

Distinguishing exoplanet companions from field stars in direct imaging using Gaia astrometry

Philipp Herz, Matthias Samland, Coryn A.L. Bailer-Jones*

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

Submitted 6 Nov. 2023, revised 29 Nov. 2023, accepted 7 Dec. 2023

ABSTRACT

Direct imaging searches for exoplanets around stars detect many spurious candidates that are in fact background field stars. To help distinguish these from genuine companions, multi-epoch astrometry can be used to identify a common proper motion with the host star. Although this is frequently done, many approaches lack an appropriate model for the motions of the background population, or do not use a statistical framework to properly quantify the results. Here we use Gaia astrometry combined with 2MASS photometry to model the parallax and proper motion distributions of field stars around exoplanet host stars as a function of candidate magnitude. We develop a likelihood-based method that compares the positions of a candidate at multiple epochs with the positions expected under both this field star model and a co-moving companion model. Our method propagates the covariances in the Gaia astrometry and the candidate positions. True companions are assumed to have long periods compared to the observational baseline, so we currently neglect orbital motion. We apply our method to a sample of 23 host stars with 263 candidates identified in the B-Star Exoplanet Abundance Study (BEAST) survey on VLT/SPHERE. We identify seven candidates in which the odds ratio favours the co-moving companion model by a factor of 100 or more. Most of these detections are based on only two or three epochs separated by less than three years, so further epochs should be obtained to reassess the companion probabilities. Our method is publicly available as an open-source python package from GitHub to use with any data.

Key words. Exoplanets - Astrometric motion - Direct imaging - open source software

1. Introduction

Young exoplanets with a favourable brightness and separation to their host star can be directly imaged. However, such exoplanets can be confused with more-distant background stars that happen to lie in the line-of-sight. A common way to distinguish these scenarios is to observe both the host star and candidate over time to look for a common proper motion and/or parallax. An example of doing this is shown in Fig. 1 (data from Janson et al. 2021 and published in Squicciarini et al. 2022). The positions of the host star and a number of exoplanet candidates were measured at two epochs, 2018 and 2021. The orange crosses show the measured change in position (relative to the host star) of candidates between these two epochs. The host star has some proper motion and parallax between 2018 and 2021: the black dashed line shows how objects with zero proper motion and parallax would move over this time period (as our view is centred on the host star). Those orange points clustered near the black point labelled 2021 are therefore consistent with being distant background stars. The candidate labelled “b”, on the other hand, has a motion more consistent with the host star and so is more likely to be a true companion.

To make this procedure quantitative we must take into account the measurement uncertainties and any covariance between them. Background stars do not have zero parallax and proper motion, so we need a proper model for their motions too. There may also be more than two epochs, so we will want to take into account all of the data within a single assessment. In the literature these confounding factors are usually not considered

(e.g. Lagrange et al. 2010; Lafrenière et al. 2011; Carson et al. 2013; Kuzuhara et al. 2013; De Rosa et al. 2015; Konopacky et al. 2016; Chauvin et al. 2017; Keppler et al. 2018; Bohn et al. 2021; Janson et al. 2021; Squicciarini et al. 2022; Franson et al. 2023; Mesa et al. 2023; Chomez et al. 2023). Such studies often compare the second epoch position with that expected for a stationary background star. If they deviate significantly the null hypothesis of being a background star is rejected and some alternative model, i.e. a companion, is implicitly accepted. This classical hypothesis testing approach does not, however, assess whether the data might be even more unlikely under the companion model.

The goal of this paper is to put this astrometric confirmation approach on a solid statistical footing. We develop a model to evaluate whether the multi-epoch motion of the candidate is more likely to be a co-moving exoplanet or a coincidental field star. Our model is based on the proper motion and parallax distributions of field stars in the same area of the sky as the candidate, and with similar magnitudes to the candidate being tested. Using the odds ratio, we compare this background model with a model in which the candidate is co-moving with the host star. We apply our method to multi-epoch measurements of candidates observed in the B-Star Exoplanet Abundance Study (BEAST) survey (Janson et al. 2021), which has an abundance of candidate objects due to observing close to the Galactic plane. A python package implementing our method is accessible via GitHub¹.

* email: calj@mpia.de

¹ github.com/herzphi/compass

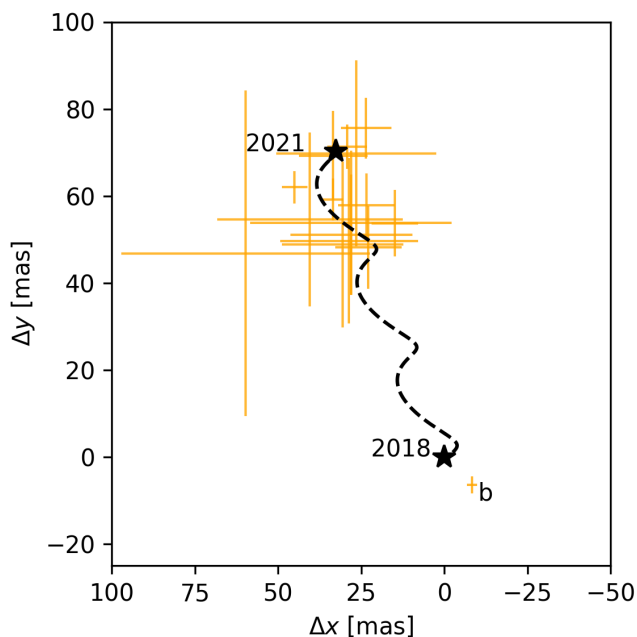


Fig. 1: Change in position of exoplanet candidates (orange crosses) relative to the star μ^2 Sco between two measurement epochs. A co-moving source should be close to the origin (labeled "2018"). A background source with zero proper motion will move according to the dashed curve (a reflection of the host star's parallax and proper motion) ending in the black star labeled "2021". The motion of μ^2 Sco b is distinct from the cloud of background stars in the field that are (through this plot) deemed not to be exoplanets. Figure adapted from Squicciarini et al. (2022).

2. Methods

We develop a probabilistic method that compares the likelihoods of position measurements of an exoplanet candidate under two models: the first assumes the candidate is a co-moving companion, the second assumes it is a field star. The former uses the proper motion and parallax of the exoplanet candidate's host star, the latter additionally uses a magnitude-dependent fit to the parallax and proper motion distributions of field stars close to the host star's line-of-sight. Our method can use an arbitrary number of epochs of astrometric observations to derive the likelihoods. We also implement a special case of the general method that neglects parallax. This is less realistic but is easier to visualize. The field star astrometric model is currently built on data from the Gaia (Gaia Collaboration et al. 2023) and 2MASS (Skrutskie et al. 2006) surveys, but could use other present or future survey.

2.1. Positional model

Consider the position of a candidate object relative to the host star as a function of time. Assuming that the candidate does not move relative to the star, then in a geocentric coordinate system this can be described as a linear motion with a superimposed parallactic motion. In a Cartesian plane projection we write this as

$$\begin{aligned}\Delta x'_i &= (x'_{a,0} - x'_{\star,0}) + (\mu_{a,x} - \mu_{\star,x}) t_i + (\varpi_a - \varpi_{\star}) s_x(t_i) \\ \Delta y'_i &= (y'_{a,0} - y'_{\star,0}) + (\mu_{a,y} - \mu_{\star,y}) t_i + (\varpi_a - \varpi_{\star}) s_y(t_i)\end{aligned}\quad (2.1)$$

where μ is the proper motion and ϖ is the parallax. Subscript a refers to the candidate object, subscript \star to the host star, subscript 0 to the true position of the candidate at a reference epoch, and subscript i to the i -th epoch relative to the reference epoch. The functions s_x and s_y are periodic phase factors for the parallax motion given by the orbital motion of the Earth around the Sun at epoch i . The primed variables denote the true position of the candidate, while the unprimed variables denote the measured position of the candidate. In this paper "true positions" denote those we would obtain in the absence of noise in our positional measurements at the epochs, but still based on some externally-provided values of parallax and proper motion. Those external values *are* considered to be noisy, and their covariances will be taken into account.

Equation 2.1 can be generalized to N measurements of the candidate's position over time $(\Delta \mathbf{x}, \Delta \mathbf{y}) = \{\Delta x_i, \Delta y_i\}$ at times $t = \{t_i\}$ with $i \in [1, N]$, where the bold font indicates a vector of measurements. The first position measurement defines the true position of the candidate in the first epoch $(\Delta x_1, \Delta y_1) = (\Delta x'_1, \Delta y'_1)$ without loss of generality.

2.2. Overall probabilistic model

Given the measured positions of the exoplanet candidate over time, we compute the likelihood of the data under two models. The first model, denoted M_c , assumes the exoplanet candidate is a co-moving companion, sharing the same proper motion and parallax as the host star. We make the simplifying assumption that the orbital period is long compared to the observational baseline and so neglect the candidate's orbital motion. This may not be justified in all cases, which we elaborate on in Section 4. The second model, denoted M_b , assumes the candidate to be a background object² with a proper motion and parallax distribution constructed from a set of background stars in a narrow field-of-view around the host star (described in Section 2.3). We then compare these two models via the odds (likelihood) ratio

$$r_{c,b} = \frac{P(\Delta \mathbf{x}, \Delta \mathbf{y} | M_c)}{P(\Delta \mathbf{x}, \Delta \mathbf{y} | M_b)} \quad (2.2)$$

which indicates which model favours the data more. This does not give the posterior probability of the model given the data, however, for which we would first need to establish model prior probabilities.

