength-dependent resolving power of R~20-100.

Unlike conventional spectroscopic data consisting of discrete flux measurements, *Gaia* XP spectra implement a continuous functional representation through basis function decomposition (De Angeli et al. 2023). Each spectrum is parameterized via 55 coefficients derived from orthonormal Gauss-Hermite functions, optimized to capture both broadband morphology and higher-frequency spectral features with minimal information loss. For analytical tractability, we transform these coefficient-based representations to wavelength-flux space using `GaiaXPy`[1], facilitating direct identification and measurement of diagnostic spectral features as described in Appendix A. This transformation yields calibrated flux values in physical units ($W\,nm^{-1}\,m^{-2}$) spanning the operational wavelength range (390–1000 nm). The low-resolution XP spectrophotometry provides sufficient spectral discrimination capability for identifying key stellar atmospheric diagnostics (Zhang et al. 2023; Li et al. 2024) while maintaining the statistical power of *Gaia*'s extensive survey volume, rendering it suited for our binary census and characterization program.

[1] http://doi.org/10.5281/zenodo.6674521

Learn the single-star sequence before searching for excess companion light.

2.2. Sample Selection Criteria

To construct our parent sample, we execute the following query on the *Gaia* DR3 database:

```
SELECT *
FROM gaiadr3.gaia_source
WHERE parallax > 1
  AND phot_bp_mean_mag < 20
  AND bp_rp BETWEEN 0 AND 5
  AND parallax_over_error > 10
  AND phot_g_mean_mag + 5*log10(abs(parallax)/100) > 3
  AND has_xp_continuous = 'true'
```

Our sample selection procedure implements multi-stage filtering protocols: First, we establish a volume-limited sample within a 1 kpc heliocentric radius (parallax $\varpi > 1$ mas) containing objects with available BP/RP (XP) continuous spectra and B-R color indices spanning $0 < \mathrm{B} - \mathrm{R} < 5.0$, encompassing F-type to M-type stars. Second, we mitigate systematic flux calibration biases that affect faint red sources in the *Gaia* photometric system by implementing a magnitude threshold of $B < 20$ mag, where these systematic effects remain below our precision requirements. Third, we excise evolved stellar populations by imposing an absolute magnitude criterion $M_G > 3.0$ mag (undereddened), removing giants and supergiants while retaining the main sequence population. We focus on main-sequence stars to ensure homogeneity in our sample, as their spectral properties are primarily determined by mass and metallicity, simplifying the forward modeling in this study. In future work, we plan to incorporate stellar age explicitly, particularly for brighter stars, by forward-modeling spectra with mass, age, and [M/H] as parameters. Fourth, we enforce an astrometric quality constraints, requiring parallax signal-to-noise ratios $\varpi/\sigma_\varpi > 10$ to ensure accurate absolute magnitude determinations. The resulting parent sample comprises 35,831,031 unique sources with *Gaia* DR3 identifiers.

2.3. Training Data

The nearby, high-S/N training set is selected to reduce binary contamination.

To build a clean training sample of main sequence stars for our data-driven emulator of XP spectra, we have implemented a multi-step filtering process targeting the main sequence within ~333 parsecs (`parallax`> 3). Our selection methodology prioritizes purity while maintaining a representative population across the observed main sequence.

We first, we applied a color-magnitude diagram (CMD) filter using PARSEC theoretical isochrones with parameters of 1 Gyr age and [M/H]= 0.2 metallicity, and a hard cut of $M_G > 3$ as shown in Figure 1. By excluding objects above this reference isochrone, we removed most of the main sequence binaries and pre-main sequence objects that would contaminate our sample. Second, we remove the white dwarf main sequence binary (WDMS) candidates selected from Li et al. (2025). As a final quality assurance measure, we required all spectra to have an average signal-to-noise ratio greater than 70. The resulting clean sample contains 746,548 main sequence stars with high quality XP spectra.

3. Data-Driven Spectral Model

In this section, we present our data-driven approach for modeling the XP spectra of both single and MSMS stars using neural networks (NN). Our methodology encompasses two interconnected components: (1) a single-star spectral emulator that maps stellar parameters directly to observed spectra in Subsection 3.1, and (2) a binary star model that combines individual stellar spectra based on fundamental stellar properties in Subsection 3.2. Our NN architecture employs a simple yet streamlined design for both single and binary star modeling. For single stars, we map photometric stellar parameters directly to spectral flux using a neural network. For binary systems, we transform physical parameters into component-specific stellar properties, and combine the resulting spectra through linear superposition.

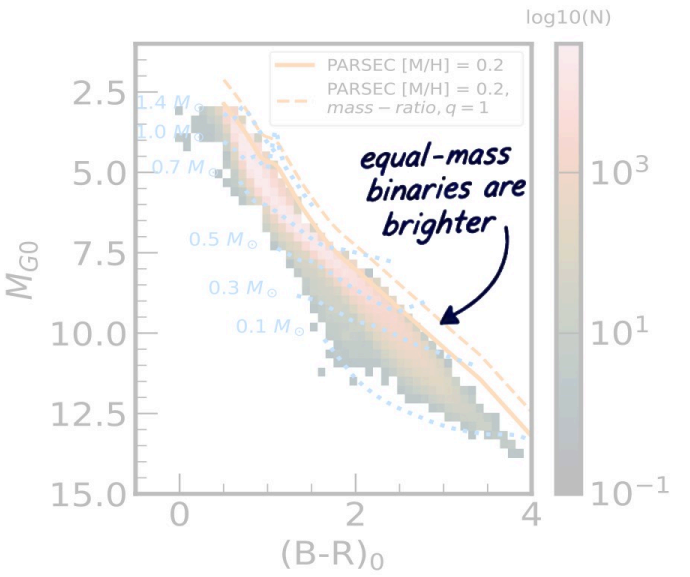


Fig. 1: Hertzsprung–Russell diagram of training sample stars. The color scale indicates the logarithm of the number of stars in each bin. The solid orange line represents a PARSEC isochrone with metallicity ([M/H] = 0.2), while the dashed orange line shows the same isochrone for equal-mass binary systems (mass ratio $q = 1$). The blue dotted lines indicate iso-mass tracks, with masses ranging from 0.1 $M_\odot$ to 1.4 $M_\odot$ as labeled.

To enhance the capability of modeling single-star XP spectra, we iteratively cleaned the training set three times, removing outliers based on fit quality metrics. This process involved evaluating the χ^2 statistic for each spectra in the training set, where high χ^2 values indicated poor agreement between observed XP spectra and the NN's predicted spectra. Outliers, defined as spectra with χ^2 values exceeding 3, were removed to ensure the training data consisted of representative single-star spectra. Three iterations of this cleaning were performed to refine the dataset progressively, each time retraining the NN to improve its accuracy in mapping stellar parameters.

We adopt photometric quantities directly related to the Hertzsprung-Russell (H-R) diagram as our stellar labels. Specifically, we use

$$\boldsymbol{\theta} = (B - R, M_G), \tag{1}$$

where $B - R$ is the dust-corrected color index, and M_G is the absolute magnitude in the Gaia G band, corrected for extinction using 3D dust maps (Green 2018; Edenhofer et al. 2024).

The de-reddening process follows the method of Li et al. (2025). We utilize the `dustmaps` package (Green 2018) to query

Colour changes spectral shape; absolute magnitude changes the brightness scale.

vary colour at fixed magnitude

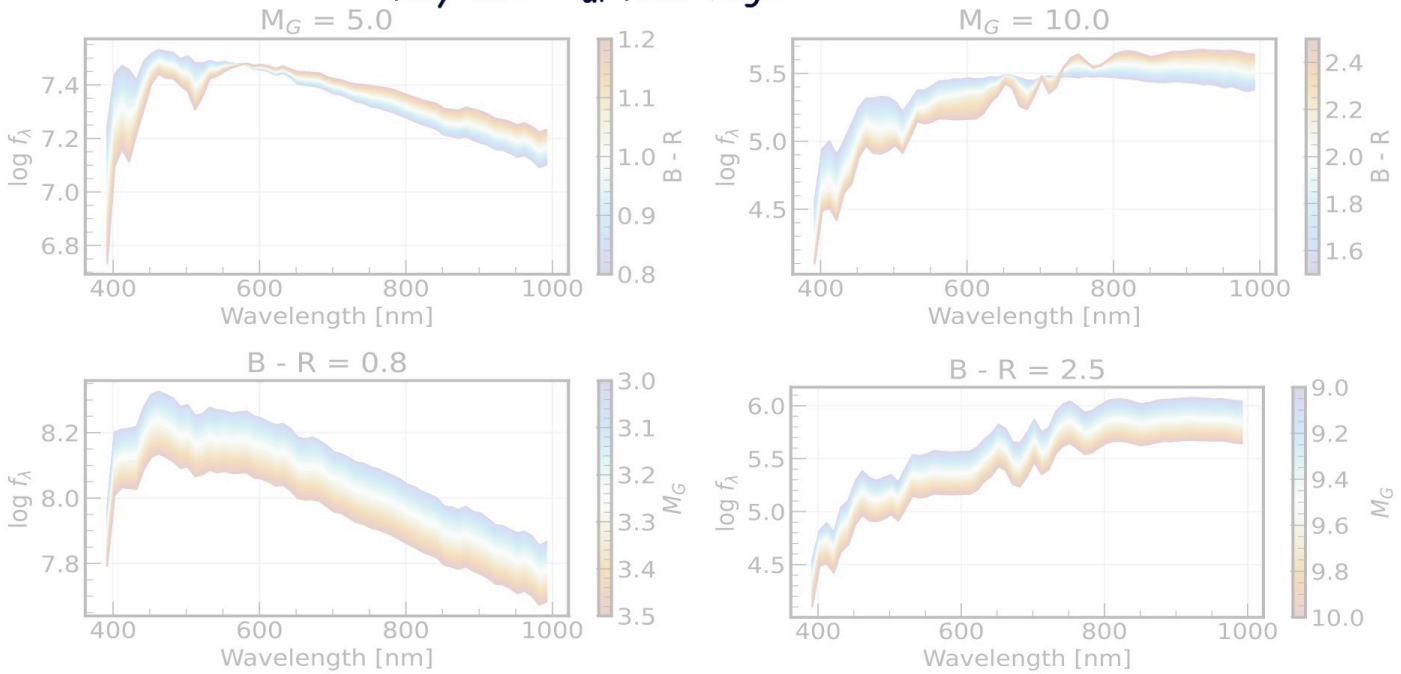


Fig. 2: Demonstration of our NN spectral emulators for single MS stars. Each panel shows the predicted XP spectra with different stellar parameters, where the x-axis is the wavelength and the y-axis is the logarithmic flux (10^{-20} W m^{-2} nm) normalised to a distance of 10 pc. Top panels: Spectra at fixed absolute G magnitude (M_G) with different B-R color. Bottom panels: Spectra at fixed B-R color with varying M_G.

vary magnitude at fixed colour

a 3D dust map and derive the extinction parameter E.[2] We adopt the average extinction curve from Zhang et al. (2023), combined with the transmission curves of the *Gaia* passbands, to compute the extinction corrections.

This choice of M_G is motivated by the fact that M_G provides a reliable and physically meaningful representation of low-mass stars ($M_G \geq 4$–5). In contrast, $\log g$ is not an ideal descriptor for low-mass stars because it becomes less sensitive to variations in luminosity in this regime. In addition, our sample spans a wide range of spectral types, from M-type to G-type stars, making it difficult to define a homogeneous set of $T_{\rm eff}$ labels. The parameters ($T_{\rm eff}$, $\log g$, [M/H]) of M-type stars are often parameterized differently from previous spectral types (e.g., Li et al. 2021; Qiu et al. 2024), leading to inconsistencies when using $T_{\rm eff}$ as a universal label. Furthermore, our approach is purely data-driven, relying on observational data rather than theoretical assumptions. By focusing on photometric quantities such as $B-R$ and M_G, we ensure that our model is both reliable and applicable to the full range of stellar types in our sample.

3.1. Spectral Emulator

The neural network architecture is designed to be as simple as possible, consisting of three hidden layers, each containing 16 neurons activated by Rectified Linear Units (ReLU). The configuration optimizes computational efficiency and spectral modeling. While XP spectra in logarithmic absolute flux space exhibit approximately linear relationships with absolute magnitude and colors, the spectral features and detailed line profiles require nonlinear transformations to accurately represent the underlying physics of stellar atmospheres.

