



Research Article

Enhanced astrometry of the rapid ASKAP continuum survey for precise localisation of fast radio bursts

Akhil Jaini¹, Adam T. Deller^{1,2}, Yuanming Wang^{1,2}, Emil Lenc³, and Marcin Glowacki^{4,5}

¹Center for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC, Australia, ²ARC Centre of Excellence for Gravitational Wave Discovery (OzGrav), Hawthorn, VIC, Australia, ³CSIRO Space and Astronomy, Epping, NSW, Australia, ⁴International Centre for Radio Astronomy Research (ICRAR), Curtin University, Bentley, WA, Australia and ⁵Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh, UK

Abstract

Fast radio bursts (FRBs) are short, intense radio signals from distant astrophysical sources, and their accurate localisation is crucial for probing their origins and utilising them as cosmological tools. This study focuses on improving the astrometric precision of FRBs discovered by the Australian Square Kilometre Array Pathfinder (ASKAP) by correcting systematic positional errors in the Rapid ASKAP Continuum Survey (RACS), which is used as a primary reference for ASKAP FRB localisation. We present a detailed methodology for refining astrometry in two RACS epochs (RACS-Low1 and RACS-Low3) through crossmatching with the Wide-field Infrared Survey Explorer (WISE) catalogue. The uncorrected RACS-Low1 and RACS-Low3 catalogues had significant astrometric offsets, with all-sky median values of $0.58''$ in RA and $-0.26''$ in Dec. (RACS-Low1) and $0.29''$ in RA and $1.24''$ in Dec. (RACS-Low3), with a substantial and direction-dependent scatter around these values. After correction, the median offset was completely eliminated, and the 68% confidence interval in the all-sky residuals was reduced to $0.2''$ or better for both surveys. By validating the corrected catalogues against other, independent radio surveys, we conclude that the individual corrected RACS source positions are accurate to a $1-\sigma$ confidence level of $0.3''$ over the bulk of the survey area, degrading slightly to $0.4''$ near the Galactic plane. This work lays the groundwork to extend our corrections to the full RACS catalogue that will enhance future radio observations, particularly for FRB studies.

Keywords: Astrometry; methods: statistical; catalogues; radio continuum: general

(Received 18 December 2024; revised 11 March 2025; accepted 14 March 2025)

1. Introduction

Fast radio bursts (FRBs) are millisecond duration single pulses of intense coherent radiation. These elusive bursts are discovered serendipitously and offer a unique window into the universe. In recent years, the localisation of the first FRBs to galaxies far beyond our own has helped to establish their extreme luminosity and provided crucial insights into their potential origins (Chatterjee et al. 2017; Bannister et al. 2019). The detection of such powerful emissions from vast distances poses significant challenges to our understanding of the processes that generate radio waves and the kinds of astrophysical objects capable of producing such short-lived pulses. The most plausible explanation for these bursts is that they originate from neutron stars (Platts et al. 2019), with theoretical and observational evidence pointing towards both magnetars (Lyubarsky 2020; Bochenek et al. 2020; Chime/Frb Collaboration et al. 2020) and high-mass X-ray binaries (Lyutikov, Barkov, & Giannios 2021) as potential sources. This suggests that multiple objects or mechanisms might be responsible for producing these enigmatic bursts. Studying FRBs is also significant because they can uniquely probe the universe's structure by detecting and characterising the highly diffuse intergalactic and circumgalactic media. These bursts traverse entire columns of plasma, reveal-

ing details about their densities, turbulence, and magnetisation. Measurements like dispersion measure (DM) and de-dispersed pulse width also provide valuable insights into the nearly invisible gas that makes up $\sim 90\%$ of baryonic matter in the universe (Macquart et al. 2020).

1.1 FRB localisation

Since the discovery of the first FRB, astronomers have identified over 2000 unique sources of FRBs (Petroff, Hessels, & Lorimer 2022), yet the precise localisation of these events remains a challenge, with just over 60 successfully linked to a host galaxy to date. Accurate astrometry is paramount in unravelling the enigmatic origins of FRBs and leveraging them as potent astrophysical and cosmological probes. For local universe galaxies with redshifts ~ 0.1 , an astrometric precision of a few arcseconds is sufficient to robustly identify a host galaxy, but for farther galaxies at redshifts ~ 1 , we need sub-arcsecond astrometry for pinpointing the exact galactic and stellar environments housing FRB sources.

Precise localisations are crucial for identifying FRB host galaxies, studying their properties and environments, and distinguishing between competing progenitor models. They also enable detailed mapping of the intergalactic medium (IGM) and its magnetic fields, enhancing our understanding of cosmic magnetism and large-scale structure. Furthermore, such localisation facilitates testing fundamental physics and cosmological principles, offering insights into dark matter, dark energy, and the early universe. Combining multi-wavelength observations with statistical

Corresponding author: Akhil Jaini; Email: ajaini@swin.edu.au

Cite this article: Jaini A, Deller A T., Wang Y, Lenc E and Glowacki M. (2025) Enhanced astrometry of the rapid ASKAP continuum survey for precise localisation of fast radio bursts. *Publications of the Astronomical Society of Australia* 42, e060, 1–20. <https://doi.org/10.1017/pasa.2025.28>

analyses of localised FRBs provides valuable constraints on their physical properties and cosmological significance (Bailes 2022).

The fleeting nature of FRBs generally limits the options available in terms of precision calibration of the observations in which they are discovered. Particularly for non-repeating bursts, accurate localisation relies heavily on measuring their positions relative to background radio sources captured in the same interferometric observation. Consequently, by improving the precision of the astrometry for these continuum sources, which are generally catalogued in large-scale telescope surveys, we can enhance the accuracy of FRB localisation.

1.2 ASKAP

The Australian Square Kilometre Array Pathfinder (ASKAP) is an interferometer which comprises of 36, 12 m diameter antennas with a maximum baseline length of 6 km. Each antenna is equipped with novel phased-array feed (PAF) receivers that, together with advanced digital systems, generate 36 simultaneous, independent beams, which collectively cover a ~ 30 sq.deg. field-of-view in the frequency range of 700–1 800 MHz (Hotan *et al.* 2021). These beams are calibrated independently and can produce separate (albeit overlapping) images, but are commonly mosaicked together to form a single large image.

ASKAP produces correlated visibilities with an integration time of 10 s and frequency resolution of 18 kHz. This data product is used for rapid surveying of the sky in continuum and spectral line processing modes. ASKAP's wide field-of-view has helped discover various transient phenomena, including flare stars, long period transients (LPTs), and, most notably, FRBs. However, the utility of ASKAP extends beyond mere detection. The interferometric capabilities of ASKAP allow for sub-arcsecond astrometry, which is crucial not only for precisely localising FRBs but also for other transient phenomena. For flare stars and LPTs, precise localisation is critical for cross-matching against crowded optical catalogues, particularly when only a single detection has been made. This ensures that the sources can be accurately identified and followed up across other wavelengths, allowing astronomers to explore their physical properties and origins comprehensively. Therefore, the results of this work, which enhance the localisation precision of ASKAP detections, are highly relevant to a broader range of data products beyond just FRBs.

1.3 FRB detection and localisation with ASKAP

ASKAP has two FRB detection pipelines which operate commensally with the standard processing. The first one uses the InCoherent Sum (ICS, started in mid-2018, Shannon *et al.* 2024) of antenna powers at a time resolution of ≈ 1 ms and frequency resolution of 1 MHz. Each beam is independently summed and dedispersed. The sensitivity of the ICS processing is proportional to $N_{\text{ant}}^{1/2}$, where N_{ant} is the number of antennas.

The second pipeline, Commensal Real-time ASKAP Fast Transients (CRAFT) COherent (CRACO, Wang *et al.* 2024) Upgrade, commenced operations in late 2023. CRACO can potentially utilise interferometric visibilities with a time resolution of ≈ 1 ms and frequency resolution of 1 MHz. It creates and dedisperses an interferometric image coherently, with sensitivity scaling proportionally to N_{ant} . Currently, CRACO operates with a time resolution of 13.8 ms, meaning its sensitivity is only comparable to ICS for FRBs that are unresolved in time. However, once CRACO

achieves its targeted 1 ms time resolution, the detection sensitivity is expected to improve by a factor of ~ 5 over ICS, significantly enhancing its capability for detecting FRBs.

Both detection systems currently operate in real time. When an FRB is detected, the system saves the voltage data to disk, which is then correlated offline, and interferometric calibration and imaging, as well as high-time resolution processing is performed as described in Scott *et al.* (2023). In addition to imaging the \sim milliseconds of time containing the FRB itself, the entire download (3–12 s) is imaged, producing a radio continuum image of the entire field with an RMS sensitivity of a few milli-Jansky. We measure an astrometric offset by comparing the positions of the sources in this image with a reference catalogue. After measuring the position of the FRB from the \sim millisecond-duration image containing it, we adjust FRB position by the astrometric offset derived from the field sources.

1.4 RACS

The Rapid ASKAP Continuum Survey (RACS) is a comprehensive all-sky survey conducted using ASKAP. It represents ASKAP's pilot survey and aims to create a shallow model of the ASKAP sky to aid in the calibration of future deep ASKAP surveys and other scientific research. RACS covers the entire southern sky below the Dec. of $+49^\circ$ across the full ASKAP band of 700–1 800 MHz (McConnell *et al.* 2020). This survey aims to provide a comprehensive model of the radio sky accessible to ASKAP to modest depth, but significantly better than previous all-sky surveys in the southern hemisphere. The RACS project encompasses the entire sky visible from the ASKAP site in Western Australia, and the initial data release includes 903 images providing information for over 3 million continuum sources, and serves as a critical reference for subsequent astronomical investigations.

RACS source lists are created for each observing run, known as an SBID (Scheduling Block Identifier). For each SBID, a mosaic image (also known as a tile image) is produced by combining data from 36 individual ASKAP beams, which overlap to cover a large area of the sky. From these tile images, source lists of detected radio sources are extracted, with key parameters like flux density, position, and source morphology. These source lists are then compiled into catalogues after undergoing further checks (as described in Hale *et al.* 2021).

Each RACS field (also referred to as scan) has a field-of-view of approximately 31 sq. deg. Each observation within RACS is made for 15 min and employs a contiguous 288 MHz band (unless mentioned otherwise) within the broader 700–1 800 MHz range, with the first data release centred at 887.5 MHz. The survey achieves an angular resolution of approximately $15''$ and aims for a point-source detection sensitivity of around 1 mJy. RACS offers significant advantages over existing radio surveys like the National Radio Astronomy Observatory's Very Large Array (NRAO VLA) Sky Survey (NVSS, Condon *et al.* 1998) and the Sydney University Molonglo Sky Survey (SUMSS, Sadler 1999), including greater depth, superior spatial resolution, sensitive wide-band coverage in the intermediate frequency regime for studying broadband spectra and transient behaviour, and publicly accessible data products such as total-intensity images, polarisation data, and a comprehensive source catalogue. RACS is also the only survey to uniformly cover the entire Southern sky at high angular resolution. Due to these advantages and overlapping sky coverage with other ASKAP data products, the positions of FRBs discovered with ASKAP are

Table 1. RACS observing details.

Epoch	Year	Central frequency (MHz)	Bandwidth (MHz)	Number of scans	Footprint	Reference
RACS-Low1	2019–2020	887.5	288	903	<i>square_6x6</i>	McConnell et al. (2020)
RACS-Low3	2023–2024	943.5	288	1 493	<i>closepack36</i>	— ^a
RACS-Mid1	2020–2021	1 367.5	144	1 493	<i>closepack36</i>	Duchesne et al. (2023)
RACS-High1	2022–2023	1 655.5	~ 200	1 493	<i>closepack36</i>	Duchesne et al. (2024)

RACS-Low3 does not have a dedicated publication but follows similar methodologies to those outlined in Hale et al. (2021) and Duchesne et al. (2024). Its documentation is available on the RACS website at: research.csiro.au/racs/home/survey/comparison/ and research.csiro.au/racs/racs-low3-dr1-raw-data/.

Table 2. Survey characteristics of VLA FIRST, VLASS, RFC and NASA WISE.

Survey	Wavelength/ Frequency	Angular resolution	Sky coverage	Number of sources	Key features
VLA FIRST	1.4 GHz	5 arcsec	~10 000 sq. deg. (Galactic caps)	~900 000	Radio survey focusing on extragalactic objects like quasars, radio galaxies
VLASS	2–4 GHz	2.5 arcsec	33 885 sq. deg. (North of declination –40 deg.)	~5 000 000	Designed for transient radio phenomena such as supernovae, black hole jets, etc.
RFC	Various (VLBI)	Milliarcsec	Entire sky	~22 000	Provides astrometric data, used for differential astrometry, phase calibrations
NASA WISE	3.4, 4.6, 12, 22 microns	6.1 - 12 arcsec	Entire sky	~563 000 000	Infrared sky survey covering asteroids, star formation, galaxy structure

corrected using the continuum source positions provided by the RACS survey.

1.4.1 RACS bands

The ASKAP observing band is split into three bands and RACS covers each of these bands (abbreviated as RACS-Low, RACS-Mid and RACS-High) in dedicated epochs, as detailed in Table 1.

1.4.2 Astrometry in RACS

RACS employs an observing setup where multiple target scans share a single common scan on a bright calibrator to provide initial phase, gain, and bandpass solutions (McConnell et al. 2020). However, any residual time- or direction-dependent errors between the calibrator and target scans can introduce astrometric errors into the initial RACS field model since no further phase calibration was applied prior to imaging and self-calibration. These errors can thus, propagate through to the final RACS catalogue positions. When compared to the International Celestial Reference Frame (ICRF), RACS sources exhibited notable systematic position offsets of magnitudes $\sim 0.6''$ in right ascension and $\sim 0.4''$ in declination (McConnell et al. 2020), and have an astrometric precision of less than $0.8''$.

In this paper, we address and correct these astrometric uncertainties, ensuring that ASKAP-discovered FRBs can be localised with high precision, thus facilitating more accurate scientific investigations.

1.5 Astrometry in reference catalogues

To test and correct the astrometry of RACS, we need to compare it with another catalogue. Ideally, this catalogue would have the same spatial resolution as RACS, be conducted at the same

radio frequency, cover the entire RACS footprint, and have negligible astrometric errors. However, no such catalogue exists. Instead, we make use of several wide-field catalogues, each with its own advantages and disadvantages. We consider four potentially useful catalogues for both the astrometric correction and validation of RACS. These are the VLA FIRST (the Very Large Array Faint Images of the Radio Sky at Twenty-cm, Becker et al. 1994), VLASS (the Very Large Array Sky Survey, Lacy et al. 2020), RFC (the Radio Fundamental Catalogue,¹ Petrov & Kovalev 2024), and NASA WISE (the National Aeronautics and Space Administration’s Wide-field Infrared Survey Explorer, Wright et al. 2010). The important features of these catalogues are detailed in Table 2.

The sources catalogued by these surveys are predominantly galaxies, which provide well-defined radio and/or infrared emissions, allowing accurate positional cross-matching. For any individual galaxy being cross-matched, the goal is to identify a corresponding source in the reference catalogues and compare the positions. Astrometry in these catalogues is, therefore, critical for accurately determining the positions of celestial objects. The FIRST survey, conducted with the VLA, has systematic astrometric errors of less than $0.05''$ and the source positions are measured to a $1\text{-}\sigma$ statistical precision of $0.5''$ at the survey threshold of 1 mJy, and much smaller for brighter sources. However, it covers only a small portion of the sky.

