

Gaia DR3 detectability of unresolved binary systems

Alfred Castro-Ginard¹, Zephyr Penoyre¹, Andrew R. Casey^{2,3}, Anthony G.A. Brown¹, Vasily Belokurov⁴, Tristan Cantat-Gaudin⁵, Ronald Drimmel⁶, Morgan Fouesneau⁵, Shourya Khanna⁶, Evgeny P. Kurbatov⁴, Adrian M. Price-Whelan⁷, Hans-Walter Rix⁵, and Richard L. Smart⁶

¹ Leiden Observatory, Leiden University, Einsteinweg 55, 2333 CC Leiden, the Netherlands
e-mail: acaastro@strw.leidenuniv.nl

² School of Physics and Astronomy, Monash University, VIC 3800, Australia

³ Centre of Excellence for Astrophysics in Three Dimensions (ASTRO-3D), Melbourne, Victoria, Australia

⁴ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom

⁵ Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

⁶ INAF - Osservatorio Astrofisico di Torino, Strada Osservatorio 20, Pino Torinese 10025 Torino, Italy

⁷ Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY 10010, USA

Received date / Accepted date

ABSTRACT

Context. *Gaia* can not individually resolve very close binary systems, however, the collected data can still be used to identify them. A powerful indicator of stellar multiplicity is the sources reported Renormalized Unit Weight Error (*ruwe*), which effectively captures the astrometric deviations from single-source solutions.

Aims. We aim to characterise the imprints left on *ruwe* caused by binarity. By flagging potential binary systems based on *ruwe*, we aim to characterise which of their properties will contribute the most to their detectability.

Methods. We develop a model to estimate *ruwe* values for observations of *Gaia* sources, based on the biases to the single-source astrometric track arising from the presence of an unseen companion. Then, using the recipes from previous *Gaia*Unlimited selection functions, we estimate the selection probability of sources with high *ruwe*, and discuss what binary properties contribute to increasing the sources' *ruwe*.

Results. We compute the maximum *ruwe* value which is compatible with single-source solutions as a function of their location on-sky. We see that binary systems selected as sources with a *ruwe* higher than this sky-varying threshold have a strong detectability window in their orbital period distribution, which peaks at periods equal to the *Gaia* observation time baseline.

Conclusions. We demonstrate how our sky-varying *ruwe* threshold provides a more complete sample of binary systems when compared to single sky-averaged values by studying the unresolved binary population in the *Gaia* Catalogue of Nearby Stars. We provide the code and tools used in this study, as well as the sky-varying *ruwe* threshold through the *Gaia*Unlimited Python package.

Key words. Galaxy: general – astrometry – methods: data analysis – methods: statistical – catalogs

1. Introduction

A large number of the stars in our Galaxy are formed in binary or multiple systems. The fraction of stars hosting at least one companion depends on the star's properties, e.g., for main sequence (MS) stars this fraction increases with stellar mass (Duchêne & Kraus 2013), with $\lesssim 50\%$ of solar-type stars being in multiple systems (Raghavan et al. 2010). It also depends on the environment where these stars are formed, finding a higher multiplicity fraction in stellar clusters or star-forming regions compared to field stars (Duchêne et al. 2018). Thus, understanding the observed properties of binary (or multiple) systems provides a better insight into star-formation processes and the evolution of stars and stellar systems (see El-Badry 2024, and references therein for a recent review).

There are several methods to unveil binarity among the current stellar surveys. Using spectroscopic data from APOGEE (Majewski et al. 2017), Price-Whelan et al. (2020) characterised the shift in the absorption lines due to the motion generated by the star's orbit in a binary system. El-Badry et al. (2021) used *Gaia* EDR3 (Gaia Collaboration et al. 2021a) astrometry to generate a catalogue of wide binary systems, using the common

proper motions and parallaxes of individually resolved components. Wallace (2024) used a combination of photometric data from *Gaia* (Gaia Collaboration et al. 2016), 2MASS (Skrutskie et al. 2006) and WISE (Wright et al. 2010) to detect the excess of light in a H-R diagram produced by the presence of a companion. However, most binary systems will remain unresolved and undetected given their intrinsic properties, distance and the limitations of our observations.

The third *Gaia* data release (*Gaia* DR3, Gaia Collaboration et al. 2023) contains around $1.5 \cdot 10^9$ sources with a full, single-star astrometric solution, but potentially contains a large number of binary systems. Penoyre et al. (2020) showed how the presence of an unresolved companion may affect the single-star solution provided by a survey such as *Gaia*. Using this effect, Belokurov et al. (2020) demonstrated how the binary properties will impact the *Gaia* DR2 (Gaia Collaboration et al. 2018) astrometric solution and how this may be indicated by an elevated value of the Renormalized Unit Weight Error (*ruwe*). With *ruwe* being an effective indicator to pinpoint stellar binarity, different *ruwe* thresholds to indicate good single-star solutions, thus flagging "bad" astrometric solutions potentially caused by unresolved binaries, have been proposed for *Gaia* DR2 (1.4, Lindegren 2018)

and *Gaia* EDR3 (1.25, Penoyre et al. 2022b). In this paper, we go a step further, and we i) estimate the maximum ruwe value compatible with single stars as a function of the position on-sky, and ii) characterise what population of binary systems is accessible through this cut on ruwe.

This paper falls within the scope of the GaiaUnlimited project¹, which aims to provide selection functions and tools for their practical usage, for the *Gaia* catalogue and its different data products. This paper thus complements the work done in the estimation of selection functions for the *Gaia* DR3 catalogue (Cantat-Gaudin et al. 2023), for any subsample drawn from the main catalogue (Castro-Ginard et al. 2023) and for the combination of *Gaia* and APOGEE (Cantat-Gaudin et al. 2024). Similarly to the previous works, the code and tools developed for this paper will be provided through the GaiaUnlimited Python package².

This paper is organised as follows. Section 2 describes how *Gaia* observes unresolved binary systems and the details of the calculation of the key metric (ruwe) to identify binarity. In Sect. 3, we explore what properties of binary systems will enhance their detectability in *Gaia*. Section 4 explores the properties of the unresolved binary systems that may be present in the *Gaia* Catalogue of Nearby Stars (GCNS, Gaia Collaboration et al. 2021b). Finally, the summary of the paper and our conclusions are presented in Sect. 5. We show some examples of how to use the GaiaUnlimited python package in the Appendix A.

2. Observation of unresolved binary systems

Gaia has continuously scanned the sky since 2014, registering the times when each source crosses the focal plane (Gaia Collaboration et al. 2016). Subsequently, all these observations undergo processing via the Astrometric Global Iterative Solution (AGIS, Lindegren et al. 2021a,b) which generates the measurements published in the catalogue. AGIS constructs an astrometric track representing a single source when a source has accumulated five valid detections. This process yields fitted astrometric parameters alongside metrics for assessing the quality of the fit, especially ruwe which we rely on in this present work (Lindegren 2018), among other criteria. In essence, ruwe serves as a metric to distinguish between good and bad AGIS solutions using a singular value. Typically, it approximates unity for well-behaved individual sources, while higher values indicate a poorer fit by AGIS.

There are several causes for a bad AGIS fit (see Sect. 2.2.1), evident in a ruwe value significantly exceeding one, which often arise from the wobbling of a source’s photocentre during the *Gaia* observations. This is the case for binary systems (or multiple stellar systems, in general), where the centre of mass (CoM) and the photocentre of the system follow distinct trajectories depending on the system’s characteristics. This effect is shown in Fig. 1, where the photocentre of a binary system (orange dotted line) oscillates around its CoM (indicated by the blue line), while the CoM itself traces the trajectory fitted by AGIS for a single source. The mismatch between the CoM and the photocentre results in an estimated ruwe of 2.07 for this system. Therefore, given that the ruwe effectively captures the presence of an unresolved companion, it serves as a valuable indicator of binarity in *Gaia* sources (Belokurov et al. 2020; Stassun & Torres 2021; Kervella et al. 2022).

