



Publication Year	2023
Acceptance in OA	2025-03-04T16:27:59Z
Title	An empirical model of the Gaia DR3 selection function
Authors	Tristan Cantat-Gaudin, Morgan Fouesneau, Hans-Walter Rix, Anthony G. A. Brown, Alfred Castro-Ginard, DRIMMEL, Ronald, David W. Hogg, Andrew R. Casey, Shourya Khanna, Semyeong Oh, Adrian M. Price Whelan, Vasily Belokurov, Andrew K. Saydjari, Gregory M. Green
Publisher's version (DOI)	10.1051/0004-6361/202244784
Handle	http://hdl.handle.net/20.500.12386/36417
Journal	ASTRONOMY & ASTROPHYSICS
Volume	669

An empirical model of the *Gaia* DR3 selection function

Tristan Cantat-Gaudin¹, Morgan Fouesneau¹, Hans-Walter Rix¹, Anthony G. A. Brown²,
Alfred Castro-Ginard², Zuzanna Kostrzewa-Rutkowska², Ronald Drimmel³, David W. Hogg^{1,4,5},
Andrew R. Casey^{7,8}, Shourya Khanna³, Semyeong Oh⁶, Adrian M. Price-Whelan⁴, Vasily Belokurov⁶,
Andrew K. Saydjari^{9,10}, and G. Green¹

¹ Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany
e-mail: cantat@mpia.de

² Leiden Observatory, Leiden University, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands

³ INAF – Osservatorio Astrofisico di Torino, Strada Osservatorio 20, Pino Torinese 10025 Torino, Italy

⁴ Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY 10010, USA

⁵ Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA

⁶ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

⁷ School of Physics & Astronomy, Monash University, Clayton 3800, Victoria, Australia

⁸ Centre of Excellence for Astrophysics in Three Dimensions (ASTRO-3D), Melbourne, Victoria, Australia

⁹ Department of Physics, Harvard University, 17 Oxford St., Cambridge, MA 02138, USA

¹⁰ Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138, USA

Received 21 August 2022 / Accepted 7 October 2022

ABSTRACT

Interpreting and modelling astronomical catalogues requires an understanding of the catalogues' completeness or selection function: what properties determine an object's probability of being included in the catalogue? Here we set out to empirically quantify the completeness of the overall catalogue of *Gaia*'s third data release (DR3). This task is not straightforward because *Gaia* is the all-sky optical survey with the highest angular resolution to date and no consistent ground truth exists to allow direct comparisons. However, well-characterised deeper imaging enables an empirical assessment of *Gaia*'s *G*-band completeness across parts of the sky. On this basis, we devised a simple analytical completeness model of *Gaia* as a function of the observed *G* magnitude and position over the sky, which accounts for both the effects of crowding and the complex *Gaia* scanning law. Our model only depends on a single quantity: the median magnitude M_{10} in a patch of the sky of catalogued sources with `astrometric_matched_transits` ≤ 10 . We note that M_{10} reflects elementary completeness decisions in the *Gaia* pipeline and is computable from the *Gaia* DR3 catalogue itself and therefore applicable across the whole sky. We calibrated our model using the Dark Energy Camera Plane Survey (DECaPS) and tested its predictions against *Hubble* Space Telescope observations of globular clusters. We found that our model predicts *Gaia*'s completeness values to a few per cent (RMS) across the sky. We make the model available as a part of the `gaiaunlimited` Python package built and maintained by the `GaiaUnlimited` project^{*}.

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

1. Introduction

Comparing model predictions to observations drawn from an astronomical catalogue requires knowledge of the selection effects and incompleteness affecting the observed list of objects. Knowing what we could not observe could be as essential as what we observed, even for simple endeavours. For instance, mapping the stellar distribution around the Sun to reconstruct our Galaxy's overall shape requires knowing the sample's limiting magnitude and whether this limit varies across the sky.

A catalogue selection function S_C describes the probability of an object being included in an astronomical catalogue. Such a function represents the combined effects of the data collection (such as detection efficiency decreasing with apparent magnitude) and data processing (such as removing sources with noisy observations). To avoid biases caused by incomplete data, astronomers commonly restrict their studies to regions of the parameter space where the sample is assumed to be complete

($S_C \sim 1$). This approach is generally substantially restrictive and could lead to a poor representation of the problem one wants to address. Instead, one needs to fold in the selection function of a catalogue. In Rix et al. (2021), we have presented a general approach to evaluating and accounting for known selection functions in modelling astronomical data.

Gaia observes the sky continuously according to a complex scanning law: a six-hour rotation around its spin axis, a 63-day day precession of the spin axis, and the annual motion of the Earth (and its Lagrange 2 point) around the Sun (Gaia Collaboration 2016). This results in an intricate pattern, covering the entire celestial sphere with, on average, ~ 15 visits per year, but with significant variations. The *Gaia* catalogue, only includes sources with at least five observations (Lindgren et al. 2018). The probability that a transit across the *Gaia* field of view leads to an observation is lower for fainter sources, mainly due to two reasons. First, the onboard source detection algorithm has a nominal faint-end threshold of $G = 20.7$ mag beyond which sources are not selected for observation; however, the onboard magnitude estimate has a precision of a few tenths of magnitude

^{*} <https://github.com/gaia-unlimited/gaiaunlimited>

(de Bruijne et al. 2015), giving sources fainter than 20.7 a non-zero observation probability. Second, in crowded areas where the density exceeds $\sim 1\,050\,000$ per square degree, *Gaia* cannot follow all transiting sources and prioritises bright objects over fainter ones (Gaia Collaboration 2016). The probability that a source benefits from five observations is therefore a complex function of the sky position (via both crowding and the scanning law) and magnitude. Using the notation introduced in Rix et al. (2021), this paper models the catalogue selection function $S^{\text{parent}}(\mathbf{q})$ of the third *Gaia* data release (DR3, Gaia Collaboration 2021), where in the present case the catalogue properties $\mathbf{q} = (\ell, b, G)$ are the sky-position (ℓ, b) and G magnitude of a source.

The most common and straightforward approach to estimating a catalogue selection function of a sample is to compare the dataset with a more complete catalogue, which often means a deeper one in terms of magnitude limit. This comparison is generally made by binning both catalogues by magnitude, colour, or sky position and computing the ratio of the source count in each bin. For instance, Rybizki et al. (2021) followed this procedure to characterise the selection function of the radial velocity sample of the second *Gaia* data release (Katz et al. 2019), and Everall & Das (2020) improved upon this version using a smooth Gaussian mixture model to solve the issue with sparsely populated bins. However, this technique is empirical and relies on an external reference of complete samples.