VLASS is a more recent survey using the VLA, at a higher frequency and a higher angular resolution. Nominally, this should lead to a statistical astrometric precision of $0.25''$ for brighter sources, but due to imperfect widefield imaging corrections, the current Quick Look (QL) images have an astrometric precision

¹Website: astrogeo.org/rfc/.

ranging from $0.5''$ to $1''$. However, it only covers a part of the RACS sky, and the higher frequency of observation means the sources are more dominated by AGN cores whereas RACS-Low tends to pick up more radio lobes. The published catalogues of VLASS are also known to have a systematic astrometric error of $\approx 0.25''$ in declination (Gordon et al. 2021).

RFC is a database of compact radio sources offering milliarc-second (mas) precision in source positions, achieved through Very Long Baseline Interferometry (VLBI). VLBI combines data from widely separated radio telescopes to form a virtual telescope with an effective aperture equal to the distance between them, enabling very high angular resolution. While RFC provides highly accurate astrometry, it only contains data on $\sim 22\,000$ sources (Petrov & Kovalev 2024). Additionally, some of these sources exhibit extended structure on larger scales (\sim arcseconds) that are resolved out by the VLBI observations. This means that the milliarcsecond positions provided by VLBI do not always represent the centroid of the larger-scale emission and may not align perfectly with the ASKAP positions.

Finally, being a space-based infrared survey, WISE is not subject to the atmospheric and ionospheric propagation effects that perturb instrumental effects as the other radio interferometric surveys in this study, and the astrometric accuracy is typically better than $0.2''$, but it is observed in a completely different band (which means WISE primarily picks up AGN cores and not radio lobes like RACS-Low).

Given these constraints, we chose to use WISE as the primary reference catalogue and while different frequency bands can cause centroid position offsets between WISE and RACS-Low for individual sources, since the Universe lacks a preferred direction, these positional offsets are statistically expected to average out across a large sample of sources within a RACS beam (as discussed in Section 2.1).

Section 2 describes the techniques employed in this project, including the modelling parameters and their explanations, the governing equations and the main assumptions. Section 3 presents the quantitative results, with separate subsections for results related to different objectives and different special cases. In Sections 4 and 5, we discuss the future scope of the project, highlighting the practical possibilities and related considerations, and end with concluding remarks.

2. Methodology

This section elucidates the techniques employed in this project. The entire modelling pipeline² was developed in Python version 3.11, incorporating relevant modules such as *astroquery* (Ginsburg et al. 2019), *astropy*, *scipy*, and *numpy* (Harris et al. 2020).

2.1 Modelling parameters for RACS-Low1

The aim of this work is to derive accurate astrometric corrections for every source in RACS-Low. Directly correcting the errors in each individual source position based on their offsets between WISE and RACS is not feasible: intrinsic differences in their centroid positions would result in new sources of error being added that would likely outweigh the corrections being made. Instead,

we developed a model to estimate the underlying astrometric corrections by averaging across as many sources as was reasonable, ensuring that we minimise introducing any additional noise.

To generate the initial offsets for RACS-Low1, we conducted a comparison of the sources within each beam of the available RACS-Low1 scans against the WISE catalogue. Out of the total 903 scans, only 805 were publicly available, with those located above Dec. $+30^\circ$ being excluded. These scans were available as a sources-only catalogue and a Gaussian components catalogue, where the latter categorised individual Gaussian-shaped components of extended sources, which is crucial for detailed studies of complex radio sources that are not well-represented by a single point source. We chose to use the Gaussian components catalogue for our work.

The different observing frequencies mean that the centroid positions in infrared (WISE) and radio (RACS-Low) may differ for any individual source. However, since the Universe lacks a preferred direction, the positional offsets caused by differences in centroid placement across the electromagnetic spectrum are expected to average incoherently over the large sample of sources within a RACS beam, thus minimally impacting the overall astrometric correction (e.g. for a field with 100 crossmatched sources each having a typical offset of $1''$, the mean residual error will be reduced to the $0.1''$ level). Nevertheless, since centroid differences could be expected to be larger on average for significantly extended sources, we implemented a filtering algorithm designed to exclude obvious extended sources. This algorithm identifies point-like sources by ensuring that the ratio of integrated flux to peak flux remains below a predetermined threshold (of 1.5, in this case), and ensuring that the sources do not have multiple Gaussian components within $30''$ of each other. This would filter out RACS sources that have extended radio lobe components or that are compact doubles, where the WISE source would not be associated with either of the RACS components. Additionally, we filtered out sources with a signal-to-noise ratio (SNR) below 6σ . Collectively, these steps ensure that only isolated, point-like sources, having a minimum SNR of 6σ , were included in the astrometric corrections. Through these measures, we effectively minimised the impact of observing frequency-based centroid discrepancies on the mean positions of sources.

In order to estimate the systematic position error in the catalogue at a given position (generally, centred on the location of an individual ASKAP beam during a given RACS scan), we must select an ensemble of sources to compare against the reference catalogue. We select all sources within a specified radius of the beam centre, where the choice of radius balances statistical precision (a larger radius yields more sources and hence higher precision) against fidelity (a smaller radius minimises contamination of the offset by the mosaicking process that combines data from adjacent beams – see Section 2.2 for more details). We tested radii between 0.5° and 2° , finding that a 0.5° radius yielded too few crossmatches, while a 2° radius introduced excessive overlap with neighbouring beams. Ultimately, a 1° radius provided the best compromise, offering sufficient crossmatches without significant overlap, and was adopted as the standard.

For each 1° beam, we applied a crossmatching threshold of $5''$, chosen to minimise false positives while accommodating typical astrometric uncertainties. With an average WISE source density of approximately 13 500 sources per square degree, this threshold was expected to perform well, as the average separation between sources is about $31''$. While this strategy worked effectively for

²GitHub repository: github.com/jainiakhil/RACS_Astrometry.

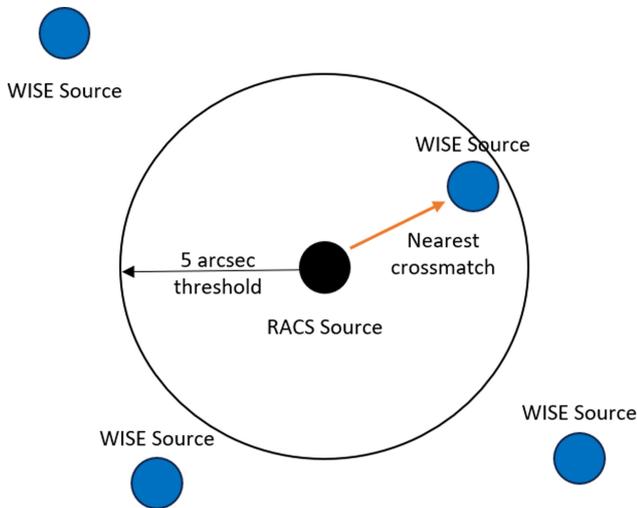


Figure 1. Graphically showing the crossmatching of a RACS source to a reference source (WISE, in this case) that is within 5" of its position. If there are more than 1 reference sources within the 5" threshold, the corresponding RACS source is not taken into consideration to avoid ambiguity in the calculations.

off-plane regions, the higher source density in the Galactic plane, combined with the prevalence of extended sources, resulted in increased false positives and decreased reliability.

The crossmatches were performed using *astropy's* built-in `match_to_catalog_sky` function (Astropy Collaboration et al. 2013, 2018, 2022). If multiple WISE sources fell within the crossmatching threshold of a single RACS source, we excluded that RACS source to avoid ambiguity in source identification (as shown graphically in Fig. 1). This ensured that each crossmatch was a unique pairing, preventing errors in astrometric corrections. Given the high source density in some regions, this step was crucial to maintain precision, and since many sources survived this cut, discarding a few at this stage did not impact the overall analysis.

Quantitatively, each 1-degree radius RACS beam contains around 200 sources. The filtering algorithm described above eliminates roughly 30–50% of these, reducing the count to about 120 sources eligible for crossmatching. Once crossmatched, an additional 30–40% of sources are typically removed, resulting in around 80 matched sources per beam. However, in the Galactic plane regions, the filtering process removes a larger fraction of sources due to increased density and complexity, resulting in fewer crossmatched sources.

This procedure was systematically applied to every RACS source within the entire sky coverage, for each beam, and for each scan, resulting in a comprehensive list of mean offsets and associated uncertainties for RA and Dec. positions.

Fig. 2 illustrates the beam-wise offsets observed in a specific RACS scan compared to WISE. These offsets range from a few milliarcseconds to several arcseconds, consistent with the results presented in McConnell et al. (2020), thereby exacerbating the localisation uncertainty of all past ASKAP FRBs.

In order to model the mean offsets as accurately as possible, we had to break the RACS dataset into smaller subsets having a common offset. We knew that the calibration process for RACS can introduce time- and direction-dependent calibration errors (as described in Section 2.7 of McConnell et al. 2020). Notably, the bandpass calibration (which is undertaken once per beam on a

very bright calibrator) can lead to a common error across all scans using the same calibration, referred to as the *scan-independent beam offset* (Δ^i), which represents the astrometric error in the direction of the calibrator at the time of the calibrator scan. Additionally, calibration terms that vary with direction and/or time (since subsequent scans can be many hours and many tens of degrees away from the bandpass calibrator) can lead to errors that are approximately common to all beams within a given scan, referred to as the *beam-independent scan offset* (Δ^j). We note that Δ is used throughout this manuscript to refer to an offset, but this can be for an individual source, the average of an ensemble of sources, or as a fitted model parameter – sub- or super-scripts are used to differentiate between these cases.

The total astrometric correction for any beam in a given scan can be defined:

$$\Delta_{modelled}^{ij} = \Delta_{modelled}^i + \Delta_{modelled}^j \tag{1}$$

$$\Delta_{residual}^{ij} = \Delta_{observed}^{ij} - \Delta_{modelled}^{ij} \tag{2}$$

where, $\Delta_{observed}^{ij}$ is the observed offset for the i -th scan and j -th beam, $\Delta_{modelled}^i$ is the *beam-independent scan offset* and $\Delta_{modelled}^j$ is the *scan-independent beam offset*, and $\Delta_{residual}^{ij}$ is the residual offset after corrections, in both RA and Dec. The offset modelling was performed using the built-in *scipy* function `optimize.minimize` (Virtanen et al. 2020). Fig. 3 provides an example of the offset modelling for a specific scan. This modelling process was subsequently applied to all RACS-Low1 scan sets.

Next, we introduced an error floor when calculating the uncertainties in RA and Dec. The uncertainty in the estimated offset for each coordinate was determined by taking a weighted sum of the observed uncertainties of both the target source and the reference source. For very bright sources, the RACS and WISE positions can be measured with very low uncertainties; however, the offset between their centroids might not be zero even for such sources. This non-zero intrinsic offset can arise due to differences in frequency coverage, resolution, and any subtle morphological mismatches between the RACS and WISE catalogues. In contrast, the *RACS position error* leads to systematic offsets that would average out over an ensemble of sources when comparing their RACS and WISE positions.

To account for this intrinsic offset and avoid underestimating the positional uncertainties for bright sources, we added a floor value in quadrature to the statistical offset uncertainty:

$$\sigma_{\Delta} = \sqrt{\sigma_t^2 + \sigma_r^2 + f^2} \tag{3}$$

where σ_t and σ_r are the catalogued uncertainties of the target and reference sources respectively, and σ_{Δ} is the final positional uncertainty of the source offset Δ after adding a floor value, f , in quadrature. The optimal floor value was established by a goodness-of-fit analysis, calculating the reduced chi-squared statistic for observed and residual offsets across several randomly selected scans using different error floor values ranging from 0" to 0.5". A floor value of 0.4" was found to yield a reduced chi-squared value closest to 1, indicating an optimal fit. This floor value of 0.4" was also validated through comparisons between RACS-Low1 and other surveys, such as FIRST and VLASS, which showed similar results. This approach ensured that any non-zero systematic contributions to the positional uncertainties were accurately captured, preventing spurious corrections that might arise from underestimating the uncertainties in the individual source positions.

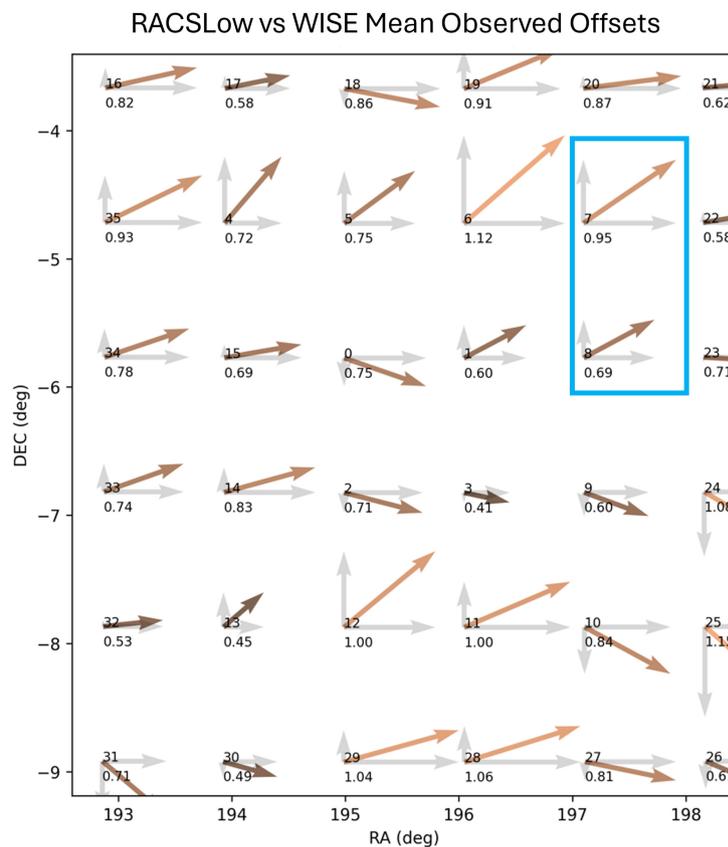
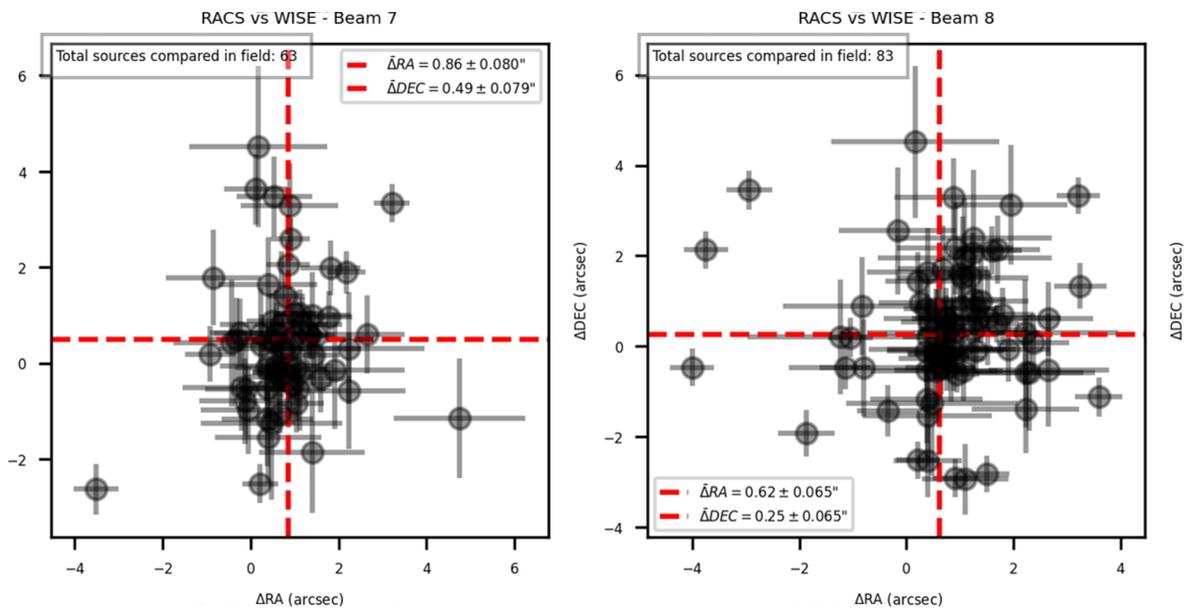


Figure 2. The mean observed offsets for a single scan are shown here. The top image shows the crossmatched offsets for each source for beams 7 and 8 of the 36 beams which are then used to calculate the mean offsets and uncertainties, with the x- and y-axes representing RA and Dec. offsets in arcsec respectively; while the bottom image shows the quiver plot showing the mean offset magnitudes and directions for all 36 beams, with the corresponding beams 7 and 8 highlighted.