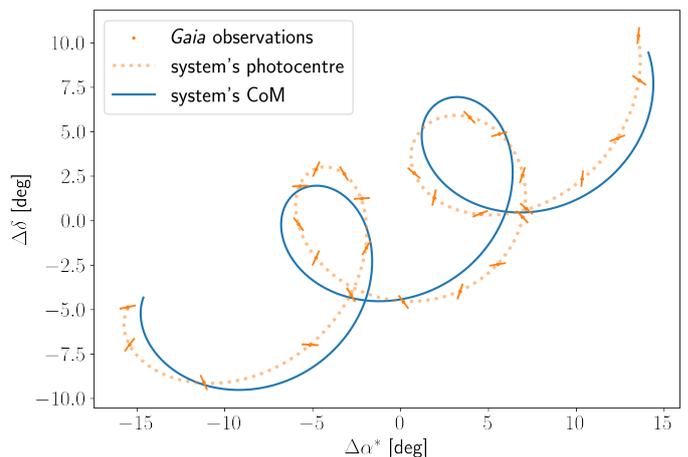


Fig. 1. Astrometric track of a binary system seen by *Gaia*. The solid blue line tracks the motion of the CoM of the system, while the dotted orange line tracks its photocentre. *Gaia* individual scans of the sources are marked with orange solid lines aligned with the observation scanning angle. The goodness-of-fit of the CoM to the *Gaia* observations for this simulated system result in a ruwe of 2.07.

2.1. ruwe for binary systems

Estimating the ruwe for unresolved binary systems involves computing the separation between the photocentre and the CoM in the along-scan direction (AL), which will depend on both the *Gaia* observations and the configuration of the binary system, while considering the astrometric error (σ_{ast}) which depends on the *Gaia* telescope configuration. As described in Lindegren (2022) (see also Sect. 4.1 in Penoyre et al. 2022a), the equations governing a five-parameter solution for noisy observations are described by

$$\begin{aligned} &\Delta\alpha^* \sin\psi + \Delta\delta \cos\psi + \\ &\Delta\varpi P_{AL} + \\ &\Delta\mu_{\alpha^*} \tau \sin\psi + \Delta\mu_{\delta} \tau \cos\psi = \delta\eta + \mathcal{N}(0, \sigma_{\text{ast}}), \end{aligned} \quad (1)$$

where ψ represents the scanning angle, $\tau = t - t_{\text{ref}}$ where t is the observation time in years and $t_{\text{ref}} = 2016.0$, P_{AL} stands for the parallax factor, $\delta\eta$ is the AL bias (derived from the projected AL separation between the CoM and the photocentre) and σ_{ast} is the AL astrometric error. This system comprises n_{obs} equations, where n_{obs} is the number of CCD observations (source transits on *Gaia*’s focal plane typically generate 9 CCD observations Gaia Collaboration et al. 2016), and is solved for $\Delta\alpha^*$, $\Delta\delta$, $\Delta\varpi$, $\Delta\mu_{\alpha^*}$ and $\Delta\mu_{\delta}$, which denote the expected biases arising from binarity (and will be zero for single-sources). From the residuals of these solutions, which quantify how much of the binary motion differs from the fitted single source astrometric solution, we can compute the χ^2 statistic, related to ruwe as

$$\text{ruwe} = \sqrt{\frac{\chi^2}{n_{\text{obs}} - 5}}, \quad (2)$$

note that the ruwe value reported in *Gaia* DR3 is a rescaled version of the original χ^2 from AGIS (see Lindegren 2018, for further details).

The details from the *Gaia* scanning law necessary to estimate a ruwe of a particular system, that is solving Eq. 1, are accessible through the *Gaia* Observation Schedule Tool (GOST³,

¹ <https://gaia-unlimited.org/>

² <https://github.com/gaia-unlimited/gaiaunlimited>

³ <https://gaia.esac.esa.int/gost/>

Fernández-Hernández et al. 2022). GOST provides observation times, parallax factors, and scanning angles linked to a given set of sky coordinates. In our case, these coordinates correspond to the pixel centres at the HEALPix level 5 resolution observed between July 25, 2014, and May 28, 2017, covering the 34-month duration of *Gaia* DR3 observations. We choose the resolution at HEALPix level 5 to limit the computational cost of the calculations, however, the procedure described here and the examples shown in Appendix A are general for any spatial resolution. On the other hand, concerning the observed system, we need to compute the projected AL separation between the system’s CoM and its photocentre at each observation time. This value is either zero (applicable for single sources, as discussed in Sect. 2.2) or computed based on the binary system configuration (as detailed in Sect. 3). Lastly, the AL astrometric error σ_{ast} is determined as a function of the G magnitude, following the relationship provided in Lindegren et al. (2021b, see their Fig. A.1). To obtain its specific numerical value, we use the function `sigma_ast` available within the `astromet` python package⁴.

2.2. ruwe threshold for single stars

Section 2.1 describes the recipe to forward model a *ruwe* value for a binary system based on its properties and location on-sky. However, to flag these binary systems we need to understand what *ruwe* values would describe a well-behaved single-source solution. For this, we characterise the maximum allowed *ruwe*, or *ruwe* threshold, which will be compatible with single sources as a function of the sky coordinates. Following our approach, we estimate *ruwe* values for a population of 10^5 simulated single sources located at the centres of each HEALPix at level 5 and observed under *Gaia* DR3. In this case, where the source’s CoM and photocentre follow identical tracks ($\delta\eta = 0$), the expected *ruwe* value will be a χ distribution scaled by $1/\sqrt{(n-5)}$, which for large n will approach a Gaussian with a dispersion that is proportional to the AL astrometric error. We assign a magnitude to each simulated single source based on a uniform distribution within $G \in (5, 20.7)$ mag, to cover the entire magnitude range. The corresponding σ_{ast} value is computed based on this magnitude. Similar to the approach by Penoyre et al. (2022a, see their Fig. 4), we fit a Student’s t-distribution to the resulting distribution of *ruwe*. From this distribution, we define the maximum *ruwe* value compatible with the star being a single source, allowing for a one-in-a-million confusion rate, i.e., solving for *ruwe* where $1 - P(\text{ruwe} < \tau_{\text{ruwe}}) = 10^{-6}$, where τ_{ruwe} is the computed *ruwe* threshold. This process is repeated 30 times to account for the random errors and to estimate an associated uncertainty ($\leq 2 \cdot 10^{-3}$) for this *ruwe* value. The final simulation-based *ruwe* threshold is the mean value of these 30 realisations.

2.2.1. Calibration with *Gaia* data

So far, we focused on the methodology to calculate the *ruwe* value for a *Gaia* observation of an isolated source, considering the presence or absence of an unseen companion. This methodology is particularly useful for modelling *ruwe* values for unresolved binary systems and deriving general conclusions for this population (see Sect. 3). However, other sources’ properties such as stellar crowding, variability or the presence of disc material around single stars (among others, Belokurov et al. 2020; Fitton et al. 2022), or attitude and geometric calibration errors and other instrumental properties, also contribute to the increase in

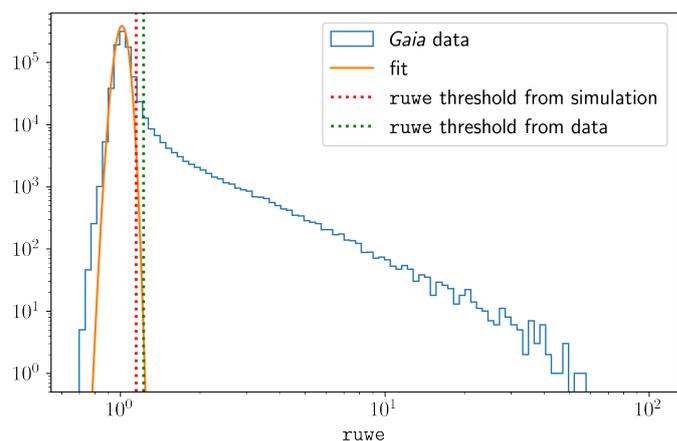


Fig. 2. Histogram of *ruwe* for the sources in the *Gaia* archive (blue line) HEALPix centred at $l = 73.58^\circ$, $b = 3.17^\circ$. The solid orange line indicates our fit of the single source population in that HEALPix. The dotted lines indicate the values of the *ruwe* threshold for single sources with (green) and without (red) crowding effects.

the sources’ *ruwe* (Lindegren 2018). In regions with particularly high stellar density, the observed single source might be influenced by neighbouring stars causing *Gaia*’s measured image to deviate from the single source PSF model. This deviation leads to difficulties in accurately determining the PSF’s centroid, resulting in an elevated *ruwe*.