There are ongoing efforts to reconstruct the *Gaia* selection function from a forward-modelling approach (Boubert et al. 2020; Boubert & Everall 2020, and subsequent papers in their Completeness of the Gaiaverse series). This approach requires modelling each step of the *Gaia* processing, from the scanning law and onboard filtering to the astrometric processing. In an upcoming paper (Castro-Ginard et al., in prep.), we will update their model, using transit data of non-variable stars to identify data-taking gaps and time-variable detection efficiencies and to extend the model to bright sources ($G \sim 1\text{--}6$ mag).

This study is part of the larger GaiaUnlimited project¹, which aims to determine the *Gaia* selection function and provide tools to the astronomical community to account for the selection effects in the *Gaia* catalogue. The present paper empirically builds an analytical model of the *Gaia* DR3 source catalogue selection function, that is to say the probability that the final catalogue contains a given source as a function of its sky position and an apparent G magnitude. We use the deep Dark Energy Camera Plane Survey (DECaPS, Schlafly et al. 2018; Saydjari et al. 2022) of the southern Galactic plane as our ‘complete’ reference to calibrate our model. We identified a simple quantity derived from the *Gaia* catalogue itself to use as a predictor of the *Gaia* completeness as a function of magnitude at any location over the sky, even outside the DECaPS footprint.

The approach and modelling are presented in Sect. 2. In Sect. 3, we verify our predictions against data which were not used in the model calibration. We discuss our model and its limitation in Sect. 4, and close with concluding remarks in Sect. 5.

2. Approach

2.1. Choice of reference datasets

We use the Dark Energy Camera Plane Survey (DECaPS) DR1 catalogue (Schlafly et al. 2018) as ground truth to calibrate our

model of the *Gaia* selection function. DECaPS is a ground-based optical and near-infrared survey of the Galactic plane using the Dark Energy Camera (DECam, Flaughner et al. 2015) mounted on the 4 m Victor M. Blanco telescope at the Cerro Tololo Inter-American Observatory (CTIO). The 2.2° diameter field of view, $0.26''/\text{pixel}$ plate scale and arcsecond seeing make these observations well-suited to resolving even the extremely crowded inner galaxy. DECaPS DR1 covers the Galactic plane with $|b| < 4^\circ$ and $5 > \ell > -120^\circ$. The survey reaches typical exposure depths of ~ 23 mag in g and r bands, and uses the crowdsourced photometric pipeline, which is specifically designed to deal with crowded fields. Using DECaPS as a reference dataset means that the present study cannot address the issue of incompleteness at angular separations that neither *Gaia* nor DECaPS can resolve (~ 1 per arcsecond for DECaPS and half an arcsecond for *Gaia*), such as stars in binary systems.

In Sect. 3.2, we verify the prediction of our model on high-density regions, using *Hubble* Space Telescope observations of the inner 3.5×3.5 arcmin 26 globular clusters collected by Sarajedini et al. (2007). The data were acquired with the Wide Field Channel of the Advanced Camera for Surveys, with photometry in the F606W and F814W filters, and are essentially complete down to magnitude 25. This data set was used in Arenou et al. (2018) to visualise the completeness of *Gaia* DR2 and the influence of crowding, but no quantitative model of completeness was proposed.

2.2. Choice of initial dependencies

This work aims to identify observable quantities that can be computed from the *Gaia* data itself, to be used as a proxy to constrain the selection function in any given field. We explored possible choices of observables by computing source count ratios between *Gaia* and DECaPS in magnitude bins in various areas on the sky. Due to the onboard resource allocation strategy prioritising bright sources, a naive expectation would be that completeness correlates with observed source density, with more populated fields being less complete. The observed source density is, in fact, a poor indicator of the true density (two areas with the same number of *Gaia* sources can differ by a factor of four in DECaPS) and thus a poor predictor of completeness. This is illustrated in the left panel of Fig. 1, and in Fig. A.1.

We also tested the following as possible indicators of completeness in a given field of view: the magnitude at which the observed luminosity function differs from an expected power law, the mode of the magnitude distribution, the magnitude of the faintest star in a given area, and the 90th percentile of magnitude. This last quantity provides a reasonable estimate of completeness in the most crowded regions but does not perform well in sparser fields. Constructing an all-sky map of the aforementioned quantities (except local source density) is also computationally very expensive, as it requires going through the entire data set of ~ 1.8 billion DR3 sources.

The best indicator of completeness we could identify is the G magnitude of the sources with the smallest number of observations. The number of observations used to compute the astrometric solution of a given source is given in the *Gaia* DR3 catalogue as `astrometric_matched_transits`². By construction, its minimum value is five because sources with

¹ <https://gaia-unlimited.org/>

² The DR3 catalogue also contains the column `matched_transits`, which counts all transits matched to a certain source even if they were not used in the construction of the catalogue. The quantity used in this study is `astrometric_matched_transits`.




Fig. 1. Completeness of *Gaia* relative to the DECaPS survey, which is taken as ‘ground truth’, in four magnitude ranges, computed in 3000 distinct patches across the DECaPS footprint. This completeness is shown as a function of *Gaia* source density (left) or M_{10} (right): *Gaia* source density is a poor predictor of completeness, while the M_{10} parameter – the median magnitude of catalogued sources with `astrometric_matched_transits` ≤ 10 in a surrounding patch of the sky – is an excellent completeness predictor. M_{10} combines the impact of source density and scanning law, as demonstrated e.g. in Fig. 2.

fewer observations were not included. In the remainder of this paper, we denote M_{10} the median magnitude of the sources with `astrometric_matched_transits` ≤ 10 in a given patch of sky. Its value is generally between 19 and 21.5 and strongly depends on how many times a given region was seen by *Gaia* and how crowded the region is. This is illustrated in Fig. 2, where the patterns introduced by stellar density and the *Gaia* scanning law are clearly visible.

We choose the median value rather than the mean because some bright sources might occasionally have a small number of matched transits, and the median is a more robust summary statistic. The model could also be calibrated on the median magnitude of sources with exactly five astrometric matched transits, but the chosen value of ten conveniently allows for a sufficient number of tracers even in sparse regions (the sparsest HEALPix level 7 contains 19 such sources), while keeping the total number of tracers manageable when building all-sky maps (210 million in *Gaia* DR3).

The variation of completeness with the M_{10} value of each investigated patch of sky is shown in Fig. 1. The completeness in any magnitude range is overall a tight function of M_{10} , although a dispersion of up to ~ 0.1 can be seen for some magnitude ranges. This dispersion effectively sets the limit of the precision one can achieve by using M_{10} as the sole predictor. The same effect is illustrated in the bottom-left panel of Fig. 3, where it can be seen that a given value of M_{10} can correspond to slightly different completeness profiles. The effect of crowding likely depends not only on the true source density but also on the magnitude distribution of the sources, and a different parameter derived from the distributions shown in the top panel of Fig. B.1 might be able to provide a second-order correction to the simple model presented in this study.