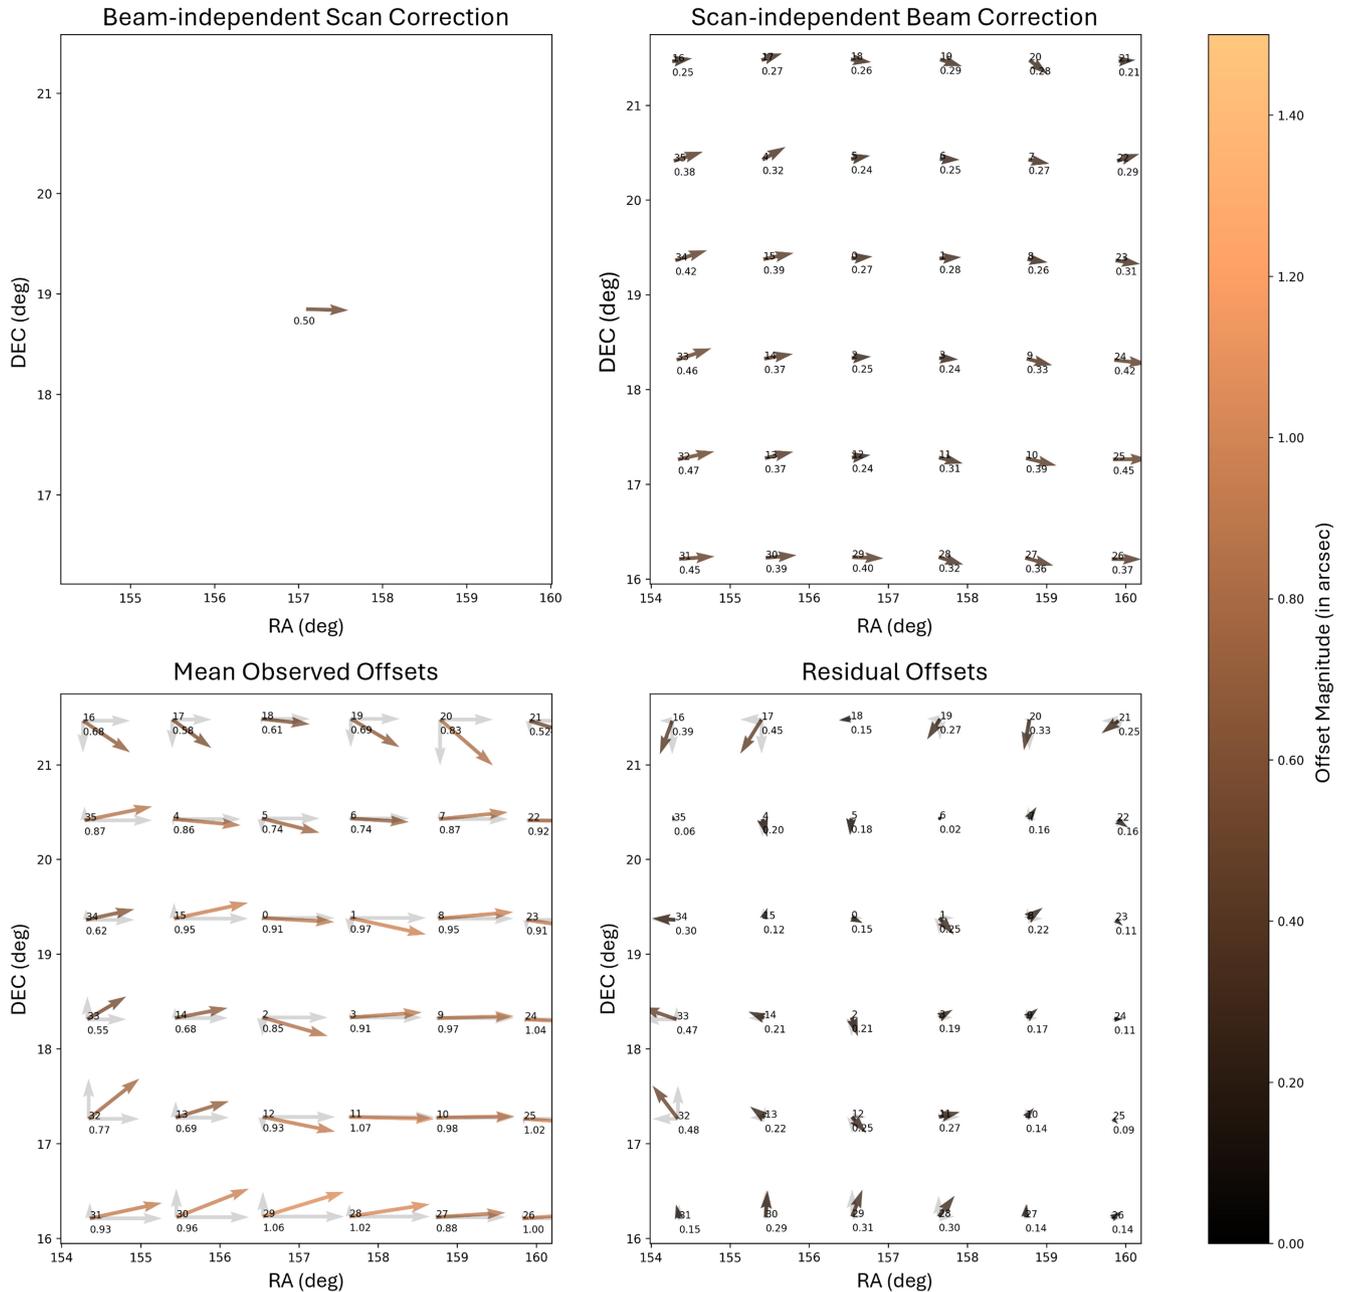


Figure 3. The observed, modelled and residual beam offsets for a particular RACS-Low1 scan (1028+18A) are shown here. The top-left plot shows the beam-independent scan offset modelled for the particular scan, the top-right plot shows the scan-independent beam offset modelled for all scans having the same bandpass calibration. The bottom-left plot shows the observed mean offsets for the scan, and the bottom-right plot shows the residual offsets after subtracting the sum of the beam-independent scan offset and the scan-independent beam offset from the observed mean offsets. As can be inferred visually and by referring to the offset magnitudes printed on the plots, the residual offsets are significantly lower than the observed offsets.

These crossmatching and modelling steps were reiterated for the residual offsets, now reducing the crossmatching threshold to $2''$. This refinement was aimed at eliminating any false positive crossmatches, further improving the precision of our modelling. By applying this more stringent threshold, we ensured that the remaining offsets reflected only genuine positional discrepancies, leading to a more accurate astrometric correction across the entire RACS sky coverage.

With approximately 600 million WISE sources in the sky compared to around 3 million RACS sources, only $\sim 0.5\%$ of WISE sources have a RACS counterpart. Assuming a uniform source distribution across the entire sky, the probability of a false crossmatch for any RACS source with a $5''$ matching radius is $\sim 3\%$. This implies that for every 100 crossmatched sources, approximately 3 are likely spurious. By further refining the crossmatching radius to $2''$, the rate of spurious matches drops to

under 0.4%, an insignificantly low number that does not meaningfully impact our corrections. However, in regions on the Galactic plane, where source density is notably higher, the probability of false matches increases substantially, which likely contributes to the reduced reliability of our corrections in those areas.

The original RACS-Low1 catalogue was obtained using the methods described in Hale et al. (2021). We generated a new catalogue with the corrected positional values by updating the RA and Dec. of each source, using positional corrections of up to three of the nearest beams within a 1-degree radius. The corrections were applied based on a weighted system, where closer corrections were given higher weight (as shown in Equation 4) in determining the final position of the source. This weighting system ensured that positional shifts were as precise as possible, minimising the impact of large, spurious offsets from more distant corrections. This method was employed to produce both the corrected source catalogue and the Gaussian component catalogue.

$$\overline{\Delta}_{final} = \frac{\sum_{i=1}^n \frac{1}{\theta_i} \Delta_i}{\sum_{i=1}^n \frac{1}{\theta_i}}, \quad \text{where } n \leq 3 \quad (4)$$

where, $\overline{\Delta}_{final}$ is the final weighted-average offset correction for a given source in the catalogue, Δ_i are the beam offset corrections nearest to the source, and θ_i is the angular separation of the i -th beam correction from the source.

The final residual offsets were compared with FIRST, VLASS and RFC as a verification check. Each of these surveys was chosen for a different reason, as previously mentioned. The results for the entire survey are consolidated in Section 3.1.

2.2 Modelling parameters for RACS-Low3

The methodology discussed in the previous Section 2.1 is further modified, improved and then extended to the new RACS-Low3 scans as well as the other RACS bands, which are detailed below.

We used the post-processed, publicly available version of the RACS-Low1 catalogue, where the scan tiles were convolved to a common resolution of 25'' and mosaicked with neighbouring tiles. While this provided uniform resolution across the sky coverage, it introduced complications for astrometric modelling. Specifically, beams located at the edges of these tiles contained mosaicked source information from adjacent tiles, leading to positional uncertainties. These edge effects caused the astrometric offsets derived from RACS-Low1 to be less reliable in regions where the mosaicking blended sources from multiple tiles, diminishing the overall efficiency of the modelling process.

In contrast, the RACS-Low3 observatory project provided us with un-mosaicked source lists for each individual beam for each scan across the entire sky coverage. This was an improvement, as we were able to directly model the true positional offsets for each beam in each scan without contamination from neighbouring scans. The un-mosaicked data ensured that the offsets derived were accurate representations of each beam's intrinsic position, unpolluted by artefacts arising from neighbouring observations. This, combined with advancements in ASKAP instrumentation and refined observation techniques between the time of RACS-Low1 and RACS-Low3, we expected would result in better overall astrometry in RACS-Low3.

RACS-Low3 consists of 1 493 unique scans in the lowest frequency band (RACS-Low). It differs from RACS-Low1 in several key aspects. Firstly, it employs a more compact beam footprint,

known as *closepack36*, as opposed to the *square_6x6* configuration used in RACS-Low1 (as detailed in Section 9.10 of Hotan et al. 2021). This updated footprint aligns RACS-Low3 observations with those of RACS-Mid1 and RACS-High1, enabling improved consistency across RACS epochs for different ASKAP bands. Additionally, RACS-Low3 uses a slightly higher observing frequency of 943.5 MHz compared to the 877.5 MHz used in RACS-Low1, where the ASKAP sensitivity is improved. In its raw state, each RACS-Low3 beam image is a square 4 × 4 degrees in size, capturing not only the main lobe of each PAF beam, but extending well out into the adjacent sidelobes. This is true for the pre-mosaicked beam images for all RACS epochs.

In our preliminary analysis of the RACS-Low3 data, we observed significantly larger and more erratic positional offsets in scans above Dec. +30° and within the Galactic plane. While the issues within the Galactic plane were anticipated due to reasons explained in Section 2.1, for regions above Dec. +30°, ASKAP's performance was likely degraded due to observations being conducted at very low elevations, where the positional offsets exceeded 7'' in some cases. These discrepancies can be attributed to three primary factors: the highly elongated point spread function (PSF) at low elevations, degraded observing conditions due to increased atmospheric distortion, and heightened solar activity impacting ionospheric stability. This resulted in missed cross-matches due to the limiting crossmatching threshold of our algorithm, which then affected the offset modelling. Better results at these declinations could be achieved by using a higher crossmatching radius. However, this approach would also introduce more false positives, reducing the reliability of the crossmatches. To address this, we plan to explore a declination-dependent cross-matching scheme as part of our future work (as discussed briefly in Section 4.3).

Therefore, a crucial modification we implemented for the RACS-Low3 modelling process was to update the code to exclude both the Galactic plane and areas above Dec. +30° from the primary modelling process, that is, scans in this declination range did not contribute to the calculation of the scan-independent beam offsets. The exclusion of these problematic regions from the main modelling pipeline allowed us to focus on the parts of the sky where offsets followed more predictable patterns. This strategic omission improved the overall robustness and accuracy of the offset models, leading to more reliable corrections for the rest of the RACS-Low3 sky coverage. We adopted a different approach for scans in these regions, using the scan-independent beam offsets generated from the unaffected regions in the previous modelling steps and applying these to model the beam-independent scan offsets for the problematic areas. This approach allowed us to still model offsets in these regions without compromising the integrity of the broader dataset.

Since producing a mosaicked, uniformly corrected all-sky catalogue from the RACS-Low3 per-beam source lists using methods similar to those outlined in previous works such as Hale et al. (2021) is beyond the immediate scope of our project, we tailored our workflow to meet the specific requirements of this study. We focused on combining the per-beam source lists, but with the updated RA and Dec. positions derived from the corrections generated for each source. This approach allowed us to refine the astrometry for the sources in RACS-Low3 in a way that directly addresses our scientific objectives. In a future work, we plan to make these corrected source lists available publicly, which can then be used to generate fully corrected RACS-Low3 catalogues.

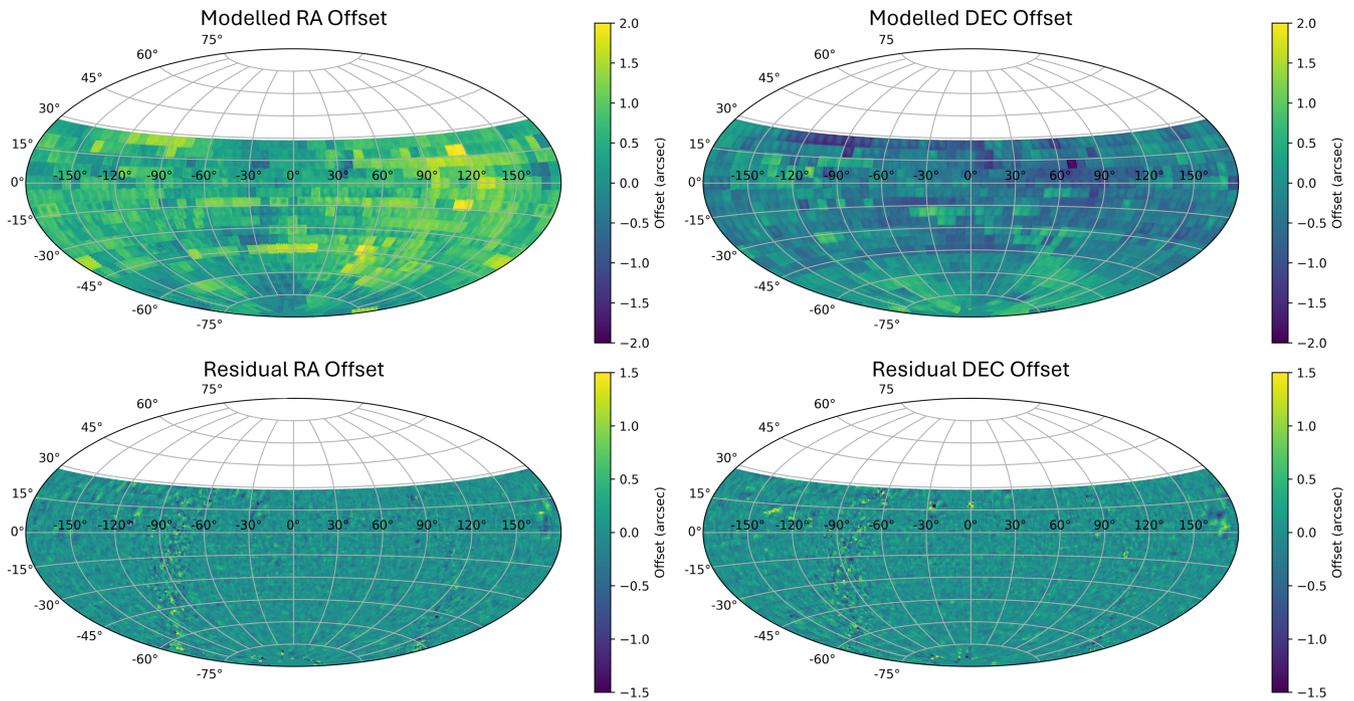


Figure 4. The modelled and residual offsets of RACS-Low1 vs WISE are shown here. The top row shows the modelled offsets and the bottom row shows the residual offsets in RA (left) and Dec. (right) for the entire sky coverage.