A detailed modelling of the previous effects on the source’s *ruwe* is out of the scope of this paper. Instead, we gauge these effects, specifically focusing on crowding, directly from the *Gaia* data employing the following approach. For HEALPix at level 5, we use `scipy`’s `curve_fit` function in Python⁵ to fit a Gaussian to the left half of the *ruwe* distribution of all the sources in the HEALPix retrieved from *Gaia* archive (left half with respect to the mode of the distribution, usually centred at one but automatically estimated from the data). Using the full Gaussian, representing a single-source population, we obtain the maximum *ruwe* allowed for single sources by applying the procedure described in Sect. 2.2. Figure 2 shows the *ruwe* distribution for the $\sim 10^6$ sources in *Gaia* corresponding to the HEALPix centred at $l = 73.58^\circ$ and $b = 3.17^\circ$. Additionally, we show the Gaussian distribution fitted to the sources centred around *ruwe* ~ 1 . In this case, the *ruwe* threshold for single sources estimated from both the simulated single-source population (see Sect. 2.2) and actual data are 1.14 and 1.22, respectively.

The data-driven approach adopted here faces limitations in its application across the whole sky due to the necessity of a minimum source count, and the *ruwe* distribution around its mode to be well represented by a Gaussian for an accurate fit. For HEALPix regions where we obtain a data-driven *ruwe* threshold, we fit a second-order polynomial to the difference of the *ruwe* thresholds, i.e., $\tau_{\text{ruwe, data}} - \tau_{\text{ruwe, simulation}}$, as a function of the logarithm of the source count within that HEALPix. With this, we derive an adjustment factor, or offset, to refine the simulation-based *ruwe* threshold accounting primarily for the influence of crowding based on the number of sources present within each HEALPix.

Figure 3 shows the estimated upper value of *ruwe* compatible with single star solutions across the sky. It shows a strong dependence on the *Gaia* scanning law, indicating that regions with a higher number of observations tend to exhibit lower *ruwe*

⁴ <https://github.com/zpenoyre/astromet.py>

⁵ <https://scipy.org/>

thresholds since the fitted AGIS solution for each source is better constrained. We also see the impact of crowding within the Galactic plane, especially toward the Galactic centre and the Magellanic clouds. These regions show elevated *ruwe* thresholds due to challenges in reliably fitting the single source track using AGIS caused by the misidentification of sources due to their high stellar density. This misidentification for a large fraction of sources in those regions will cause the single source *ruwe* distribution to be skewed towards *ruwe* larger than one, resulting in a higher threshold value after our calibration.

The *ruwe* threshold variation across different sky coordinates presented in our work represents an improvement compared to prior studies that employed a singular *ruwe* value across the whole sky. For instance, in *Gaia* DR2, Lindegren (2018) suggested a *ruwe* threshold of 1.4 to characterise good *Gaia* solutions, with higher values potentially indicating binarity or other influencing factors. For *Gaia* EDR3, Penoyre et al. (2022a) investigated the imprints left in the *ruwe* distribution attributed to nearby binary systems (within 100 pc) and recommended a threshold of 1.25 for identifying potential binarity (the authors' discussion focuses around *luwe*, which stands for local-normalised *uwe*). These different *ruwe* thresholds suggested in different studies highlight the sensitivity of *ruwe* to different observational factors, their relation to specific *Gaia* data releases, e.g., DR2 or DR3, and the characteristics of the observed regions.

2.2.2. Selection probability of sources with high *ruwe*

We compute the selection probability for sources in *Gaia* DR3 with a value for *ruwe* higher than the estimated *ruwe* threshold, using the approach described in Castro-Ginard et al. (2023) to estimate selection functions for *Gaia* subsamples. We note here that for a universe consisting only of single stars without any photocentre motion, and for perfectly behaved and calibrated *Gaia* instruments, by construction the probability to select sources above the *ruwe* threshold would be 10^{-6} , independent of sky position or source brightness. In reality (as shown below) the selection probability is much higher and this can indicate, among others, the presence of unseen companions to the selected sources.

Within each HEALPix at level 5, we determine the expected selection probability as

$$E(p) = \frac{k + 1}{n + 2}, \quad (3)$$

where k represents the count of sources exceeding the estimated sky-varying *ruwe* threshold (shown Fig. 3) and n is the total number of sources within that HEALPix. Figure 4 shows the estimated selection probability across the sky, revealing that the selection of these sources, which in some regions may translate to the detectability of binary systems defined as sources with large *ruwe*, correlates with the *Gaia* scanning law. We see a higher selection probability in crowded areas, e.g., the Galactic plane, particularly towards the Galactic centre, and the Magellanic Clouds or the hot pixels corresponding to the cores of globular clusters. We can't interpret these high-selection probability regions as being highly populated by binary systems. As previously discussed, several reasons in both the sources or instrument properties can contribute to an increased *ruwe* value. Therefore, these high-selection probability regions may be seen as regions with an important contribution of factors other than binarity to increase the sources' *ruwe*, tracing regions of high contamination when binary systems are selected solely based on

the latter. In this paper, our modelling primarily focuses on the increase in *ruwe* caused by a bad single-source fit by AGIS due to the presence of an unresolved companion, which is typically the predominant contribution across most of the sky, i.e. Galactic latitudes above the plane and avoiding crowded regions. Consequently, selecting binary systems based on *ruwe* within highly crowded areas will include a higher level of contamination or interference from other contributing factors, even with our sky-varying *ruwe* threshold taking crowding into account.

3. Properties of detected binary systems

The approach presented in Sect. 2 allows us to estimate *ruwe* values for any source observed by *Gaia*, including unresolved binary systems, based on their properties and location on-sky. With this approach, and considering a specific population of binary systems, we can characterise what properties will enhance the probability of detection for unresolved binary systems in *Gaia* DR3, when detected as systems whose estimated *ruwe* is higher than the threshold shown in Fig. 3.

3.1. Simulated binary population

We simulate a population of binary systems to predict their detectability as a function of the system's parameters. This simulated population is not intended to mirror a realistic representation of binary systems distributed throughout the Galaxy. Instead, the simulation aims to cover a broad range of stellar binary characteristics, e.g. mass and light ratios or orbital properties, general enough to allow for an exploration of the effect of the binary parameters on the probability of being selected based on a *ruwe* threshold which varies across the sky.

We generate around one million MS binaries to study the detectability variations of different simulated parameters. First, we retrieve the stellar parameters from the PARSEC⁶ isochrones (Bressan et al. 2012) for a uniform grid of different ages. We isolate the isochrones' MS and take the corresponding masses as the mass of the binary primary component. We assign a mass ratio (q) following a uniform distribution from 0 to 1 to each of the primary masses to build a binary system. These systems are uniformly placed at varying distances ranging from 10pc to 5kpc. We calculate *Gaia*'s G magnitude from the mass and distance distributions, considering no extinction, and restrict our analysis to systems within *Gaia*'s observational limits, i.e., $G < 20.7$ mag. The remaining parameters for the simulated binaries follow the distributions described in Table 1 and are similar to those used in Penoyre et al. (2022a, their Table 1). The semi-major axis of the binary orbit is computed using the masses of both components and the orbital period. Finally, since we consider MS-MS binaries, the luminosity ratio is computed from the mass ratio as $l = q^{3.5}$.

3.2. Selection probability and binary system properties

We distribute the generated population of binary systems across various sky locations, to evaluate the fraction of these systems that would be detected using *ruwe* (estimated as in Sect. 2.1). Figure 5 shows this detection fraction for the different binary configurations considered at the pixel centred at $l = 15.3^\circ$ and $b = -14.4^\circ$, with yellow pixels representing a nearly complete detection while bluer pixels indicating a detection fraction closer to zero within that specific configuration.