2.3. Source count ratios relative to DECaPS

We study the completeness in 1085 patches of size 18×7.2 arcmin across the DECaPS DR1 footprint, sampling a wide

range of source densities. The size of the patches was chosen to allow us to avoid gaps in the coverage of the DECaPS DR1 data, mainly present near the Galactic centre. The distribution of those patches is shown in Fig. C.1. Patches containing at least 10 000 *Gaia* sources are further divided into two, four, or eight bins to provide a finer spatial resolution in densest areas, for a total of 2906 individual regions. We spatially match the *Gaia* data to DECaPS with a 1 arcsecond radius. When several DECaPS sources are present within this radius, we consider the best match to be the source whose r magnitude is closest to the *Gaia* source’s G . Matches with a magnitude difference larger than 1 mag are discarded. Despite *Gaia* having a finer angular resolution (0.4–0.5 arcsec, [Gaia Collaboration 2018](#)) than DECaPS (~ 1 arcsec, [Schlafly et al. 2018](#)), we find that less than 0.5% of *Gaia* sources have no DECaPS counterpart, indicating that the *Gaia* completeness is mostly limited by crowding rather than resolution. The colour and magnitude of these missing sources seem to be a random subset of the *Gaia* data. They appear to follow lines of constant declination on the sky, which suggests that they correspond to an instrumental effect of the Dark Energy Camera (e.g. bleeding trails caused by the presence of bright stars) rather than spurious *Gaia* detections. A small fraction of the *Gaia* sources ($\sim 0.3\%$, see Fig. D.1) lack a G -band magnitude. Since the present study investigates the completeness of the *Gaia* catalogue as a function of magnitude, our procedure treats these sources as if they were missing from the *Gaia* catalogue.

We estimate the G magnitude of the missing *Gaia* sources from their DECaPS ($r, r-i$) photometry. The conversion is performed by fitting a linear relation with the form $G = a \times r + b \times (r-i) + c$ in each patch of sky separately to account for the fact that photometric transformations are extinction-dependent. We then compute the fraction of DECaPS sources with a *Gaia* counterpart in bins of G magnitude of width 0.2 mag, from $G = 15$ to 23. The completeness as a function of magnitude is shown in Fig. 3, colour-coded by the value of M_{10} for each region. In




Fig. 2. Map of the parameter M_{10} in the direction of the Galactic centre; M_{10} is the median G magnitude (here in HEALPix regions of level 10) of *Gaia* sources with `astrometric_matched_transits` ≤ 10 , reflecting the outcomes for faint sources of the *Gaia* pipeline completeness decisions. The complex pattern results from the combination of the *Gaia* scanning law and stellar density, which in turn depends on Galactic structure and dust distribution. Prominent patches with bright M_{10} can be seen. Several globular clusters are also visible, for instance the prominent M 22 near $(\ell, b) = (10, -7.5)$.

the densest regions, the completeness reaches 50% at $G \sim 19$, while in sparse regions that benefited from large numbers of observations, the *Gaia* catalogue appears essentially complete down to $G \sim 20.5$. For a given value of M_{10} , the magnitude at which a 90% completeness is reached can vary by up to 0.4 mag, although with a typical scatter of less than 0.2 mag (comparable with the width of our magnitude binning).

2.4. Fitting the model

We model the completeness curve computed in each region with a sigmoid function. To capture the change of slope from dense to sparse regions as well as the slight asymmetry of the curve, we define a generalised sigmoid with the baroque but flexible analytic form:

$$S(G | M_{10}) = 1 - 0.5 \times \left[\tanh\left(\frac{x(M_{10}) - G}{y(M_{10})}\right) + 1 \right]^{z(M_{10})}, \quad (1)$$

where x is the magnitude of the inflexion point, y controls how steeply the completeness drops at the inflexion point (smaller values correspond to a steep decrease), and z describes the skewness of the function ($z < 1$ means it is flatter at bright magnitudes). The effect of varying these three parameters is illustrated in Fig. E.1.

We initially fit the triplet (x, y, z) independently in each of the 2906 patches. The generalised sigmoid defined in Eq. (1) allows the fit to reproduce the observed completeness of all patches with residuals smaller than 2%. Unfortunately, we cannot compute and provide an all-sky map of these three parameters because they can only be derived directly where DECaPS data is available. Instead, we investigate and model the relation between the three parameters and M_{10} .

The parameter x (which roughly sets the magnitude of 50% completeness) scales almost linearly with M_{10} . The parameters y and z (describing the slope and skewness of the curve) mostly follow two regimes, remaining roughly constant when $M_{10} < 20.5$, then increasing in value which results in a steeper, less asymmetric shape of the completeness curve for higher values of M_{10} . The high- M_{10} regime corresponds to the flux-limited selection function, where transiting sources are only granted a detection window if the *Gaia* sky mapper measures a magnitude $G_{\text{onboard}} < 20.7$ for that particular transit (Gaia Collaboration 2016). The low- M_{10} regime corresponds to areas on the sky where crowding plays a major role in the selection function.

To capture the variation of x , y , and z with M_{10} through these two regimes, we model them as a broken slope relation, with the same location M_{break} of the break for all three:

$$x(M_{10}) = \begin{cases} a_x M_{10} + b_x & \text{if } M_{10} < M_{\text{break}} \\ c_x M_{10} + (a_x - c_x) M_{\text{break}} + b_x & \text{otherwise} \end{cases} \quad (2)$$




Fig. 3. *Gaia*'s catalogue completeness $S(G | M_{10})$ as a function of G magnitude for a given completeness parameter M_{10} . *Top left*: empirically determined *Gaia* G completeness derived from the comparison with DECaPS photometry in thirteen DECaPS patches that are each colour-coded by their M_{10} . *Top right*: model completeness, from Eq. (1), for the corresponding M_{10} values. *Bottom left*: empirically determined completeness derived from the comparison with DECaPS photometry in 105 DECaPS patches with $M_{10}=19.4, 19.8, 20.2, 20.6$, and 21 (within 0.01 mag). The dispersion at a given M_{10} corresponds to the scatter seen in Fig. 1 (right panel) and the residuals in Fig. 5. *Bottom right*: model completeness for these five values of M_{10} .

$$y(M_{10}) = \begin{cases} a_y M_{10} + b_y & \text{if } M_{10} < M_{\text{break}} \\ c_y M_{10} + (a_y - c_y) M_{\text{break}} + b_y & \text{otherwise} \end{cases} \quad (3)$$

$$z(M_{10}) = \begin{cases} a_z M_{10} + b_z & \text{if } M_{10} < M_{\text{break}} \\ c_z M_{10} + (a_z - c_z) M_{\text{break}} + b_z & \text{otherwise} \end{cases} \quad (4)$$