3. Results

3.1 RACS-Low1 corrections

The observed and residual RACS–WISE offsets after corrections are shown in Aitoff projection plots in Fig. 4 for ~ 1.7 million crossmatched sources in 28 980 beams. The residual offsets in both RA and Dec. remain uniformly small across the entire sky coverage, with a few exceptions. These outliers are mainly located in the Galactic plane or in scans affected by very bright sources that contaminate the entire sky coverage or by observational data that is irreparably flawed (a few examples are discussed in Section 4.1). The median estimated uncertainty for both RA and Dec. is $0.12''$, and is calculated using a weighted mean of the documented positional errors from both RACS-Low1 and WISE. However, this method assumes that the positional correction is uniform across the entire sky coverage and that the RACS-Low1 and WISE source centroids align perfectly for every source, neither of which will be strictly true. As a result, while these values suggest that most RACS sources will have positional uncertainties below $0.3''$, they underestimate the true positional uncertainties. This is further discussed in Section 3.2.3.

To quantify these findings further, Fig. 5 provides histogram plots of the same parameters, with the median and 68% confidence interval values indicated. As the histograms illustrate, the mean bias is completely eliminated in both RA and Dec. while the stochastic variation is also greatly reduced; the 68% confidence intervals for the residual offset are reduced from $0.58''^{+0.43}_{-0.44}$ to $0.00''^{+0.16}_{-0.17}$ in RA and from $0.26''^{+0.45}_{-0.40}$ to $0.00''^{+0.17}_{-0.17}$ in Dec.

These results demonstrate that we have significantly improved the astrometry of RACS-Low1. To better quantify the residual errors and determine the appropriate uncertainty to assign to the

corrected RACS catalogues, we now compare the corrected RACS positions with other suitable reference catalogues.

3.2 Verification of corrections of RACS-Low1 using other catalogues

We performed cross-verification of the RACS-Low1 astrometric corrections using comparisons with FIRST, VLASS, and RFC, and the results are tabulated in Table 3.

3.2.1 RACS-Low1 vs FIRST

We selected the VLA-FIRST catalogue for verification checks despite its limited coverage of the RACS sky because it is the closest in both radio frequency and angular resolution of the available options. Therefore, while we cannot use FIRST to assess any spatial variations in the residual RACS-Low1 error, it is the most suitable catalogue to evaluate the performance accurately in the regions it is available. By crossmatching RACS-Low1 sources with FIRST in the same beam-wise bins used for WISE, we generated histograms of the offsets in RA and Dec. Only ~ 5000 RACS-Low1 beams overlap with FIRST, with each beam having an average of 50 crossmatched sources. These histograms, shown in Fig. 6, include offset distributions both before and after the corrections. Clear improvements can be seen from this figure and Table 3, and the histograms also exhibit a more Gaussian distribution, which is the expected result of well-behaved astrometric corrections. These results will factor in to our estimate of the residual RACS-Low1 uncertainties in the off-plane region. Finally, we note that since the RACS–FIRST offsets are calculated for ensembles of sources across a small but finite patch of sky, these values may still slightly underestimate the true uncertainty for any individual RACS-Low source.

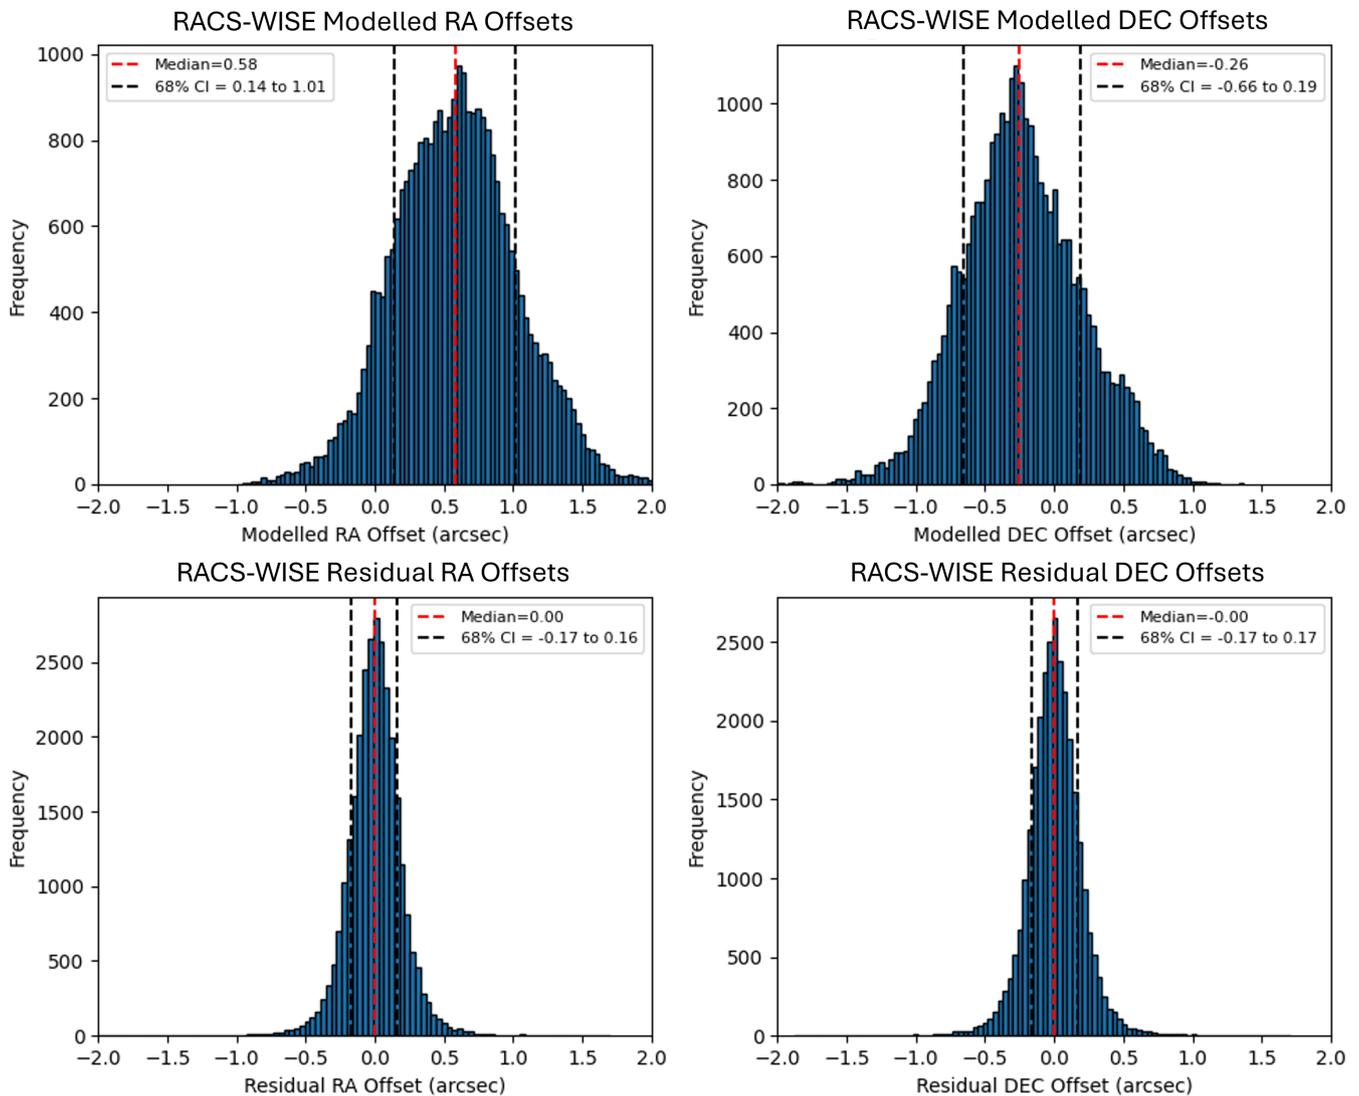


Figure 5. The modelled (top row) and residual (bottom row) offsets for RACS-Low1 vs WISE. The median RA offset (left) improves from $0.58''$ to $0.00''$ and the median Dec. offset (right) decreases from $0.26''$ to $0.00''$. The 68% confidence intervals also narrow down significantly.

3.2.2 RACS-Low1 vs VLASS

Next, we compare RACS-Low1 with VLASS, that has a wider sky coverage than FIRST (down to a declination of -40°). The higher observation frequency of VLASS will mean that for any sources that exhibit frequency-dependent structure, the centroid offsets from RACS-Low1 will be larger than was the case for FIRST. Moreover, the VLASS catalogue available on VizieR (Ochsenbein, Bauer, & Marcout 2000), which we used for verification, has known inherent all-sky offset of approximately $0.25''$ in declination (as documented in Section 3.3 of Gordon *et al.* 2021), which we corrected prior to our cross-matching. However, as can be seen in Figure 8 of the same article, there's a clear direction dependence on the Dec. offset, but it is unclear to what extent the spatial variations of VLASS residual errors have been mitigated by this single all-sky correction we employed (meaning that our single all-sky correction is likely to leave some residual spatially-dependent errors). $\sim 20\,000$ RACS-Low1 beams overlap with VLASS, with each

beam having an average of 50 crossmatched sources. The median offsets and 68% confidence intervals are shown in Fig. 7 and Table 3.

We further divide the dataset into on-plane and off-plane regions, to separate the more complex Galactic Plane region between Galactic Latitudes of -10° and $+10^\circ$, where we know the reliability of our model is lower. As can be seen in Table 3, results are similar to FIRST in the off-plane regions. In the Galactic plane region, while the median RA and Dec. values are mostly similar, we see the histogram spread worsening in both RA and Dec., which shows that our model performs worse in these regions, as expected. As with the comparison to FIRST, the fact that ensemble source averages are used to derive these offsets means that the true residual RACS position uncertainty may be higher.

3.2.3 RACS-Low1 vs RFC

Finally, we use RFC which is a VLBI all-sky catalogue with high astrometric precision. However, as it is restricted to relatively bright sources that are compact on mas scales, the source density is

Table 3. Astrometric offset comparison for corrected RACS-Low1 with FIRST, VLASS, and RFC.

Survey comparison:	Total sources	Median RA offset	68% confidence interval	Median Dec. offset	68% confidence interval
Corrected RACS-Low1 vs	Crossmatched	(arcseconds)	(RA, arcseconds)	(arcseconds)	(Dec., arcseconds)
FIRST	~ 250 000	0.02	-0.18-0.23	0.00	-0.21-0.22
VLASS (Total)	~ 1 000 000	0.04	-0.18-0.27	-0.05	-0.30-0.21
VLASS (Off-Plane)	~ 900 000	0.03	-0.18-0.26	-0.05	-0.29-0.20
VLASS (On-Plane)	~ 100 000	0.08	-0.17-0.37	-0.06	-0.36-0.23
RFC (Total)	~ 12 000	0.04	-0.25-0.34	-0.03	-0.34-0.25
RFC (Off-Plane)	~ 11 000	0.04	-0.25-0.33	-0.03	-0.33-0.25
RFC (On-Plane)	~ 1 000	0.03	-0.29-0.41	-0.05	-0.43-0.26

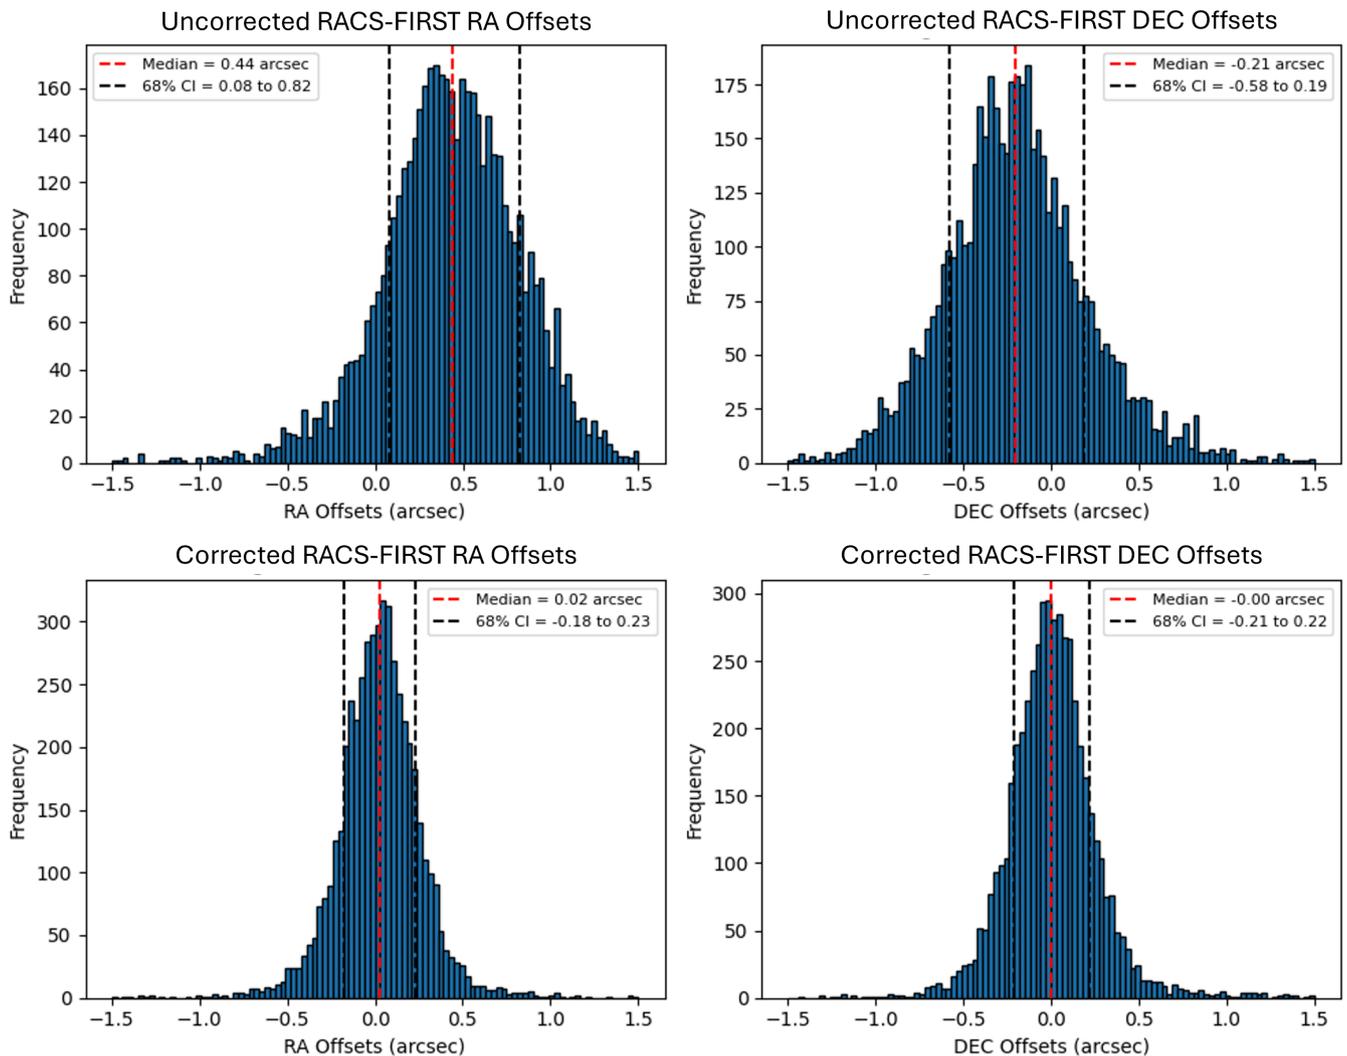


Figure 6. The uncorrected (top row) and corrected offsets (bottom row) in RA (left) and Dec. (right) for RACS-Low1 vs FIRST. The median offset values in RA and Dec. reduce from 0.44" and -0.21" to 0.02" and 0.00" respectively.