⁶ <http://stev.oapd.inaf.it/cgi-bin/cmd>

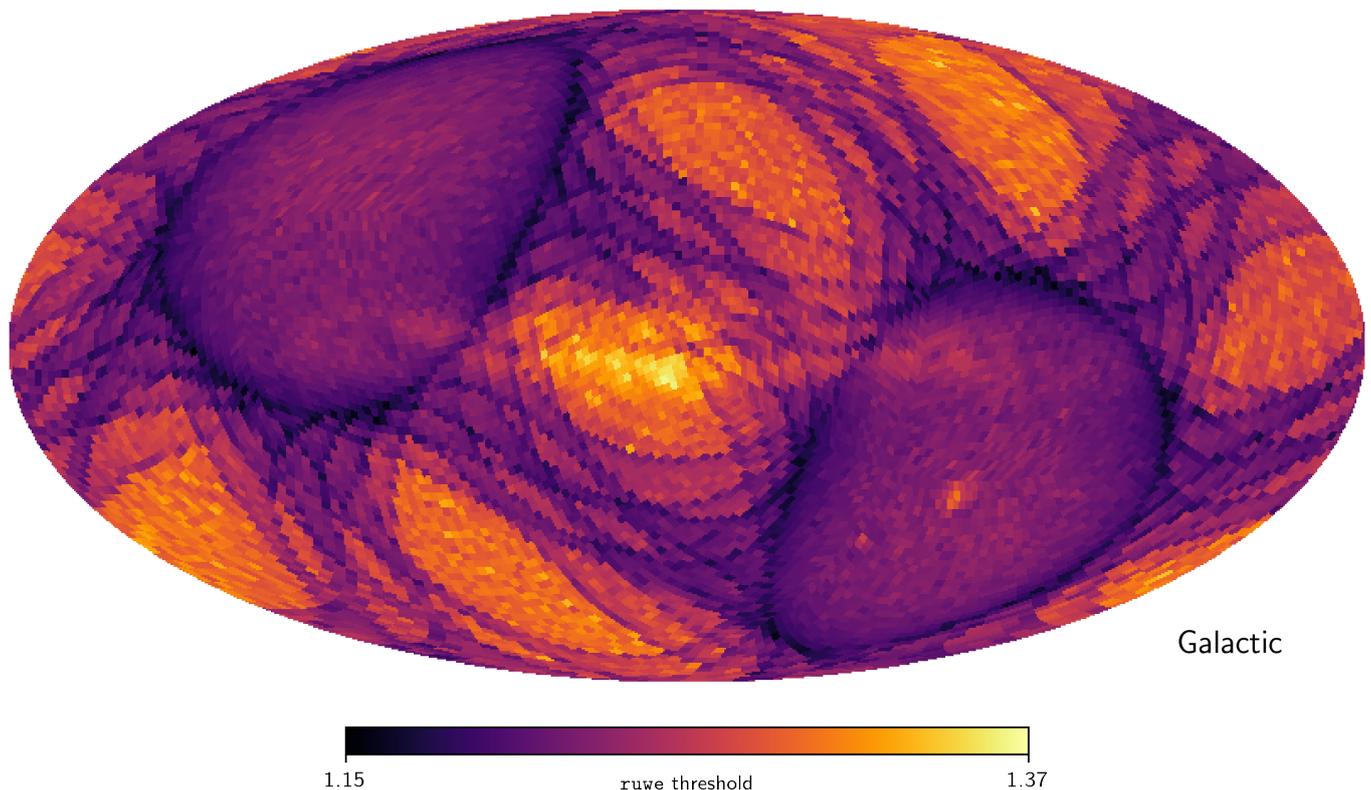


Fig. 3. Sky map of the sky-varying *ruwe* threshold at the resolution of HEALPix level 5. Sources above these *ruwe* threshold indicate potential binary systems. These *ruwe* values show a strong influence of the *Gaia* scanning law, while also highlighting regions with high source density.

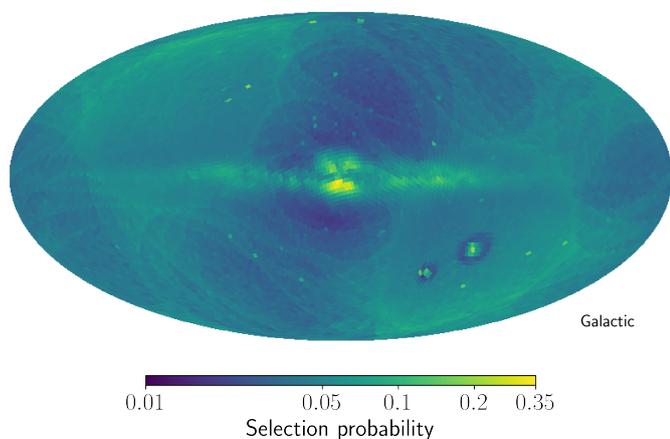


Fig. 4. Selection probability for sources with a *ruwe* higher than the threshold represented in Fig. 3. The regions with higher selection probability, e.g., cores of globular clusters, the Galactic plane and the Magellanic Clouds, contain sources with a *ruwe* increased due to a combination of binarity, source crowding and other effects. These regions may indicate a higher contamination rate when selecting binaries solely based on *ruwe*. The dark halos around the Magellanic Clouds are due to an excess of faint sources with low *ruwe* in that region.

As expected, the binary detectability peaks at a mass ratio of $q = 0.5$. No systems are detected when $q = 0$ (equivalent to single sources), or when $q = 1$ where the CoM and the photocentre trace the same trajectory, which means the value of *ruwe* is not changed.

The detection probability shows a peak at an orbital period aligning with the observational baseline of the survey, which is

Parameter	Distribution
Distance [pc]	$U(10, 5000)$
$\log(\text{Period [yr]})$	$U(-2.56, 2)$
Mass ratio	$U(0, 1)$
Eccentricity	$U(0, 1)$
Initial phase [rad]	$U(0, 2\pi)$
Azimuthal view angle (ϕ) [rad]	$U(0, 2\pi)$
$\cos(\text{Polar view angle } (\theta) \text{ [rad]})$	$U(-1, 1)$
Planar projection angle (ω) [rad]	$U(0, 2\pi)$

Table 1. Distributions for the simulated binary system population.

34 months for *Gaia* DR3. This probability gradually declines for longer periods, given that no observable orbital motion is evident within this observation time (Penoyre et al. 2020). Furthermore, we observe notable declines in the fraction of detected systems when the orbital period coincides harmonically with a year, prominently visible in Fig. 5 at periods of 1 year and 1/2 year. This decrease in detectability arises due to a confusion of the orbital motion with a parallax shift, leading to a situation where a seemingly good astrometric fit (resulting in low *ruwe*) corresponds to an erroneous parallax. These features in the period distribution are also visible in the astrometric binaries in *Gaia* DR3 (Halbwachs et al. 2023, also selected based on *ruwe*). We observed that the strength of this effect, however, depends on the direction of the observed binary systems (details of the *Gaia* scanning law).

From the first column in Fig. 5, we see that our ability to detect binary systems is monotonically decreasing with distance,

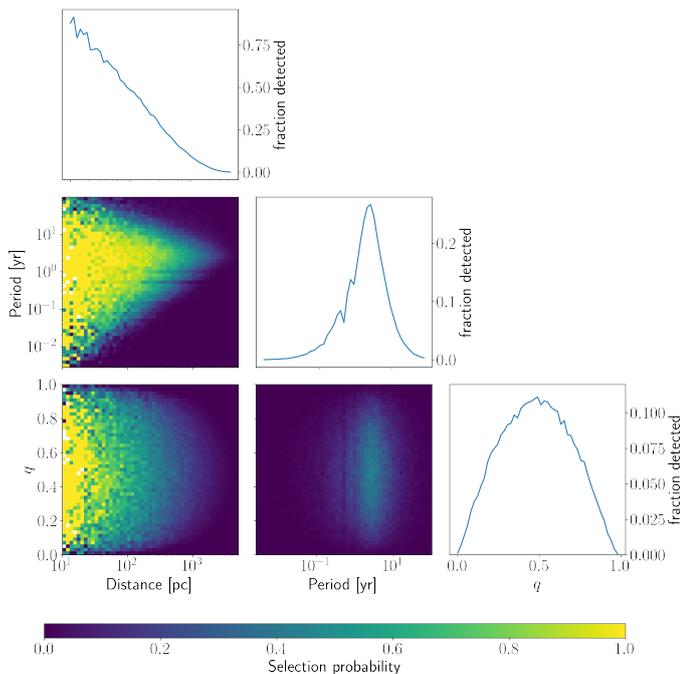


Fig. 5. Heatmap of binary systems' properties, i.e., distance, period and mass ratios, that will enhance their detection. The detectability of these systems monotonically decreases with distance, reaching a detection fraction of 50% at around 100 pc. There is a strong detectability window around periods of three years, corresponding to the *Gaia* observation time baseline. Systems with mass ratios of zero, i.e., single sources, or one show no sign of binarity on *ruwe* since there is no wobbling of their CoM around the photocentre.