Let us first consider the denominator in Eq. 2.2, the likelihood under the background model. This can be computed as a combination of two probability density functions (PDFs). The first of these is the probability of the noisy position measurements given the true positions, $P(\Delta \mathbf{x}, \Delta \mathbf{y} | \Delta \mathbf{x}', \Delta \mathbf{y}')$, which reflects the noise in the determination of the centroid of a point-spread-function on the detector. We assume this to be a Gaussian with mean $(\Delta \mathbf{x}', \Delta \mathbf{y}')$ and a covariance matrix Λ that reflects the accuracy of, and correlations in, the measurements. The second term is $P(\Delta \mathbf{x}', \Delta \mathbf{y}' | t, M_b)$. This represents the spread in possible "true" positions of a background star at specific times arising from the spread in the proper motion and parallax of the background star population. We assume this term to be a Gaussian distribution with a mean given by the right side of Eq. 2.1. The terms in that equation, as well as the covariance matrix of the Gaussian (which we denote Λ'), come from our fit to the background star population (Section 2.3). These two PDFs we then

² We use the term "background objects" and "field stars" interchangeably, as most field stars are more distant than the host star.

combine via a marginalization to give the required likelihood

$$P(\Delta\mathbf{x}, \Delta\mathbf{y} | \mathbf{t}, M_b) = \int_{\Delta\mathbf{x}', \Delta\mathbf{y}'} P(\Delta\mathbf{x}, \Delta\mathbf{y} | \Delta\mathbf{x}', \Delta\mathbf{y}') \times P(\Delta\mathbf{x}', \Delta\mathbf{y}' | \mathbf{t}, M_b) d\Delta\mathbf{x}' d\Delta\mathbf{y}'. \quad (2.3)$$

This shows that the background model likelihood is a convolution of the two PDFs. Given that these are both Gaussian, the result of the convolution is also a Gaussian, with mean $(\Delta\mathbf{x}', \Delta\mathbf{y}')$ and covariance matrix $(\Lambda + \Lambda')$. More details are provided in Appendix A.

We take the same approach to computing the likelihood of the data under the co-moving companion model M_c , the numerator in Eq. 2.2. Here, however, the relative position of the exoplanet candidate is assumed to be constant, meaning that the relative position measurements have zero variance and so $P(\Delta\mathbf{x}', \Delta\mathbf{y}' | \mathbf{t}, M_c)$ is simply a delta function. The convolution in Eq. 2.3 is then trivial, resulting in $P(\Delta\mathbf{x}, \Delta\mathbf{y} | \mathbf{t}, M_c)$ being a Gaussian with mean $(\Delta\mathbf{x}', \Delta\mathbf{y}')$ and covariance matrix Λ .

Now that we have expressions for both likelihoods in terms of the measured relative positions and their covariances, and in terms of the parallax and proper motion distribution of the background model, we can compute the odds ratio in Eq. 2.2 and decide which model better explains the data.

The above general method we refer to as the proper motion and parallax covariance method. If the parallax is negligible, or if the observations are separated by almost exactly a year, then we can set the parallax in Eq. 2.1 to zero. If we also have only two epochs, then instead of using the two position measurements directly, we can convert them into a single proper motion. We call this the "proper motion-only" method. In this special case, the "true" proper motion $\Delta\boldsymbol{\mu}' = (\Delta\mu'_x, \Delta\mu'_y)$ relative to the host star can be written

$$\Delta\mu'_x = \frac{\Delta x'_2 - \Delta x'_1}{t_2 - t_1} \quad (2.4)$$

which we compare to the measured relative proper motions from the data (same equation without the accents). The likelihoods for the two models are then

$$P(\Delta\boldsymbol{\mu} | M_c) = \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \text{Var}(\Delta\mu_x) & \text{Cov}(\Delta\mu_x, \Delta\mu_y) \\ \text{Cov}(\Delta\mu_x, \Delta\mu_y) & \text{Var}(\Delta\mu_y) \end{bmatrix}\right) \quad (2.5)$$

and

$$P(\Delta\boldsymbol{\mu} | M_b) = \int P(\Delta\boldsymbol{\mu} | \Delta\boldsymbol{\mu}') P(\Delta\boldsymbol{\mu}' | M_b) d\Delta\boldsymbol{\mu}'. \quad (2.6)$$

Both equations are Gaussian distributions. In the co-moving companion model the likelihood is centered at zero with a covariance matrix given by the uncertainties in position measurements. In the background model the likelihood is centered at the proper motion $\Delta\boldsymbol{\mu}'$ of the background star distribution relative to the host star with a covariance matrix given by the convolution of the proper motions of this background star population with the measurement uncertainties.

A more complete derivation of both the full model and this simpler (but less representative) proper motion-only model is given in Appendix A.³

³ More details can be found at http://mpia.de/homes/calj/astrometric_companion_model.pdf.




Fig. 2: K_S -band magnitude distribution of stellar objects in Gaia DR3 in the area of the sky within 0.3° of the 23 BEAST stars. Those in red have measured 2MASS K_S -band photometry. Those in black stripes have K_S predicted by a colour transformation.

2.3. Proper motion and parallax distribution model for field stars

In order to assess whether a candidate is astrometrically consistent with a given field star population (i.e. to evaluate the likelihood $P(\Delta\mathbf{x}, \Delta\mathbf{y} | \mathbf{t}, M_b)$), we have to create a model for the parallax and proper motion distributions of the population. Ideally this model would be a function of the true distance and velocity of the star, but we of course do not have this information for an arbitrary candidate. However, we know that distance and velocity – and therefore parallax and proper motion – depend on the measurable direction in the Galaxy (e.g. Bailer-Jones 2023). We can also condition our model on other relevant measurements, the most obvious being magnitude, as this contains some information about distance and stellar population. Ideally we would also use colour (as a crude proxy of mass and age and thus having velocity dependence), but this is often not available for many exoplanet surveys.

We build our model empirically using Gaia DR3 (Gaia Collaboration et al. 2023). We select stars that are near to the candidate in Galactic coordinates and that have similar brightness. The latter can be hard to achieve because Gaia observes in the optical and is not as deep as infrared exoplanet surveys. To address this we use a positional cross-match of Gaia with 2MASS (Skrutskie et al. 2006) to assign infrared magnitudes to Gaia stars where possible. For those Gaia objects that do not have a 2MASS counterpart, we assign synthetic K_S magnitudes using the Gaia colour–colour transformation based on the G -band magnitude and $G_{BP} - G_{RP}$ colour from Riello et al. (2021). While this transformation is not perfect, it significantly increases our field star sample size. Figure 2 shows in red all bright field stars covered by 2MASS within 0.3° of the 23 target stars of the BEAST. The distribution of stars for which we calculated synthetic K_S -Band magnitudes via the colour transformations are shown with black stripes.

Using our sample, we then build a smooth model of the astrometric distributions as a function of magnitude. This will allow us to evaluate (if necessary by extrapolation) the astrometric model at the candidate's brightness. To build the model we first group the data into magnitude bins, then fit a two-dimensional Gaussian distribution to the proper motions and parallaxes in each bin. An example of one magnitude bin is shown in Fig. 3. We then fit a linear function to the mean values of the 2D Gaus-




Fig. 3: Proper motion distribution of stellar objects with K_S -magnitude between 18 and 19 in a 0.3° sky area around HIP 82545 (μ^2 Sco). The elliptical contours are the boundaries that encompass 50%, 90% and 99% of the stellar objects.

sian as a function of magnitude (Fig. 4), excluding points that lie outside the 10th to 90th percentile range of the magnitudes to achieve a more robust fit. This fit smooths out the variations and allows us to extrapolate the model to fainter candidates than are in Gaia. The standard deviations as a function of magnitude are fit as an exponential

$$\sigma_{\text{fit}}(m) = \sigma_{\text{min}} + a \exp(-b(m - m_0)) \quad (2.7)$$

with $a > 0$. This fit has a minimum value of σ_{min} , a fixed constant. The magnitude m_0 is the mean of the observed magnitudes. An example of this fit is shown in Fig. 5. A linear fit is used instead if the exponential fit residuals are larger than those of a linear fit. In making this linear fit we prevent $\sigma_{\text{fit}}(m)$ from becoming smaller than 1 mas/yr.

We did not identify any definitive and consistent trend in the correlation between parallax and proper motion. We therefore opted to fit a constant value instead.

2.4. Verification of the method using simulations

A verification of our model would ideally be based on a large set of real data in which the true nature of each candidate is known, but this is not available. Here we use simulations to assess the reliability of the odds ratio. We simulate 1000 trajectories of a co-moving companion and 1000 trajectories of background model objects.

Starting with the real star μ^2 Sco, we construct a background model for proper motion and parallax. The initial position of the candidate relative to μ^2 Sco is randomly generated, and the position at subsequent epochs (one year apart) is given by

$$x_{i+1} = x_i + (\mu_x + \mathcal{N}(0, \sigma = 3\text{mas/yr})) \Delta t \quad (2.8)$$

and similarly for y . This equation represents either a co-moving companion with a zero (relative) proper motion ($\mu_x = \mu_y = 0$),




Fig. 4: The variation of the mean of the proper motion distribution in field stars as function on stellar magnitude. Each bin includes 200 stellar objects in a 0.3° sky area around HIP 82545 (μ^2 Sco). Each point is the mean value of the 2D Gaussian fit in each magnitude bin. Only those points lying within the 10th–90th percentile range of magnitudes are used for the linear fit.

or a relative proper motion of the background model to μ^2 Sco ($\mu_x = \mu_{RA}(m) - \mu_{x,\mu^2 \text{ Sco}}$, where $\mu_{RA}(m)$ is a fixed value defined by the brightness of the candidate), both with a normal distribution \mathcal{N} to simulate the propagation of the proper motion uncertainty. Here we choose a standard deviation of 3mas/yr. We simulate 1000 paths the candidate could take under the companion model M_c using Eq.2.8 with $\mu = 0$ and another 1000 paths under the background object model M_b using a single fixed proper motion a background object with a similar brightness of the candidate would have. The spread seen in Fig.6 thus arises from the noise we add at each epoch. Our method distinguishes between those paths via a logarithmic odds ratios greater than zero (companion model favoured, shown in blue) and smaller than zero (background model favoured, shown in black). We find that all 2000 trajectories are correctly identified as coming from the model from which they were generated.

3. Results

The B-star Exoplanet Abundance Study (Janson et al. 2021) employs the VLT/SPHERE Beuzit et al. (2019) extreme adaptive optics instrument to conduct a direct imaging survey of 85 B-type stars in the young (5–20 Myr) and relatively nearby (120–150 pc) Scorpius-Centaurus (Sco-Cen) region (Fig. 7). The targets have similar distance and G -Band magnitude, but vary in colour. By targeting young B-type stars, their exoplanets should be relatively bright in the near-infrared. Before this survey, B-stars had not been systematically surveyed for exoplanets. Radial velocity surveys, which are more sensitive to close-in planets, reveal comparatively few planets around massive stars (Reffert et al. 2015). Thus BEAST addresses the question of whether massive stars can form massive planets at larger separations.