The resulting hierarchical model has 10 free hyperparameters (three for each of x , y , and z , plus the location of the break). We add a final free parameter σ representing the noise on the observed completeness profiles. The noise is assumed to be Gaussian and constant with magnitude. This is a rough approximation, and here σ acts like a nuisance parameter rather than a model for the noise. The corresponding log-likelihood is:

$$-n \log \sigma - \frac{1}{2\sigma^2} \sum_{i=1}^n (\text{obs}_i - \text{pred}_i)^2 \quad (5)$$

where n is the total number of data points: the source count ratios in 40 magnitude bins $\times$ 2906 patches. We maximise the log-likelihood with the Markov chain Monte-Carlo sampler *emcee* (Foreman-Mackey et al. 2013) and explore the parameter space with 32 walkers for 10 000 steps each. We impose that a_x , c_x , and c_z must be positive, c_y must be negative, σ must be between 0 and 1, and M_{break} between 19 and 21. The priors on the other parameters are left unbounded. The sampling takes about two hours on an 8-core laptop. We discard the first 1000 iterations as burn-in³. We provide the median of posterior samples for each parameter in Table 1. The final relation between the parameters of the sigmoid and M_{10} are shown in Fig. 4.

³ The resulting chain is 70–90 times longer (depending on the parameter) than the autocorrelation time estimated by *emcee*.




Fig. 4. Relation between the parameters of the sigmoid and M_{10} in Eq. (1), for 500 samples from the MCMC chain: the inflection point, $x(M_{10})$ – Eq. (2), on the *left*; the inverse slope, $y(M_{10})$ – Eq. (3), at *center*; and the skewness, $z(M_{10})$ – Eq. (4), on the *right*. The magnitude of 50% completeness, $x(M_{10})$ is well approximated by M_{10} . Fields with fainter M_{10} have a steeper (smaller $y(M_{10})$) and more symmetrical ($z(M_{10})$ towards 1) selection function. Examples of resulting sigmoids are shown in Fig. 3 for a range of M_{10} values.




Fig. 5. Mean completeness residuals, i.e. model predicted minus observed completeness, as a function of magnitude, for patches in two different ranges of M_{10} . The shaded areas correspond to the 16th to 84th percentile and 5th to 95th percentile intervals. Expectedly, the residuals are highest near 50% completeness and smaller in the highly complete or dramatically incomplete regime.

Table 1. Parameters for the selection function $S(G | M_{10})$ in Eq. (1).

Parameter	Median	Uncertainty	
a_x	0.985	+0.002	−0.001
b_x	0.649	+0.030	−0.031
c_x	0.693	+0.001	−0.002
a_y	−0.004	+0.003	−0.003
b_y	0.223	+0.060	−0.051
c_y	−0.093	+0.002	−0.002
a_z	0.006	+0.003	−0.003
b_z	0.034	+0.068	−0.060
c_z	0.351	+0.004	−0.004
M_{break}	20.519	+0.002	−0.001
σ	0.02060	+0.00004	−0.00004

Notes. The parameters for Eq. (1) are defined via Eqs. (2)–(4), and were obtained via MCMC. We adopt the median of the posterior chain as the best value, and state the 16th to 84th percentile confidence interval. Here, σ is a nuisance parameter in the fitting procedure. The corner plot of the full posterior chain is shown in Fig. F.1.

Figure 5 shows the mean and dispersion of the residuals, for two ranges of M_{10} . The prediction is most precise where the model predicts completeness of 0 or 100%. Where the prediction is least precise, the dispersion of the residuals reaches about 5%.




Fig. 6. Summary of the workflow used to build the selection function $S(G | M_{10})$ (Eq. (1)) as a function of sky position and magnitude G .

Figure 6 summarises the workflow and how we use the hyperparameters to predict the completeness as a magnitude function at any sky position.

3. Testing the model

3.1. With more DECaPS data

We verify the prediction of our model by applying it to regions of the DECaPS footprint that were not used for the fitting step.




Fig. 7. Comparison of the direct, empirical and model-predicted completeness maps, illustrated at $G \sim 21$. *Top left*: map of the direct completeness estimate, i.e. the ratio of source densities in *Gaia* and DECaPS in the magnitude range $20.9 < G < 21.1$. *Top right*: map of the quantity M_{10} used to predict the model completeness. *Bottom left*: completeness at $G = 21$ predicted from the M_{10} map and the model of Eq. (1). *Bottom right*: map of the difference between the predicted and observed completenesses. We note that using more external information, the model-predicted completeness map (*bottom left*) is effectively a de-noised version of the empirical completeness map (*top left*).

The $2^\circ \times 2^\circ$ field of view shown in Fig. 7 was chosen to straddle the boundary between a stripe which received a large number of visits and an adjacent area with much fewer transits. The M_{10} parameter was mapped by computing the median G magnitude of stars with ten or fewer `astrometric_matched_transits` in spatial bins of 2.4×2.4 arcmin. The main diagonal feature, splitting the field of view in two, is due to the *Gaia* scanning law. The finer structure is shaped by patchy dust extinction.

The model-predicted completeness map (bottom left of Fig. 7) obtained from the M_{10} map (top right) is less noisy than the map obtained directly from source count ratios (top left). However, as discussed in Sect. 4.1, the use of M_{10} as the sole predictor of completeness can lead to local biases of a few per cent (within the amplitude of the residuals shown in Fig. 5 and the bottom left panel of Fig. 3). In Fig. 7, this leads to slightly overestimating the completeness of the most complete area (with $M_{10} > 21.1$).

3.2. HST data of globular clusters

The cores of globular clusters (GCs) are among the most challenging regions for *Gaia* due to their high densities. In some particularly dense clusters, the completeness at $G = 18$ is close to zero.

The data (presented in Sect. 2.1) contains observations of the inner 3 arcmin of 26 GCs. We split each field of view into a core (inner 1.5 arcmin) and a surrounding area. Since these objects are very dense and the field of view is smaller than the spatial binning we used in our model calibration, fifteen cluster regions have values of M_{10} which are locally smaller than the range (19.11–21.23) covered by the calibration set (computed on a coarser spatial resolution). We show in Fig. 8 that our model can predict the completeness even when extrapolated to these crowded, low- M_{10} fields. The extrapolation only seems to fail by 20% in the most extreme case of crowding, which corresponds to the inner 1.5 arcmin of Omega Centauri ($M_{10} \sim 15$).