reaching 50% of systems detected at around 100pc. To capture the variation of this behaviour across different regions of the sky, we estimate the *ruwe* of the simulated binary population described in Sect. 3.1 in different pixels from the north to the south ecliptic poles at ecliptic longitude $\lambda = 0^\circ$, and from $\lambda = 0^\circ$ to $\lambda = 360^\circ$ at ecliptic latitude $\beta = 0^\circ$, as shown in Fig. 6. We use ecliptic coordinates since the *Gaia* scanning law is symmetric in this coordinate system, and the strongest contribution to the variation of the *ruwe* threshold across the sky (see Sect. 2.2). Also, we focus on periods from one to four years which is the window for optimal detectability, since most of the selected binary systems will be in this period range (see discussion in Sect. 4). Figure 6 shows how the variation of the distance where 50% of the simulated systems are detected (where the median *ruwe* is equal to the *ruwe* threshold shown in Fig. 3). We see that in the ecliptic poles, and up to $|\beta| = 45^\circ$, where *Gaia* has uniformly scanned the sky with a large number of observations, we can detect binary systems at larger distances (up to around 1 kpc) due to a more constrained astrometric solution, thus resulting in a lower *ruwe*. For the $|\beta| < 45^\circ$, on the other hand, the distance where we detect at least 50% of the simulated binary systems can decrease down to 800pc. Along the λ direction, the distance reached shows a more irregular behaviour since this region captures larger variations in the *Gaia* scanning law. However, the distances where we can detect systems can vary ~ 200 pc depending on the scanning configurations. If we include in the analysis the whole range of simulated periods, the distances where we detect 50% of the simulated systems range from around 60 up to 150pc, approximately, depending on the (λ, β) direction.

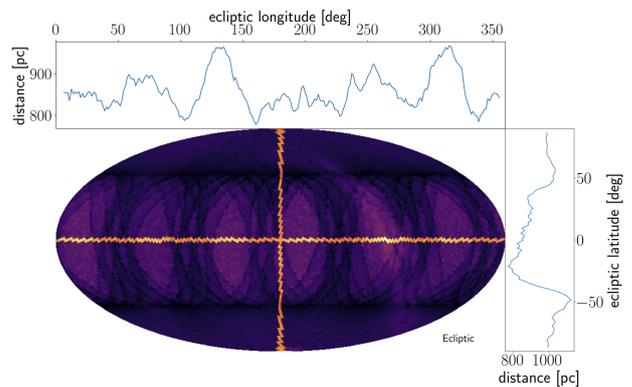


Fig. 6. Central panel: Same as in Fig. 3, in ecliptic coordinates. The pixels highlighted in the vertical and horizontal directions show the regions where we estimate the maximum distance where we detect half of the simulated systems, which is shown in the right and top panels, respectively. Top panel: Distance where 50% of the systems are detected, for the highlighted pixels along the ecliptic $\lambda = 0^\circ$. Right panel: Same as the top panel, for ecliptic $\beta = 0^\circ$.

4. The GCNS binary population through GUMS

Gaia Collaboration et al. (2021b) produced a catalogue of 331 312 objects within 100 pc which was released together with *Gaia* EDR3. The catalogue includes a list of probable members to the two closest open clusters, i.e., the Hyades and Coma Berenices, white dwarfs (WDs) and resolved binary systems. It also includes numerous sources with *ruwe* significantly higher than one, which may be potential unresolved binary systems that may have gone unnoticed. We find 74 785 potential unresolved binary systems with a catalogued *ruwe* higher than the threshold for single sources we computed in Sect. 2.2 (see Fig. 3). Based on our approach, we can select more binary system candidates compared to the selection based on *ruwe* higher than 1.4 and 1.25, which provide 55 754 and 70 848, respectively. In Fig. 7 we show how our selection of unresolved binary systems is distributed in a colour-magnitude diagram (probability estimated as in Eq. 3 with respect to the colour and magnitude bins). We observe a high selection probability in the brighter MS, which monotonically decreases towards faint magnitudes. A similar trend is observed in the WD sequence ($10 \lesssim M_G \lesssim 15$ and $0 \lesssim G_{BP} - G_{RP} \lesssim 1$), where the selection probability is higher in the brighter sequence. The region between the MS and the WD sequence, which shows a selection probability unexpectedly high, corresponds to sources with "poor astrometric solutions" in the GCNS catalogue (see Section 2.1 in Gaia Collaboration et al. 2021b, for further details). Although this region also contains MS-WD binaries, the *ruwe* for the sources in this region is probably also increased for reasons other than binarity. Therefore, this region may represent a highly contaminated region in our analysis.

To understand the completeness and contamination of this sample, and the types of systems that our method selects among the unresolved binary system candidates we use the *Gaia* Universe Model Snapshot (GUMS, Robin et al. 2012). GUMS provides, through the *Gaia* Object Generator (GOG, Luri, X. et al. 2014), astrometry, photometry and astrophysical parameters of sources as if they were observed by *Gaia*. GUMS only provides around 100 000 sources within 100 pc, therefore, to increase the statistical significance of the analysis we increase the volume to analyse to 200 pc, resulting in 882 107 sources of which 243 377 are true binary systems.

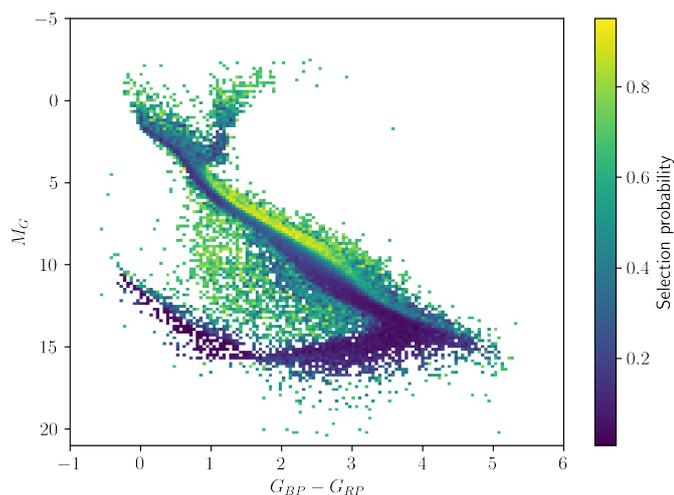


Fig. 7. Selection probability of the GCNS sources with *ruwe* larger than our threshold shown in Fig. 3 across the colour-magnitude diagram. For the MS and WD sequence, we see a gradient on the selection probability for bins at the same $G_{BP} - G_{RP}$ colour. The region between the MS and the WD sequence, which contains the MS-WD binaries, has a high concentration of sources labelled with poor astrometric solutions in the GCNS

We forward model a *ruwe* value for these GUMS sources following the recipe in Sect. 2.1 (see also Appendixy A for instructions on how to use our *GaiaUnlimited* python package). Then, we detect these binary systems based on *ruwe* as we have done for the GCNS. We can recover 51 524 true binary systems with our variable *ruwe* approach, which still results in a more complete sample when compared to *ruwe* selections higher than 1.4 and 1.25, resulting in 43 398 and 50 431, respectively. The contamination of our sample is one single source selected as a binary system, as expected by the design of our approach (see Sect. 2.2). This contamination is zero in the case of *ruwe* > 1.4, due to the more restrictive criteria, and 20 in the case of *ruwe* > 1.25. Therefore, our approach provides a sample of unresolved binary systems with the best compromise between completeness and contamination when compared to other criteria to select this type of object.

Working with simulated data has the advantage that we can access all the true simulated parameters. In this case, we can study the population of binary systems depending on the evolutionary stages of their two components. We have identified five main evolutionary stages present in GUMS, i.e., Giant stars, WD, and high-, low- and mid-MS stars defined as stars with $M_G < 4$, $M_G > 10$ and in-between these cuts, respectively. Figure 8 shows the distribution of mass and light ratios for GUMS binaries within 200 pc - and the subset which are unresolved and have detectable excess *ruwe*. This shows the range of important (though often not observable) physical properties we might expect from a range of binary systems. MS stars, especially those that are cooler and of lower mass, follow a relatively tight distribution (akin to our previous heuristic of $l \propto q^{3.5}$), above this (and thus closer in light ratio for a given mass ratio) we see binaries containing stars the mass and temperature of our Sun and higher, which are significantly less tightly clustered but follow a similar trend. Systems containing Giants generally fall below these populations (as the Giant primary is effectively over-luminous for its mass). Systems containing WDs are quite distinct in this space - having generally much more mass compared to their luminosity

(except for WD-WD binaries which cluster very close to $q \approx 1$ and $l \approx 1$).