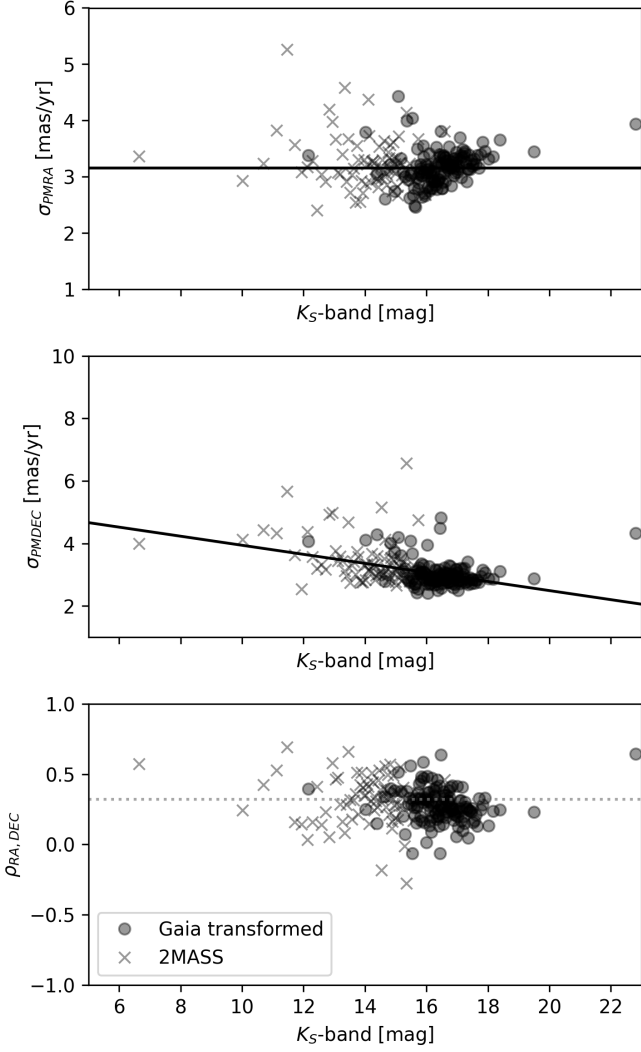


Fig. 5: The variation of the standard deviation and correlation of the proper motion distributions for the field stars around HIP 82545 (μ^2 Sco) as a function of magnitude. Each point comes from a 2D Gaussian fit over a narrow magnitude bin (e.g. as shown in Fig. 3). Each bin includes 200 stars in a 0.3° sky area around the target star. Crosses denote fits using stars with 2MASS magnitudes, and circles denote those with magnitudes computed from the Gaia colour transformation.

Certain formation scenarios, like disk instability, might primarily take place in the outermost areas of extensive disks surrounding massive stars (Helled et al. 2014). Sco-Cen is located close to the Galactic plane, so observations often include many spurious candidates in the field-of-view. This makes the survey a prime target for testing our method of distinguishing bound companions from faint field stars.

Of the 85 stars in BEAST, 23 have candidates with at least two epochs of observation (as of 2022-04-01). These 23 stars have in total 263 candidates. The projected separations between the candidates and their host stars range from 49 au to 1457 au (Fig. 8).

We apply both our proper motion-only model and our proper motion and parallax covariance model to all 263 candidates. The results are shown in Table 1, sorted by descending logarithmic




Fig. 6: Simulated test of our method by propagating the positions of 2000 simulated candidates to μ^2 Sco over four epochs. Half are propagated according to the co-moving model, which just adds zero mean noise at each epoch. The other half are propagated according to the proper motion of background stars, plus noise. The colours denote the logarithmic odds ratio that our method computes for each trajectory. All 2000 are correctly identified.




Fig. 7: Colour-magnitude diagram of all BEAST stars, except for η Cen at $(B_P - R_P, G) = (4.75, 2.25)$ mag.




Fig. 8: Distribution of the separations of all candidate from all targets identified in BEAST. This assumes the candidates are at the same distance as their respective target stars.

Table 1: The 20 candidates with logarithmic odds ratio greater than zero under the parallax and proper motion model, sorted by decreasing $\log_{10} r_{cb}(\mu, \varpi)$ (Eq. 2.2). Those above 1.0 we identify as likely companions (in bold face). The column "Candidate ID" contains unique identifier strings assigned to each candidate by Janson et al. (2021) in most cases, and/or the commonly accepted companion name ("B" referring to a very probable stellar-mass companion). The full table including results for 263 candidates (including six candidates published in Viswanath et al. 2023 that did not have follow-up when our list was compiled) is available online.

Host Star	Host Star	Candidate ID	K_s -band (mag)	N_{obs}	Δt_{obs} (years)	$\log_{10} r_{cb}(\mu)$	$\log_{10} r_{cb}(\mu, \varpi)$
HIP 71865	b Cen	b (1oh0dfaj)	17.17	3	20.9	15.78	≥ 300
HIP 61257	HD 109195	"B" (58pwu3dv)	12.81	3	2.8	1.81	≥ 300
HIP 81208	HD 149274	B	13.63	2	2.7	-2.71	91.66
HIP 82545	μ^2 Sco	b (snojiu5b)	16.08	2	3.1	8.00	51.36
HIP 81208	HD 149274	C	12.56	2	2.7	2.29	42.58
HIP 52742	HD 93563	x9ld1uh0	12.03	2	1.1	1.57	12.56
HIP 76048	HD 138221	vslvj1zp	14.50	2	0.2	1.42	2.00
HIP 60009	ζ Cru	rtolbhl8	20.49	2	1.0	0.07	0.32
HIP 62327	HD 110956	nlikloif	19.64	2	1.0	0.08	0.30
HIP 69011	HD 123247	962kt2vw	18.16	2	0.2	0.81	0.26
HIP 60009	ζ Cru	ebmx222k	19.97	2	1.0	0.05	0.18
HIP 74100	HD 133937	ldj61i8l	19.97	2	1.0	0.08	0.18
HIP 52742	HD 93563	mm9uw3ha	19.49	2	1.1	0.01	0.16
HIP 60009	ζ Cru	qhsttxkw	20.14	2	1.0	0.03	0.15
HIP 60009	ζ Cru	n41zi1pt	19.51	2	1.0	0.04	0.15
HIP 60009	ζ Cru	4lh805ve	18.85	2	1.0	0.06	0.14
HIP 60009	ζ Cru	a7214wz3	19.39	2	1.0	0.09	0.12
HIP 60009	ζ Cru	5o2rkr75	20.11	2	1.0	0.05	0.10
HIP 60009	ζ Cru	d9oco7ww	19.01	2	1.0	0.17	0.07
HIP 60009	ζ Cru	gtartoxk	19.23	2	1.0	0.12	0.06




Fig. 9: Proper motion model of the background stars for μ^2 Sco for the right ascension direction (top) and the declination direction (bottom) in the ICRS reference system. This is similar to Fig.4, but now includes the BEAST candidates shown in orange and transformed to ICRS.

odds ratio $\log_{10} r_{cb}(\mu, \varpi)$, i.e. the most likely real companions

appear first. In what follows we focus on the results using the parallax and proper motion model in the final column. Visualizations of all 20 candidates are shown in Fig.13a, Fig.13b, Figs. B.1, B.2 and B.3. The complete table of candidate astrometry and photometry from the BEAST survey will be published in Delorme et al. (in preparation).

Most interesting are the seven candidates that have an odds ratio greater than ten (i.e. $\log_{10} r_{cb}(\mu, \varpi) > 1$). Two of these, b Cen b and μ^2 Sco b, are well-established as exoplanets (Janson et al. 2021; Squicciarini et al. 2022). We further infer that the candidates of HIP 81208 (two candidates), HIP 61257, HIP 52742, and HIP 76048 are likely true companions (but not necessarily in the exoplanet mass regime). The remaining candidates in Table 1 are still formally favoured by the co-moving companion model, but not by much. These will require more epochs or other observations in order to determine their nature.

We now examine the results on some of the individual, high odds ratio candidates, which will also serve to illustrate how our method works.

3.1. Results on individual candidates

μ^2 Sco b The astrometric field star model for the target μ^2 Sco is shown in Fig. 9. The black line is our fit to the background stars (described in Section 2.3), which can be compared to the various BEAST candidates in this field in orange. Nearly all of the candidates agree with the field star model, showing the large degree of field star contamination that can be present in exoplanet searches. Just one does not agree; this is the exoplanet μ^2 Sco b. The PDF over the proper motion (i.e. the likelihood) for this candidate under the background model is shown by the black curves/contours in Fig. 10. The likelihood for the compan-

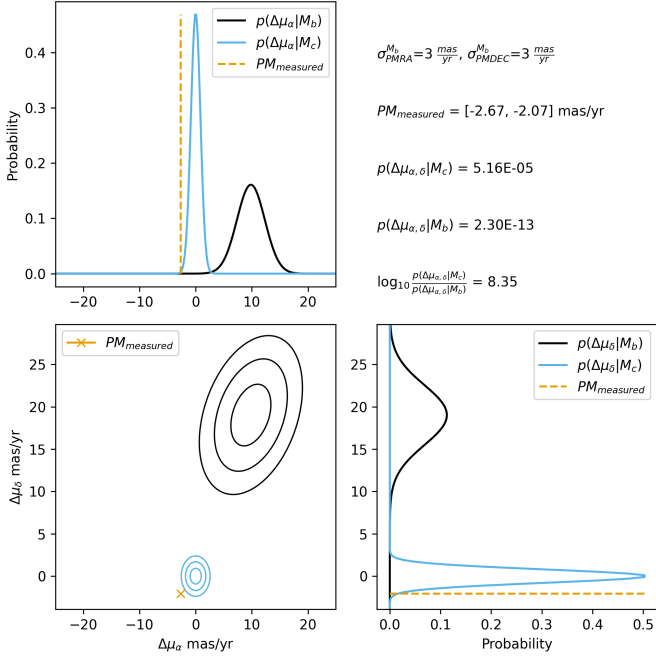


Fig. 10: Demonstration of the proper motion-only method for the known exoplanet μ^2 Sco b. The likelihood under the background model is shown in black and the likelihood under the co-moving companion model is shown in blue. The bottom left panel shows these as two-dimensional (Gaussian) distributions. The other two panels show the one-dimensional marginal distributions. The measured relative proper motion is shown in orange: the uncertainties in this is not shown because it is included in the two likelihoods (see Eq. 2.5 and 2.6).