As an example, we map the completeness prediction of globular cluster NGC 1261 in Fig. 9. Our model correctly identifies the regions of 100% and 0% completeness. The intermediate regions appear as a ring of noise on the residuals map (bottom-right panel of Fig. 9), due to the small spatial binning and the narrow magnitude range.

3.3. Comparison to the Gaiaverse model

We compare our predictions with the model of [Everall & Boubert \(2022, hereafter EB22\)](#), which itself is the *Gaia* DR3 update of the model developed by [Boubert & Everall \(2020\)](#)




Fig. 8. Completeness of *Gaia* relative to HST as a function of M_{10} in five chosen magnitude ranges, for an extreme crowding regime: the core and outskirts of 26 globular clusters. The lines show the expected completeness according to our model. Remarkably, our model (Eq. (1)) provides unbiased completeness predictions based on *Gaia* information alone.

for DR2. Their predictions do not rely on comparisons to reference data but on a model of the *Gaia* scanning law and of the detection efficiency as a function of magnitude. We show the all-sky map of completeness at $G = 21$ predicted by both models in Fig. 10.

Boubert & Everall (2020) point out that modelling the effects of crowding is a complex task, as one can only know the observed density of sources while crowding depends on the true density (including those missing from the catalogue). For this reason, the effect of globular clusters or high-density regions near the Galactic centre is more clearly visible in our model, which naturally accounts for crowding via the M_{10} parameter.

The most striking difference between both models is that EB22 predict much higher completeness even in non-crowded regions, with an essentially 100% complete catalogue at $G = 21$ across most of the sky, while our model predicts that a 100% completeness at this magnitude is only achieved in the regions most favoured by the scanning law. This is supported by comparisons to the DECaPS data (including those used to calibrate the model), which show most of the Galactic plane is only 60–80% complete at this magnitude.

A likely explanation for this discrepancy is that EB22 overestimates faint source detection probabilities. Figure 7 from EB22 shows that for sources with $G = 21$, the reconstructed detection probability is $\sim 30\%$, which translates to a 90% probability of having five detections after 25 scans (the number for the least visited regions in DR3) and 99.5% probability after 40 scans (the median number for the whole sky). The EB22 model estimates detection efficiencies from photometric time series published in the *Gaia* archive for variable stars. This sample is likely to be biased towards stars with high-quality measurements and low photometric errors. A more realistic estimate of the range of detection efficiencies in a given region can be obtained from individual transit data of all *Gaia* sources. This data is not publicly available and will be used in a future publication (Castro-Ginard et al., in prep.)

4. Discussion

4.1. Limits and potential improvements

With its three parameters, the generalised sigmoid functional form defined in Eq. (1) turns out to be sufficiently flexible to approximate the observed *Gaia*-to-DECaPS count ratio as a

function of magnitude in any region to within two per cent. Given the limited sky coverage of our ground truth catalogue (DECaPS, which only covers $\sim 7\%$ of the celestial sphere), we need to predict the values of Eq. (1)'s three parameters, using the quantity M_{10} defined in Sect. 2.2 and computed from the *Gaia* data itself. The scatter observed in the right panel of Fig. 1 varies with M_{10} and with G magnitude and illustrates the limitations of predicting selection function $S(G | M_{10})$ using M_{10} as the sole predictor of completeness.

We could not identify a single quantity providing a better precision than M_{10} , which encodes the combined effect of the scanning law and crowding. Nonetheless, it may be possible to establish second-order corrections based on other *Gaia*-derived quantities. The detection probability as a function of magnitude likely depends on the magnitude distribution of sources in a given field of view, not just on their total number. We investigated the residuals of our model but only found hints of additional correlations between completeness and the total number of scans, or completeness and observed source density, in some restricted ranges of M_{10} and G magnitude. Establishing the right functional form and choice of dependency for such ad-hoc corrections would be a difficult task. One might be tempted to follow a machine-learning approach and let the machine determine the most relevant predictors of the *Gaia*-to-DECaPS count ratio. This would, however, incur the likely risk of overfitting, as many of the correlations in the reference data cannot reliably be generalised to the entire sky unless they can be supported by some understanding of the instrumental pipeline.

A potential improvement to the M_{10} proxy could be to characterise the entire distribution of `astrometric_matched_transits` with magnitude G (illustrated in Fig. B.1 for five chosen patches) rather than just its value at the faint end. The slope and shape of the drop in the number of matched transits may contain information on the level and type of crowding affecting the observations. Another direction of improvement to explore is to see whether sky areas with a broader range of scanning angles are more likely to be complete since the sources missed by *Gaia* are more likely to be different at each visit (Gaia Collaboration 2016; Pancino et al. 2017). The dispersion in scanning angles over a given area could therefore be an additional parameter in the empirical description of the selection function. This quantity is available for each source in the *Gaia* catalogue as `scan_direction_strength_k2`, but testing its validity as a secondary predictor of completeness would be difficult in the present context because the regions with the densest clustering of scanning angles are located outside the DECaPS footprint (see e.g. Fig. 1a in Everall et al. 2021). Using the DECaPS DR2 release (Saydjari et al. 2022), which more than doubles the survey area, could mitigate this problem in future work.

Finally, the model constructed in this study assumes that the (x, y, z) parameters of the sigmoid are related to M_{10} via a broken-slope relation, with a total of ten free hyperparameters. A more complex model (for instance, with more breaks) would decrease the residuals shown in Fig. 5 (and smooth out the kink near $G \sim 20.2$ in its right panel), but unless the increase in complexity can be justified by some knowledge of the instrumental behaviour of *Gaia*, a simpler model is more likely to be valid outside the DECaPS footprint.

4.2. Dust extinction makes *Gaia* more complete

A natural but perhaps counter-intuitive effect of interstellar extinction is to increase completeness as a function of apparent




Fig. 9. Comparison of the direct (empirical) and model-predicted completeness maps, illustrated for the globular cluster NGC 1261. *Top left:* map ratio of the number of sources in *Gaia* and HST in the magnitude range $20.9 < G < 21.1$. *Top right:* map of M_{10} used to predict the completeness. *Bottom left:* predicted completeness at $G = 21$. *Bottom right:* map of the difference between the predicted and observed completeness: the variance is largest in the intermediate completeness regime (see Fig. 5), producing a ring-like structure in the residuals map.

magnitude. Foreground sources of a given apparent magnitude G are more easily detected when projected against a ‘dark’ background.

Of course, dust extinction still reduces the probability of *Gaia* catalogue membership for sources of a given set of physical properties and distance. Given that the selection function must be phrased in terms of observables, modelling sources as a function of distance and absolute magnitude requires a 3D extinction map.