The strong dependence of light ratio on mass ratio leads to many unresolved binary systems being indistinguishable from single objects from their luminosity (and also in general from their spectra) as even moderate mass ratios, e.g. $q = 0.5$, lead to much smaller light ratios ($l < 0.1$). Thus the light ratio, though crucial for fully characterising binary systems, is rarely measurable.

Comparing the full population to the detectable we can get a feeling for which kinds of systems are easier or more difficult to pick out from astrometry alone. We can see a low detection fraction for systems with Δ close to zero (as the astrometric motion is negligible) - which includes both 'twin' systems ($q \sim l \sim 1$) as well as faint MS-MS and MS-Giant binaries (with $q \lesssim 0.1$). The detection fraction is also low for systems containing WDs where the second component is similarly or less luminous (top left of each panel, because these are dim enough that the astrometric error is high enough to dominate over astrometric motion).

Detectability functions with period and distance

Figure 9 shows the fraction of binaries that are detected as a function of period, expanding on Figure 15 of Penoyre et al. (2022a). This is a very strong 'window function' limiting us to only binaries for which we see a significant fraction of the orbit ($P \lesssim 10$ years, with the fraction dropping as P^{-2}) and those whose orbits are wide enough that the astrometric motion is significant compared to the astrometric error ($P \gtrsim 1$ month, with the fraction increasing proportional to $a \propto P^{\frac{3}{2}}$, see Penoyre et al. 2020, 2022a; Andrew et al. 2022, for more detail).

Even for a sample out to 200 pc (which will be dominated by the number of systems close to that distance limit), we detect more than 80% of systems at the optimal period of 34 months (*Gaia* DR3's time-baseline). For a closer limiting distance that fraction increases mostly due to the larger apparent magnitude and parallax, both leading to the astrometric signal being more dominant over astrometric noise.

Separating systems by the type of the primary we see essentially the same curve, with brighter primaries more easily detectable (again because of the reduced astrometric noise). Low mass primaries are significantly harder to detect (as the orbit is smaller for the same period) and so are WD primary systems (as these are necessarily low luminosity). The most massive primaries, especially the Giants, are even more strongly peaked in period, particularly at the low period end. This is mostly a general astronomical bias rather than an astrometric one, since these are both rarer (thus they don't appear in significant numbers at small distances) and smaller period orbits can be unresolvable at larger distances) and more limited in their minimum period by their large stellar radii and the requirement that the stars do not merge. Comparing instead the detectability based on the type of the secondary we now see that WD companions give significantly larger signals (because of their large q and low l , giving large Δ). In general, it is not a bad assumption that each of the curves shown here is a simple rescaling of the distribution across all stars, with the small caveat that it must saturate at a detection fraction of 1. Figure 10 shows a similar detectability function but now as a function of distance. All populations generally have lower detection fractions at increasing distances, where they are both smaller orbits on-sky and dimmer, thus the astrometric error can increase.

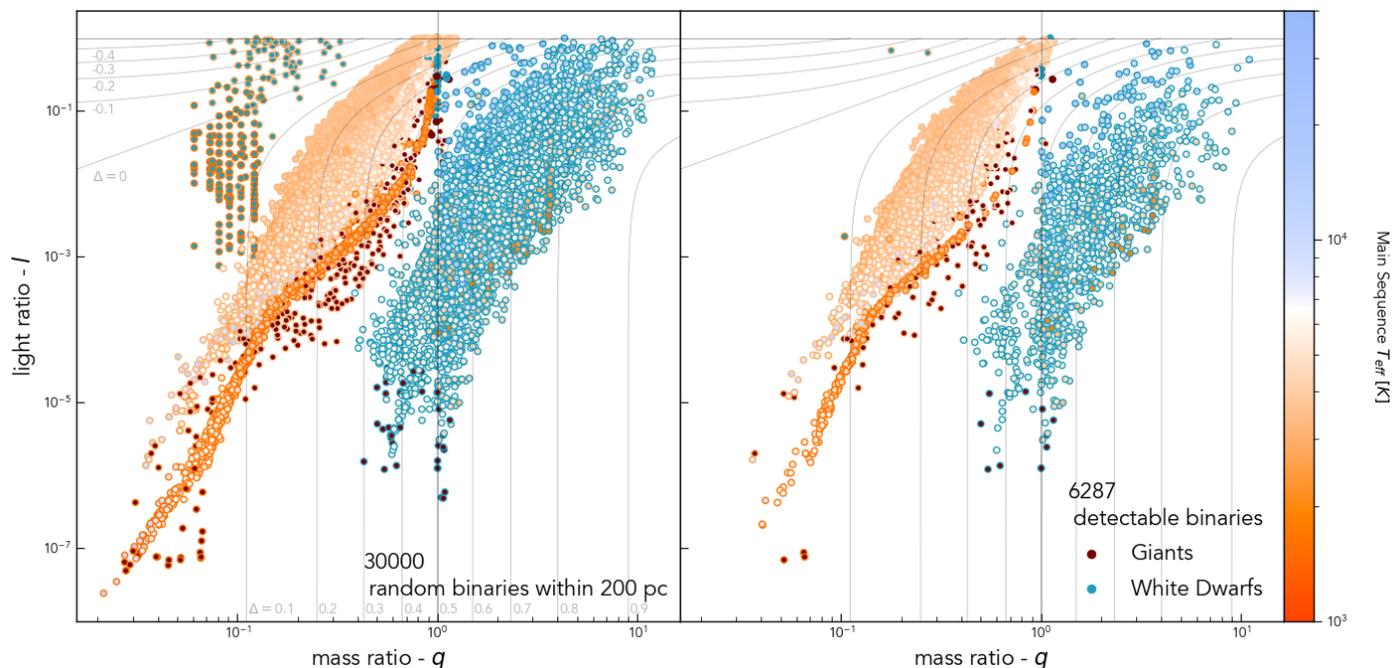


Fig. 8. Mass and light ratios of a random subset of 30,000 (of 242,485) binaries from GUMS within 200 pc. The central colour of each point shows the type of the primary, and the outline shows the type of the secondary. MS stars are given a colour that maps approximately to their visual colour based on their effective temperature. Giants ($M_G < 4$ and $G_{BP} - G_{RP} > 0.9$) are shown in dark red and WDs ($M_G - 3(G_{BP} - G_{RP}) < 9.5$) in blue. The left panel shows all random binaries and the right panel only those which are unresolved and whose $ruwe$ is greater than the threshold for their region of the sky. Because we choose the primary to be the brighter component ($l < 1$ always) the populations are effectively reflected through $l = 1$ and $q \rightarrow \frac{1}{q}$ - thus the population of systems containing WDs appears at both low and high q depending on the relative brightness of the other component. Lines of constant $\Delta = \frac{l-q}{(1+l)(1+q)}$ are shown - describing the relative offset of the centre of light from the CoM (with positive values corresponding to systems where the photocentre is closer to the secondary).

Evidently, both the period and distance functions apply to the data at once, so we start by showing the detectable distance fraction split by period. For ideal periods (close to the peak at 34 months), the detectable fraction is close to 1 at small distances and still near 80% at the edge of our sample - as shown in Fig. 6 this will only drop below 50% for sources ~ 1 kpc and further. Longer-period systems are harder to detect but follow the same kind of behaviour, whilst short-period systems drop off non-linearly as the population approaches the astrometric noise floor of *Gaia*. Separating by primary type we see that dim primaries (WDs and low-MS stars) also drop off quickly with distance for the same reason, and more massive brighter systems follow the general trend. Finally, separating by secondary type, we again see that WD companions are the easiest to detect, but the slope of the detection fraction with distance does not depend on the secondary type.

For all of these systems, the limiting factor with distance is the ratio of the magnitude of the astrometric motion to the astrometric error, which will necessarily decrease with distance. For populations for which no sources remain in the (approximately) constant regime of astrometric error ($G \gtrsim 13$), this ratio drops more rapidly and thus the curve steepens. We would expect this steepening to apply to all types of sources eventually, modulated by their inherent luminosity, and thus the brightest Giants may still be detectable as binaries at large (> 1 kpc) distances. This estimation does not account for extinction and crowding increasing with distance, both of which will cause extra astrometric noise and may lower the distance at which binary detection is achievable, especially in crowded and obscured regions such as the Galactic disc, Galactic centre and the Magellanic clouds.