ion model is shown by the blue curves/contours, and the measured proper motion is shown as the orange point/line. In this example we see that the measurement is far more consistent with the companion model than with the background model. This is a rather clear-cut case even by visual inspection; many other cases in appendix B are more ambiguous.

b Cen b This candidate, which was identified as an exoplanet by Janson et al. (2021), has three epochs in our analysis. Figure 11 shows the change in positions of the candidate under the background model M_b (black), the (unchanged) position of the candidate under the co-moving companion model M_c (blue), and the measured positions of the candidate (orange). The first observation of this candidate was at the epoch 2000.4 (circle) based on archival images by Shatsky & Tokovinin (2002). In the relative reference frame to the host star, a co-moving companion would not move from its 2000.4 position in our model, and so is not shown for other epochs. The next observation of the candidate took place in 2019.2 with BEAST. Based on the background proper motion and parallax distribution, as well as the host star's proper motion, a background object in the same area of the sky as the host star would have moved about 850 mas since 2000.4. However, the figure shows that by 2019.2 the candidate has moved much less (orange point). Another epoch at 2021.3 confirms that the motion of the candidate is much closer to the host star's motion than it is to that of the background population. The background star hypothesis is clearly ruled out in this case. But because the measured positions lie near the 99% con-

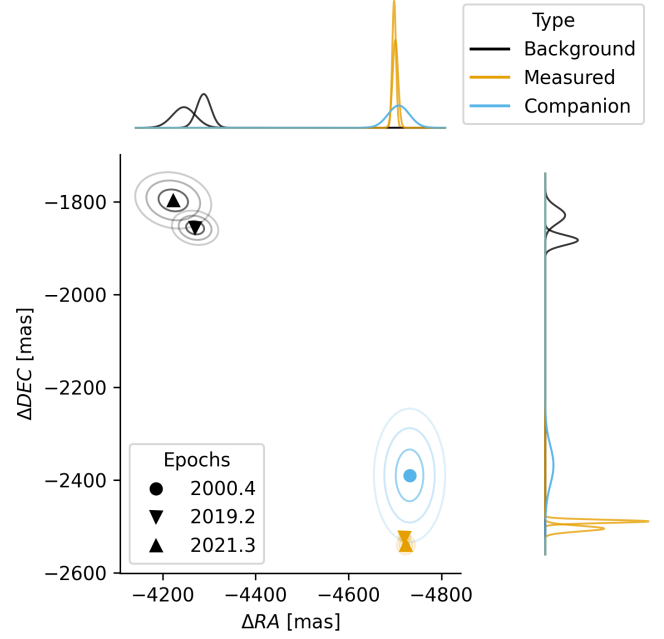


Fig. 11: Visualization of the predicted positions of the candidate companion b of the star HIP 71865 (b Cen) under the proper motion and parallax model. A co-moving companion would remain at the position of the first epoch (blue circle) because orbital motion is not included. A field star with the modelled proper motion and parallax of nearby (mostly background) stars would be measured at the two later epochs at the two positions shown by the black triangles. The actual measured change in positions of the candidate are shown as orange triangles. The fact that these are much nearer to the blue distribution means this is likely to be a true companion, something that is properly quantified by our method. The contour lines show 50%, 90%, and 99% of the enclosed probability, reflecting the propagated uncertainty in the parallaxes, proper motions, and BEAST position measurements. The marginal likelihoods are shown on both axes. This visualization does not show the covariances between the measurements at different epochs, which are nonetheless taken into account by our method (see Eq. A.4).

tour of our co-moving companion model, some residual motion compared to the co-moving companion model likely remains. This can be explained by orbital motion, which is not taken into account, over the 21-year observational baseline (Janson et al. 2021).

HIP 61257 "B" This highly probable candidate of HIP 61257 can be further examined by including non-BEAST data. Kouwenhoven et al. (2005) discuss a potential companion around HIP 61257, but they ultimately classify it as a background object. This potential companion has a K_S -band magnitude of 12.43 mag and a separation of 5540 mas, which coincides with the magnitude and separation of one of our candidates with logarithmic odds ratio greater than zero in both modelling frameworks. The astrometric motion plot (Fig. 12) shows that this candidate is co-moving with the host star. As the astrometric data from Kouwenhoven et al. (2005) are reported without separation and position angle uncertainties, we adopted an uncertainty of 10mas in the relative right ascension and declination. This com-

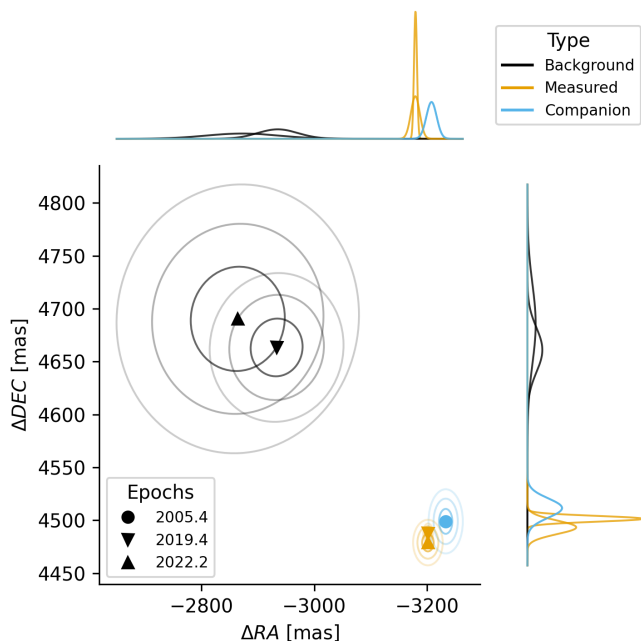


Fig. 12: The results of our model applied to the position measurements of HIP 61257 "B" from Kouwenhoven et al. (2005) and Janson et al. (2021) over a 17-year baseline. See the caption to Figure 11 for a description.

panion was also discussed by Gratton et al. (2023), who identified it to be a low-mass star ($0.083 \pm 0.01 M_{\odot}$) based on its Gaia and K-band magnitude. The logarithmic odds ratios without the archival data are $\log r_{cb}(\mu, \varpi) = 12.91$ and $\log r_{cb}(\mu) = 2.45$. We therefore confirm this candidate as a bona-fide binary companion based on its astrometry.

HIP 81208 B and C These two candidates are identified as co-moving companions by Viswanath et al. (2023). They identify HIP 81208 B as a 67 MJup object, mostly likely a brown dwarf, and HIP 81208 C as a $0.135 M_{\odot}$ low-mass star. Our results for these are shown in Figs. 13a and 13b (using astrometry and photometry from Table C.1 in Viswanath et al. 2023). Both candidates are favoured by the co-moving companion model with our more sophisticated parallax and proper motion method. The proper motion only model rejects HIP 81208 B but accepts HIP 81208 C. This system was recently analyzed further and found to be a gravitationally-bound hierarchical quadruple system comprised of low mass objects with a newly discovered companion to the C component (Cb, Chomez et al. 2023).

HIP 52742 x91d1uh0 and mm9uw3ha Two candidates with magnitudes of $K_S = 12.03$ mag and $K_S = 19.49$ mag are formally favoured by our companion model. The brighter one we identify as a companion with high significance (log odds ratio 12.6). Gratton et al. (2023) also identified this as a companion: Adopting an age of 82.5 Myr (obtained from assuming membership of a comoving group outside Sco-Cen; Janson et al. 2021) they estimated a mass of $0.51 M_{\odot}$ at a projected separation of 176 au. In a search for astrometric acceleration from a comparison of Hipparcos and Gaia DR3 proper motions, Brandt (2021) found marginally significant evidence for an acceleration of this star, which might be evidence of this companion. The fainter

candidate companion we identify to this star has a log odds ratio of just 0.16, which is not significant.

HIP 76048 vs1vj1zp We identify one potential companion to this star with an odds ratio of 100 in our analysis, albeit with a very short baseline of just 0.2 years. Brandt (2021) found no significant evidence of acceleration from their Hipparcos–Gaia DR3 proper motion study.

4. Discussion

4.1. Interpretation

Our analysis favours the co-moving companion model for five candidates from the BEAST data set accessible for this study, as well as two candidates in the HIP 81208 system that have a second epoch from Viswanath et al. (2023). Of these seven candidates, two are confirmed exoplanets: b Cen AB b (Janson et al. 2021) and μ^2 Sco b (Squicciarini et al. 2022). HIP 61257 "B" is very likely a stellar-mass companion (Gratton et al. 2023). The two candidates of HIP 81208 are discussed as being a brown dwarf and a low-mass stellar companion by Viswanath et al. (2023). The remaining two companions of the host stars HIP 52742 and HIP 76048 – with a measurement baseline of only 0.2 years – are unconfirmed at the time of writing.

Most of the targets in our analysis have only two observation epochs with small temporal baselines (of order one year), so many candidates show little motion relative to the host star between the epochs. It is just these cases where ad hoc approaches to assessing companionship are inconclusive, and our statistical framework is most useful. For many of these short baselines the co-moving model is not favoured. Longer temporal baselines make it easier to distinguish between the models, especially as the targets of direct imaging campaigns tend to be nearby stars with large proper motion.

4.2. Model assumptions

Our model does not currently take into account orbital motion, so the co-moving companion model may not be favoured even if the candidate is a true companion with significant orbital motion over the observational baseline. Orbital periods for directly-imaged planets tend to be long, on the order of centuries or even millennia. In those cases, neglecting the orbital motion will often not affect our analysis. In contrast, planets with relatively short orbits like β Pic b and c (23.6 and 3.3 years respectively, Lacour et al. 2021) can show significant orbital motion.

If orbital motion can be directly observed then the candidate is very likely to be confirmed as a companion (e.g. Marois et al. 2008). This would be a relatively straight forward extension to our modelling approach by including a model with path curvature. Plausible paths (priors) could be generated from a Keplerian model (e.g. using Blunt et al. 2020) and then marginalizing over them to determine the model likelihood. An alternative approach is to show that a candidate's motion is consistent with that of field stars while also showing that the range of orbits that can explain the data are beyond the escape velocity of the system (Nielsen et al. 2017; Wagner et al. 2022).