The neural network defines a function f that maps from stellar parameters to the predicted XP spectrum:

$$f(\lambda; \boldsymbol{\theta}, \boldsymbol{w}) = \mathbf{F}(\lambda), \tag{2}$$

where $\mathbf{F}(\lambda)$ represents the flux at wavelength λ, $\boldsymbol{\theta}$ are the stellar parameters, and $\boldsymbol{w}$ denotes the trainable parameters of the network.

We train the network using the Adam optimizer with a learning rate of 10^{-4} and a batch size of 16,384. The loss function is defined as the chi-squared (χ^2) statistic, which measures the discrepancy between the observed and predicted spectra:

$$\begin{aligned}\mathcal{L} &= \frac{1}{N}\sum_{i=1}^{N}\chi_i^2 \\ &= \frac{1}{N}\sum_{i=1}^{N}(\mathbf{F}_i^{\rm obs}(\lambda) - f(\lambda;\boldsymbol{\theta}_i,\boldsymbol{w}))^T\mathbf{C}_i^{-1}(\mathbf{F}_i^{\rm obs}(\lambda) - f(\lambda;\boldsymbol{\theta}_i,\boldsymbol{w})),\end{aligned} \tag{3}$$

where $\mathbf{F}_i^{\rm obs}(\lambda)$ represents the observed spectrum of the i-th star, $\mathbf{C}_i$ is the covariance matrix for the i-th observation, and N is the total number of stars in the training set.

3.2. Binary Spectral Model

3.2.1. Binary System Parameterization

For modeling binary star systems, we perform stellar parameterization based on stellar properties.

[2] The dust map provides E in units of mag per parsec, which can be converted to extinction at any wavelength by applying the appropriate extinction curve.

Build a binary spectrum by adding the two components in linear flux.

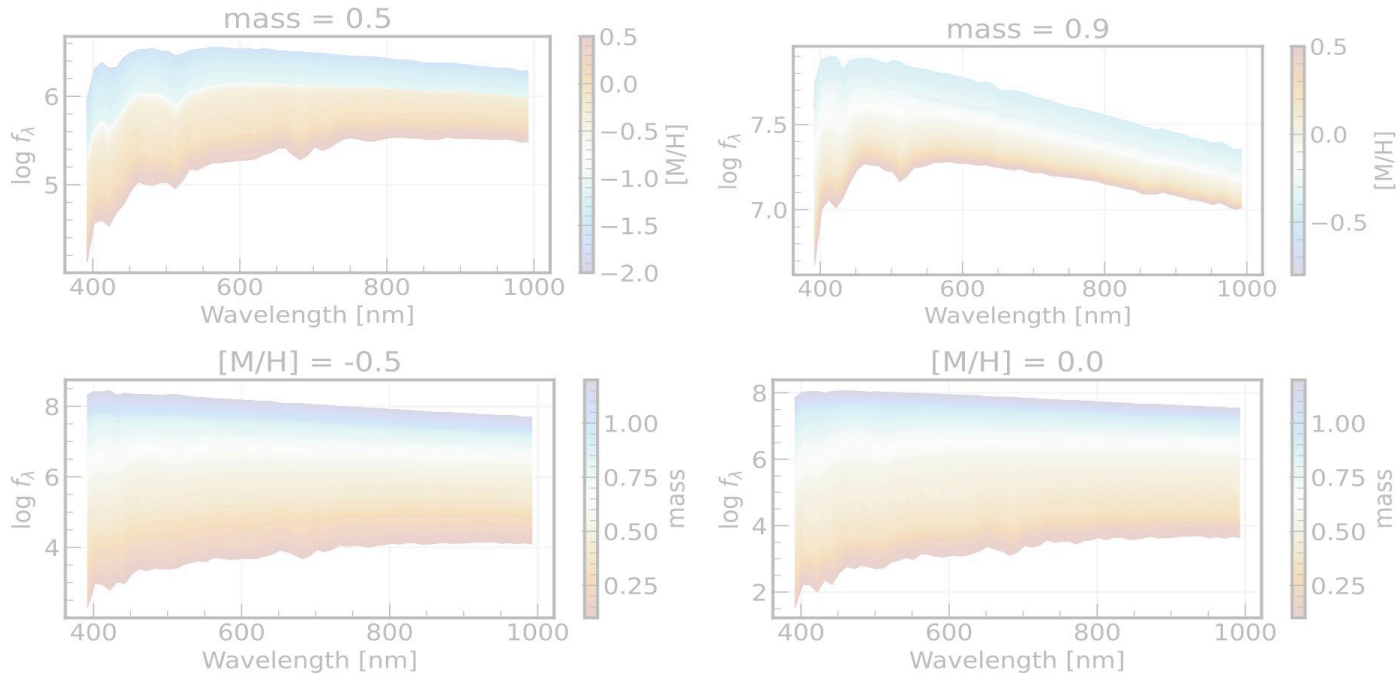


Fig. 3: Demonstration of NN spectral emulators for single MS stars. Top panels: Spectra at fixed stellar mass with different $[M/H]_{phot}$. Bottom panels: Spectra at fixed $[M/H]_{phot}$ with varying stellar mass.

Each binary system is characterized by three primary parameters:

$$\boldsymbol{\phi} = (m_1, [\mathrm{M/H}]_{\mathrm{photo}}, q), \tag{4}$$

where m_1 represents the mass of the primary star, $[M/H]_{phot}$ is the metallicity (assumed identical for both components), and $q = m_2/m_1$ is the mass ratio, constrained to $0 < q \leq 1$. This parameterization captures the key properties of binary systems while minimizing the number of free parameters.

3.2.2. Mapping Stellar Parameters to the H-R Diagram

To translate the physical parameters into observable quantities on the H-R diagram, we use the PARSEC stellar evolution models (Bressan et al. 2012; Chen et al. 2014) with ages of 5 Gyr. We implement this mapping using a neural network function $\mathcal{M}$ that captures the nonlinear transformations:

$$\boldsymbol{\theta} = \mathcal{M}(m_1, [\mathrm{M/H}]_{\mathrm{photo}}), \tag{5}$$

where $\boldsymbol{\theta} = (B{-}R, M_G)$ represents the position in the H-R diagram expressed as dust-corrected color index and absolute magnitude.

Both stars share metallicity and distance; q sets the secondary mass.

The photometric properties of each component in a binary system are computed separately:

$$\boldsymbol{\theta}_1 = \mathcal{M}(m_1, [\mathrm{M/H}]_{\mathrm{photo}}), \tag{6}$$

$$\boldsymbol{\theta}_2 = \mathcal{M}(q, m_1, [\mathrm{M/H}]_{\mathrm{photo}}). \tag{7}$$

Here, $[M/H]_{photo}$ denotes the photometric metallicity derived from broadband colors. The photometric metallicity reflects the integrated effect of metal content on the stellar atmosphere's opacity and emergent flux, which may differ from element-specific spectroscopic determinations. Full details of this mapping implementation and its validation are provided in Appendix B.

3.2.3. Combined Binary Spectrum

Once the H-R diagram positions of both components are determined, we predict the individual spectra using the same mapping function f from Section 3.1:

$$\mathbf{F}_1(\lambda) = f(\boldsymbol{\theta}_1), \tag{8}$$

$$\mathbf{F}_2(\lambda) = f(\boldsymbol{\theta}_2). \tag{9}$$

The combined spectrum of the unresolved binary system is then computed as the sum of the individual component spectra:

$$\mathbf{F}_{\mathrm{binary}}(\lambda; m_1, [\mathrm{M/H}]_{\mathrm{photo}}, q) = \mathbf{F}_1(\lambda) + \mathbf{F}_2(\lambda). \tag{10}$$

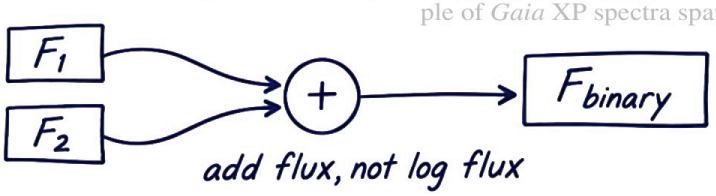
This additive model assumes the flux contributions from the two stars are linearly superimposed, which is a valid approximation for non-interacting binary systems.

3.2.4. Extinction Modeling

We model extinction as a wavelength-dependent attenuation of stellar flux following standard extinction laws. The extinction is characterized by a parameter E, which scales the universal extinction curve $R(\lambda)$. The observed flux is related to the intrinsic flux by:

$$\mathbf{F}_{\mathrm{obs}}(\lambda) = \mathbf{F}_{\mathrm{intrinsic}}(\lambda) \cdot \exp(-E \cdot R(\lambda)), \tag{11}$$

where $R(\lambda)$ represents the relative extinction at each wavelength. For our analysis, we adopt the average extinction curve derived by Zhang et al. (2023), which was determined from a large sample of *Gaia* XP spectra spanning diverse Galactic environments.



At fixed distance, companion light raises the combined spectrum above the primary alone.

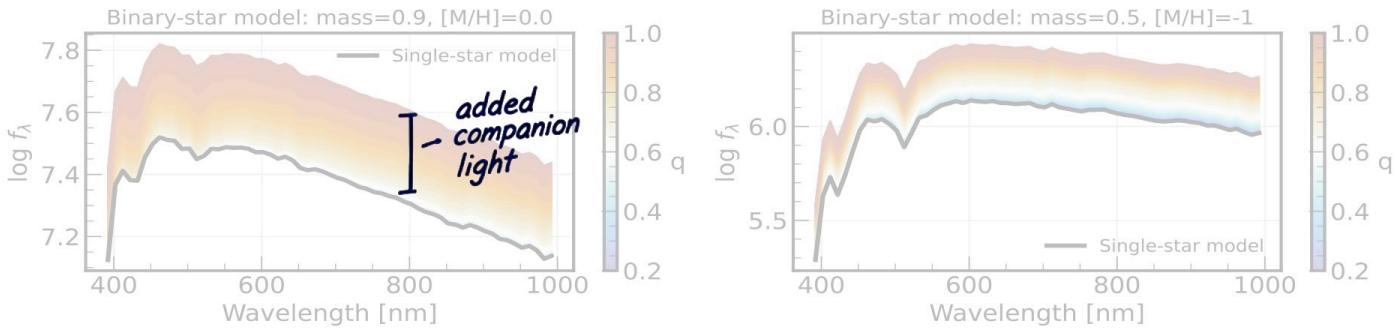


Fig. 4: Comparison of single-star and binary-star XP spectral models for systems with varying mass ratios. Black lines represent the single-star model, while the color gradient represents binary-star models with varying mass ratios ($q = m_2/m_1$), ranging from $q = 0.2$ (blue) to $q = 1.0$ (red). The binary models incorporate consistent $[M/H]_{phot}$, and distance for both components.

3.2.5. Parallax Error Propagation

We normalize all XP spectra to the flux at 10 pc using the parallax from Gaia DR3. Since the parallax measurement has uncertainties that cannot be neglected, we perform error propagation as follows:

Parallax uncertainty must also enter the flux covariance.

$$f_{10} = f_{\rm obs} \left(\frac{100}{\varpi}\right)^2 \tag{12}$$

where f_{10} is the flux normalized to 10 pc, $f_{\rm obs}$ is the observed flux at the star's distance, and ϖ is the parallax in milliarcseconds (mas). The error propagation process accounts for the uncertainty in parallax measurements when normalizing XP flux spectra to a standard distance of 10 pc.

We propagate the errors using two components:

$$C_{f_{10}} = g_\varpi^2 C_{f_{\rm obs}} + J J^T \sigma_\varpi^2, \tag{13}$$

where g_ϖ is defined as the scaling factor:

$$g_\varpi = \left(\frac{100}{\varpi}\right)^2; \tag{14}$$

The first term scales the original covariance matrix ($C_{f_{\rm obs}}$) by the square of the scaling factor, while the second variance term accounts for parallax uncertainty. This second term is computed using the Jacobian (J) of the scaling function with respect to parallax, multiplied by the square of the parallax error (σ_ϖ):

$$J = \frac{\partial g_\varpi}{\partial \varpi} f_{\rm obs} = -\frac{20000}{\varpi^3} f_{\rm obs} \tag{15}$$

The resulting covariance matrix ($C_{f_{10}}$) incorporates both the scaled original uncertainties and the additional uncertainty introduced by the parallax measurement. This error propagation ensures that measurement uncertainties are properly accounted for in our model fitting and parameter inference.