much lower than FIRST or VLASS. Due to this limitation, we deviate from the beam-wise approach used earlier and directly compare RACS-Low1 sources with RFC sources. Given this change in approach and RFC's higher frequency and resolution, the differences in centroid positions for extended sources would prevent

per-source centroid offsets from averaging out, thereby negatively affecting our verification of RACS-Low1 corrections. As RFC does not have any information about the extent of the sources, we first crossmatch it with VLASS, in the region it is available, and filter out all the sources that correspond to extended sources in VLASS.

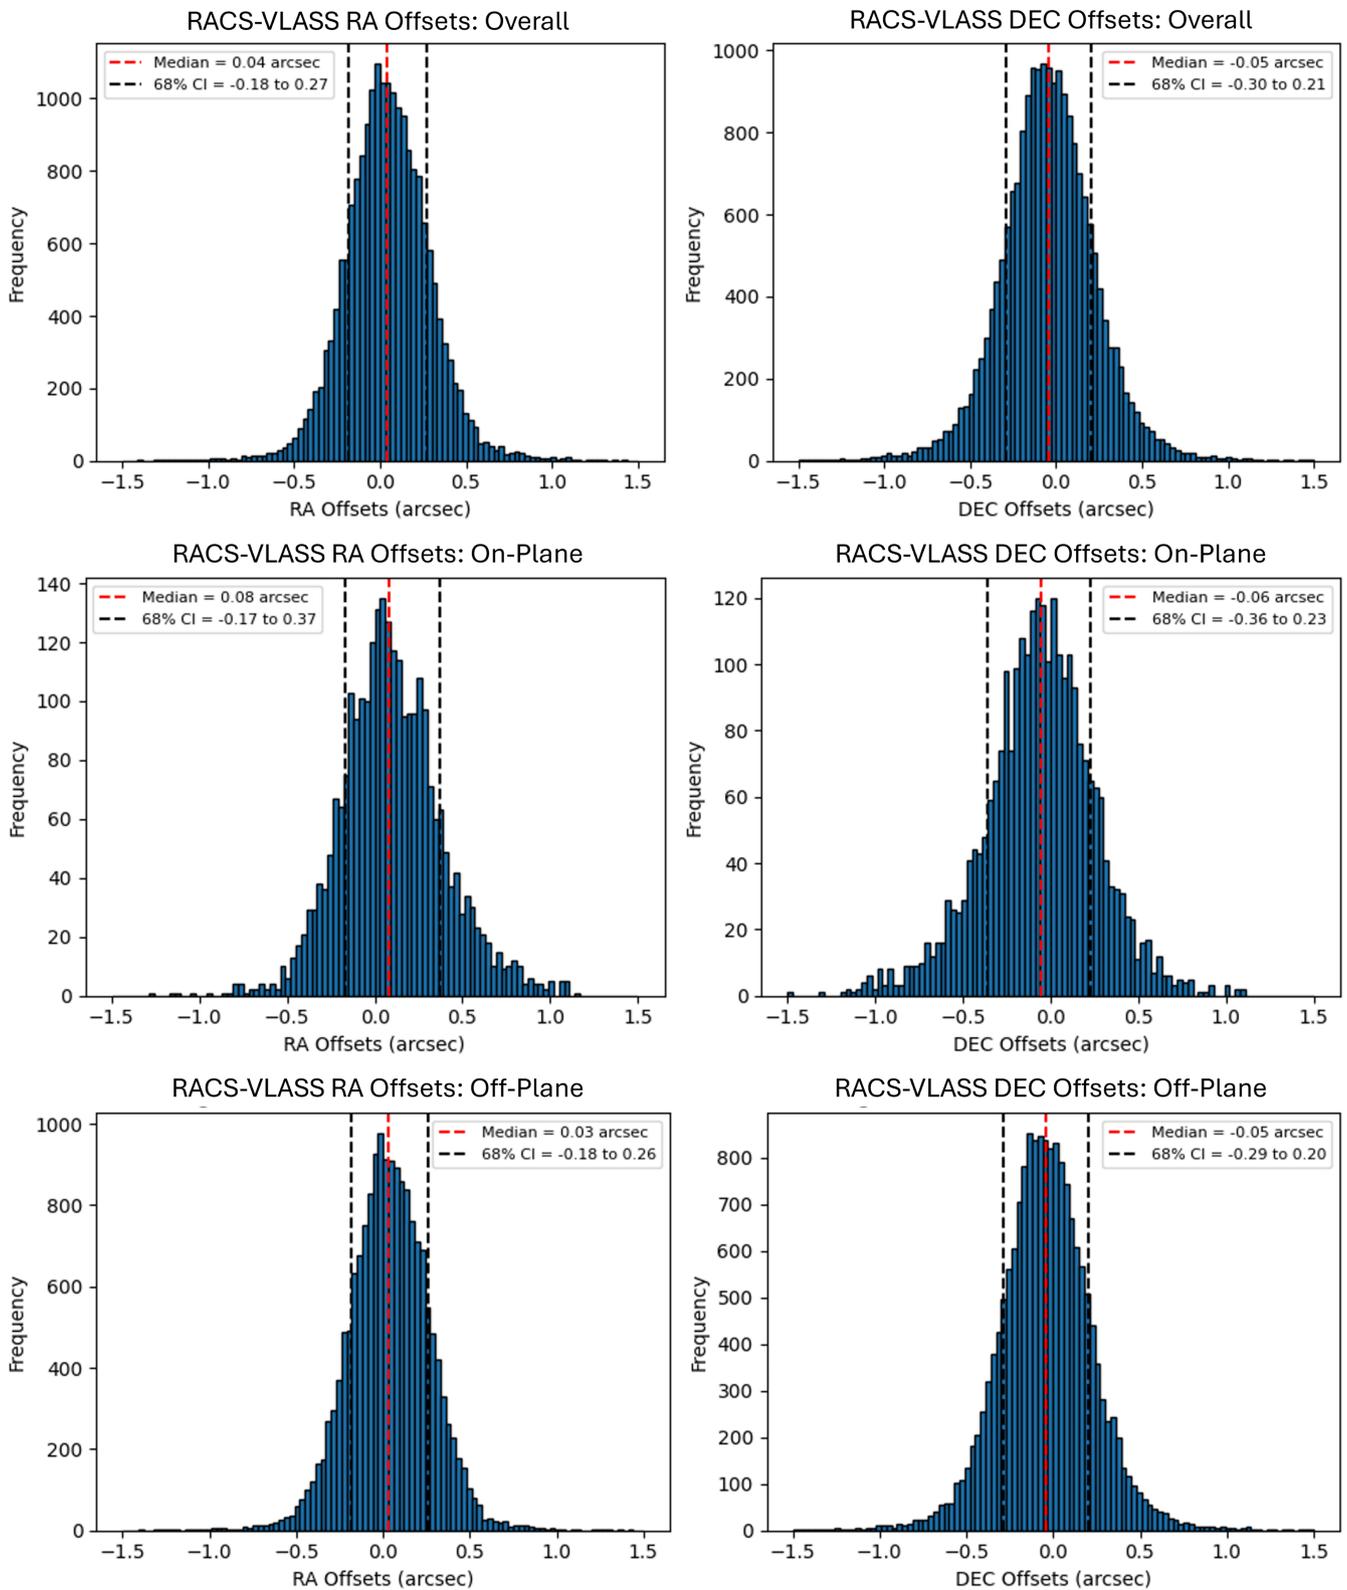


Figure 7. The corrected offsets in RA (left) and Dec. (right) for RACS-Low1 vs VLASS, with the overall histograms in the first row, the Galactic plane regions in the middle row, and the off-plane regions in the last row.

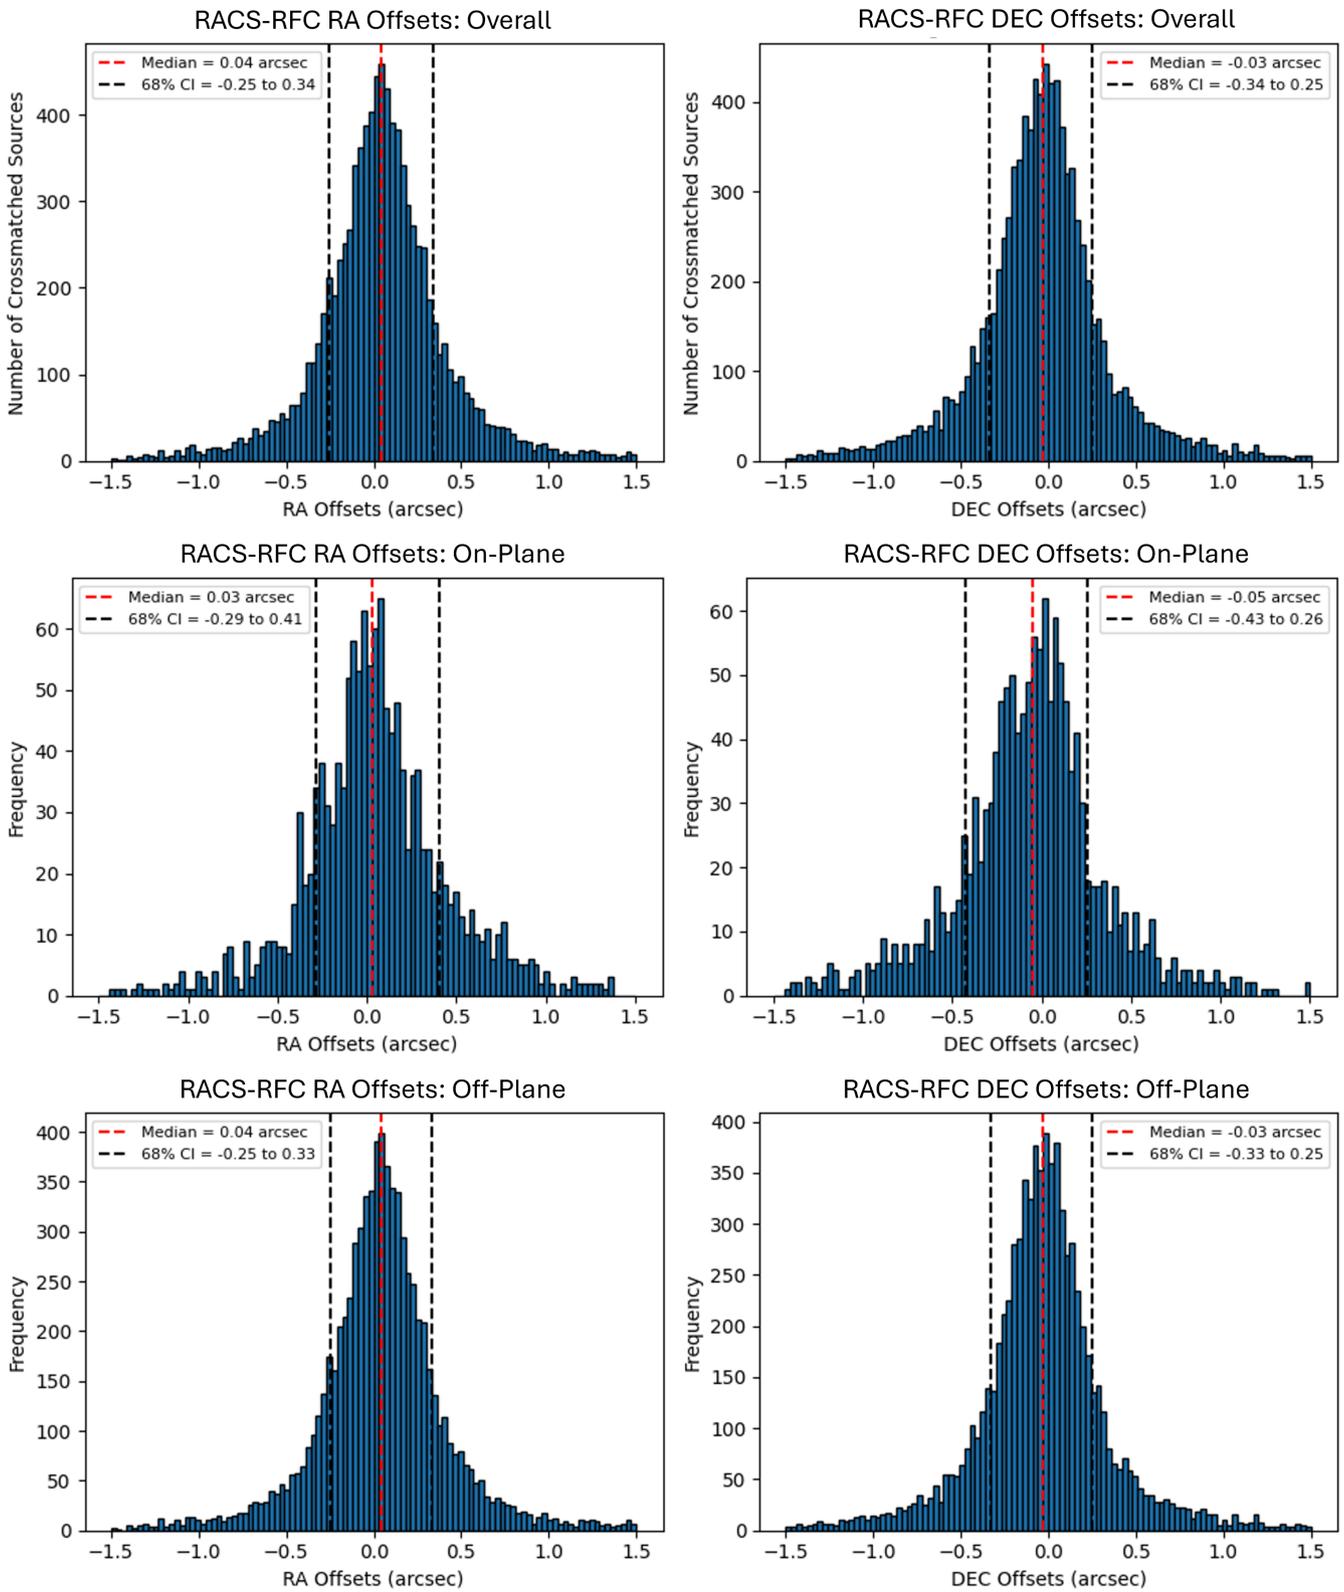


Figure 8. The corrected offsets in RA (left) and Dec. (right) for RACS-Low1 vs RFC, with the overall histograms in the first row, the Galactic plane region in the middle row, and the off-plane region in the last row. Of the ~ 18 000 RFC sources in the RACS-Low1 footprint, after the extended sources are filtered out, ~ 12 000 sources have a valid crossmatch to RACS-Low1 sources.