We build our field star astrometric model using Gaia parallaxes and proper motions. On account of the limited depth of Gaia, we then have to extrapolate our model to the fainter magnitudes of our candidates. This can introduce biases, and thus incorrectly favour or disfavour the background model.

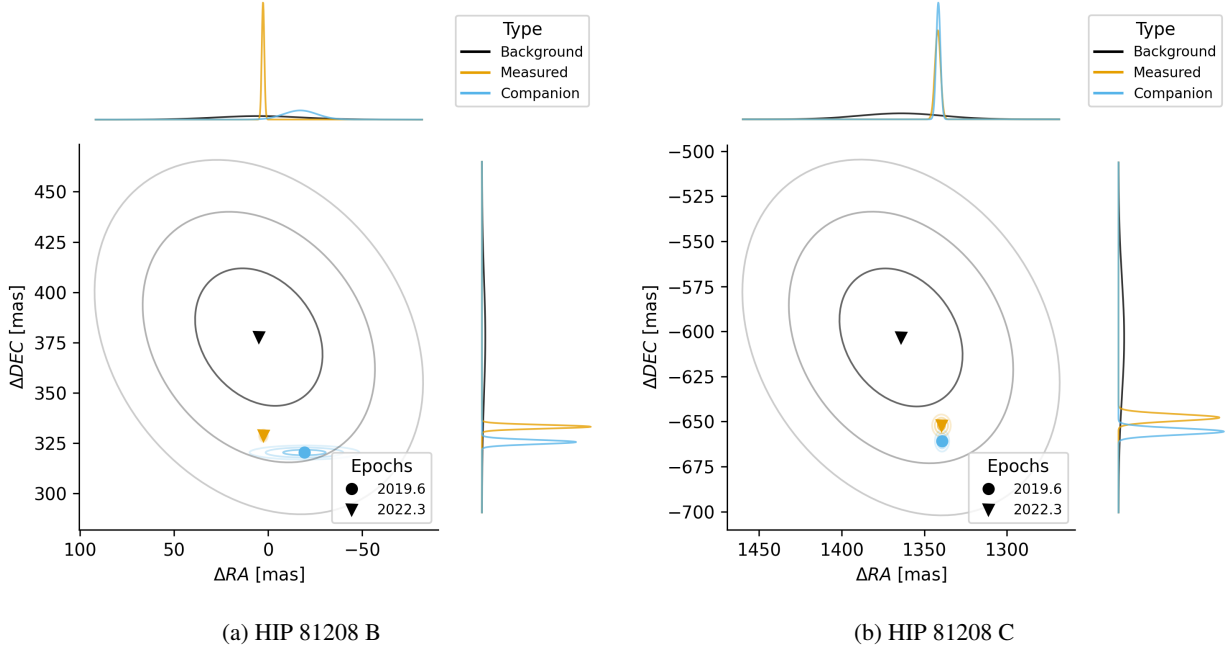


Fig. 13: Astrometric motion of the two candidates of HIP 81208 with an observation baseline of three years. Viswanath et al. (2023) report both objects as co-moving companions based on their proper motion analysis. Our analysis supports this claim.




Fig. 14: Difference between the K_S magnitude calculated from Gaia via colour transformations and the measured 2MASS K_S magnitudes for 934 023 objects from Gaia with 2MASS counterparts in the vicinity of μ^2 Sco.

Another drawback of Gaia is that it observes in the optical, whereas direct imaging surveys for planets are currently done mostly in the near-infrared. To obtain the necessary infrared magnitudes of the Gaia sources we had to crossmatch to an infrared survey. We chose 2MASS because it is all-sky. However, it is not very deep, so for many field stars we instead had to predict the K_S -band magnitude using colour transformations from Gaia. We used the transformations of Riello et al. (2021) for objects in the colour range $-0.5 < G_{BP} - G_{RP} < 2.5$. The accuracy of this for brighter stars around μ^2 Sco is demonstrated in Fig. 14. The 2MASS objects only cover a K_S -band range from 10 to 18 mag while the transformed photometry extends to 21 mag. Some ex-

oplanet surveys are also conducted at shorter wavelengths, such as J and H , for which colour transformations are more reliable as they are closer to the observed Gaia bands. Ideal for our purposes, of course, would be an infrared version of Gaia.

As explained in section 2.3, in our background model the parallax and proper motion distributions depend only on direction and magnitude. A further improvement would be to add dependence on the colour, if that is provided by the imaging survey. We could then also use the measured colour to infer something about the intrinsic properties of the object, e.g. its spectral type (e.g. Parviainen et al. 2019). While this may help determine whether or not it is a low-mass object, it would not tell us whether it is gravitationally bound.

In computing the odds ratio we only consider kinematics. We do not take into account the number density of background stars or the angular separation between the candidate and target star. Yet for a given separation, the less dense the background, the more likely the candidate is a genuine companion (e.g. Tamura 2016; Squicciarini et al. 2022). This information could be incorporated as an additional multiplicative odds ratio, although it requires a model (or measurement) of the stellar density at the faint magnitudes at which we conduct our exoplanet survey.

Finally, we emphasise that we report our results as a ratio of likelihoods of two models, where each likelihood is the probability of the data given the model. To convert each likelihood into a posterior probability of the model given the data, we would need to adopt a prior probability for the model, and know that our models are exhaustive. The latter is not yet the case, as we have neglected orbital motion, for example. The prior could incorporate the direction-dependent variation of the background star number density.

5. Summary

This work has introduced a statistical method that uses multi-epoch astrometry of an imaged exoplanet candidate to compare

a co-moving companion model with a chance-aligned field star model. It puts what is commonly referred to as the “common proper motion test” on a probabilistic footing.

Our statistical model enables a quantitative analysis of an arbitrary number of epochs, a task that cannot be achieved effectively through visual inspection. We consider the proper motion and parallax of the host star and the candidate and evaluate the likelihoods under two different models for the candidate: one in which it is a co-moving companion with negligible orbital motion, the other in which it is a member of the field star population. For the latter we build a probabilistic model of the distribution of the proper motions and parallaxes of field stars as a function of magnitude, using a fit to Gaia data in the field of each target star.

We applied our method to a sample of 263 candidates around 23 stars from the B-Star Exoplanet Abundance Study (BEAST). We first developed a purely proper motion based method, which we then extended to take into account the parallax. This model accommodates the covariance in the astrometry both between the measurements and across multiple epochs, for both Gaia astrometry and the direct measurements. We identify seven candidates as co-moving companions. Five of these have been identified as real companions in the literature, including the two exoplanets μ^2 Sco b and b Cen(AB) b. The remaining two candidates are priority targets for further investigation.

Our modelling approach is publicly available as an open-source Python package on GitHub, allowing for easy evaluation and visualization of existing and new data. While this work presents an improvement over current practices, there is scope for further improvement. Of particular importance is the inclusion of exoplanet orbital motion in the companion model, the incorporation of stellar number densities, and discriminating field stars from exoplanets based on their spectral information.

Acknowledgements. We are grateful for the helpful comments during the refereeing process which improve the paper. The observations this study is based on were acquired at the ESO VLT telescope (program 1101.C-025). This publication makes use of data from the European Space Agency (ESA) mission Gaia (<http://www.cosmos.esa.int/gaia>), processed by the Gaia Data Processing and Analysis Consortium (DPAC, <http://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 Multi-lateral Agreement. This publication also makes use of data from the Two Micron All Sky Survey, which is a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center/California Institute of Technology, funded by the National Aeronautics and Space Administration and the National Science Foundation.

References

- Bailer-Jones, C. A. L. 2023, *AJ*, 166, 269
 Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, 631, A155
 Blunt, S., Wang, J. J., Angelo, I., et al. 2020, *AJ*, 159, 89
 Bohn, A. J., Ginski, C., Kenworthy, M. A., et al. 2021, *A&A*, 648, A73
 Brandt, T. D. 2021, *ApJS*, 254, 42
 Carson, J., Thalmann, C., Janson, M., et al. 2013, *ApJ*, 763, L32
 Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017, *A&A*, 605, L9
 Chomez, A., Squicciarini, V., Lagrange, A. M., et al. 2023, *A&A*, 676, L10
 De Rosa, R. J., Nielsen, E. L., Blunt, S. C., et al. 2015, *ApJ*, 814, L3
 Franon, K., Bowler, B. P., Zhou, Y., et al. 2023, *ApJ*, 950, L19
 Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, *A&A*, 674, A1
 Gratton, R., Squicciarini, V., Nascimbeni, V., et al. 2023, *A&A*, 678, A93
 Helled, R., Bodenheimer, P., Podolak, M., et al. 2014, in *Protostars and Planets VI*, ed. H. Beuther, R. S. Klessen, C. P. Dullemond, & T. Henning, 643–665
 Janson, M., Gratton, R., Rodet, L., et al. 2021, *Nature*, 600, 231
 Janson, M., Squicciarini, V., Delorme, P., et al. 2021, *Astronomy and Astrophysics*, 646, A164
 Keppler, M., Benisty, M., Müller, A., et al. 2018, *A&A*, 617, A44
 Konopacky, Q. M., Rameau, J., Duchêne, G., et al. 2016, *ApJ*, 829, L4

- Kouwenhoven, M. B. N., Brown, A. G. A., Zinnecker, H., Kaper, L., & Portegies Zwart, S. F. 2005, *A&A*, 430, 137
 Kuzuhara, M., Tamura, M., Kudo, T., et al. 2013, *ApJ*, 774, 11
 Lacour, S., Wang, J. J., Rodet, L., et al. 2021, *A&A*, 654, L2
 Lafrenière, D., Jayawardhana, R., Janson, M., et al. 2011, *ApJ*, 730, 42
 Lagrange, A. M., Bonnefoy, M., Chauvin, G., et al. 2010, *Science*, 329, 57
 Marois, C., Macintosh, B., Barman, T., et al. 2008, *Science*, 322, 1348
 Mesa, D., Gratton, R., Kervella, P., et al. 2023, *A&A*, 672, A93
 Nielsen, E. L., Rosa, R. J. D., Rameau, J., et al. 2017, *AJ*, 154, 218
 Parviainen, H., Tingley, B., Deeg, H. J., et al. 2019, *Astronomy & Astrophysics*, 630, A89
 Reffert, S., Bergmann, C., Quirrenbach, A., Trifonov, T., & Küstler, A. 2015, *A&A*, 574, A116
 Riello, M., Angeli, F. D., Evans, D. W., et al. 2021, *Astronomy and Astrophysics*, 649, A3
 Shatsky, N. & Tokovinin, A. 2002, *A&A*, 382, 92
 Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, 131, 1163
 Squicciarini, V., Gratton, R., Janson, M., et al. 2022, *Astronomy and Astrophysics*, 664, A9
 Tamura, M. 2016, *Proceedings of the Japan Academy, Series B*, 92, 45
 Viswanath, G., Janson, M., Gratton, R., et al. 2023, *A&A*, 675, A54
 Wagner, K., Apai, D., Kasper, M., McClure, M., & Robberto, M. 2022, *AJ*, 163, 80