3.3. Maximum Likelihood Parameter Estimation

We employ a likelihood-based parameter inference of both single and binary stellar systems with a framework below.

3.3.1. Likelihood Formulation

Assuming Gaussian uncertainties in the observed fluxes, we define the likelihood function as:

$$\mathcal{L}(\mathbf{\Phi}) \propto \exp\left(-\frac{1}{2}\chi^2(\mathbf{\Phi})\right), \tag{16}$$

where $\mathbf{\Phi}$ represents either single-star parameters $(\boldsymbol{\theta}, E)$ or binary parameters $(m_1, [M/H]_{\rm photo}, q, E)$. The χ^2 statistic is defined as:

$$\chi^2(\mathbf{\Phi}) = \Delta\mathbf{F}^T C_f^{-1} \Delta\mathbf{F}, \tag{17}$$

where $\Delta\mathbf{F}$ is the vector of residuals between observed and model-predicted fluxes, and C_f is the covariance matrix of flux uncertainties.

3.3.2. Modified Likelihood with Prior Information

To incorporate prior information about extinction, we maximize a modified likelihood function:

$$\ln \mathcal{L}_{\rm mod}(\mathbf{\Phi}) = -\frac{1}{2}\chi^2(\mathbf{\Phi}) - \frac{1}{2}\lambda_p \cdot (E - E_{\rm ref})^2, \tag{18}$$

where the second term functions as a penalty encouraging the extinction parameter E to remain close to the reference value $E_{\rm ref}$ from the extinction map (Edenhofer et al. 2024). The parameter λ_p, set to 0.05, controls the strength of the extinction constraint in the modified likelihood function. This term acts as a regularization, anchoring the fitted extinction E to the reference value $E_{\rm ref}$ from the dust map (Edenhofer et al. 2024).

3.3.3. Optimization Procedure

For parameter estimation, we employ the Trust Region Reflective (TRF) algorithm (Coleman & Li 1994), which is effective for constrained optimization problems. The TRF method constructs a local quadratic approximation of the objective function within a trust region, ensuring stability and avoiding issues such as overstepping or divergence that can occur with other optimization methods.

The parameter space exploration is bounded by physical constraints:

- Primary mass $m_1 \in [0.1,\ 1.5]\ M_\odot$
- Photometric metallicity $[M/H]_{\rm phot} \in [-2.5,\ 0.7]$ dex
- Mass ratio q strictly satisfies $0.1 < q < 1.0$

The optimal parameters are obtained by solving:

$$\hat{\mathbf{\Phi}} = \arg\min_{\mathbf{\Phi}} (-\ln \mathcal{L}_{\rm mod}(\mathbf{\Phi})) \tag{19}$$

In our implementation, we compute the residual vector and extinction penalty separately, then combine them into an augmented residual vector. This allows us to treat the extinction penalty as an additional "observation" in our least-squares formulation while rigorously enforcing the parameter boundaries through the TRF algorithm's constraint-handling capabilities.

4. Validation

To validate our binary parameter recovery method, we constructed a suite of mock binary systems designed to span the full range of expected stellar parameters. Since we already have the single star emulator as described in section 3, we will first generate stellar labels and then use the single star emulator to mock the single and binary XP spectra. We generated 10,000 synthetic binary systems using a uniform random distribution for primary masses (m_1) ranging from 0.1 to 1.3 $M_\odot$ and $[M/H]_{\rm phot}$ ranging from −2.0 to +0.5 dex. Secondary masses (m_2) were similarly drawn from a uniform distribution between 0.1 and 1.3 $M_\odot$. We then performed random pairing of primary and secondary stars, followed by reordering when necessary to ensure $m_1 > m_2$ in all cases, resulting in mass ratios ($q = m_2/m_1$) spanning the full physical range from 0.1 to 1.0. For each synthetic binary, we generate mock *Gaia* XP spectra from our emulator with Gaussian noise properties based on the Gaia XP error model, incorporating the appropriate signal-to-noise. These mock spectra were then processed through our parameter inference pipeline, allowing us to quantify biases and scatter in the recovery of primary mass, metallicity, and mass ratio across the entire parameter space, as shown in Figure 5.

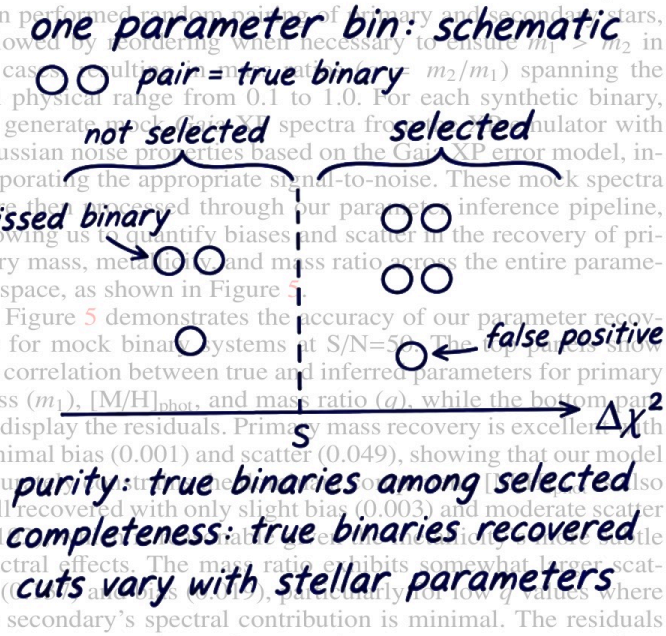


Figure 5 demonstrates the accuracy of our parameter recovery for mock binary systems at S/N=50. The top panels show the correlation between true and inferred parameters for primary mass (m_1), $[M/H]_{\rm phot}$, and mass ratio (q), while the bottom panels display the residuals. Primary mass recovery is excellent, with minimal bias (0.001) and scatter (0.049), showing that our model accurately captures the dominant spectral features. Metallicity is also well recovered with only slight bias (0.003) and moderate scatter (0.1[illegible] dex), indicating sensitivity to subtle spectral effects. The mass ratio exhibits somewhat larger scatter ([illegible] dex), particularly for low q values where the secondary's spectral contribution is minimal. The residuals show no systematic trends across the parameter ranges, confirming our binary model's robustness for systems with moderate to high signal-to-noise ratios.

The best-fit masses show larger scatter for stars more massive than 1 $M_\odot$, as these stars evolve on the HR-diagram during their main-sequence lifetime. This highlights that we cannot neglect the effects of stellar ages for more massive stars. In future work, we will further refine our approach by explicitly modeling colors and magnitudes as functions of mass, age, and metallicity to better account for evolutionary effects.

4.1. Classification Methodology

To identify binary systems, we compare the goodness-of-fit between single-star and binary models using the χ^2 statistic. A source is classified as a binary if:

$$\Delta\chi^2 = \chi^2_{\rm single} - \chi^2_{\rm binary} > S, \tag{20}$$

where S represents the optimal threshold for binary classification.

The threshold S is calibrated to balance completeness and purity, with an emphasis on minimizing false positives. We optimize using the F_β-score with $\beta = 0.5$:

$$F_\beta = (1+\beta^2) \cdot \frac{\text{Purity} \cdot \text{Completeness}}{(\beta^2 \cdot \text{Purity}) + \text{Completeness}}. \tag{21}$$

Setting $\beta < 1$ explicitly prioritizes purity over completeness, reducing contamination from single stars incorrectly classified as binaries.

4.2. Parameter-Dependent Thresholds

We determine optimal thresholds within bins of primary mass (m_1) and $[M/H]_{\rm phot}$ and mass ratio (q) rather than applying a uniform threshold across all stellar types. The optimal threshold S varies across these parameters as shown in Table C.1, reflecting how binary detectability depends on stellar properties.

For validation, we generate a mock dataset of 100,000 spectra (50,000 single and 50,000 binary) using our spectral emulators. We randomly sample 10,000 stars with primary masses ranging from 0.1 to 1.5 $M_\odot$, photometric metallicities ($[M/H]_{\rm photo}$) from -2.0 to +0.5 dex, and for binary systems, mass ratios (q) from 0.1 to 1.0. Each spectra is then simulated at five different signal-to-noise ratios (S/N = 5, 10, 20, 50, 100) to assess classification performance across varying data quality conditions. We add Gaussian noise to the synthetic fluxes according to these S/N levels to realistically mimic observational uncertainties.

For each mock spectrum, we perform both single-star and binary model fits to obtain best-fit parameters $\mathbf{\Phi}_{\rm single}$ and $\mathbf{\Phi}_{\rm binary}$, then calculate the χ^2 values and their difference ($\Delta\chi^2$). Within each bin of primary mass, metallicity, and mass ratio, we determine the threshold S that maximizes the F_β-score. Figure 6 shows the performance metrics as a function of S for all mock dataset. The optimal point is set to maximize the F_β value, which in this case is 2.37. This approach ensures our classification accounts for the varying sensitivity of the $\Delta\chi^2$ statistic across different regions of parameter space.

Binaries with nearly equal-mass components typically require lower thresholds for identification as they produce more spectral deviations from single-star models. Conversely, systems with low mass ratios require more stringent thresholds to avoid false positives,. The resulting parameter-dependent threshold map in Table C.1 enables binary classification across the entire dataset while balancing purity and completeness.

4.3. Verification with Binary Catalogs

Recovery rates in known binary catalogues are not the purity of the full new sample.

Table 1: Fraction of Systems labeled as binaries by adopted threshold

Dataset	Astrometric/Orbital Binary	SB1	SB2	single
Gaia NSS	67%	81%	93%	–
APOGEE	–	52%	85%	34%

The *Gaia* DR3 Non-Single Star (NSS) catalog (Gaia Collaboration et al. 2023b) identified 813,687 multiple systems, including 169,227 astrometric orbital solutions, 220,372 spectroscopic solutions, 87,073 eclipsing binaries, and 338,215 acceleration

Stricter cuts trade recovered binaries for a cleaner selected sample.

These mock tests assess recovery under the assumed spectral model.

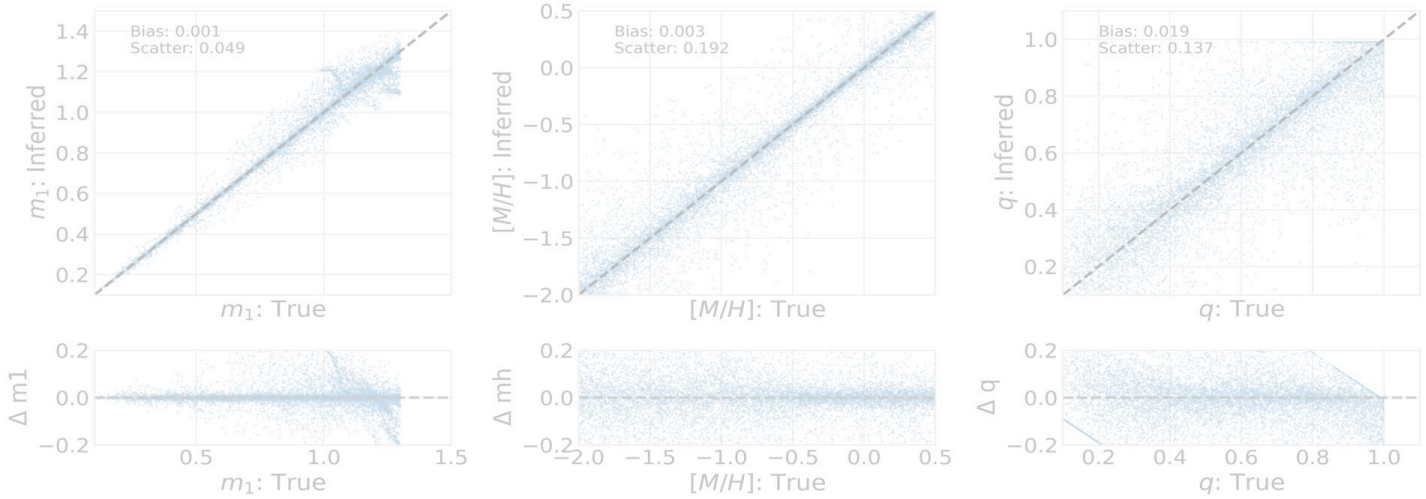


Fig. 5: Validation of parameter recovery for the MSMS mock data. Top row: Comparison between true and inferred values for primary mass (m_1, left), $[M/H]_{phot}$([M/H], middle), and mass ratio (q, right). The black dashed lines show the one-to-one relation, while blue points represent individual systems in our validation sample. Inset text shows the median bias and scatter for each parameter. Bottom row: Residuals (inferred minus true values) as a function of the true parameter values, highlighting the consistency of our parameter recovery across the full range of stellar properties. Primary masses are recovered with high precision (scatter of 0.049 $M_\odot$), while $[M/H]_{phot}$ show broader but still reliable recovery (scatter of 0.192 dex). Mass ratios are determined with good precision (scatter of 0.137) with a small positive bias (0.019) at low q values because of the detection sensitivity limitations. The hard limits in the right panel result from enforcing $q < 1$.