We then crossmatch the remaining point-like RFC sources with the corrected RACS-Low1 catalogue, producing histograms as shown in Fig. 8, with the median offsets and 68% confidence intervals listed in Table 3. To assess model performance in different regions, we again divided the dataset into on-plane and off-plane regions. In the Galactic plane region, the median RA and Dec. offsets are still within the margin of error, though the 68% confidence intervals are noticeably broader. The histograms for the off-plane region align with the overall dataset, as expected, once again confirming the model's consistency outside the more complex Galactic plane region.

We further divided the dataset into different declination ranges to assess whether the performance of our astrometric model varied with declination. Across these ranges, the median RA offsets were found to vary between 0.00'' and 0.10'', while the median Dec. offsets ranged from 0.00'' to 0.06''. The 68% confidence intervals for both RA and Dec. remained consistent, varying within $\pm 0.40''$ from the median. However, we see that the model's performance slightly deteriorates at Dec. above $+10^\circ$ and below Dec. -60° , likely due to challenges related to lower elevation observations by ASKAP in these regions.

Since these values are measured using individual sources, any true centroid position differences contribute directly to the measured offset, meaning that these offsets should overestimate the true residual uncertainty in the corrected RACS positions (unlike the checks made using FIRST and VLASS). This provides a useful upper bound, complementing the beam-averaged approach with VLASS and FIRST.

The comparison with FIRST provided 68% confidence intervals of up to $\pm 0.23''$ in RA and $\pm 0.22''$ in Dec. The comparison with VLASS yielded 68% confidence intervals of $\pm 0.37''$ in RA and $\pm 0.36''$ in Dec. for on-plane regions, while off-plane regions showed 68% confidence intervals of $\pm 0.26''$ in RA and $\pm 0.29''$ in Dec. Similarly, the comparison with RFC gave us 68% confidence intervals of up to $\pm 0.41''$ in RA and $\pm 0.43''$ in Dec. for on-plane regions, and $\pm 0.33''$ for both RA and Dec. in off-plane regions. After carefully evaluating the underestimates from this methodology, and overestimates from single-source matches with RFC, we recommend assuming $1\text{-}\sigma$ systematic uncertainties of $\pm 0.30''$ for both RA and Dec. in regions outside the Galactic plane is suitable for RACS-Low1 corrections. For Galactic plane regions, where the model encounters more significant challenges due to complex source environments and fewer point sources, slightly higher $1\text{-}\sigma$ systematic uncertainties of $\pm 0.40''$ in both RA and Dec. are more appropriate.

The residual distributions are not completely Gaussian, with a (small) excess of fields at large errors ($>3\text{-}\sigma$). Specifically, $\sim 1.2\%$ of FIRST fields in RA and $\sim 2.1\%$ in Dec. (all of which are off-plane) have a residual error greater than $0.9''$ ($3\text{-}\sigma$), compared to the 0.3% predicted for a Gaussian distribution. These deviations are dominated by a handful of fields in the *excluded regions* (as discussed in Section 4.1), where the RACS-Low1 data proved impossible to fully correct. For the vast bulk of the sky away from the Galactic plane, the post-correction uncertainty we prescribe should be appropriate: in the off-plane region, residual errors exceeding $0.6''$ ($2\text{-}\sigma$) are seen in just 3.0% of FIRST fields for RA, and 4.3% of FIRST fields for Dec. Comparable results are seen when comparing against VLASS in the off-plane region ($\sim 2.0\%$ fields in RA and $\sim 4.0\%$ in Dec. exceeded $0.6''$ residual offsets.) In the on-plane region, where we employ a more conservative estimate of $0.4''$ for the $1\text{-}\sigma$ uncertainty, no excess of

VLASS field residuals is seen at either the $2\text{-}\sigma$ ($\sim 2.3\%$ of fields in RA and $\sim 2.6\%$ of fields in Dec.) or $3\text{-}\sigma$ ($\sim 0.2\%$ of fields in RA and $\sim 0.1\%$ in Dec.) level. Accordingly, while our recommended post-correction uncertainty for RACS-Low1 astrometry should be usable for almost the entire survey, we advise caution when using RACS-Low1 data from the *excluded regions*, all of which are detailed in Section 4.1.

3.3 RACS-Low3 corrections

After completing the corrections for RACS-Low1, we applied an upgraded workflow to correct the 1 493 un-mosaicked RACS-Low3 scans (53 748 beams) using the same WISE data. The modelled and residual RACS-WISE offsets for RACS-Low3 are represented in Aitoff projection plots in Fig. 9.

Several key points arise from preliminary inspection. First, the RACS-Low3 scans exhibit significantly denser coverage and are more uniformly distributed, extending over the entire sky below Dec. $+49^\circ$. However, the pre-correction offsets in both RA and Dec. are considerably larger than those in RACS-Low1 even in regions common to both these surveys. For regions above Dec. $+30^\circ$, which are unique to RACS-Low3, pre-correction offsets sometimes exceed $7''$. Post-correction, these offsets are markedly reduced across most of the sky, demonstrating the effectiveness of our model. The model's performance is notably weaker in the Galactic plane, as observed with RACS-Low1, as well as in regions above Dec. $+30^\circ$, due to the inherent limitations of ASKAP at low elevations, such as elongated PSFs and atmospheric distortions.

We recognise that the performance of ASKAP degrades significantly above Dec. $+30^\circ$, so we first calculate the results for scans below Dec. $+30^\circ$, shown quantitatively in Fig. 10. As the histograms illustrate again, the mean bias is fully eliminated in both RA and Dec. while the stochastic variation is also significantly reduced; the 68% confidence intervals for the residual offset are reduced from $0.37''_{-1.13}^{+1.19}$ to $0.00''_{-0.18}^{+0.18}$ in RA and from $1.19''_{-1.24}^{+1.51}$ to $0.00''_{-0.18}^{+0.18}$ in Dec. This shows that the astrometric accuracy as estimated by the WISE residuals, appears consistent between RACS-Low1 and RACS-Low3 in these regions. This implies that, somewhat surprisingly, the mosaicking in the RACS-Low1 catalogue did not have a significantly detrimental effect on our ability to correct the catalogue positions.

When calculating the results for the entire RACS-Low3 sky, including the regions above Dec. $+30^\circ$, we similarly find that the mean bias is completely eliminated and the stochastic variation is also reduced; the 68% confidence intervals for the residual offset are reduced from $0.29''_{-1.19}^{+1.22}$ to $0.00''_{-0.20}^{+0.20}$ in RA and from $1.24''_{-1.29}^{+1.76}$ to $0.00''_{-0.21}^{+0.21}$ in Dec. These results are noticeably worse because of the poor corrections in regions above Dec. $+30^\circ$, due to reasons discussed above.

3.4 Verification of corrections of RACS-Low3 using other catalogues

We performed cross-verification of the RACS-Low3 astrometric corrections using comparisons with FIRST, VLASS, and RFC, collapsing the data into a single analysis for simplicity, as shown in Fig. 11. Crossmatching RACS-Low3 with FIRST resulted in median offsets in RA and Dec. improving from $0.19''$ and $1.13''$

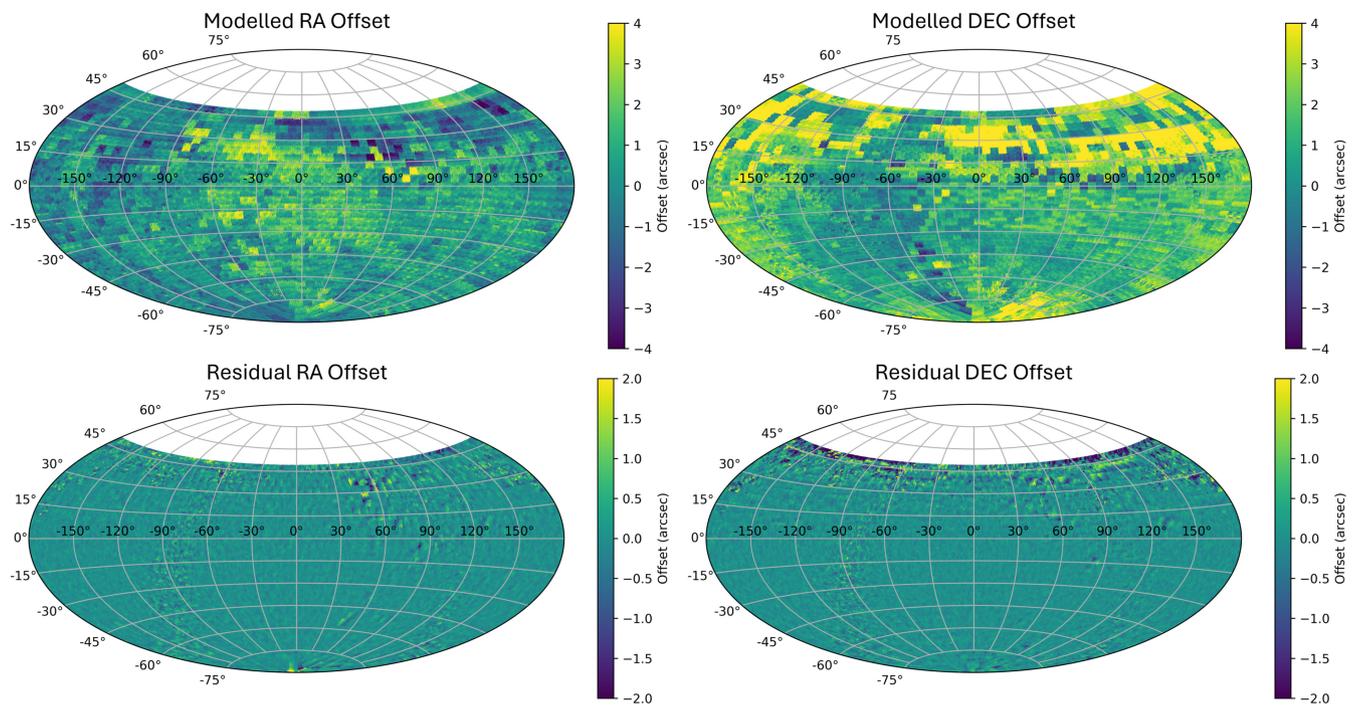


Figure 9. The modelled and residual offsets of RACS-Low3 vs WISE. The top row shows the observed offsets and the bottom row shows the residual offsets in RA (left) and Dec. (right) for the entire sky coverage.

(uncorrected) to $0.02''$ and $0.06''$ (corrected). Owing to the limited coverage of FIRST, only $\sim 9\,000$ RACS-Low3 beams are included, with each beam having an average of 50 crossmatched sources.

For the comparison of RACS-Low3 with VLASS, we used the latest catalogues available on CIRADA,³ the official repository for VLASS data. These updated catalogues include corrections for the inherent Dec. offset observed in previous versions of VLASS. This produced median RA and Dec. offsets of $0.01''$ and $0.11''$, respectively. $\sim 42\,000$ beams of RACS-Low3 have overlapping sky coverage with VLASS, with each beam having an average of 50 crossmatched sources. There were persistent outliers in Dec. offsets, primarily within the Galactic plane and Dec. above $+20^\circ$, which contributed to larger Dec. offsets and indicate a continued challenge for accurate astrometric correction in these regions.

Crossmatching with the high-precision RFC catalogue yielded median RA and Dec. offsets of $0.00''$ and $0.06''$ for $\sim 14\,000$ crossmatched sources. Although the RA offsets met expectations, the Dec. offsets exhibited a larger median than anticipated, mostly because of the poor corrections in regions above a Dec. of $+20^\circ$. These results suggest that while the corrections improved the overall astrometric accuracy of RACS-Low3, further refinement is necessary, particularly in challenging regions like the Galactic plane and higher declinations.

3.5 RACS-Low3 vs RACS-Low1

As a final sanity check, we crossmatched the corrected sources in RACS-Low3 with those in RACS-Low1. The histograms in Fig. 12

show the median offsets and their 68% confidence intervals of $-0.03_{-0.28}^{+0.27}$ in RA and $0.09_{-0.23}^{+0.27}$ in Dec. This comparison demonstrates that the corrected sources in RACS-Low3 are, within the margin of error, in the same positions as the corrected sources in RACS-Low1. However, in a few outlier regions (as evidenced in the residual plots shown in Fig. 13), our model does not perform optimally due to the significantly lower number of point sources in each beam.

Following these cross-verifications, while our model improves the positional accuracy of RACS-Low3, the larger number of beams that are less effectively corrected – particularly in regions within the Galactic Plane and at northern declinations above Dec. $+30^\circ$ – indicates that the overall reliability of the final RACS-Low3 corrections is lower compared to RACS-Low1. This highlights the need for further refinement of the model, especially in areas with fewer point sources, to ensure optimal accuracy across the entire RACS-Low3 sky coverage. Given these results, we propose not to use RACS-Low3 for further processing at this stage and to continue exclusively with RACS-Low1.

4. Discussion

The encouraging outcome of our present focused exploration to correct the positional offsets in RACS-Low, paves way to proceed to the next step to generate a catalogue of corrected positions of all continuum sources across the entire ASKAP frequency band. Learning from the challenges we faced during our present endeavour, we aim to develop a more efficient and adaptive pipeline to model the offsets in RACS-Mid and RACS-High, as well as any future ASKAP survey.

³Website: cirada.ca/vlascatalogueql0.