Appendix A: Mathematical description of the multi-epoch probabilistic model

We describe here in more detail the probabilistic model introduced in section 2.2. In high-contrast imaging, the set of positions of candidates are generally determined relative to the host star, i.e.

$$\begin{aligned}\Delta\mathbf{x}' &= \mathbf{x}'_a - \mathbf{x}'_\star \\ \Delta\mathbf{y}' &= \mathbf{y}'_a - \mathbf{y}'_\star\end{aligned}\quad (\text{A.1})$$

where the subscript a refers to the candidate and \star to the host star. The accents on the coordinates \mathbf{x}'_a and \mathbf{x}'_\star indicate that they are “true”, i.e. without positional measurement uncertainties. The noisy, observed versions are \mathbf{x}_a and \mathbf{x}_\star at the (noise-free) times t . The times we use in our models are measured relative to the Gaia epoch of observation of the host star, because we use Gaia data to build our proper motion and parallax distribution models for the field stars.

As the true positions are nonetheless predictions based on noisy parallaxes and proper motions (and propagated in time using Eq. 2.1), we write the probability density functions of these true positions as

$$P(\Delta\mathbf{x}', \Delta\mathbf{y}' \mid t, M(\varpi, \mu)) \quad (\text{A.2})$$

where M can be M_c or M_b and the PDF can be written as Gaussians with the variances and covariances between the two positions $\Delta\mathbf{x}'$ and $\Delta\mathbf{y}'$. For the companion model, these (co)variances are derived from the Gaia data for the host star. For the background model, these (co)variances come from our fit to the distribution of the Gaia parallaxes and proper motions of a set of field stars, as described in section 2.3. The likelihoods for both of these model (which appear in the odds ratio of Eq. 2.2) are obtained by marginalizing over the unknown true positions ($\Delta\mathbf{x}', \Delta\mathbf{y}'$) predicted by the model, as follows

$$\begin{aligned}P(\Delta\mathbf{x}, \Delta\mathbf{y} \mid t, M(\varpi, \mu)) \\ &= \int P(\Delta\mathbf{x}, \Delta\mathbf{y}, \Delta\mathbf{x}', \Delta\mathbf{y}' \mid t, M(\varpi, \mu)) d\Delta\mathbf{x}' d\Delta\mathbf{y}' \\ &= \int P(\Delta\mathbf{x}, \Delta\mathbf{y} \mid \Delta\mathbf{x}', \Delta\mathbf{y}') \times \\ &\quad P(\Delta\mathbf{x}', \Delta\mathbf{y}' \mid t, M(\varpi, \mu)) d\Delta\mathbf{x}' d\Delta\mathbf{y}' .\end{aligned}\quad (\text{A.3})$$

$P(\Delta\mathbf{x}, \Delta\mathbf{y} \mid \Delta\mathbf{x}', \Delta\mathbf{y}')$ is the likelihood of the measured positions given the true positions, and so reflects the noise in the determination of the image centroid. We take this to be a Gaussian with mean $(\Delta\mathbf{x}', \Delta\mathbf{y}')$ and $2N$ -dimensional Gaussian Λ , where N is the number of measurement epochs. Equation A.3 is a convolution of two Gaussians, which results in another Gaussian.

In the *background model* M_b , $P(\Delta\mathbf{x}', \Delta\mathbf{y}' \mid t_i, M_b)$ is a Gaussian with mean $(\Delta\mathbf{x}', \Delta\mathbf{y}')$ and covariance matrix

$$\Lambda' = \begin{bmatrix} \text{Var}(\Delta x'_1) & \text{Cov}(\Delta x'_1, \Delta y'_1) & \text{Cov}(\Delta x'_1, \Delta x'_2) & \text{Cov}(\Delta x'_1, \Delta y'_2) & \dots \\ & \text{Var}(\Delta y'_1) & \text{Cov}(\Delta y'_1, \Delta x'_2) & \text{Cov}(\Delta y'_1, \Delta y'_2) & \dots \\ & & \text{Var}(\Delta x'_2) & \text{Cov}(\Delta x'_2, \Delta y'_2) & \dots \\ & & & \text{Var}(\Delta y'_2) & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad (\text{A.4})$$

where terms are given in the order $(\Delta x'_1, \Delta y'_1, \Delta x'_2, \Delta y'_2, \dots)$. The symmetry of the covariance matrix fills the lower triangle of

Eq. A.4. The individual terms of Λ' are computed from Eq. 2.1 to be

$$\begin{aligned}\text{Var}(\Delta x'_i) &= t_i^2 [\text{Var}(\mu_{b,x}) + \text{Var}(\mu_{\star,x})] \\ &\quad + s_x(t_i)^2 [\text{Var}(\varpi_b) + \text{Var}(\varpi_\star)] \\ &\quad + 2t_i s_x(t_i) [\text{Cov}(\varpi_b, \mu_{b,x}) + \text{Cov}(\varpi_\star, \mu_{\star,x})] \\ \text{Var}(\Delta y'_i) &= t_i^2 [\text{Var}(\mu_{b,y}) + \text{Var}(\mu_{\star,y})] \\ &\quad + s_y(t_i)^2 [\text{Var}(\varpi_b) + \text{Var}(\varpi_\star)] \\ &\quad + 2t_i s_y(t_i) [\text{Cov}(\varpi_b, \mu_{b,y}) + \text{Cov}(\varpi_\star, \mu_{\star,y})] \\ \text{Cov}(\Delta x'_i, \Delta y'_i) &= t_i^2 [\text{Cov}(\mu_{b,x}, \mu_{b,y}) + \text{Cov}(\mu_{\star,x}, \mu_{\star,y})] \\ &\quad + s_x(t_i) s_y(t_i) [\text{Var}(\varpi_b) + \text{Var}(\varpi_\star)] \\ &\quad + t_i s_y(t_i) [\text{Cov}(\varpi_b, \mu_{b,x}) + \text{Cov}(\varpi_\star, \mu_{\star,x})] \\ &\quad + t_i s_x(t_i) [\text{Cov}(\varpi_b, \mu_{b,y}) + \text{Cov}(\varpi_\star, \mu_{\star,y})] \\ \text{Cov}(\Delta x'_i, \Delta x'_j) &= t_i t_j [\text{Var}(\mu_{b,x}) + \text{Var}(\mu_{\star,x})] \\ &\quad + t_i s_x(t_j) [\text{Cov}(\varpi_b, \mu_{b,x}) + \text{Cov}(\varpi_\star, \mu_{\star,x})] \\ &\quad + t_j s_x(t_i) [\text{Cov}(\varpi_b, \mu_{b,x}) + \text{Cov}(\varpi_\star, \mu_{\star,x})] \\ &\quad + s_x(t_i) s_x(t_j) [\text{Var}(\varpi_b) + \text{Var}(\varpi_\star)] \\ \text{Cov}(\Delta y'_i, \Delta y'_j) &= t_i t_j [\text{Var}(\mu_{b,y}) + \text{Var}(\mu_{\star,y})] \\ &\quad + t_i s_y(t_j) [\text{Cov}(\varpi_b, \mu_{b,y}) + \text{Cov}(\varpi_\star, \mu_{\star,y})] \\ &\quad + t_j s_y(t_i) [\text{Cov}(\varpi_b, \mu_{b,y}) + \text{Cov}(\varpi_\star, \mu_{\star,y})] \\ &\quad + s_y(t_i) s_y(t_j) [\text{Var}(\varpi_b) + \text{Var}(\varpi_\star)] \\ \text{Cov}(\Delta x'_i, \Delta y'_j) &= t_i t_j [\text{Cov}(\mu_{b,x}, \mu_{b,y}) + \text{Cov}(\mu_{\star,x}, \mu_{\star,y})] \\ &\quad + t_i s_y(t_j) [\text{Cov}(\varpi_b, \mu_{b,x}) + \text{Cov}(\varpi_\star, \mu_{\star,x})] \\ &\quad + t_j s_x(t_i) [\text{Cov}(\varpi_b, \mu_{b,y}) + \text{Cov}(\varpi_\star, \mu_{\star,y})] \\ &\quad + s_x(t_i) s_y(t_j) [\text{Var}(\varpi_b) + \text{Var}(\varpi_\star)] .\end{aligned}\quad (\text{A.5})$$

From the convolution in Eq. A.3, we see that $P(\Delta\mathbf{x}, \Delta\mathbf{y} \mid t, M_b)$ is a $2N$ -dimensional Gaussian with mean $(\Delta\mathbf{x}', \Delta\mathbf{y}')$ and covariance matrix $\Lambda + \Lambda'$. Note that equations A.5 assume that there is no covariance between the field star proper motion and parallax and the host star proper motion and parallax. This is not strictly true, because both the background model and the host star data are drawn from Gaia, and Gaia shows correlations between sources separated by small angles on the sky. But in practice our background model describes the distribution of an ensemble of stars, and so is less affected by Gaia’s correlations between individual sources.

In the *co-moving companion model* M_c , the position of the companion relative to the star is constant, leading to a zero covariance for the second term under the integral in Eq. A.3 making $P(\Delta\mathbf{x}', \Delta\mathbf{y}' \mid t, M_c)$ a delta function. $P(\Delta\mathbf{x}, \Delta\mathbf{y} \mid t, M_c)$ is therefore a $2N$ -dimensional Gaussian with mean $(\Delta\mathbf{x}', \Delta\mathbf{y}')$ and covariance Λ .