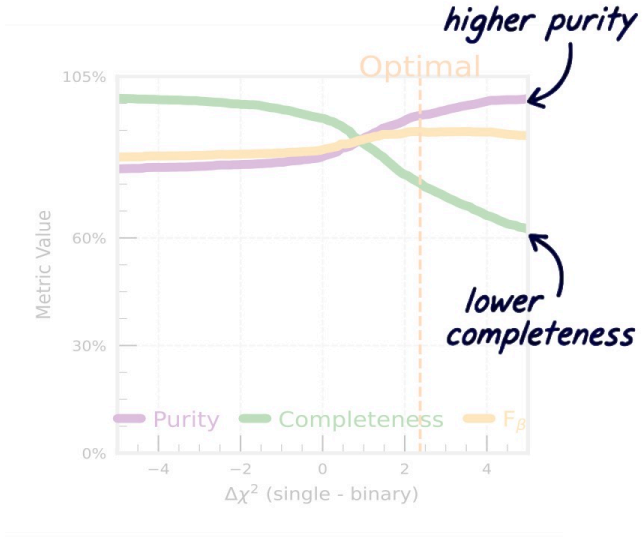


Fig. 6: Performance metrics as a function of the difference between χ^2 single and χ^2 binary values. The plot shows purity (precision, purple line), completeness (recall, green line), and F_β score (orange line) across different threshold values, where $\beta = 0.5$. The vertical dashed line is the χ^2 threshold that maximum the F_β score. Beyond this threshold, completeness decreases most rapidly, followed by F_β score, while purity remains high.

solutions (Gaia Collaboration et al. 2023a). Spectroscopic binaries are most detectable at periods of 1-1000 days, limited by stellar rotation at short periods and the DR3 baseline at longer periods. The NSS catalog represents more than orders of magnitude increase in the number of well-characterized binary systems compared to pre-*Gaia*, catalogs (Gaia Collaboration et al. 2023b).

Our verification as shown in Fig 8 demonstrates that the vast majority of sources in the NSS catalog show preference for binary-star models when fitting their XP spectra, providing independent validation of our methodology. This correlation is particularly pronounced for double-lined spectroscopic binaries (SB2 and SB2C types), where the χ^2 improvement for binary models over single-star fits consistently exceeds our detection threshold. The stronger signal for SB2 systems is expected, as these binaries typically have more comparable flux ratios between components, making their composite spectral signatures more distinctive.

As shown in Table 1, for *Gaia* NSS catalog, our methodology correctly classifies 67% of astrometric/orbital binaries, 81% of SB1, and 93% of SB2 systems as binaries based on our parameter-dependent thresholds in subsection 4.2. The APOGEE binary sample (El-Badry et al. 2018), derived from ∼20,000 main-sequence star spectra in APOGEE DR13, includes ∼2,500 binaries and ∼200 triples identified through spectral fitting with composite models, alongside ∼700 velocity-variable systems. For APOGEE, the correct classification rates are 52% for SB1 and 85% for SB2 systems, confirming the robustness of our spectral fitting approach. Additionally, only 34% of APOGEE's single-star sample is classified as binaries in our XP fitting, further validating our results, as some of these "single" stars may be long-period q≈1 binaries indistinguishable from single stars in earlier models due to limited parallax precision.

The astrometric-only binaries (Orbital) show a more varied response in our analysis, with detection rates correlating strongly with the magnitude difference between components. These patterns across different NSS solution types provide insight into the sensitivity limits and complementary nature of spectrophotometric binary detection compared to spectroscopic and dynamic

A two-star model can explain spectral residuals that a single-star fit leaves behind.

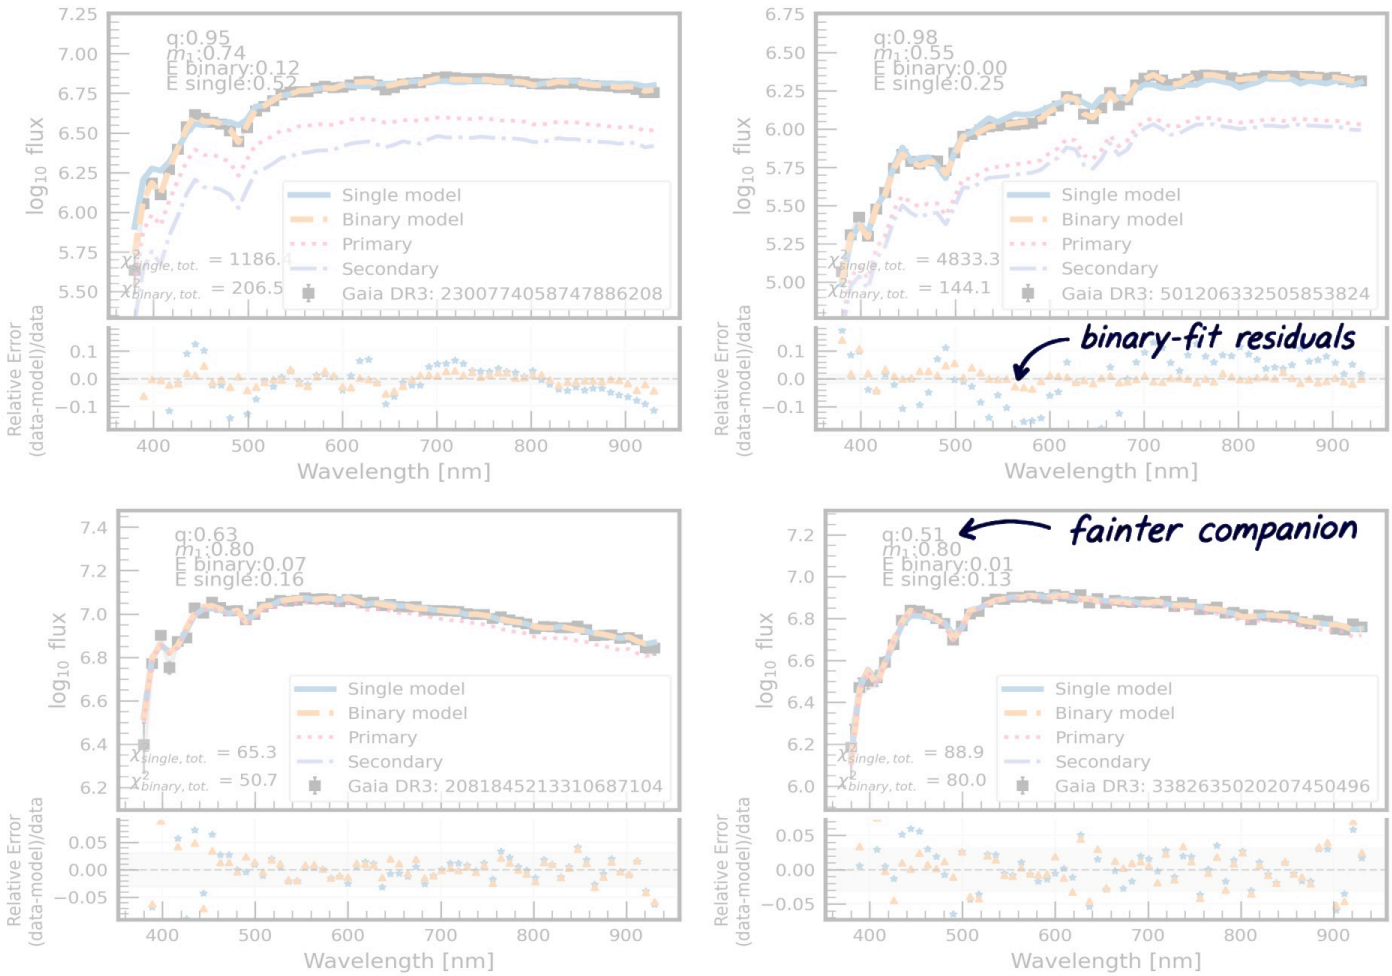


Fig. 7: *Gaia* XP spectra and best-fitting models for four MSMS binaries. In each panel, the black line represents the observed flux. The blue solid line shows the best-fit single-star model, while the orange dashed line represents the best-fit binary model. The pink dotted line corresponds to the primary component of the binary model, and the dark blue dotdash line shows the secondary component. For each system shown, the binary model is preferred as it achieves a lower χ^2 value than the single-star model. Top row: Example NSS SB2 binary. Bottom row: Example NSS astrometric binaries.

approaches. To improve usability, we also provide the cross-matched catalog between our binary fits and the NSS catalog in Appendix C.3.

We further validate the mass estimates derived from our XP fitting method by comparing them with those provided in the NSS catalog, specifically the `gaiadr3.binary_masses` table (Gaia Collaboration et al. 2023a). This table contains pre-computed masses for the primary (m_1) and secondary (m_2) components of binary systems. The NSS mass estimates are derived using a variety of observational techniques, primarily combining high-precision astrometric, spectroscopic, and, in some cases, photometric (e.g., eclipsing binary) data. We adopt the binary mass values from all available solutions, including the solution type, such as 'Orbital+SB2' (astrometric orbit combined with double-lined spectroscopy), 'EclipsingSpectro(SB2)' (eclipsing binary photometry with SB2 spectroscopy), or 'SB1+M1' (single-lined spectroscopy with primary mass from isochrone fitting). These methods leverage Kepler's Third Law and, where applicable, radial velocity semi-amplitudes and orbital inclinations to resolve individual masses. As shown in Fig. 9, the m_1 and m_2 values derived from our XP fitting method exhibit strong consistency with the NSS catalog's mass estimates. This agreement validates the robustness of our approach, particularly given the NSS catalog's reliance on diverse and complementary observational data.

4.4. Validation with Dark Companion Systems

The identification of single stars represents the inverse challenge of binary detection. Achieving high purity in binary star identification inevitably comes at the cost of both binary sample completeness and single star sample purity. The chi-squared comparison method requires careful calibration, with threshold settings that vary as demonstrated in Section 5.2. To validate our classification approach, we sought benchmark "single-star-like" systems to verify whether our method correctly categorizes them. Such validation would be demonstrated by these systems showing minimal improvement in fit quality when modeled as binaries rather than single stars. For this purpose, we utilized two specialized samples:

The gain is stronger when both components contribute substantial light.

A dark companion can hide behind a single-star-like spectrum.