4.3. Arguments of the selection function

In this study, we expressed the *Gaia* source catalogue selection function as a function of magnitude and position (G, ℓ, b). We find no evidence that this fundamental *Gaia* selection function depends significantly on the source colour: in a given part of the sky, two sources with the same G magnitude appear to have equal probabilities to be included in *Gaia*, regardless of their colour. This result is not surprising, because the *Gaia* sky mapper and the astrometric instruments on board the spacecraft operate in the G band.

We point out that due to strong correlations between the observables, investigating the chromaticity of the selection function is a much more complex task than simply expressing detection rates as a function of colour. For astrophysical reasons, red stars tend to be intrinsically fainter than blue stars. Interstellar extinction acts in the same direction, making sources appear both fainter and redder. On the other hand, areas of the sky heavily obscured by dust are redder but also more complete due to the background being less crowded in the magnitude range where *Gaia* operates.

4.4. Selection function for subsets of the *Gaia* catalogue

This paper only addresses the completeness of the sample of *Gaia* catalogue entries with a published position and G magnitude, establishing the selection function noted $S_C^{\text{parent}}(G, \ell, b)$ in the notation of Rix et al. (2021). In practice, most users will be interested in comparing other *Gaia* quantities with theoretical models, such as observed G_{BP} and G_{RP} fluxes, parallaxes, proper motions, or more advanced data products provided by the *Gaia* pipelines such as astrophysical parameters




Fig. 10. Global comparison of the completeness maps predicted by the ab initio completeness model (EB22, *top panel*), and our empirically derived M_{10} -based completeness model (*bottom panel*). The overall morphology of the two maps is similar, but our empirical completeness model implies far greater incompleteness (at $G = 21$), especially in the regions of high source densities.

(Gaia Collaboration 2022). It is not clear whether the approach used in this study is suitable for selecting further, more restricted subsets of the *Gaia* data.

First, M_{10} might not be a good predictor of the completeness of *Gaia* subsets, say stars with spectra from the radial velocity spectrometer (RVS; Gaia Collaboration 2016), because different instruments on board the spacecraft have different crowding limits: 1 050 100 sources per square degree for the astrometric instrument, 750 000 for the BP/RP spectrographs, and 35 000 for RVS. It may, however, be possible to construct equivalent quantities to characterise particular subsets, for example an equivalent of M_{10} for *Gaia* sources with BP/RP. Second, it is not clear that

the generalised sigmoid function (Eq. (1)) is a good functional form for the selection functions of various *Gaia* subsets. Third, the selection function of some subsets will depend on more than just G magnitude and sky position. For instance, Everall & Das (2020) and Rybizki et al. (2021) express the RVS and $\text{ruwe} < 1.4$ completeness as functions of $(G - G_{\text{RP}})$.

For thinking about the construction of more complex selection functions, we refer the reader to Rix et al. (2021), who provide recommendations on how to construct the sample function $S_C^{\text{sample}}(\mathbf{q})$ of a given subset of the *Gaia* data selected on attributes $\mathbf{q}$. In general, an overall selection function can be approximated as a multiplication reflecting the different Boolean

steps in the sample selection:

$$S_C(\mathbf{q}, G, \ell, b) = S_C^{\text{sample}}(\mathbf{q}) \times S_C^{\text{parent}}(G, \ell, b). \quad (6)$$

4.5. Non-stellar sources and extragalactic objects

The present study does not explicitly consider non-stellar sources and our model is calibrated on the DECaPS data, for which most sources are stars. At high Galactic latitudes, the density of galaxies is higher than that of stars. The completeness for extended sources is especially difficult to estimate with comparisons to reference data sets, as some deep data sets (e.g. DECaLS, Dey et al. 2019) list different parts of resolved galaxies as distinct catalogue entries, where *Gaia* processed the extended source as a single object (and vice versa). Further characterisation of the completeness of the *Gaia* catalogue could also be performed with the Pan-STARRS data (Chambers et al. 2016), which reach magnitudes of 22.5–23 in the *r* band and cover ~30 000 square degrees, mostly off the Galactic plane. For the specific case of M 31, the completeness of *Gaia* and the *Gaia* Andromeda Photometric Survey (providing epoch photometry for all sources including non-variables Evans et al. 2022), the Panchromatic *Hubble* Andromeda Treasury data (Dalcanton et al. 2012) covering 0.5 square degrees of the Andromeda galaxy’s disc would also be a valuable data set.

5. Summary and conclusion

This study is part of a paper series by the GaiaUnlimited project that aims to characterise the *Gaia* selection function and provide the astronomical community with corresponding data and tools. This paper presents an analytical model of the *Gaia* DR3 completeness as a function of observed *G* magnitude and position on the sky. Our model depends on a single quantity which is derived from the *Gaia* data itself: the median magnitude M_{10} in a patch of the sky of catalogued sources with `astrometric_matched_transits` ≤ 10. The quantity M_{10} reflects the elementary processes and decisions made by the *Gaia* pipeline to turn observations into the published *Gaia* catalogue and naturally accounts for the effects of crowding and the *Gaia* scanning law.

As ground truth, we rely on the DECaPS survey, which is deeper than *Gaia* and whose pipeline is optimised for high-density fields, to calibrate our model. We test our predictions against DECaPS and *Hubble* Space Telescope observations of globular clusters. Our model predicts the observed completeness with a precision of up to a few per cent. We make the model available as a Python package through the GaiaUnlimited web page, along with documentation and tutorials. The present model only provides a selection function for the *Gaia* DR3 entries with a published *G* magnitude and sky position. Sub-samples of this catalogue will be characterised in upcoming GaiaUnlimited publications.

Acknowledgments. 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. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular, the institutions participating in the *Gaia* Multilateral Agreement. A.R.C. is supported in part by the Australian Research Council through a Discovery Early Career Researcher Award (DE190100656) and through Discovery Project DP210100018. DECaPS data were retrieved with the `astro-datalab` (<https://github.com/astro-datalab/datalab/>) Python package, and the *Gaia* data with `astroquery` (Ginsburg et al. 2019). This work also made use of the Python packages `astropy` (Astropy Collaboration 2018), `scipy` (Virtanen et al. 2020), `astroML` (Vanderplas et al. 2012), `MWplot` (<https://pypi.org/project/mw-plot/>), `numpy` (Harris et al. 2020), `plotly` (Inc. 2015), `healpy` (<http://healpix.sourceforge.net>; Górski et al. 2005; Zonca et al. 2019), `pandas` (McKinney et al. 2010), and `matplotlib` (Hunter 2007). T.C.G. acknowledges an extensive use of TOPCAT (Taylor 2005) and Jupyter notebooks (Kluyver et al. 2016).