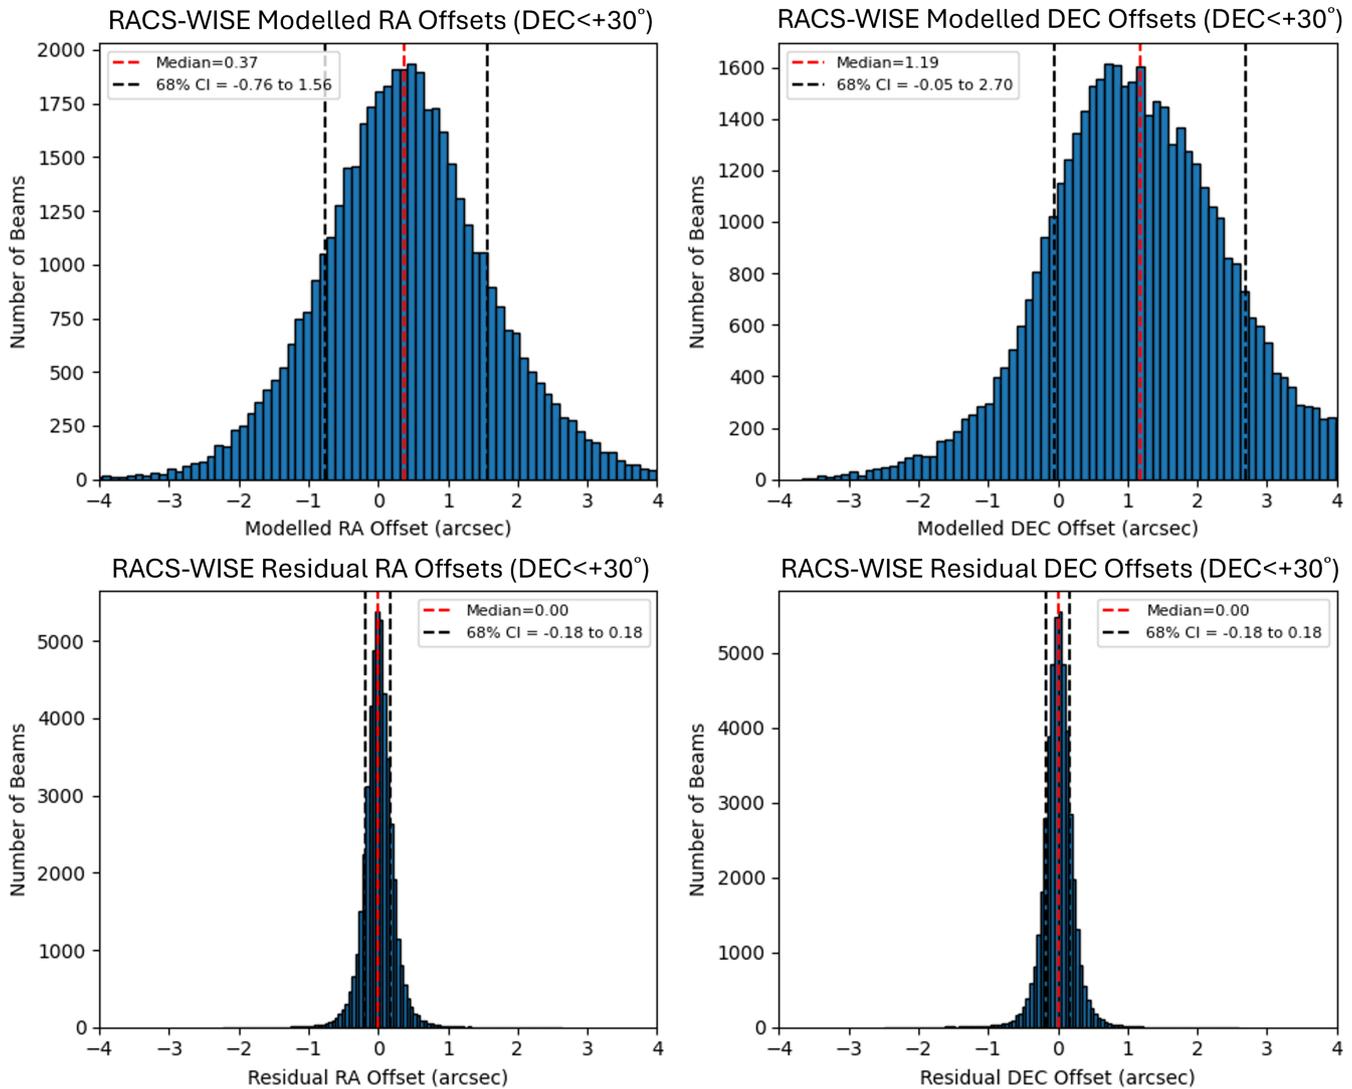


Figure 10. The modelled (top row) and residual (bottom row) offsets for RACS-Low3 vs WISE for scans below Dec. +30°, to have a fair comparison between the corrections for RACS-Low3 and RACS-Low1. The median RA offset (left) improves from 0.37'' to 0.00'' and the median Dec. offset (left) decreases from 1.19'' to 0.00''. The 68% confidence intervals also get significantly smaller.

4.1 Accuracy of corrections and excluded regions

A quick analysis of the residual offsets of RACS-Low1, as depicted in Fig. 4, highlights several regions where the astrometric corrections are less reliable than expected. Notably, the Galactic plane exhibits poor corrections, which is further elaborated in Section 2.1. Additionally, certain *excluded regions* present significantly flawed corrections, as also seen in the RACS-Low1–VLASS offsets in the top row of Fig. 13.

For instance, Hercules A, a very bright radio source located at RA $16^{\text{h}}51^{\text{m}}$ and Dec. $+04^{\circ}59'$, corrupted the original data, likely due to its sidelobes, which shift the centroids of real sources by layering noise over the top, leading to unreliable corrections within a 3° radius. Similarly, the Crab Nebula at RA $05^{\text{h}}34^{\text{m}}$ and Dec. $+22^{\circ}01'$ also caused contamination, resulting in poor corrections within a similar radius. Additionally, regions at RA $22^{\text{h}}12^{\text{m}}$ and Dec. $+16^{\circ}50'$, at RA $11^{\text{h}}20^{\text{m}}$ and Dec. $+05^{\circ}00'$, and at RA $00^{\text{h}}00^{\text{m}}$ and Dec. $+17^{\circ}00'$, show significant astrometric errors despite the

absence of known sources or external factors. In these areas, corrections are highly unreliable, and uncertainties are estimated at $\pm 0.50''$ for RA and $\pm 1.00''$ for Dec.

Given the magnitude of these uncertainties, we suggest that for radio crossmatching/astrometric purposes, it would be preferable to not use RACS-Low1 in these regions. Instead, alternate catalogues like VLASS or RACS-Low3, which shows much more reliable corrections (as evidenced by the residual plots in Fig. 9 and RACS-Low3–VLASS offsets in the bottom row of Fig. 13), should be considered.

4.2 Present use of the RACS-Low1 corrections

The corrected RACS-Low1 catalogues are currently being utilised in the CRAFT Effortless Localisation and Enhanced Burst Inspection pipeline (CELEBI, Scott *et al.* 2023) to refine the positions of ASKAP-discovered FRBs. A comprehensive review of this

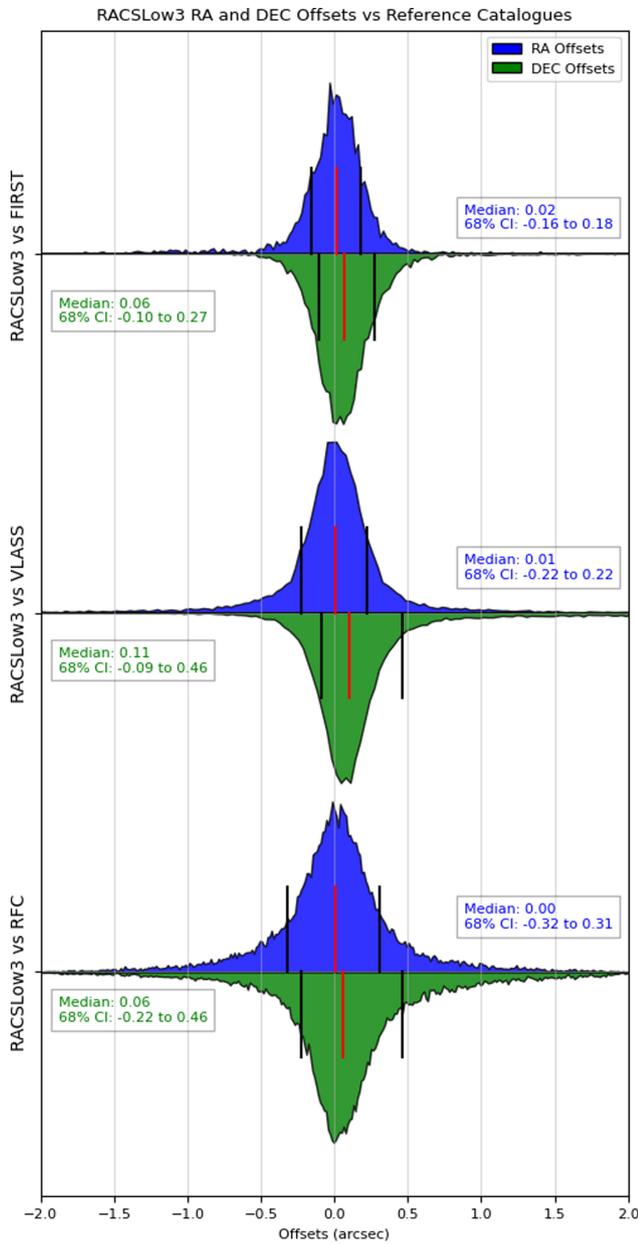


Figure 11. The corrected offsets in RA and Dec. for RACS-Low3 vs reference catalogues, with FIRST comparisons in the top row, VLASS in the middle row and RFC in the bottom row.

integration, along with other updates to the pipeline, will be discussed in Glowacki & Dial (in preparation). As an example, Fig. 14 shows the positional offsets of FRB20230718A (a detailed review of the methodology is presented in Glowacki et al. 2024) before and after using the corrected RACS-Low1 catalogue in CELEBI. The change in the source position can be visualised from the DECAM⁴ image of the FRB field. Additionally, a list of FRBs discovered using the ICS system, with their positions updated using the corrected RACS-Low1 catalogue, is presented in Shannon et al. (2024).

⁴Website: noirlab.edu/public/programs/ctio/victor-blanco-4m-telescope/decam/.

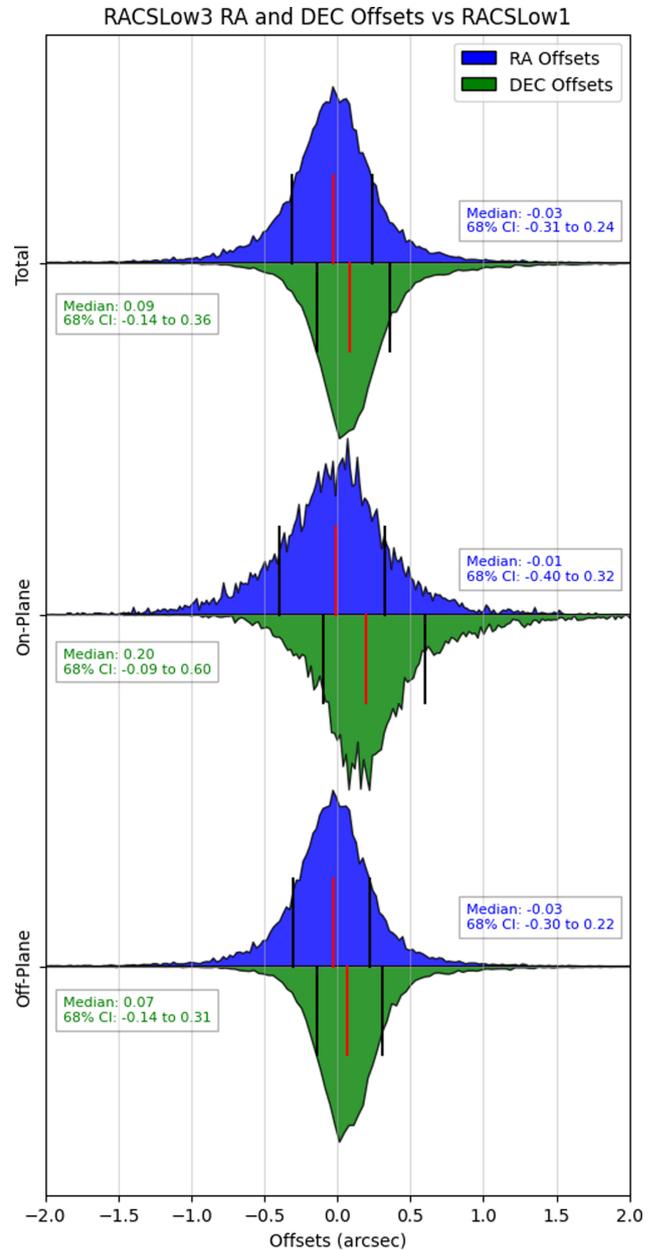


Figure 12. The offsets in RA and Dec. for RACS-Low3 vs RACS-Low1 post corrections, with the overall histograms in the top row, the on-plane region in the middle row, and the off-plane region in the bottom row. The overall histograms only include regions below Dec. +30°, to compare common data available for these surveys, which corresponds to ~46 000 RACS-Low3 beams, with each beam having an average of 50 crossmatched sources.

4.3 Future plans

Building on the astrometric improvements with RACS-Low1 and RACS-Low3, we plan to extend our corrections to the full RACS catalogue, including RACS-Mid and RACS-High. By applying similar crossmatching techniques with catalogues like WISE, FIRST, VLASS, as well as the corrected RACS-Low1, we aim to improve astrometric accuracy across all ASKAP frequencies. As with RACS-Low3, we will use the un-mosaicked per-beam source lists for both RACS-Mid and RACS-High to avoid the blending of astrometric errors that occur in the tile imaging process.

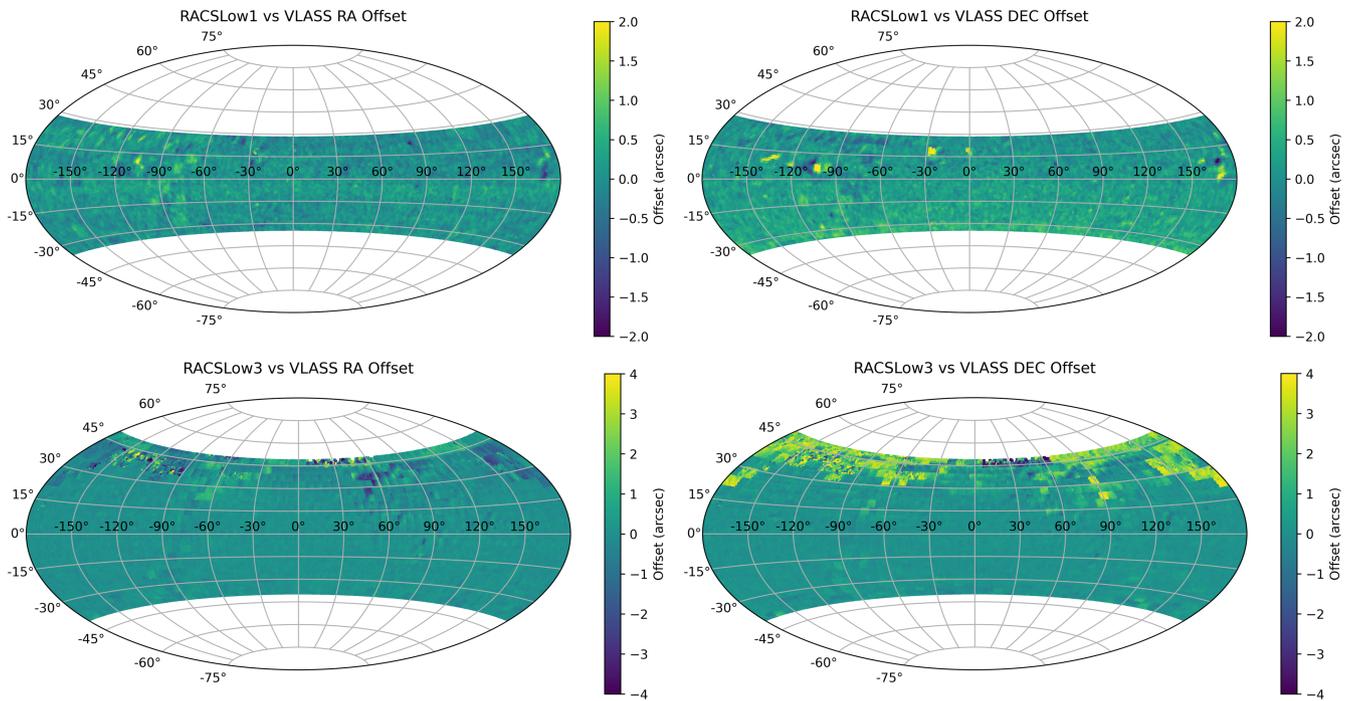


Figure 13. The top row shows the RA (left) and Dec. (right) offsets of RACS-Low1 vs VCLASS and the bottom row shows the RA (left) and Dec. (right) offsets of RACS-Low3 vs VCLASS. The underperforming regions in both RACS-Low1 and RACS-Low3 can be clearly visualised from these plots.