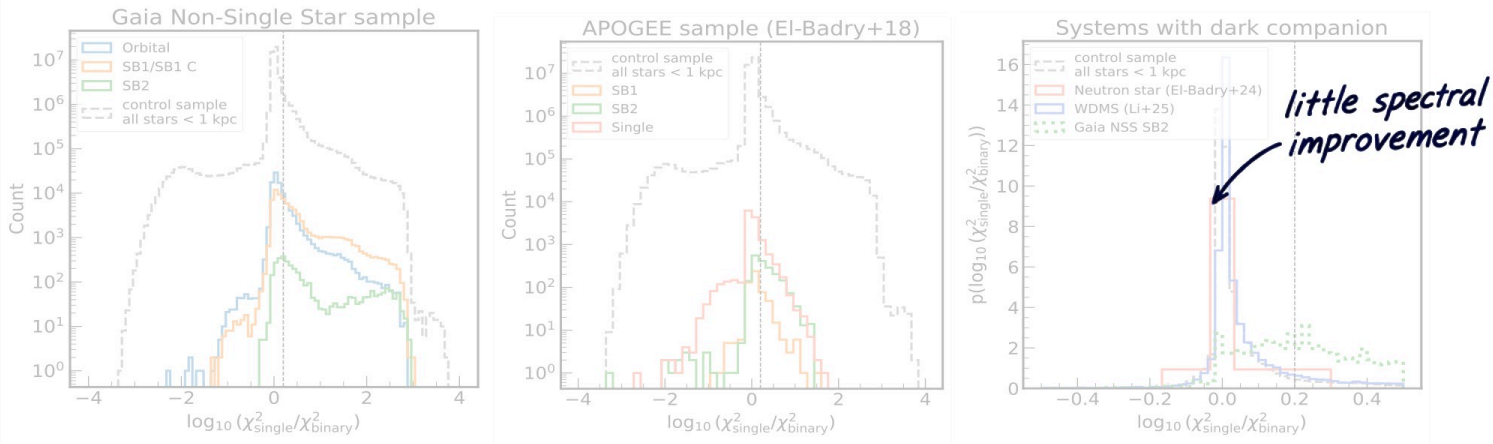


Fig. 8: Distribution of the logarithmic χ^2 improvement, $\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary})$, showing the statistical evidence favouring binary over single-star spectral models. Left panel: Distributions for Gaia DR3 Non-Single Star (NSS) sources: systems with orbital solutions (blue), single-lined spectroscopic binaries (SB1/SB1C; orange), and double-lined spectroscopic binaries (SB2; green). Right panel: Distributions for specialized populations with likely dark companions: neutron star binaries (El-Badry et al. 2024) (red) and white dwarf-main sequence binaries (WDMS) (Li et al. 2025) (blue). The Gaia DR3 NSS SB2 sample (green dotted line) is shown for reference. Common elements in each panel: The distribution for all stars within 1 kpc (black dashed line) acts as a control representing the general population. A vertical dashed line marks 0.2 for reference. Higher values of the χ^2 improvement indicate stronger evidence for binarity; known binaries typically show significant positive values.

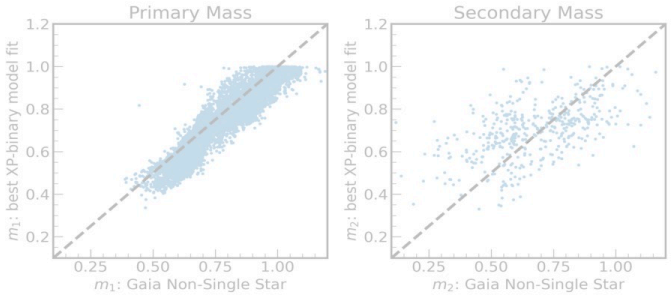


Fig. 9: Comparison of primary (left) and secondary (right) masses derived from the XP binary model (m_1, m_2) against those from the *Gaia* Non-Single Star catalog (m_1, m_2). The dashed line represents the one-to-one relation.

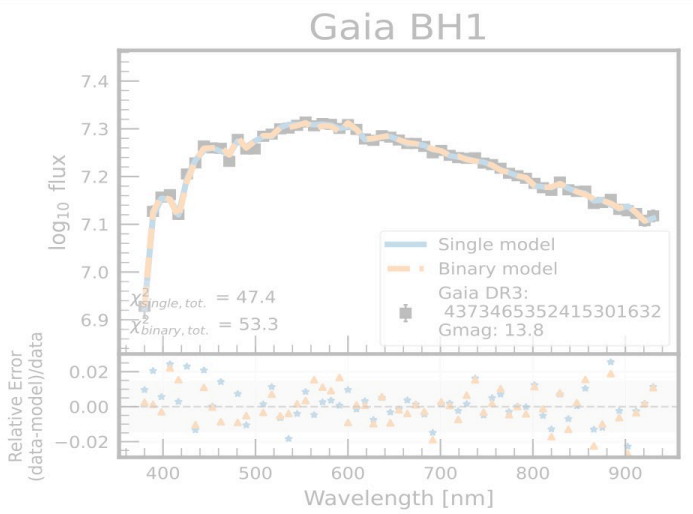


Fig. 10: *Gaia* XP spectra and best-fitting models for *Gaia* BH1. The black line represents the observed flux. The blue solid line shows the best fit single-star model, while the orange dashed line represents the best fit binary model. Bottom panel: Residuals (Δ log flux) for both models.

First, we examined the neutron star companion sample from El-Badry et al. (2024), which contains 21 astrometric binaries with solar-type stars orbiting dark companions with masses near 1.4 $M_\odot$. Since these companions contribute virtually no light to the systems, their spectra should closely resemble those of single stars despite their binary nature. Second, we analyzed the white dwarf-main sequence (WDMS) binary sample from Li et al. (2025), which identifies approximately 30,000 WDMS systems from XP spectra. In most of these systems, the main sequence primary is at least 10 times brighter than the white dwarf secondary (Li et al. 2025), though some were flagged as MSMS binaries using the same chi-squared method.

The right panel in Figure 8 compares the chi-squared distribution of these samples against a control sample consisting of all *Gaia* sources within 1 kpc. The results demonstrate that for the neutron star sample, the binary model provides minimal improvement over the single-star model ($\log_{10}(\chi^2_{\rm MS}) - \log_{10}(\chi^2_{\rm MSMS}) < 0.2$). Similarly, 84% of the WDMS sample shows no notable χ^2 improvement with the binary model, although some systems in this sample could potentially be MSMS binaries with white dwarfs as tertiary companions. This validation confirms our method appropriately classifies systems with dark or faint companions as effectively single stars from a spectroscopic perspective, despite their astrometric binary nature.

single-star fit already sufficient

We also validate the χ^2 difference for *Gaia* Black Hole 1 (BH1) as shown in Figure 10, we see that the total $\chi^2_{\rm single}$ is 47.3, $\chi^2_{\rm binary}$ is 53.3, which also proves our method can successfully

No luminous secondary signature does not mean there is no companion.

The golden sample is a stricter subset of the binary-candidate catalogue.

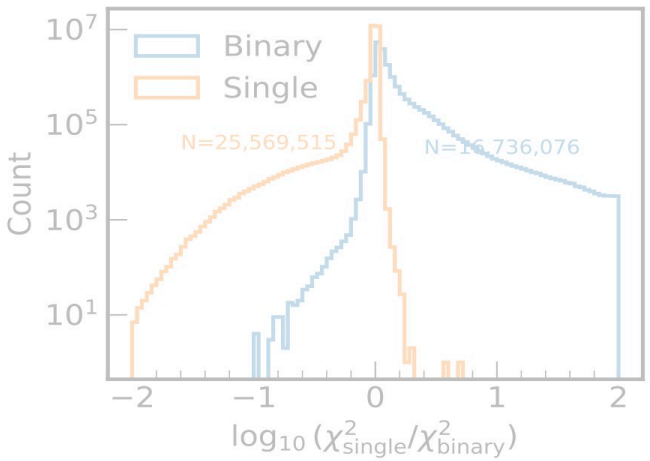


Fig. 11: Distribution of $\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary})$ for a subset of *Gaia* DR3 XP spectra, illustrating the statistical separation between single-star and binary models. The histogram shows the frequency of sources as a function of $\Delta\chi^2$, with the classifiation based on an optimal threshold S (calibrated to maximize the F_β-score with $\beta = 0.5$) used to classify binary systems.

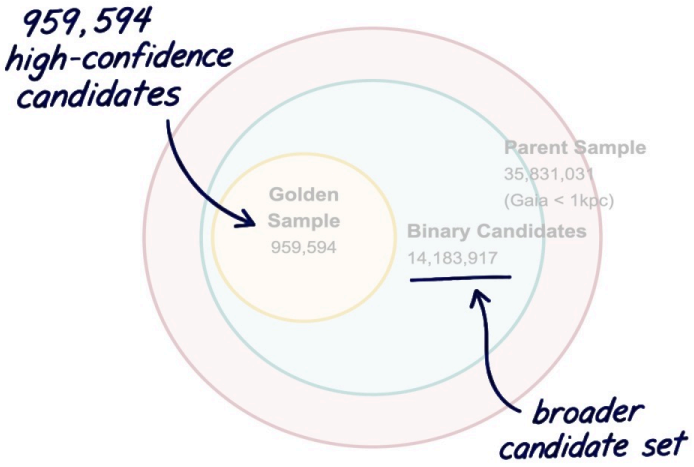


Fig. 12: Venn diagram showing the hierarchical relationship between our sample selections. The parent sample consists of 35 million stars within 1 kpc from Gaia DR3. From this population, we identify 14 million binary candidates based on spectral model fitting criteria. The golden sample is identified through selection criteria in Table 2.

identify it as a single star / dark companion. This validation confirms that our method appropriately classifies systems with dark or faint companions as effectively single stars from a spectroscopic perspective.

5. The MSMS Binary Catalog

5.1. Data Products

We have compiled and published the main-sequence (MSMS) binary catalog based on *Gaia* DR3 XP spectra. This catalog contains 14 million binary candidates identified from our parent sample of 35,831,031 stars within 1 kpc of the Sun. The binary candidates were selected through a combination of spectroscopic modeling techniques that compare the χ^2 values between single-star and binary-star model fits.

The selected-candidate fraction is not a completeness-corrected binary fraction.

For each entry, we include the best-fit results from both single-star and binary spectral models, allowing users to evaluate the binary classification (12). The binary model parameters include primary mass, system metallicity, mass ratio, and derived photometric properties such as colors and absolute magnitudes for both components.

The full MSMS Binary Catalog dataset, along with NN weights enabling users to forward model XP single/binary spectra according to their requirements, is now available for public access at https://doi.org/10.5281/zenodo.15166184 Li (2025).

5.2. Binary Candidate Selection

To identify binary systems from our stellar sample, we employ a statistical approach based on the χ^2 difference between single-star and binary-star model fits. Our method leverages the model structure described in Section 3, extending it to accommodate binary systems by combining the fluxes of two stars with different masses but sharing the same $[\mathrm{M/H}]_{\rm phot}$, and distance. This approach is particularly sensitive to binaries with high mass ratios, which were identified as outliers during the self-cleaning process described in Section 3.

We begin with a parent sample of 35 million stars within 1 kpc. For each star, we perform two separate model fits: one assuming that the object is a single star and another assuming it is a binary system. The chi-square difference ($\Delta\chi^2$) between these two models provides a statistical measure of how much the data favors the binary hypothesis over the single-star hypothesis.

To establish appropriate thresholds for binary classification, we divide our sample into bins based on primary mass ($[\mathrm{M/H}]_{\rm phot}$) and mass ratio ($q$), as shown in Table C.1. These thresholds are calibrated to balance completeness and contamination in our binary identification, accounting for the varying sensitivity of our method across different stellar parameters.

A star is classified as a binary candidate if the χ^2 difference exceeds the threshold corresponding to its mass bin:

$$\Delta\chi^2 = \chi^2_{\rm single} - \chi^2_{\rm binary} > \mathrm{threshold}(m_1, [\mathrm{M/H}]_{\rm phot}, q) \quad (22)$$

Fig. 11 shows the distribution of $\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary})$ values for our parent sample. Binary candidates systematically exhibit larger chi-squared differences, with the distribution showing a clear bimodality.

Using the thresholds defined in each m_1-$[\mathrm{M/H}]_{\rm phot}$-$q$ bin, we identify 14,183,917 binary candidates from our parent sample, corresponding to an overall binary fraction of approximately 39%. This is consistent with expected binary fractions from previous studies (e.g. Gao et al. 2014; Liu 2019; Moe et al. 2019), although the exact fraction varies with stellar mass and metallicity environment.

5.3. Golden Binary Sample

To facilitate studies requiring higher confidence binary identifications, we defined a "golden sample" of 959,594 binary systems, representing the most reliable binary candidates. These golden sample members satisfy stringent selection criteria and quality cut detailed in Table 2.

The golden cuts favour reliable fits and main-sequence pairs within the model's useful range.