References

- Arenou, F., Luri, X., Babusiaux, C., et al. 2018, *A&A*, **616**, A17
- Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, *AJ*, **156**, 123
- Boubert, D., & Everall, A. 2020, *MNRAS*, **497**, 4246
- Boubert, D., Everall, A., & Holl, B. 2020, *MNRAS*, **497**, 1826
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints [arXiv:1612.05560]
- Dalcanton, J. J., Williams, B. F., Lang, D., et al. 2012, *ApJS*, **200**, 18
- de Bruijne, J. H. J., Allen, M., Azaz, S., et al. 2015, *A&A*, **576**, A74
- Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *AJ*, **157**, 168
- Evans, D. W., Eyer, L., Busso, G., et al. 2022, *A&A*, accepted, <https://doi.org/10.1051/0004-6361/202244204>
- Everall, A., & Boubert, D. 2022, *MNRAS*, **509**, 6205
- Everall, A., & Das, P. 2020, *MNRAS*, **493**, 2042
- Everall, A., Boubert, D., Koposov, S. E., Smith, L., & Holl, B. 2021, *MNRAS*, **502**, 1908
- Flaugher, B., Diehl, H. T., Honscheid, K., et al. 2015, *AJ*, **150**, 150
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Gaia Collaboration (Prusti, T., et al.) 2016, *A&A*, **595**, A1
- Gaia Collaboration (Brown, A. G. A., et al.) 2018, *A&A*, **616**, A1
- Gaia Collaboration (Brown, A. G. A., et al.) 2021, *A&A*, **649**, A1
- Gaia Collaboration (Creevey, O. L., et al.) 2022, *A&A*, accepted, <https://doi.org/10.1051/0004-6361/202243800>
- Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, **157**, 98
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Katz, D., Sartoretti, P., Cropper, M., et al. 2019, *A&A*, **622**, A205
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, eds. F. Loizides, & B. Schmidt (IOS Press), 87
- Lindgren, L., Hernández, J., Bombrun, A., et al. 2018, *A&A*, **616**, A2
- McKinney, W., van der Walt, S., Millman, J. 2010, in *Proceedings of the 9th Python in Science Conference*, 445, 51
- Pancino, E., Bellazzini, M., Giuffrida, G., & Marinoni, S. 2017, *MNRAS*, **467**, 412
- Plotly Technologies Inc. 2015, Collaborative data science (Montréal: Plotly Technologies Inc.)
- Rix, H.-W., Hogg, D. W., Boubert, D., et al. 2021, *AJ*, **162**, 142
- Rybicki, J., Rix, H.-W., Demleitner, M., Bailer-Jones, C. A. L., & Cooper, W. J. 2021, *MNRAS*, **500**, 397
- Sarajedini, A., Bedin, L. R., Chaboyer, B., et al. 2007, *AJ*, **133**, 1658
- Saydjari, A. K., Schlafly, E. F., Lang, D., et al. 2022, *ApJS*, submitted, [arXiv:2206.11909]
- Schlafly, E. F., Green, G. M., Lang, D., et al. 2018, *ApJS*, **234**, 39
- Taylor, M. B. 2005, in *ASP Conf. Ser.*, **347**, 29
- Vanderplas, J., Connolly, A., Ivezić, Ž., & Gray, A. 2012, in *Conference on Intelligent Data Understanding (CIDU)*, 47
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, **17**, 261
- Zonca, A., Singer, L., Lenz, D., et al. 2019, *J. Open Source Softw.*, **4**, 1298

Appendix A: Observed *Gaia* source density and true source density

The crowding limit of the astrometric instrument on board the *Gaia* spacecraft is 1,050,000 sources per square degrees (Gaia Collaboration 2016). In denser regions, the processing pipeline is not able to track all sources transiting through the focal plane and prioritises bright sources. Since the set of discarded faint stars can be different every time *Gaia* scans a given region, the final source density in the *Gaia* catalogue locally reaches values as high as ~ 1.4 million sources per square degree, but the crowding issue artificially sets an upper limit on the catalogue source density, as shown in the left panel of Fig. A.1.

The right panel of Fig. A.1 shows that the parameter M_{10} scales almost linearly with true source density (approximated by the density of DECaPS sources brighter than $r=23$), while the *Gaia* source density is essentially flat for $M_{10} < 20.5$ and never reaches above 1.5 million sources per square degree.

Appendix B: Number of matched transits with magnitude

For the sources brighter than $G \sim 18$, the number of transits matched to a given source is directly related to the number of times a region of the sky was observed by *Gaia*. The number of matched transits decreases for fainter sources, which are not observed at every transits due to the on-board processing prioritising bright sources (top panel of Fig. B.1) until it reaches a minimum value of five. Sources with `astrometric_matched_transits` < 5 are not included in the *Gaia* catalogue. The G magnitude at which this minimum value is reached varies across the sky, as it strongly depends on the level of crowding and on how many times the region was seen by *Gaia*.

In this paper, we note M_{10} the median magnitude of the sources with `astrometric_matched_transits` ≤ 10 , which quantifies the typical magnitude of the sources. This quantity can be computed for any patch of sky and is a good predictor of completeness at all magnitudes.

Appendix C: Location of our reference DECaPS fields

Figure C.1 shows the location of the patches of sky within the DECaPS footprint which were used to calibrated our completeness model.

Appendix D: Sources without a G magnitude

This paper presents a selection function for the sample of *Gaia* DR3 sources with published positions and magnitude G . A small fraction (variable across the sky but on average 0.3% of the *Gaia* DR3 catalogue) do not have an associated G magnitude. We illustrate their distribution in Fig. D.1.

Appendix E: Generalised Sigmoid

Figure E.1 illustrates the generalised sigmoid function $S(G)$ defined in Equation 1 for various choices of its parameters x , y , z .

Appendix F: Model fitting

The data we fit in this paper are the source count ratios of *Gaia* to DECaPS, computed in 40 magnitude bins ($G=15$ to 23 in step of 0.2 mag) in 2906 patches on the sky, representing a total of $n = 40 \times 2906 = 116240$ data points obs_i . The model presented in Sect. 2.4 contains ten parameters $\Theta = (a_x, b_x, c_x, a_y, b_y, c_y, a_z, b_z, c_z, M_{break})$ and predicts the completeness for a given value of M_{10} at a given magnitude G .