Appendix B: Candidates with positive log odds ratios

Figures B.1, B.2, and B.3 show the modelling results for the 18 BEAST-only candidates that have logarithmic odds ratio $\log r_{icb} > 0$. The results for the two targets with additional data are shown in Figure 13.

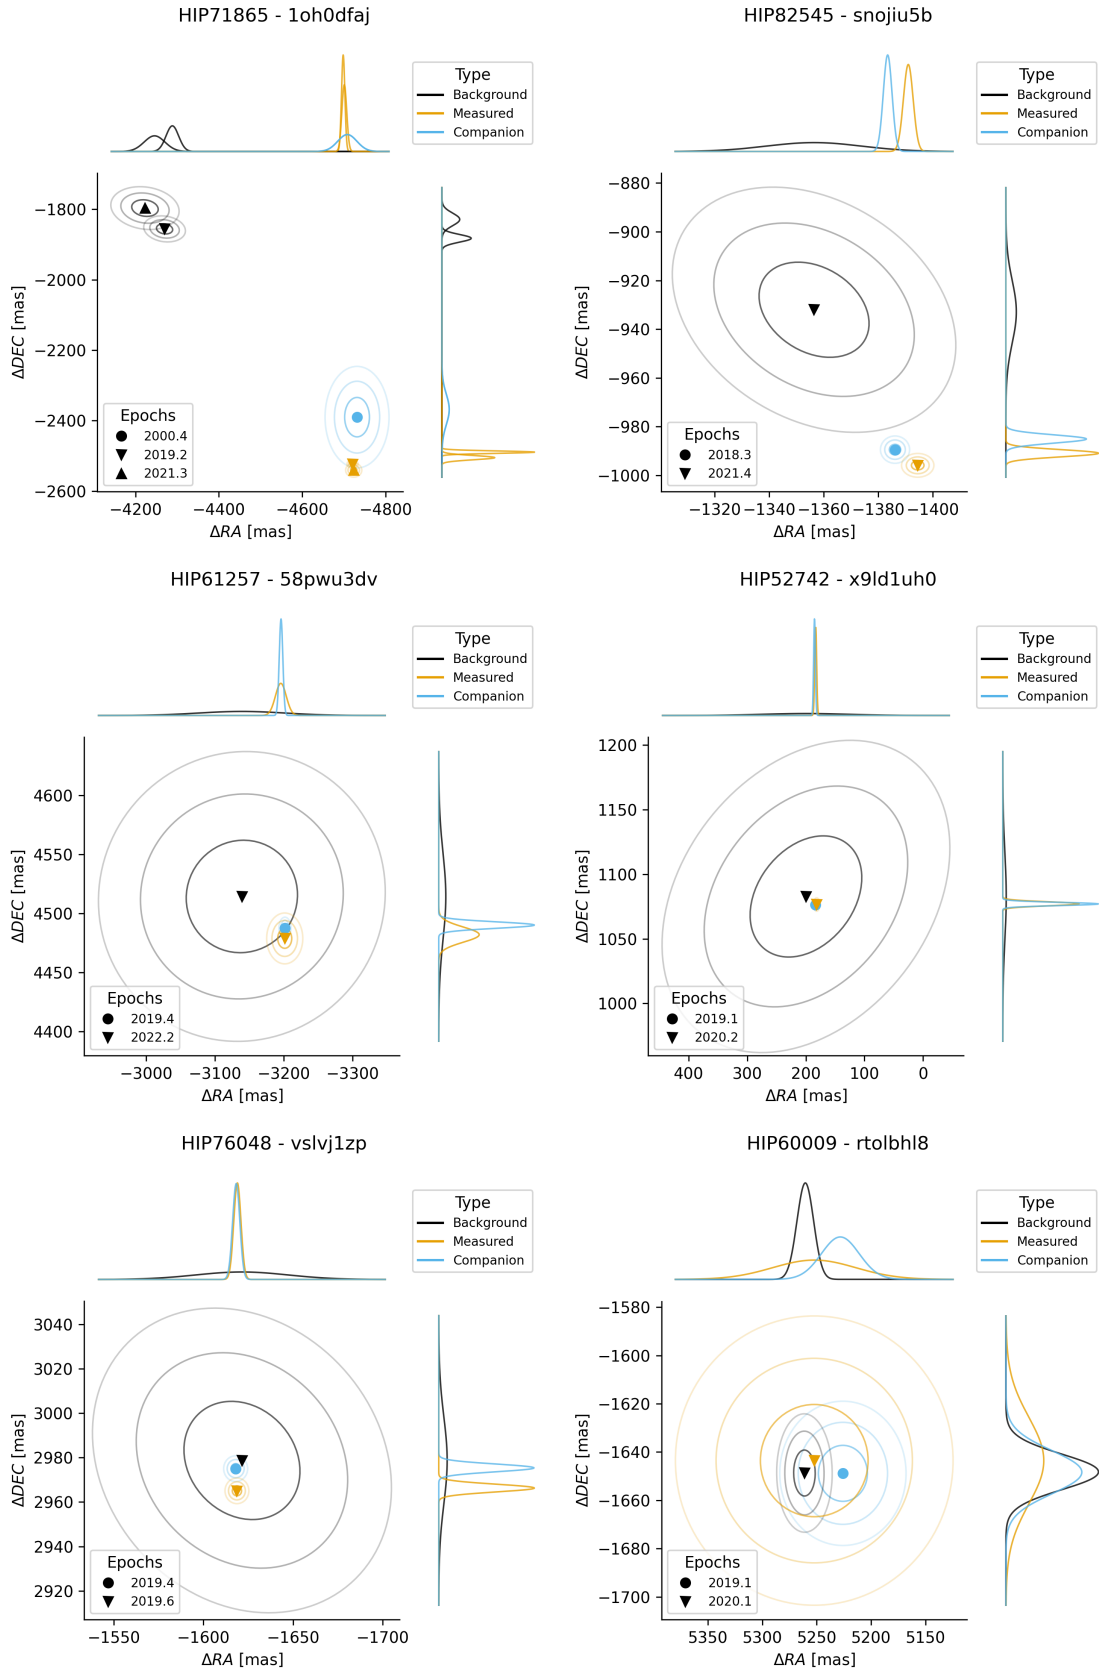


Fig. B.1: Candidates from BEAST with an logarithmic odds ratio $\log r_{tcb} > 0$. The schematics of the plots are explained in Figure 11.