Table 2: Selection criteria for the golden binary sample

Criterion	Description
predicted_class = 1	Binary classification that passes the χ^2 threshold in subsection 5.2
$\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary}) > 0.2$	χ^2 improvement for binary fit
$\chi^2_{\rm single}/{\rm Dof} < 3 \vee \chi^2_{\rm binary}/{\rm Dof} < 3$	Good fit quality (either single or binary model)
fit_status_single $\geq 2 \vee$ fit_status_binary ≥ 2	Successful convergence (either single or binary model)
$m_1 < 1.0\ M_\odot$	Primary mass cut
$[{\rm M/H}]_{\rm phot} > -1$	$[{\rm M/H}]_{\rm phot}$ cut
$M_G \geq 4.5$	Absolute magnitude cut
E_binary < 0.2	Extinction cut
$0.4 \leq q \leq 1.0$	mass-ratio cut
$0.2 \leq$ flux_ratio_mean ≤ 1.0	flux-ratio cut
flag_extrap_m2 == False	secondary mass cut ($m_2 > 0.1$)

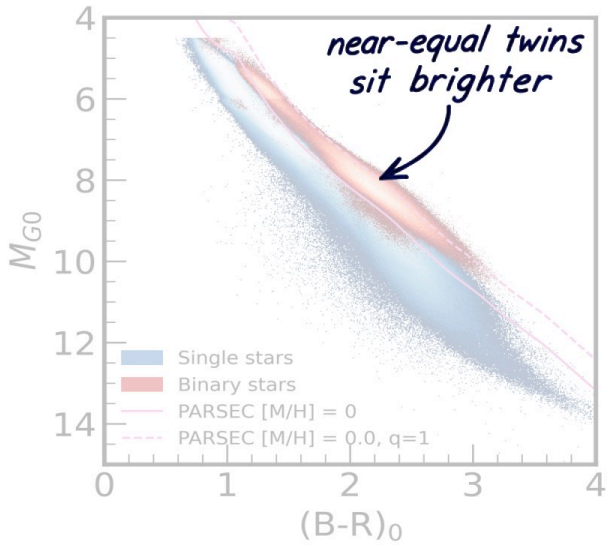


Fig. 13: **HR diagrams for golden binary sample of mass-ratio** > 0.9**.** The background blue shows the density distribution of selected single stars in, with the scale bar indicating the logarithmic number of stars per bin. The pink solid and dashed lines show PARSEC isochrones (Bressan et al. 2012) for [M/H]=0, and $q = 1$ binary configurations.

The "golden sample" of 959,594 systems, selected from an initial catalog of 35,831,031 systems (Table C.2), represents the highest confidence binary candidates identified in this work. These systems were chosen to ensure robust binary identification and focus on well-characterized main-sequence pairs suitable for detailed analysis.

These golden sample members satisfy stringent criteria including a significant chi-square improvement with $\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary}) > 0.2$, good fit quality for both single or binary models with $\chi^2 < 3 \times 61$ (3 per degree of freedom), successful convergence of model fits, primary masses restricted to $M_1 < 1.0\ M_\odot$, $[{\rm M/H}]_{\rm phot}$ from the binary fit > -1, absolute G-band magnitude $M_G \geq 4.5$ to focus primarily on main sequence stars, and extinction $E_{\rm binary} < 0.2$ to ensure the extinction isn't too severe. Additionally, systems must have reliable mass ratios between $0.4 \leq q \leq 1.0$, and flux ratios between 0.2 and 1.0. These cuts are based on practical considerations from mock tests and are not overly stringent, though they may exclude some binaries.

Fig. 12 illustrates the relationship between our parent sample, binary candidates, and golden binary sample using a Venn diagram, highlighting the nested nature of these increasingly selective subsets. Fig. 13 shows HR diagrams for our binary systems across different mass ratio ranges. The distinct positions of equal-mass binaries (right panel) compared to lower mass-ratio systems (left panel) in the color-magnitude diagram support the effectiveness of our classification approach, with nearly equal-mass binaries appearing systematically brighter than their single-star counterparts of the same color. Our catalog includes the stellar parameters of both the primary and the inferred secondary companion for all binary candidates, along with the $\Delta\chi^2$ value and classification confidence metrics. The MSMS Binary Catalog represents an extensieve uniform samples of binary systems to date, enabling population studies across different Galactic environments and investigating the dependence of binary fraction on stellar and environmental parameters.

These cuts improve confidence by sacrificing some completeness.

6. Discussion and Conclusions

We discuss binary star systems identified from *Gaia* DR3 XP spectra in the following sections. In Section 6.1, we first validate our spectral fitting results with enhanced parallax error treatment, demonstrating the robustness of our methodology. We then examine potential caveats and limitations of our binary analysis in Section 6.2.

6.1. Parallax Errors

Two factors contribute to the underestimation of uncertainties for true binaries in our analysis. First, treating binaries as single stars introduces errors in parallax estimates (Gaia Collaboration et al. 2023a). We validate the fitting results using parallax solutions derived from binary orbit assumptions. Second, binaries exhibit poor astrometric fits due to orbital motion, leading to underestimated uncertainties in Gaia's astrometric solutions, including parallax. We apply a simulation-driven method (El-Badry 2025) to correct the parallax error and compare the results with our findings.

6.1.1. Comparison with Parallax from Binary Solutions

The NSS catalog fits complex models to describe the behavior of non-single stars, accounting for orbital motion in binary sys-

Dust, stellar evolution and low-mass extrapolation limit how literally to read the fitted parameters.

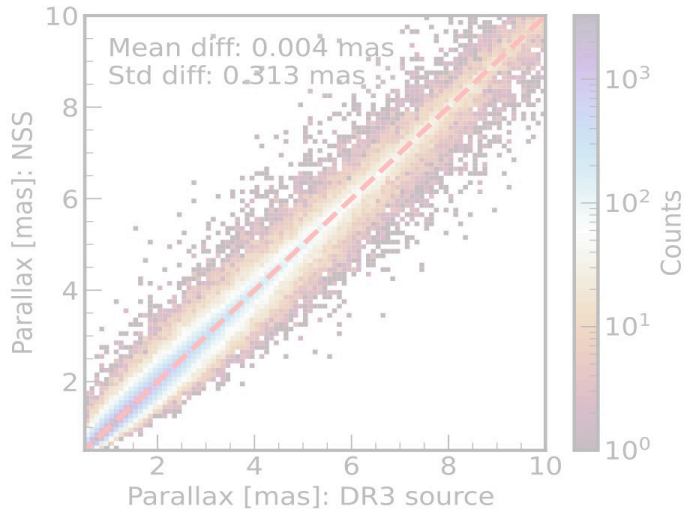


Fig. 14: Comparison of parallax from *Gaia* DR3 source catalog and NSS calalog. The dashed red line represents the one-to-one relation, and the color scale denotes the number density.

tems (Gaia Collaboration et al. 2023a). It includes solutions for astrometric, spectroscopic, and eclipsing binaries, providing parallax estimates derived from these models. The comparison between DR3 source and NSS parallax solutions (Fig. 14) reveals a mean difference of 0.004 mas with an RMS scatter of 0.3 mas, indicating that source parallaxes slightly underestimate the true values of the NSS binary systems. We further compare the χ^2 improvement using different parallax sources in Figure 15, where the right panel demonstrates that the difference between NSS parallax and source parallax is negligible. This suggests that for sources without NSS orbital solutions, using the DR3 source parallax for XP fitting may be a reasonable approach.

6.1.2. Inflating Parallax Uncertainties

To address underestimation, we adopt the inflation model from El-Badry (2025) to adjust uncertainties based on the Renormalized Unit Weight Error (RUWE), a metric of astrometric fit quality where higher values indicate poorer fits. Many binaries with close companions have elevated RUWE values due to photocenter oscillations from orbital motion, which single-star models cannot accommodate (Gaia Collaboration et al. 2023a). For binaries with orbital periods of a few years, photocentric acceleration further increases RUWE. The model of El-Badry (2025) adjusts reported uncertainties as a function of RUWE, parallax, and apparent magnitude, transforming error distributions into Gaussian forms.

We compare the χ^2 improvement ($\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary})$) in Figure 15. For astrometric (orbital) binaries, the difference before and after the parallax inflation is negligible. However, for SB1 systems, we observe modest differences: For binaries with $\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary})< 0.5$, the XP fit incorporating parallax error inflation demonstrates better performance than without inflation. Specifically, binaries previously showing $\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary})$ between 0 and 0.2 shift to the 0.2-0.5 bin after applying parallax error inflation, indicating improved classification capability. For binaries with $\log_{10}(\chi^2_{\rm single}/\chi^2_{\rm binary})> 0.5$, parallax error inflation produces no significant difference in results.

6.2. Caveats and Limitations

6.2.1. Extinction Modeling

Our extinction modeling relies on an average extinction curve derived from Gaia XP spectra (Zhang et al. 2023). This approach, while effective for population-level analysis, may not fully capture the variation in extinction properties across different Galactic environments. Recent studies by Zhang & Green (2025) demonstrate that $R(V)$, which characterizes the slope of the extinction curve in optical wavelengths and is related to the typical size of dust grains, can vary depending on the line of sight. This variation could introduce systematic uncertainties in our derived stellar parameters, particularly in regions with anomalous dust properties or high extinction values. Future work could incorporate spatially-dependent extinction curves to further refine the binary parameter estimates.

6.2.2. Low-Mass Secondary Extrapolation

For binary systems with secondary masses ($m_1 \times q$) below 0.1 $M_\odot$, our parameter estimates should be treated with caution. The stellar structure models we employ (Bressan et al. 2012) have a lower mass limit of 0.1 $M_\odot$, meaning that binary fits with low-mass secondaries rely on extrapolation from our NN training set. We provide a quality flag `flag_m2_extrap` in our catalog to identify these systems.

Secondaries below 0.1 solar masses fall outside the calibrated stellar-model range.

6.2.3. Evolutionary Effects for Higher-Mass Stars

For stars more massive than approximately 0.8 $M_\odot$, which corresponds to the highest mass that remains on the main sequence after 12.5 Gyr at [Fe/H] = -2 (Paxton et al. 2011; Choi et al. 2016), evolutionary effects become increasingly important. These stars likely have evolved from their zero-age main sequence positions in the color-magnitude diagram, making our modeling from mass and metallicity to color-absolute magnitude less accurate without accounting for age variations. Current parameter estimates for these higher-mass systems do not incorporate age as a variable, potentially introducing systematic biases in the inferred mass ratios and metallicities. A future refinement of this work should consider stellar ages and initial masses of binary systems when modeling their spectral and photometric properties.

6.3. Summary

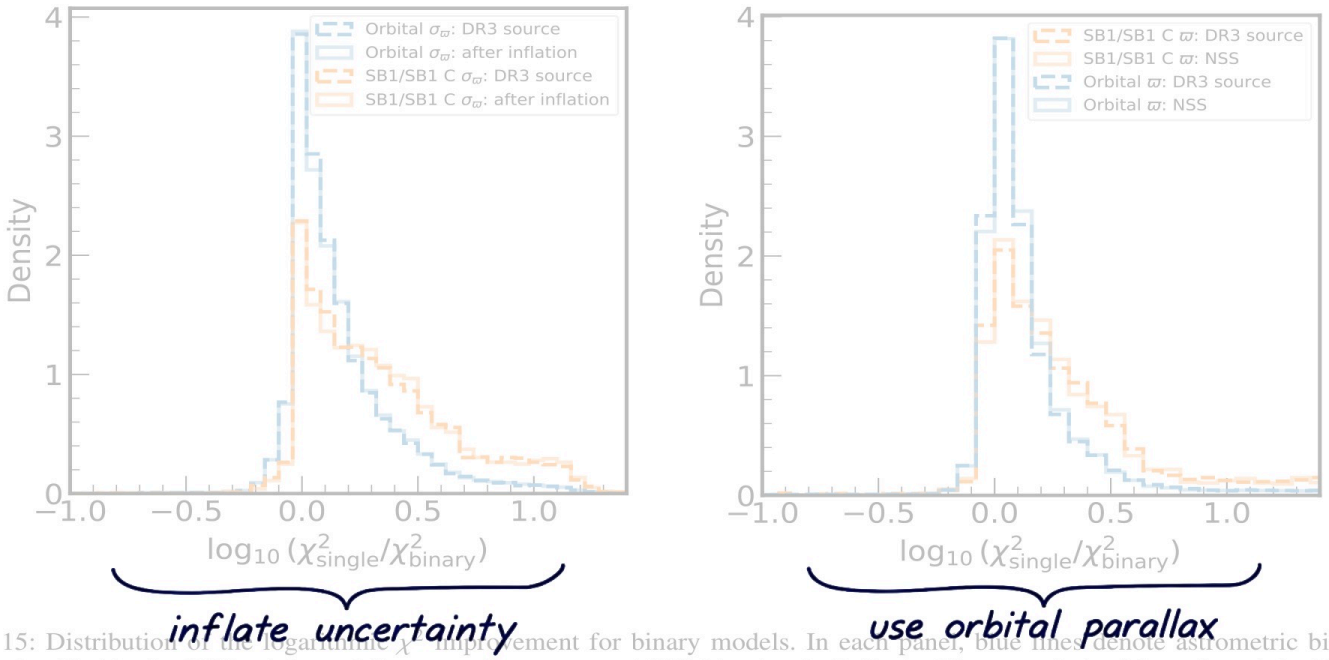
This study presents the main-sequence binary (MSMS) Catalog, identifying ~14 million binary star candidates from Gaia Data Release 3 (DR3) BP/RP spectra. The analysis uses spectrophotometric data for 35 million stars within 1 kpc of the Sun, selected from 220 million Gaia DR3 sources.