Assuming that the data (observed count ratios obs_i) are affected by Gaussian noise, the likelihood of the model (predicting completeness $pred_i$) is:

$$\mathcal{L}(\Theta) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi}\sigma^2} e^{-\frac{(obs_i - pred_i)^2}{2\sigma^2}} \quad (F.1)$$

The corresponding log-likelihood is:

$$\begin{aligned} \log \mathcal{L}(\Theta) &= \sum_{i=1}^n \log \frac{1}{\sqrt{2\pi}\sigma^2} e^{-\frac{(obs_i - pred_i)^2}{2\sigma^2}} \\ &= n \log \frac{1}{\sqrt{2\pi}\sigma^2} - \frac{1}{2\sigma^2} \sum_{i=1}^n (obs_i - pred_i)^2 \\ &= n \log \sqrt{2\pi} - n \log \sigma - \frac{1}{2\sigma^2} \sum_{i=1}^n (obs_i - pred_i)^2 \end{aligned} \quad (F.2)$$

which for a fixed value of σ simplifies to ordinary least squares. Here we keep σ as a free parameter, and maximise the quantity $\log \mathcal{L}(\Theta, \sigma)$. We adopt flat priors on all parameters, only enforcing that σ must be positive. The corner plot of the posterior chain is shown in Fig. F.1. We choose the median of each parameter chain as the best value for each parameter. These values are reported in table 1.

In practice, the noise on the source count ratios varies with M_{10} and G , is unlikely to be Gaussian, and is truncated since the observed completeness is bounded between 0 and 1 by design. The parameter σ must therefore be considered as a nuisance parameter rather than a full model of the noise.




Fig. A.1. Left: distribution of *Gaia* DR3 sources densities computed in 196,608 HEALPix regions level 7 across the whole sky. Crowding artificially sets an upper limit of the catalogue source density. Right: source density as a function of M_{10} in *Gaia* DR3 (black) and DECaPS with $r < 23$ (cyan) in the 2906 patches we use to calibrate our model (see Fig. C.1). Outside the sparsest regions, the *Gaia* source density provides almost no information on the true source density and level of crowding.




Fig. B.1. Distribution of number of matched transits with magnitude and sky position. Top: number of `astrometric_matched_transits` against magnitude for sources in five selected regions. Bottom: magnitude distribution of the sources with `astrometric_matched_transits` ≤ 10 , and their median magnitude M_{10} .




Fig. C.1. Location of the 1085 patches of sky (some of which are further divided into smaller patches for a total of 2906) used to compute *Gaia*-to-DECaPS source ratios and calibrate our model. Their observed (*Gaia*) source densities range from ~ 4000 to ~ 1.4 million sources per square degree. The background map is the integrated *Gaia* flux map (ESA/*Gaia* DPAC).




Fig. D.1. Sources without a listed *G* magnitude. Left: sky distribution of 100,000 randomly selected *Gaia* sources without a *G* magnitude. The area displayed focuses on the inner Milky Way, but such sources are found all across the sky. Right: fraction of sources (out of a random sample of 3 million *Gaia* sources) without a *G* magnitude, computed in bins of Galactic latitude. The fraction is close to 0.3% over the entire sky but is significantly higher in regions close to the Galactic plane.




Fig. E.1. Plots of the generalised sigmoid function defined in equation 1 obtained varying the x , y , and z parameters.




Fig. F.1. Corner plot of the posterior chains of the ten parameters ($a_x, b_x, c_x, a_y, b_y, c_y, a_z, b_z, c_z, M_{break}$) constraining the selection function $S(G|M_{10})$, and the nuisance parameter σ .

Appendix G: Example use of our Python selection function package

Here we show two examples of maps created with the model presented in this study, obtained through our Python package `gaiaunlimited`. The package is hosted on Github⁴. Further examples, including the construction of higher-resolution maps, are available as a Jupyter notebook tutorial⁵. This section is not intended as documentation for the `gaiaunlimited` package. Complete documentation and tutorials can be found through the Github repository and will be kept updated with every release of the code.

Appendix G.1. All-sky map

The following snippet of code produces the all-sky completeness map shown in Fig. G.1.

```
from gaiaunlimited.selectionfunctions import DR3SelectionFunctionTCG_hpx7
dr3sf = DR3SelectionFunctionTCG_hpx7()

from gaiaunlimited.utils import get_healpix_centers
coords_of_centers = get_healpix_centers(7)
import numpy as np
gmag = np.ones_like(coords_of_centers) * 21
print(f'Computing the completeness for {len(coords_of_centers)} points.')
completeness = dr3sf.query(coords_of_centers, gmag)

import healpy as hp
hp.mollview(completeness, coord=['Celestial', 'Galactic'], min=0, max=1, title='completeness at G=21')
```

⁴ <https://github.com/gaia-unlimited/gaiaunlimited>

⁵ <https://github.com/gaia-unlimited/gaiaunlimited/blob/main/docs/dr3-empirical-completeness.ipynb>




Fig. G.1. Figure created by the code provided in Sect. G.1, using our gaiaunlimited package.

Appendix G.2. Magnitude of 99% completeness

The following snippet of code produces the map indicating the magnitude corresponding completeness of 99%, shown in Fig. G.2.

```
from gaiaunlimited.selectionfunctions import DR3SelectionFunctionTCG_hpx7, m10_to_completeness
DR3SelectionFunctionTCG_hpx7()
m10_values = dr3sf.m10map[:,2] #retrieve the value of M_10 in each healpixel

S = 0.99 #we want to compute the magnitude of 99% completeness
g = np.arange(17,22,0.1)
idx = np.array([ (np.abs(m10_to_completeness(g,m)-S)).argmin() for m in m10_values ])
hp.mollview(g[idx],coord=['Celestial','Galactic'],nest=True,min=17,max=21,cmap='inferno',title='Magnitude at
99% completeness')
```


G magnitude at 99% completeness

Fig. G.2. Figure created by the code provided in Sect. G.2, using our *gaiainlimited* package.

Appendix G.3. Grid of coordinates

The following snippet of code produces Fig. G.3.

```
from gaiainlimited.selectionfunctions import DR3SelectionFunctionTCG_hpx7
dr3sf = DR3SelectionFunctionTCG_hpx7()

import numpy as np
from astropy import units as u
from astropy.coordinates import SkyCoord

l = np.linspace(-10, 10, 50)
b = np.linspace(-10, 10, 50)
l, b = [c.flatten() for c in np.meshgrid(l, b)]
coords = SkyCoord(frame="galactic", l=l*u.degree, b=b*u.degree)

import matplotlib.pyplot as plt
plt.figure(figsize=(8,4))
for i,G in enumerate([20.5,21]):
    completeness = dr3sf.query(coords,G*np.ones_like(coords))
    plt.subplot(1,2,i+1)
    plt.scatter(l, b, c=completeness, s=12, marker='s', vmin=0, vmax=1)
    plt.xlabel('$\ell$ (degrees)'); plt.ylabel('$b$ (degrees)')
    plt.title(f'completeness at G={G}')
    plt.xlim(10,-10); plt.ylim(-10,10)
    plt.minorticks_on()
plt.tight_layout()
```




Fig. G.3. Figure created by the code provided in Sect. G.3, using our `gaiaunlimited` package.