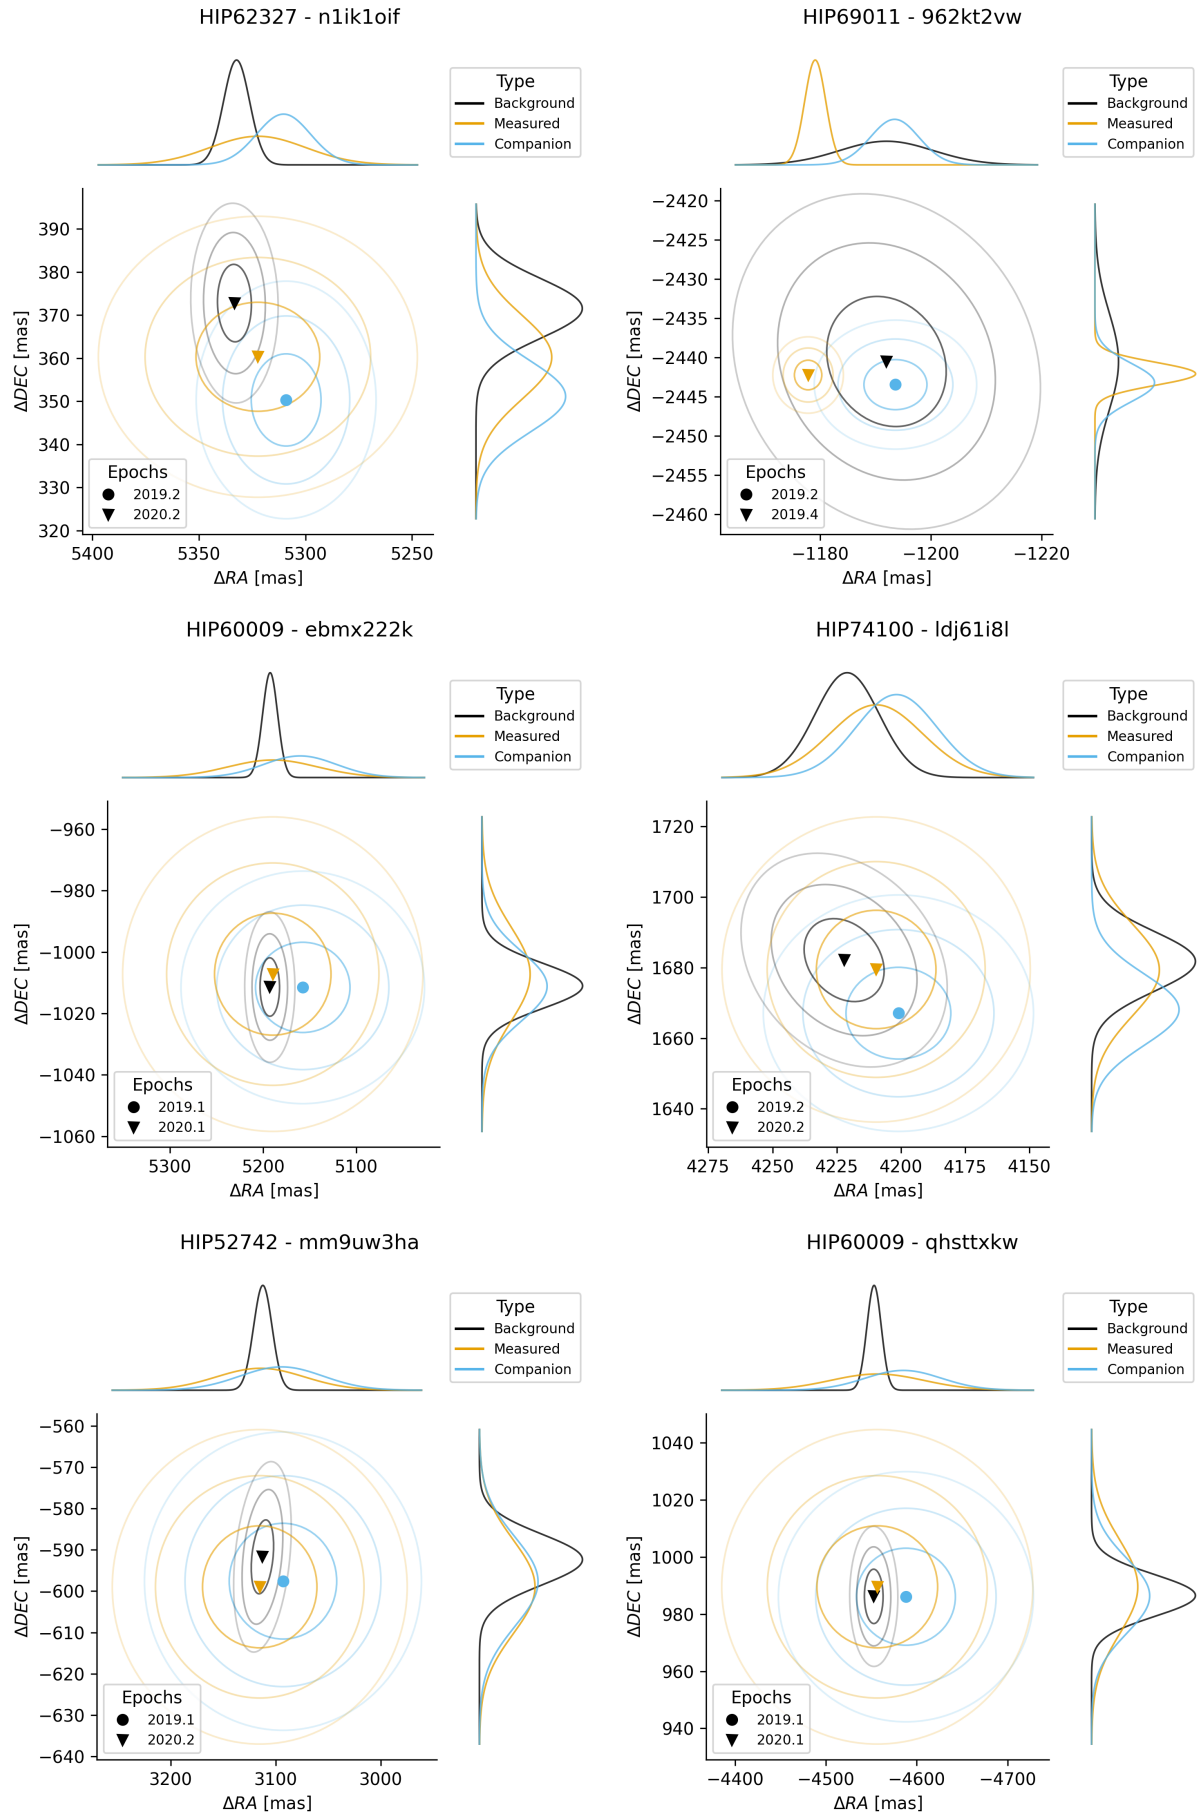


Fig. B.2: Candidates from BEAST with an logarithmic odds ratio $\log r_{tcb} > 0$. The schematics of the plots are explained in Figure 11.

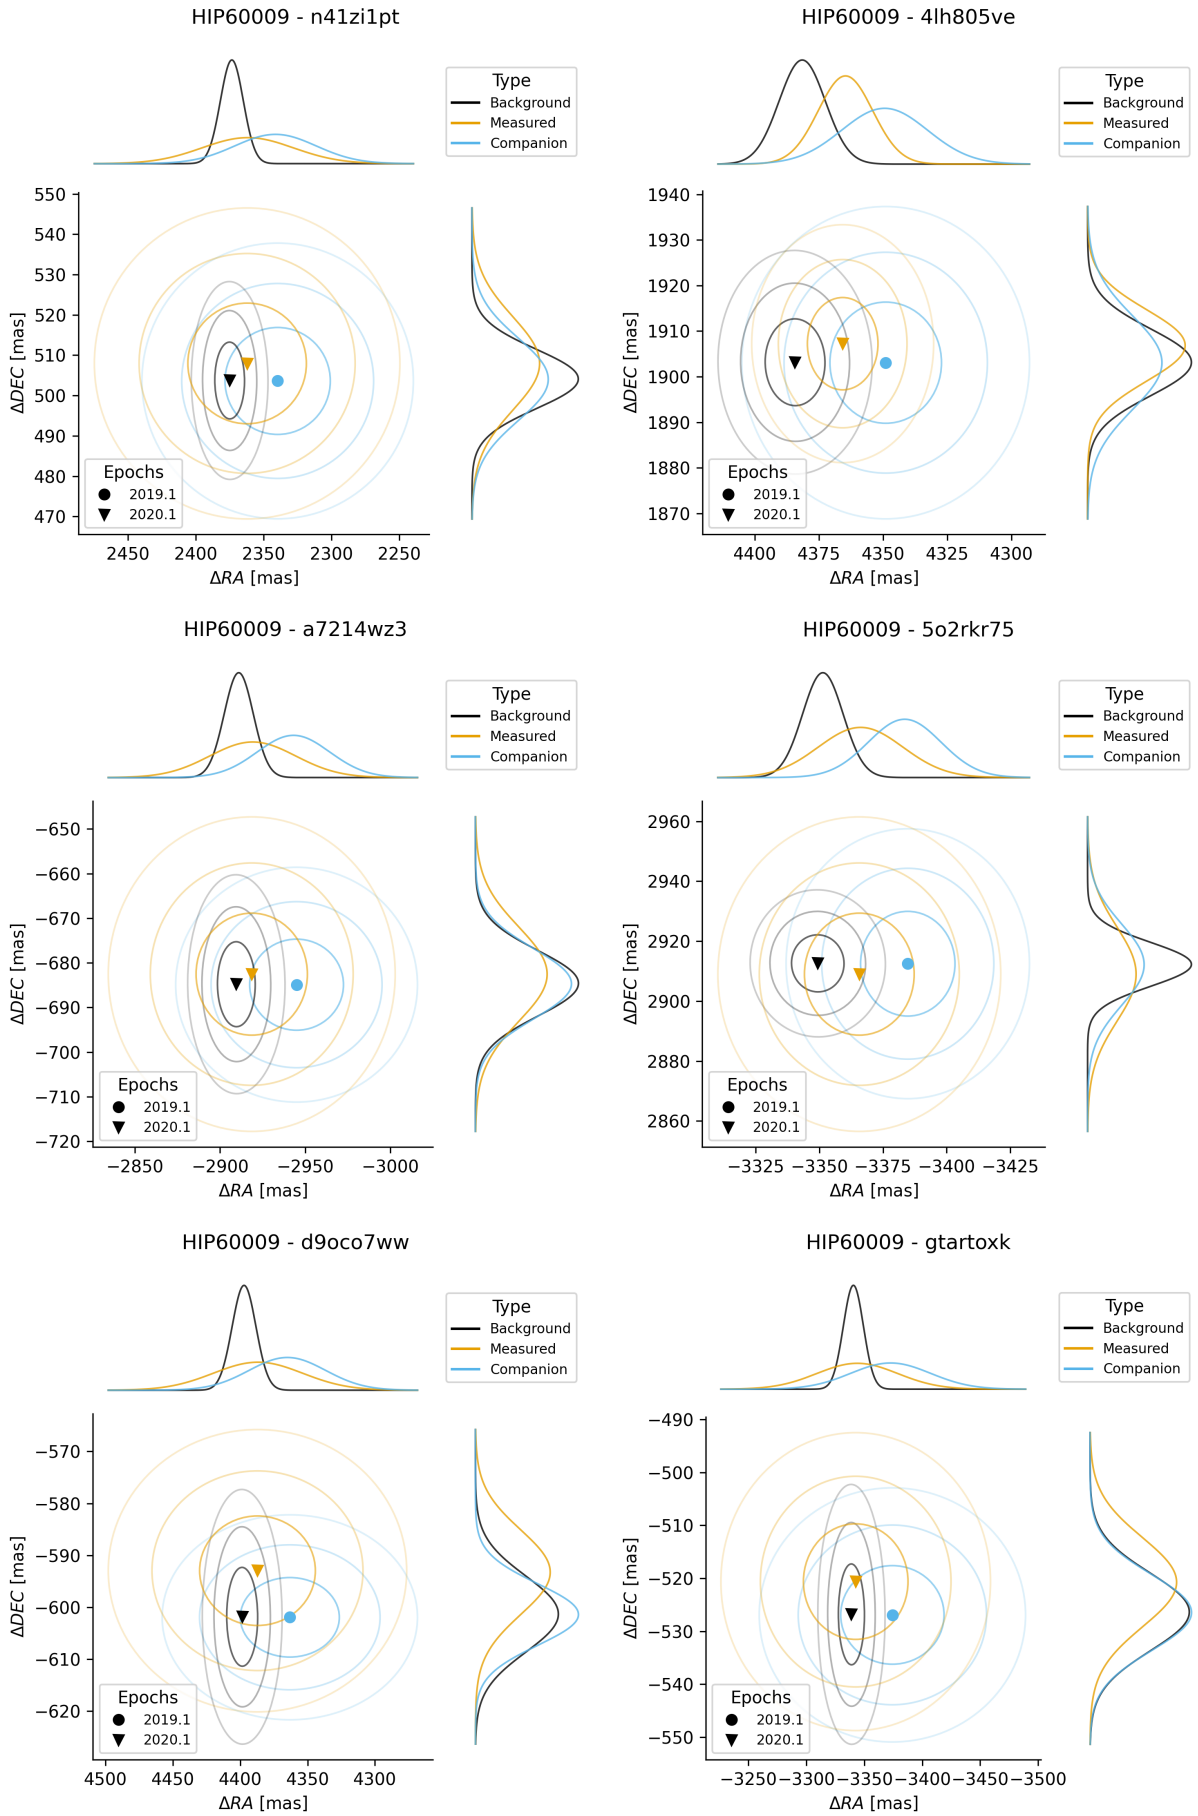


Fig. B.3: Candidates from BEAST with an logarithmic odds ratio $\log r_{tcb} > 0$. The schematics of the plots are explained in Figure 11.