5. Summary and conclusions

In this study, we have focused on the capability to detect unresolved binary systems in *Gaia* DR3, using $ruwe$ as a key metric. We developed a forward model for predicting $ruwe$ values, which depend on the sky coordinates and a set of astrophysical properties of the observed stellar system (single or binary). This model allowed us to determine the upper limit of $ruwe$ that will describe good *Gaia* measurements for single sources. Consequently, we established a $ruwe$ threshold that varies across different regions of the sky (Fig. 3), where $ruwe$ values exceeding this threshold will potentially indicate the presence of an unseen companion. Then, using the approach described in Castro-Ginard et al. (2023) to estimate the selection function of subsamples of *Gaia* data, we estimate the selection probability of a *Gaia* source with a $ruwe$ above our sky-dependent threshold. We note that this selection probability will trace areas where binary systems can be more, or less, easily detected across most of the sky, except in the very crowded regions. In these regions, Fig 4 will trace contamination when selecting binary systems solely based on $ruwe$, since this will be increased due to contributions other than binarity.

We simulate a broad range of binary systems to describe what types of systems, e.g., mass and luminosity ratios, period or distances, will lead to a $ruwe$ above the threshold and therefore make these systems detectable using this approach. The probability of detecting these systems decreases with increasing distance. Specifically, about 50% of the simulated binary systems were detectable at a distance of approximately 100 pc. However, for binary systems with orbital periods as long as the temporal baseline of the survey, this detectability range extends signifi-

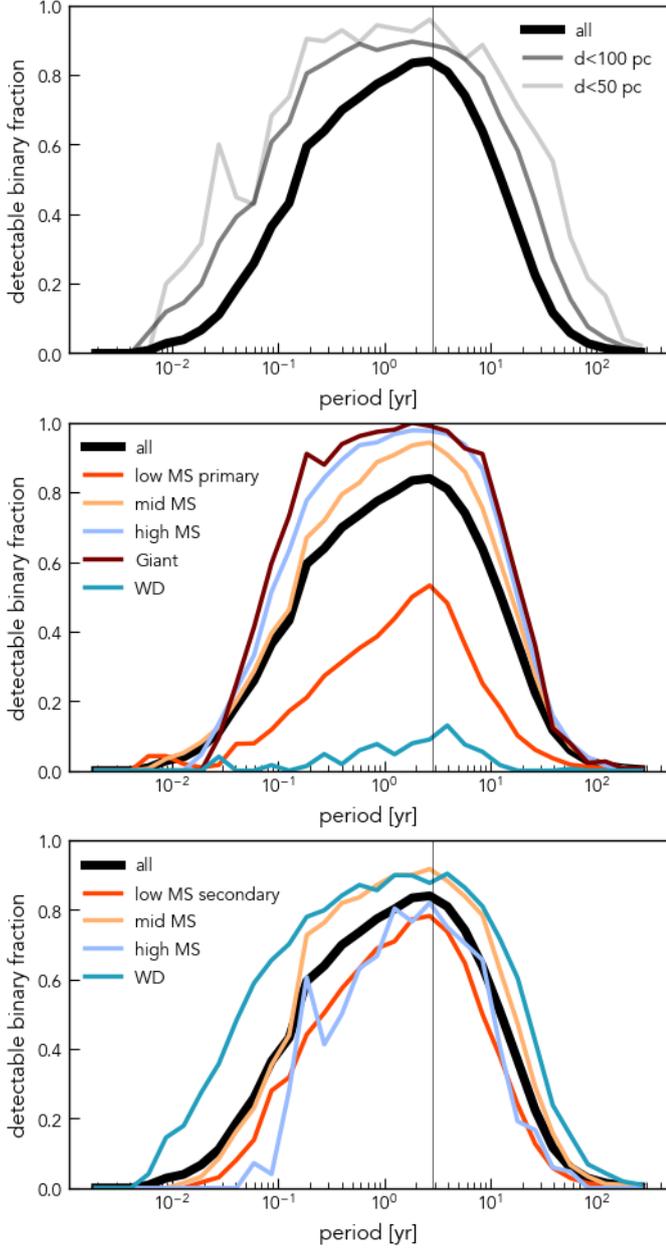


Fig. 9. The fraction of detectable binaries as a function of period, comparing the full population (within 200 pc) to subsets. Top panel: detectability within some distance range. Middle panel: detectability dependent on the type of the primary. Bottom panel: detectability dependent on the type of the secondary. The vertical line at 34 months shows the time baseline of *Gaia* DR3 (and is the optimal period for binary detection). MS stars are split into three subgroups - high ($M_G < 4$), low ($M_G > 10$) and mid (in-between these two cuts). Around 15% of MS primaries fall into the high bin, and the same for the low. Close to half of the secondaries are low, and the vast majority of the rest are mid (the number of high-mass MS secondaries is small, and Giant secondaries are negligible)

cantly, allowing for the detection of half of these systems at distances up to about 1 kpc. We also see that the imprints of binarity on *ruwe* are weaker for systems with orbital periods around 1 or 1/2 years (due to a confusion with a parallax shift, Penoyre et al. 2020) and mass ratios going to zero (single source) or one (no apparent wobbling of the photocentre around the CoM). Therefore, systems with these characteristics are less likely to be iden-

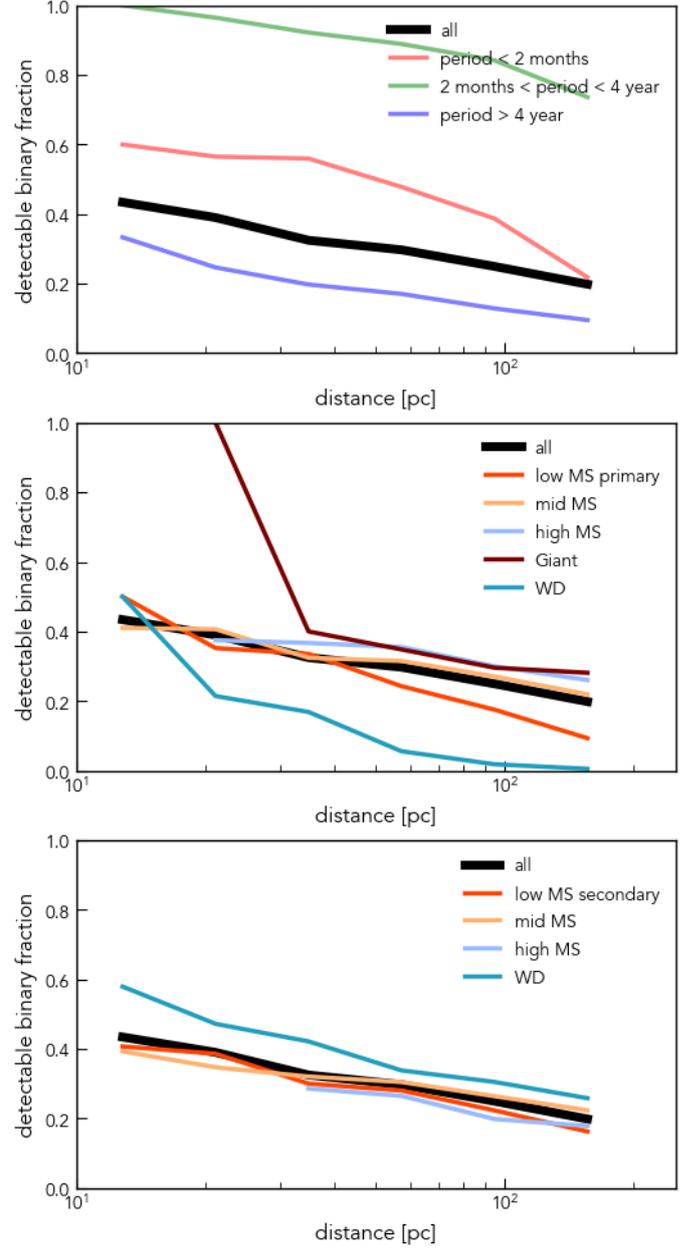


Fig. 10. Similar to Fig. 9 but separated by the distance of the binary. We show the function for all systems in the sample and for subsamples separated by binary period (top), type of primary (middle) and type of secondary (bottom).

tified as binaries in *Gaia* when relying exclusively on *ruwe* criteria.