Binaries are identified when the binary model yields a lower χ^2 than the single-star model, with a parameter-dependent χ^2 improvement threshold for mass ratios $0.5 \leq q \leq 1.0$. A subset of approximately one million systems ("golden sample") is defined using a stricter chi-squared improvement, achieving 93% completeness of SB2s validated using the *Gaia* non-single star catalog.

This approach overcomes limitations of spectroscopic methods, which rely on radial velocity shifts and favor short-period binaries (periods < 1000 days). By comparing spectral fits, the method detects binaries across mass ratios $0.5 \leq q \leq 1.0$. Validation with synthetic spectra confirms recovery of primary mass and$[{\rm M/H}]_{\rm phot}$with errors < 10%. The 14 million candidates re-

Higher-mass stars need age information that this fixed-age mapping does not include.

Changing the parallax treatment has a modest effect on these validation distributions.



inflate uncertainty

use orbital parallax

Fig. 15: Distribution of the logarithmic χ^2 improvement for binary models. In each panel, blue lines denote astrometric binary stars classified by the NSS catalog, while orange lines represent SB1 binaries. Left: Dashed lines are derived from the original DR3 source catalog parallax errors, while solid lines show results after error inflation following the model of El-Badry (2025). Right: The dashed lines use parallaxes from the DR3 source catalog, while solid lines use parallaxes from the NSS catalog, which incorporate binary orbital solutions in their parallax estimates.

veals a binary fraction of 39% among main-sequence stars, with mass- and metallicity-dependent trends in the mass-ratio distribution.

Acknowledgements. JL thanks Coryn Bailer-Jones for constructive discussion and insights. We acknowledge support from the European Research Council through ERC Advanced Grant No. 101054731. This work 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. YST is supported by the National Science Foundation under Grant AST-2406729.

References

Aznar Cuadrado, R. & Jeffery, C. S. 2001, A&A, 368, 994
Badenes, C., Mazzola, C., Thompson, T. A., et al. 2018, ApJ, 854, 147
Bressan, A., Marigo, P., Girardi, L., et al. 2012, MNRAS, 427, 127
Castro-Ginard, A., Penoyre, Z., Casey, A. R., et al. 2024, A&A, 688, A1
Chen, Y., Girardi, L., Bressan, A., et al. 2014, MNRAS, 444, 2525
Choi, J., Dotter, A., Conroy, C., et al. 2016, ApJ, 823, 102
Coleman, T. F. & Li, Y. 1994, Cornell University Technical Report
De Angeli, F., Weiler, M., Montegriffo, P., et al. 2023, A&A, 674, A2
Duchêne, G. & Kraus, A. 2013, Annual Review of Astronomy and Astrophysics, 51, 269
Edenhofer, G., Zucker, C., Frank, P., et al. 2024, A&A, 685, A82
El-Badry, K. 2024, Gaia's Binary Star Renaissance
El-Badry, K. 2025, How to use Gaia parallaxes for stars with poor astrometric fits
El-Badry, K., Rix, H.-W., & Heintz, T. M. 2021, MNRAS, 506, 2269
El-Badry, K., Rix, H.-W., Latham, D. W., et al. 2024, The Open Journal of Astrophysics, 7 [arXiv:2405.00089]
El-Badry, K., Ting, Y.-S., Rix, H.-W., et al. 2018, MNRAS, 476, 528
Gaia Collaboration, Arenou, F., Babusiaux, C., et al. 2023a, A&A, 674, A34
Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1
Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023b, A&A, 674, A1
Gao, S., Liu, C., Zhang, X., et al. 2014, ApJ, 788, L37
Green, G. 2018, The Journal of Open Source Software, 3, 695
Holl, B., Sozzetti, A., Sahlmann, J., et al. 2023, A&A, 674, A10
Kollmeier, J. A., Zasowski, G., Rix, H.-W., et al. 2017, arXiv e-prints, arXiv:1711.03234
Li, C.-q., Shi, J.-r., Yan, H.-l., et al. 2021, ApJS, 256, 31
Li, J. 2025, Dataset for MSMS binary: Millions of Main-Sequence Binary Stars from Gaia BP/RP Spectra
Li, J., Liu, C., Zhang, B., et al. 2021, ApJS, 253, 45
Li, J., Ting, Y.-S., Rix, H.-W., et al. 2025, arXiv e-prints, arXiv:2501.14494
Li, J., Wong, K. W. K., Hogg, D. W., Rix, H.-W., & Chandra, V. 2024, ApJS, 272, 2
Liu, C. 2019, MNRAS
Liu, C., Fu, J., Shi, J., et al. 2020, arXiv e-prints, arXiv:2005.07210
Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, AJ, 154, 94
Moe, M., Kratter, K. M., & Badenes, C. 2019, ApJ, 875, 61
Paxton, B., Bildsten, L., Dotter, A., et al. 2011, ApJS, 192, 3
Penoyre, Z., Belokurov, V., & Evans, N. W. 2022a, MNRAS, 513, 5270
Penoyre, Z., Belokurov, V., & Evans, N. W. 2022b, MNRAS, 513, 2437
Price-Whelan, A. M., Hogg, D. W., Rix, H.-W., et al. 2020, ApJ, 895, 2
Qiu, D., Li, J., Zhang, B., et al. 2024, MNRAS, 527, 11866
Raghavan, D., McAlister, H. A., Henry, T. J., et al. 2010, ApJS, 190, 1
Rix, H.-W., Ting, Y.-S., Sippel, A., et al. 2019, Bulletin of the American Astronomical Society, 51, 104
Seeburger, R., Rix, H.-W., El-Badry, K., Xiang, M., & Fouesneau, M. 2024, MNRAS, 530, 1935
Zhang, B., Jing, Y.-J., Yang, F., et al. 2022, ApJS, 258, 26
Zhang, X. & Green, G. M. 2025, Science, 387, 1209
Zhang, X., Green, G. M., & Rix, H.-W. 2023, MNRAS, 524, 1855

Correlated spectral errors must be accounted for before judging fit quality.

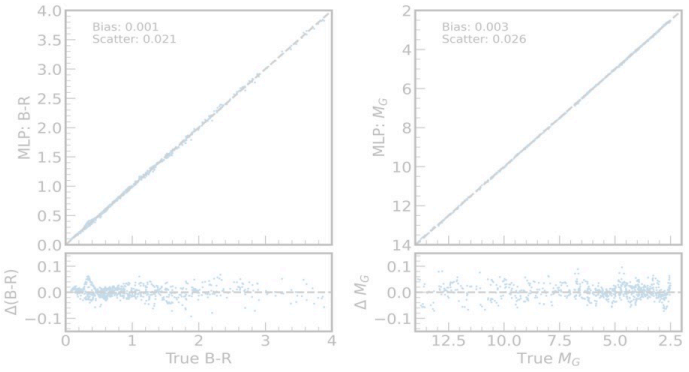


Fig. .1: Validation of the neural network-based mapping function $\mathcal{M}$ that translates stellar physical parameters to H-R diagram coordinates. Left panels: Comparison between true and predicted B-R color, showing good agreement with minimal bias (0.001) and scatter (0.021). Right panels: Validation of absolute G-magnitude (M_G) predictions, displaying negligible bias (0.003) and low scatter (0.026). The bottom panels show residuals (predicted minus true values) across the parameter range, demonstrating no systematic trends.

Appendix A: Sampled XP spectra

To accurately process XP spectra, we must account for covariances between observed fluxes across wavelengths. Our sampling range of 392–992 nm and step size of 10 nm were chosen to trim noisy spectral edges and ensure positive-definite covariance matrices (Zhang et al. 2023).

The χ^2 statistic between predicted and observed fluxes is calculated as:

$$\chi^2 = \Delta \boldsymbol{f}^T C_{\boldsymbol{f}_{\rm obs}}^{-1} \Delta \boldsymbol{f} = |L\Delta \boldsymbol{f}|^2 , \tag{A.1}$$

where $\Delta \boldsymbol{f} \equiv \boldsymbol{f}_{\rm pred} - \boldsymbol{f}_{\rm obs}$ represents the residuals between predicted and observed fluxes.

To ensure numerical stability in this calculation, we diagonalize each covariance matrix using its eigendecomposition:

$$C_{\boldsymbol{f}_{\rm obs}} = UDU^T , \tag{A.2}$$

where U is orthonormal and D is diagonal. We impose a minimum value of 10^{-9} (in units of $10^{-36}\,\mathrm{W^2\,m^{-4}\,nm^{-2}}$) on the diagonal elements of D (yielding $\hat{D}$) and define $L \equiv \hat{D}^{-1/2}U^T$.

The covariance matrix $\mathbf{C}_{\boldsymbol{f}_{\rm obs}}$ incorporates several uncertainty components:

$$\mathbf{C}_{\boldsymbol{f}_{\rm obs}} = \mathbf{C}_{\boldsymbol{f}_{\rm gaia}} + \mathrm{diag}\left(0.005\boldsymbol{f}_{\rm gaia}\right)^2 + 0.005^2 \boldsymbol{f}_{\rm gaia}\boldsymbol{f}_{\rm gaia}^T + 0.001^2 \boldsymbol{f}_{\rm BP}\boldsymbol{f}_{\rm BP}^T + 0.001^2 \boldsymbol{f}_{\rm RP}\boldsymbol{f}_{\rm RP}^T, \tag{A.3}$$

where $\boldsymbol{f}_{\rm gaia}$ represents the sampled flux, and $\boldsymbol{f}_{BP}$ and $\boldsymbol{f}_{RP}$ are zero-padded vectors containing the sampled BP and RP spectra, respectively.

The terms in this expression account for:

1. The sample-space covariance matrix from *Gaia* DR3, denoted as $\mathbf{C}_{\boldsymbol{f}_{\rm gaia}}$
2. A 0.5% independent flux uncertainty at each wavelength
3. A 0.5% uncertainty in the zero-point calibration of the overall spectral flux
4. A 0.1% uncertainty in the zero-point calibration of the BP spectrum
5. A 0.1% uncertainty in the zero-point calibration of the RP spectrum

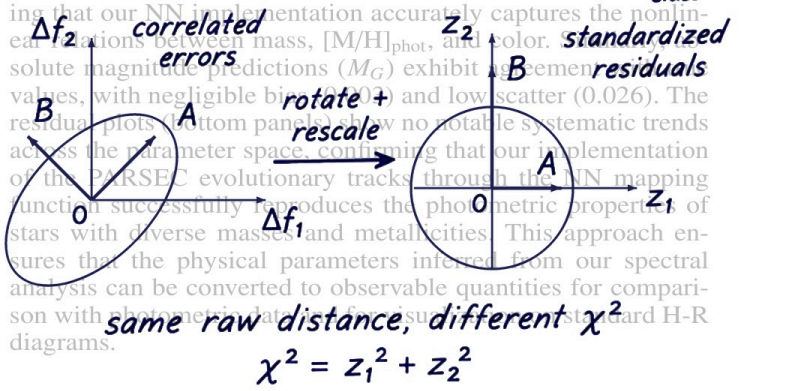
Following Zhang et al. (2023), we inflate the original covariance matrix to avoid dominance by small uncertainties reported in *Gaia* DR3. The zero-point correction involves three components: the overall spectral zero-point and the relative zero-points for BP and RP spectra.

Appendix B: From physical parameters to CMD

The mapping function $\mathcal{M}$ that translates stellar physical parameters to H-R diagram coordinates is implemented as a neural network with three hidden layers, maintaining consistency with the spectral emulator architecture described in Section 3.1.