Finally, we use the GUMS simulation within 200 pc and apply our forward model for *ruwe*, to provide insight into the unresolved binary population in the GCNS. Our sky-varying *ruwe* criteria for identifying potential binary systems outperforms the widely used threshold value of 1.4, both in terms of the completeness and the purity of the resulting sample. In the GUMS sample, we find that the most numerous binary systems, when both components are MS stars, are not the easiest to detect in *Gaia*. The stronger imprints in *ruwe*, thus higher detectability, are generated by binary systems with a luminous primary component (high-MS, mid-MS or Giant star). On the other hand, binary systems with a dim primary component (WD or low-MS

stars) will have a low detection probability. For instance, only about 1% of WD-WD binary systems simulated in GUMS were detected.

This study provides a new definition of what imprints on *ruwe* we can expect from binarity, and it is adapted to the systems' location on-sky. However, most of the existing binary systems in our Galaxy will still go unnoticed due to the CoM wobbling around the photocentre being dominated by the astrometric noise of *Gaia*, therefore resulting in a *ruwe* value below the threshold. The tools we have developed provide a valuable resource to detect binary system candidates in *Gaia* data and characterise the properties of the population from which the selected systems are drawn. Following the example of previous *GaiaUnlimited* selection functions, we make the code and tools generated in this work publicly available through the *GaiaUnlimited* Python package.

Acknowledgements. We thank David W. Hogg for his contributions to the *GaiaUnlimited* project. We thank Douglas Boubert, Andrew Everall and Tereza Jerabkova for their help in preparing the GUMS binary catalogue. This work is a result of the *GaiaUnlimited* project, which has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 101004110. The *GaiaUnlimited* project was started at the 2019 Santa Barbara *Gaia* Sprint, hosted by the Kavli Institute for Theoretical Physics at the University of California, Santa Barbara. ZP acknowledges that this project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101002511 - VEGA P). This work has made use of results from the European Space Agency (ESA) space mission *Gaia*, the data from which were processed by the *Gaia Data Processing and Analysis Consortium* (DPAC). Funding for the DPAC has been provided by national institutions, in particular, the institutions participating in the *Gaia* Multilateral Agreement. The *Gaia* mission website is <http://www.cosmos.esa.int/gaia>. The authors are current or past members of the ESA *Gaia* mission team and of the *Gaia* DPAC. This work made use of *Astropy*⁷ a community-developed core Python package and an ecosystem of tools and resources for astronomy (Astropy Collaboration et al. 2022).

References

- Andrew, S., Penoyre, Z., Belokurov, V., Evans, N. W., & Oh, S. 2022, *MNRAS*, 516, 3661
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167
- Belokurov, V., Penoyre, Z., Oh, S., et al. 2020, *MNRAS*, 496, 1922
- Bressan, A., Marigo, P., Girardi, L., et al. 2012, *MNRAS*, 427, 127
- Cantat-Gaudin, T., Fouesneau, M., Rix, H.-W., et al. 2023, *A&A*, 669, A55
- Cantat-Gaudin, T., Fouesneau, M., Rix, H.-W., et al. 2024, arXiv e-prints, arXiv:2401.05023
- Castro-Ginard, A., Brown, A. G. A., Kostrzewa-Rutkowska, Z., et al. 2023, *A&A*, 677, A37
- Duchêne, G. & Kraus, A. 2013, *ARA&A*, 51, 269
- Duchêne, G., Lacour, S., Moraux, E., Goodwin, S., & Bouvier, J. 2018, *MNRAS*, 478, 1825
- El-Badry, K. 2024, arXiv e-prints, arXiv:2403.12146
- El-Badry, K., Rix, H.-W., & Heintz, T. M. 2021, *MNRAS*, 506, 2269
- Fernández-Hernández, J., Joliet, E., & Garralda, N. 2022, *GOST Software User Manual*, https://gaia.esac.esa.int/gost/docs/gost_software_user_manual.pdf
- Fitton, S., Tofflemire, B. M., & Kraus, A. L. 2022, *Research Notes of the American Astronomical Society*, 6, 18
- Gaia* Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, 616, A1
- Gaia* Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021a, *A&A*, 649, A1
- Gaia* Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, 595, A1
- Gaia* Collaboration, Smart, R. L., Sarro, L. M., et al. 2021b, *A&A*, 649, A6
- Gaia* Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, *A&A*, 674, A1
- Halbwachs, J.-L., Pourbaix, D., Arenou, F., et al. 2023, *A&A*, 674, A9
- Kervella, P., Arenou, F., & Thévenin, F. 2022, *A&A*, 657, A7
- Lindgren, L. 2018, *gAIA-C3-TN-LU-LL-124*
- Lindgren, L. 2022, *gAIA-C3-TN-LU-LL-136*
- Lindgren, L., Bastian, U., Biermann, M., et al. 2021a, *A&A*, 649, A4
- Lindgren, L., Klioner, S. A., Hernández, J., et al. 2021b, *A&A*, 649, A2
- Luri, X., Palmer, M., Arenou, F., et al. 2014, *A&A*, 566, A119
- Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *AJ*, 154, 94
- Penoyre, Z., Belokurov, V., & Evans, N. W. 2022a, *MNRAS*, 513, 2437
- Penoyre, Z., Belokurov, V., & Evans, N. W. 2022b, *MNRAS*, 513, 5270
- Penoyre, Z., Belokurov, V., Wyn Evans, N., Everall, A., & Koposov, S. E. 2020, *MNRAS*, 495, 321
- Price-Whelan, A. M., Hogg, D. W., Rix, H.-W., et al. 2020, *ApJ*, 895, 2
- Raghavan, D., McAlister, H. A., Henry, T. J., et al. 2010, *ApJS*, 190, 1
- Robin, A. C., Luri, X., Reylé, C., et al. 2012, *A&A*, 543, A100
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, 131, 1163
- Stassun, K. G. & Torres, G. 2021, *ApJ*, 907, L33
- Wallace, A. L. 2024, *MNRAS*, 527, 8718
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, 140, 1868

⁷ <http://www.astropy.org>

Appendix A: Using the GaiaUnlimited python package

We provide the code we used to estimate the selection function of binary systems as a new tool in the GaiaUnlimited python package⁸. The necessary data, i.e., all-sky scanning law details at a resolution of HEALPix level 5, to generate the plots in the paper are also included within the package.

Listing 1 shows how to use the new BinarySystemsSelectionFunction class. Similarly to the previous GaiaUnlimited selection functions, it takes an astropy coordinate object and returns the queried quantity, e.g., sky-varying ruwe threshold (to produce Fig 3) or selection probabilities of sources with high ruwe (computed as in Castro-Ginard et al. 2023, produces Fig. 4) or the *Gaia* scanning law details.

```

1 from gaiaunlimited.selectionfunctions import binaries
2
3 SF = binaries.BinarySystemsSelectionFunction()
4
5 from astropy import units as u
6 from astropy.coordinates import SkyCoord
7
8 coord = SkyCoord(ra = ra*u.degree, dec = dec*u.degree, frame = 'icrs')
9
10 # Generates the selection probabilities based on RUWE.
11 # Function to generate Fig. 4.
12 probability, variance = SF.query(coord, return_variance = True)
13
14 # Returns the RUWE threshold above which a source is considered a potential binary system.
15 # Function to generate Fig. 3.
16 ruwe = SF.query_RUWE(coord, crowding = True)

```

Listing 1. Example of code to initialise and query the selection function of binary systems.

We also provide our forward model to compute a ruwe value from the binary system properties, and the *Gaia* scanning law in the SimulateGaiaSource class. Listing 2 shows how to initialise the simulation of a system at a given sky location, providing the period, eccentricity and initial phase of the system, as well as the *Gaia* scanning law details (observation times, scanning angles and parallax factors, provided by GOST). The AL measurements and errors are computed as a function of the other system properties. Finally, a ruwe value is computed from these simulated observations. Full documentation of the package is available through the GaiaUnlimited documentation website⁹.

```

1 from gaiaunlimited.utils import SimulateGaiaSource
2
3 source = SimulateGaiaSource(ra, dec, period, eccentricity, initial_phase)
4
5 al_pos, al_err = source.observe(Gmag,
6                               parallax,
7                               semimajor_axis,
8                               mass_ratio,
9                               luminosity_ratio,
10                              phi,
11                              theta,
12                              omega,)
13
14 ruwe = source.unit_weight_error(al_pos, al_err)

```

Listing 2. Example of code to simulate a *Gaia* observation as a function of the source parameters (single or binary), and estimate its ruwe.

⁸ <https://github.com/gaia-unlimited/gaiaunlimited>

⁹ <https://gaiaunlimited.readthedocs.io/en/latest/>