Figure .1 validates our stellar parameter-to-observable mapping function $\mathcal{M}$, showing notable accuracy in translating physical parameters to H-R diagram positions. The correlation between true and predicted color indices (B-R) shows high precision with minimal bias (0.001) and scatter (0.021), demonstrating that our NN implementation accurately captures the nonlinear relations between mass, $[\mathrm{M/H}]_{\rm phot}$, and color. Similarly, absolute magnitude predictions (M_G) exhibit agreement with true values, with negligible bias (0.003) and low scatter (0.026). The residual plots (bottom panels) show no notable systematic trends across the parameter space, confirming that our implementation of the PARSEC evolutionary tracks through the NN mapping function successfully reproduces the photometric properties of stars with diverse masses and metallicities. This approach ensures that the physical parameters inferred from our spectral analysis can be converted to observable quantities for comparison with photometric data, facilitating studies in the standard H-R diagrams.

The small upper plots test the parameter-to-colour mapping, not binary classification.



Appendix C: The Catalog

Appendix C.1: Binary candidate selection

We provide both the single and binary fitting results for ∼ 35 million sources in *Gaia* database. For each source, we include the optimal parameters from both single-star and binary-star models, along with quality flags and derived quantities. The catalog contains 14,183,917 binary candidates, with a subset of 959,594 sources meeting our stricter criteria for the "golden sample." For each source, we provide the primary and secondary stellar parameters (masses, metallicities), the mass ratio, flux ratios in different bands, and quality metrics including χ^2 values and fit convergence flags. The catalog also includes *Gaia* parallax and extinction estimates from both our model fits and external dust maps. Users can filter the catalog based on the provided quality flags to select samples optimized for either completeness or purity depending on their scientific objectives.

Appendix C.2: Quality flags

The catalog includes **quality flags** to assess the reliability and convergence status of single- and binary-model fits. The `fit_status_single` and `fit_status_binary` codes indicate termination reasons: `-1` for invalid inputs, `0` for exceeding function evaluations, and 1–4 for convergence via `gtol`, `ftol`, `xtol`, or multiple tolerances. The Boolean `fit_success_single` and `fit_success_binary` flags denote overall success (typically `True` for status codes ≥ 1, subject to additional criteria like physically plausible parameters). The `flag_extrap_m2` highlights

Mock-calibrated purity and completeness vary across stellar-parameter bins.

Table C.1: Binary Classification Thresholds and Performance Metrics by Parameter Bins

Parameter Bins			Threshold	Performance Metrics		
Stellar Mass	$[M/H]_{phot}$	Mass Ratio q		F_β	Purity	Completeness
0.1-0.5	−2.0–−1.0	0.1-0.5	4.75	80.1%	96.0%	48.1%
0.1-0.5	−2.0–−1.0	0.5-0.9	1.27	92.5%	95.5%	82.2%
0.1-0.5	−2.0–−1.0	0.9-1.0	−0.77	88.8%	88.1%	91.6%
0.1-0.5	−1.0–−0.5	0.1-0.5	5.09	73.3%	95.8%	37.8%
0.1-0.5	−1.0–−0.5	0.5-0.9	1.61	92.8%	95.7%	82.9%
0.1-0.5	−1.0–−0.5	0.9-1.0	−1.91	83.5%	80.8%	96.3%
0.1-0.5	−0.5-0.0	0.1-0.5	4.42	78.5%	94.3%	47.0%
0.1-0.5	−0.5-0.0	0.5-0.9	2.12	93.8%	95.8%	86.6%
0.1-0.5	−0.5-0.0	0.9-1.0	3.73	87.0%	97.3%	61.2%
0.1-0.5	0.0-0.5	0.1-0.5	5.11	85.3%	99.2%	54.6%
0.1-0.5	0.0-0.5	0.5-0.9	5.35	96.8%	99.3%	87.9%
0.1-0.5	0.0-0.5	0.9-1.0	6.11	97.0%	99.4%	88.4%
0.5-1.0	−2.0–−1.0	0.1-0.5	3.47	72.3%	83.8%	46.6%
0.5-1.0	−2.0–−1.0	0.5-0.9	1.82	87.6%	92.2%	73.2%
0.5-1.0	−2.0–−1.0	0.9-1.0	–	–	–	–
0.5-1.0	−1.0–−0.5	0.1-0.5	3.54	74.1%	89.8%	43.6%
0.5-1.0	−1.0–−0.5	0.5-0.9	4.04	89.2%	96.8%	68.0%
0.5-1.0	−1.0–−0.5	0.9-1.0	−7.11	80.4%	76.7%	99.6%
0.5-1.0	−0.5-0.0	0.1-0.5	3.80	72.2%	91.3%	39.4%
0.5-1.0	−0.5-0.0	0.5-0.9	2.32	89.7%	94.1%	75.7%
0.5-1.0	−0.5-0.0	0.9-1.0	2.95	74.9%	89.3%	45.5%
0.5-1.0	0.0-0.5	0.1-0.5	3.92	77.8%	92.5%	47.6%
0.5-1.0	0.0-0.5	0.5-0.9	4.40	94.7%	97.9%	83.7%
0.5-1.0	0.0-0.5	0.9-1.0	8.29	97.7%	99.7%	90.2%
1.0-1.5	−2.0–−1.0	0.1-0.5	3.12	55.4%	67.0%	32.8%
1.0-1.5	−2.0–−1.0	0.5-0.9	3.86	89.5%	95.4%	71.7%
1.0-1.5	−2.0–−1.0	0.9-1.0	4.58	96.2%	98.3%	88.4%
1.0-1.5	−1.0–−0.5	0.1-0.5	2.03	56.8%	66.8%	35.5%
1.0-1.5	−1.0–−0.5	0.5-0.9	4.32	89.7%	95.5%	72.4%
1.0-1.5	−1.0–−0.5	0.9-1.0	45.54	91.6%	100.0%	68.5%
1.0-1.5	−0.5-0.0	0.1-0.5	1.90	63.3%	76.0%	37.9%
1.0-1.5	−0.5-0.0	0.5-0.9	1.90	87.8%	92.8%	72.2%
1.0-1.5	−0.5-0.0	0.9-1.0	4.41	85.9%	96.7%	59.4%
1.0-1.5	0.0-0.5	0.1-0.5	2.91	63.4%	84.9%	31.5%
1.0-1.5	0.0-0.5	0.5-0.9	3.38	89.1%	96.9%	67.6%
1.0-1.5	0.0-0.5	0.9-1.0	3.34	82.9%	90.9%	61.2%
Overall			–	83.6%	91.6%	66.0%

mock purity mock completeness

cases where the binary fit yields a secondary mass $< 0.1\,M_\odot$, signaling potential extrapolation beyond the model's calibrated range. These flags collectively help users gauge the robustness of fitted parameters and identify fits requiring scrutiny.

The `least_squares` optimizer in `Scipy` uses three key tolerance parameters to determine convergence:

gtol **Gradient tolerance**
Optimization stops when the maximum gradient magnitude falls below this threshold.
Corresponds to quality flag: `fit_status_* = 1`

ftol **Function tolerance**
Stops when relative improvements in the cost function become negligible.
Corresponds to quality flag: `fit_status_* = 2`

xtol **Parameter tolerance**
Terminates when parameter adjustments become sufficiently small.
Corresponds to quality flag: `fit_status_* = 3`

both `ftol + xtol`
Stopping condition met for both function and parameter tolerances.
Corresponds to quality flag: `fit_status_* = 4`

Typical default values are 10^{-8} for all tolerances. Lower values enforce stricter convergence.

Appendix C.3: Cross-match with Gaia Non-Single Star catalog

In addition to our primary binary star catalog, we provide a supplementary dataset cross-matching our identified binary systems with the Gaia DR3 Non-Single Star (NSS) catalog. Table C.3a summarizes the key NSS parameters included in this cross-matched catalog. This cross-matched catalog is particularly valuable for systems with both spectrophotometric signatures (identified through our neural network approach) and dynamical evidence (from the Gaia NSS solutions), allowing for more robust mass ratio determinations and complete orbital characterization.

High purity does not imply that most binaries are recovered.

Table C.2: Parameter Descriptions for the MSMS Binary Catalog (Parent Sample, Binary Candidates, and Golden Sample)

Parameter	Type	Description
source_id	Int	Gaia DR3 source identifier
		Fundamental parameters
ra	Double	Right ascension by *Gaia* DR3 (deg)
dec	Double	Declination by *Gaia* DR3 (deg)
bp_rp	Double	Gaia DR3 B-R colors
phot_g_mean_mag	Double	Gaia DR3 G-band magnitude
parallax	Double	Gaia DR3 parallax (mas)
parallax_error	Double	Uncertainty in parallax measurement
radial_velocity	Double	Gaia DR3 radial velocity (km/s)
radial_velocity_error	Double	Uncertainty in Gaia DR3 radial velocity (km/s)
E_edhf	Double	Extinction from Bayestar19 dust map
A_G_edhf, A_B_edhf, A_R_edhf	Double	Extinction in G, B, and R bands from Edenhofer et al. (2024)'s dustmap
		Best-fit results for single XP spectral model
mass_single	Double	Stellar mass from single-model ($M_\odot$)
e_mass_single	Double	Uncertainty in stellar mass from single-model ($M_\odot$)
mh_single	Double	Metallicity from single-model solution [M/H]
e_mh_single	Double	Uncertainty in metallicity from single-model solution
E_single	Double	Extinction from single-model
e_E_single	Double	Uncertainty in extinction from single-model
chi2_mass_mh	Double	χ^2 value of the single-model fit
		Best-fit results for binary XP spectral model
m1_binary	Double	Primary star mass ($M_\odot$)
e_m1_binary	Double	Uncertainty in primary star mass ($M_\odot$)
m2_binary	Double	Secondary star mass ($M_\odot$)
mh_binary	Double	System photometric metallicity [M/H]
e_mh_binary	Double	Uncertainty in system photometric metallicity
q_binary	Double	Mass ratio (m_2/m_1)
e_q_binary	Double	Uncertainty in mass ratio
E_binary	Double	Extinction from binary-model
e_E_binary	Double	Uncertainty in extinction from binary-model
chi2_mass_mh_binary	Double	χ^2 value of the binary-model fit
flux_ratio_mean	Double	Mean XP flux ratio between primary and secondary components
flux_ratio_std	Double	Standard deviation of the XP flux ratio between primary and secondary components
b_r_prim	Double	B-R color of the primary component
b_r_second	Double	B-R color of the secondary component
MG_prim	Double	Absolute G magnitude of the primary component
MG_second	Double	Absolute G magnitude of the secondary component
		Quality flags
fit_status_single	int	Status code for single-model fit convergence
fit_status_binary	int	Status code for binary-model fit convergence
fit_success_single	Boolean	Whether the single-model fit was successful
fit_success_binary	Boolean	Whether the binary-model fit was successful
flag_extrap_m2	Boolean	Whether secondary mass from the binary fit lower than 0.1 $M_\odot$

Table C.3: Binary Catalog Information

(a) Selected Parameters from the Gaia DR3 Non-Single Star (NSS) Catalog

Parameter	**Description**
source_id	Gaia DR3 source identifier
nss_solution_type	Binary solution type
	Radial Velocity Parameters
radial_velocity	Gaia radial velocity (km/s)
radial_velocity_error	Uncertainty in the radial velocity measurement
rv_chisq_pvalue	p-value of the chi-square test for radial velocity variability
rv_expected_sig_to_noise	Expected signal-to-noise ratio for radial velocity measurements
rv_nb_transits	Number of radial velocity transits
rv_renormalised_gof	Renormalised goodness of fit for radial velocity solutions
	Orbital Parameters
period	Orbital period of the binary system (days)
eccentricity	Orbital eccentricity
semi_amplitude_primary	Semi-amplitude of the radial velocity curve for the primary component (km/s)
semi_amplitude_secondary	Semi-amplitude of the radial velocity curve for the secondary component (km/s)
	Stellar Parameters
m1	Mass of the primary component (solar masses)
m2	Mass of the secondary component (solar masses)
fluxratio	Flux ratio of the secondary to the primary in the G-band (F_2/F_1)

Note: For additional details, see the Gaia DR3 documentation: https://gea.esac.esa.int/archive/documentation/GDR3/Gaia_archive/chap_datamodel/sec_dm_non--single_stars_tables/