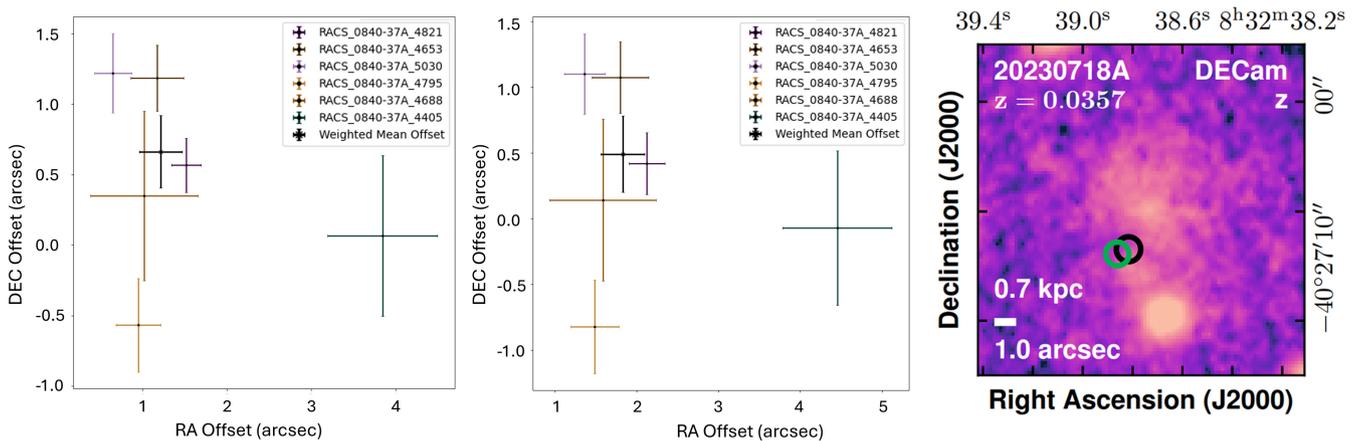


Figure 14. An example of determining the positional offsets of FRB20230718A while using the uncorrected (left) and the corrected (middle) RACS-Low1 catalogues with CELEBI. After updating the RACS catalogue, the estimated offset correction based on the field sources changed from $1.21'' \pm 0.26''$, $0.66'' \pm 0.26''$ to $1.75'' \pm 0.27''$, $0.49'' \pm 0.29''$ in RA and Dec. respectively. After including the assumed $1\text{-}\sigma$ systematic uncertainty of $0.4''$ (as the Galactic latitude of this source is -0.37°) and the statistical FRB position uncertainty, the final position of the FRB is updated from $08^{\text{h}}32^{\text{m}}38^{\text{s}}.82 \pm 0.54''$ to $08^{\text{h}}32^{\text{m}}38^{\text{s}}.86 \pm 0.52''$ in RA and from $-40^{\circ}27'06.78'' \pm 0.56''$ to $-40^{\circ}27'06.95'' \pm 0.53''$ in Dec. after corrections. The effect of this change on the FRB's localisation within the host galaxy is shown in the DECcam image of the host (right), where the black error ellipse indicates the localisation with the uncorrected RACS-Low1 catalogue, and the green ellipse shows the improved localisation achieved with the corrected version (Shannon et al. 2024).

Learning from our experiences in correcting RACS-Low1 and RACS-Low3, we will optimise our correction algorithm for future iterations. By first crossmatching with the corrected RACS-Low catalogues, we can apply larger crossmatching radii in regions with large positional offsets without significantly increasing the risk of false-positive matches, by avoiding the need to use excessively large radii directly with the dense WISE catalogue, where source confusion is more likely. This will help us generate enhanced

corrections not only in the Galactic plane but also for areas at higher declinations.

We also plan to investigate the impact of ionospheric stability on positional offsets in the different RACS epochs, particularly for observations conducted at low elevations. Given the heightened solar activity during some of these observations, ionospheric fluctuations may have introduced additional astrometric distortions. A potential direction for future work is to analyse time-dependent

trends, such as diurnal or seasonal variations in positional errors, to assess whether ionospheric conditions contribute systematically to the observed discrepancies. This could inform more refined corrections, especially for fields observed under varying ionospheric conditions.

This unified, corrected RACS catalogue will serve as a reliable foundation for precise localisation for declinations below $+20^\circ$, the first of its kind at this resolution in the Southern Hemisphere, offering over 5 times the astrometric accuracy compared to the previous metric (SUMSS). It will complement VLASS, which already provides a reliable localisation catalogue in the Northern Hemisphere. Together, these resources will enhance future radio observations, particularly for FRB studies, by offering comprehensive astrometric accuracy across both hemispheres.

5. Conclusion

In this paper, we developed and applied a robust astrometric correction model for the RACS-Low1 and RACS-Low3 radio continuum surveys conducted with ASKAP. By crossmatching RACS sources with the WISE catalogue, we were able to produce accurate positional corrections for both RACS-Low1 and RACS-Low3, significantly improving their astrometric precision. For RACS-Low1 corrections with WISE, our model demonstrated substantial improvements, with median RA and Dec. offsets reducing to within $0.00''$ and the 68% confidence intervals were calculated to within $\pm 0.30''$ all across the RACS-Low1 sky coverage, except the Galactic plane, where they were slightly worse at $\pm 0.40''$. This consolidation of the corrected offsets allowed for a significant enhancement in the astrometric accuracy of the RACS catalogue, which in turn improved the precision of FRB localisations.

We validated our correction model through cross-comparisons with other high-precision catalogues such as FIRST, VLASS, and RFC. These comparisons confirmed the effectiveness of our model, with residual offsets for RACS-Low1 well within acceptable limits. However, the model showed some limitations in regions with sparse point sources, such as the Galactic plane, where the model performed less effectively.

RACS-Low3, being an unprocessed catalogue at the time of this study, posed additional challenges, but we successfully implemented our correction model using a similar methodology. While RACS-Low3 initially presented larger offsets than RACS-Low1, especially above Dec. $+30^\circ$, the corrected positions showed marked improvements across most of the sky. However, due to the larger number of beams in RACS-Low3 that were corrected less optimally, the final astrometric accuracy was less reliable than that of RACS-Low1. Cross-verifications between RACS-Low1 and RACS-Low3 further demonstrated the reliability of the corrected positions, although certain outlier regions remained where the model did not perform as well, largely due to the limited number of point sources available for crossmatching or inherent limitations of the ASKAP telescope.

We also plan to release the corrected RACS-Low1 catalogues and RACS-Low3 per-beam source lists publicly, along with all the corrections applied, making them available for use by the broader scientific community. These corrections will not only mitigate the inherent astrometric uncertainties in RACS-Low epochs but also provide a framework for future improvements in radio surveys, ensuring that ASKAP continues to deliver cutting-edge data for the study of radio transients, and a wide range of other astronomical phenomena. This work paves the way for future astrometric

corrections in other ASKAP surveys, including RACS-Mid and RACS-High, where we plan to develop a more adaptive pipeline that addresses the challenges encountered in this study.

By leveraging these meticulous corrections, we enable robust associations between FRBs and sources in optical catalogues, facilitating the identification of their multiwavelength counterparts. This precision is critical for pinpointing host galaxies, determining their redshifts, and uncovering details about their environments. Beyond FRBs, improved astrometry supports the study of other transient phenomena. For instance, accurate localisation aids in identifying optical counterparts of flare stars, enabling spectral type classification and further characterisation of their physical properties. Similarly, for LPTs, better positional precision can reveal potential optical companions or help constrain their nature.

Acknowledgement. The authors thank Simon C.-C. Ho for comments on the manuscript. AJ, ATD, and YW acknowledge the support of the Australian Research Council (ARC) grant DP220102305. MG is supported by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (DP210102103), and through UK STFC Grant ST/Y001117/1. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission. We acknowledge the use of OpenAI's ChatGPT-4o (<https://chatgpt.com/>) for providing suggestions on refining sentence structure in this manuscript. This scientific work uses data obtained from Inyarrimanha Ilgari Bundara/the Murchison Radio-astronomy Observatory. We acknowledge the Wajarri Yamaji People as the Traditional Owners and native title holders of the Observatory site. CSIRO's ASKAP radio telescope is part of the Australia Telescope National Facility (<https://ror.org/05qajvd42>). Operation of ASKAP is funded by the Australian Government with support from the National Collaborative Research Infrastructure Strategy. ASKAP uses the resources of the Pawsey Supercomputing Research Centre. Establishment of ASKAP, Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory and the Pawsey Supercomputing Research Centre are initiatives of the Australian Government, with support from the Government of Western Australia and the Science and Industry Endowment Fund. This paper includes archived data obtained through the CSIRO ASKAP Science Data Archive, CASDA (<https://data.csiro.au>). A part of this work was performed on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR programme receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government. This research has made use of the VizieR catalogue access tool, CDS, Strasbourg, France Ochsenbein 1996. The original description of the VizieR service was published in Ochsenbein, Bauer, & Marcout (2000). The Canadian Initiative for Radio Astronomy Data Analysis (CIRADA) is funded by a grant from the Canada Foundation for Innovation 2017 Innovation Fund (Project 35999) and by the Provinces of Ontario, British Columbia, Alberta, Manitoba and Quebec, in collaboration with the National Research Council of Canada, the US National Radio Astronomy Observatory, and Australia's Commonwealth Scientific and Industrial Research Organisation.

Data availability statement. The updated RACS-Low1 sources-only and Gaussian components catalogues, which include the astrometric corrections applied in this work, are made available in CSV format through the PASA Datastore: <https://doi.org/10.57891/6ksa-md85>. Similarly, the corrected RACS-Low3 datasets, based on the same methodology, are provided as beam-wise, unmosaicked source lists in VOTable (VOT) format through the PASA Datastore: <https://doi.org/10.57891/z4gk-at64>. These datasets are intended to support the broader astronomical community in conducting high-precision localisation and cross-matching studies. Further details, including file structure and usage notes, are provided within the dataset descriptions on the respective PASA Datastore pages.

References

- ASTROPY Collaboration, et al. 2018, *AJ*, 156, 123. <https://doi.org/10.3847/1538-3881/aabc4f>. arXiv: 1801.02634 [astro-ph.IM].
- ASTROPY Collaboration, et al. 2022, *ApJ*, 935, 167. <https://doi.org/10.3847/1538-4357/ac7c74>. arXiv: 2206.14220 [astro-ph.IM].
- ASTROPY Collaboration, et al. 2013, *A&A*, 558, A33. <https://doi.org/10.1051/0004-6361/201322068>. arXiv: 1307.6212 [astro-ph.IM].
- Bailes, M. 2022, *Sci*, 378, abj3043. <https://doi.org/10.1126/science.abj3043>. arXiv: 2211.06048 [astro-ph.HE].
- Bannister, K. W., et al. 2019, *Sci*, 365, 565. <https://doi.org/10.1126/science.aaw5903>. arXiv: 1906.11476 [astro-ph.HE].
- Becker, R. H., White, R. L., & Helfand, D. J. 1994, in *Astronomical Data Analysis Software and Systems III*, Vol. 61, Astronomical Society of the Pacific Conference Series, ed. D. R. Crabtree, R. J. Hanisch, & J. Barnes, 165.
- Bochenek, C. D., Ravi, V., Belov, K. V., Hallinan, G., Kocz, J., Kulkarni, S. R., & McKenna, D. L. 2020, *Natur*, 587, 59. <https://doi.org/10.1038/s41586-020-2872-x>. arXiv: 2005.10828 [astro-ph.HE].
- Chatterjee, S., et al. 2017, *Natur*, 541, 58. <https://doi.org/10.1038/nature20797>. arXiv: 1701.01098 [astro-ph.HE].
- Chime/Frb Collaboration, et al. 2020, *Natur*, 582, 351. <https://doi.org/10.1038/s41586-020-2398-2>. arXiv: 2001.10275 [astro-ph.HE].
- Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., & Broderick, J. J. 1998, *AJ*, 115, 1693. <https://doi.org/10.1086/300337>.
- Duchesne, S. W., et al. 2024, *PASA*, 41, e003. <https://doi.org/10.1017/pasa.2023.60>. arXiv: 2311.12369 [astro-ph.GA].
- Duchesne, S. W., et al. 2023, *PASA*, 40, e034. <https://doi.org/10.1017/pasa.2023.31>. arXiv: 2306.07194 [astro-ph.IM].
- Ginsburg, A., et al. 2019, *AJ*, 157, 98. <https://doi.org/10.3847/1538-3881/aafc33>. arXiv: 1901.04520 [astro-ph.IM].
- Glowacki, M., et al. 2024, *ApJ*, 962, L13. <https://doi.org/10.3847/2041-8213/ad1f62>. arXiv: 2311.16808 [astro-ph.GA].
- Gordon, Y. A., et al. 2021, *ApJS*, 255, 30. <https://doi.org/10.3847/1538-4365/ac05c0>. arXiv: 2102.11753 [astro-ph.GA].
- Hale, C. L., et al. 2021, *PASA*, 38, e058. <https://doi.org/10.1017/pasa.2021.47>. arXiv: 2109.00956 [astro-ph.GA].
- Harris, C. R., et al. 2020, *Natur*, 585, 357. <https://doi.org/10.1038/s41586-020-2649-2>.
- Hotan, A. W., et al. 2021, *PASA*, 38, e009. <https://doi.org/10.1017/pasa.2021.1>. arXiv: 2102.01870 [astro-ph.IM].
- Lacy, M., et al. 2020, *PASP*, 132, 035001. <https://doi.org/10.1088/1538-3873/ab63eb>. arXiv: 1907.01981 [astro-ph.IM].
- Lyubarsky, Y. 2020, *ApJ*, 897, 1. <https://doi.org/10.3847/1538-4357/ab97b5>. arXiv: 2001.02007 [astro-ph.HE].
- Lyutikov, M., Barkov, M. V., & Giannios, D. 2020, *ApJ*, 893, L39. <https://doi.org/10.3847/2041-8213/ab87a4>. arXiv: 2002.01920 [astro-ph.HE].
- Macquart, J.-P., et al. 2020, *Natur*, 581, 391. <https://doi.org/10.1038/s41586-020-2300-2>. arXiv: 2005.13161 [astro-ph.CO].
- McConnell, D., et al. 2020, *PASA*, 37, e048. <https://doi.org/10.1017/pasa.2020.41>. arXiv: 2012.00747 [astro-ph.IM].
- Ochsenbein, F., Bauer, P., & Marcout, J. 2000, *A&AS*, 143, 23. <https://doi.org/10.1051/aas:2000169>. arXiv: astro-ph/0002122 [astro-ph].
- Ochsenbein, F. 1996, The vizier database of astronomical catalogues. <https://doi.org/10.26093/CDS/VIZIER>. <https://vizier.cds.unistra.fr>.
- Petroff, E., Hessels, J. W. T., & Lorimer, D. R. 2022, *A&ARv*, 30, 2. <https://doi.org/10.1007/s00159-022-00139-w>. arXiv: 2107.10113 [astro-ph.HE].
- Petrov, L., & Kovalev, Y. 2024, arXiv e-prints (October): arXiv:2410.11794. <https://doi.org/10.48550/arXiv.2410.11794>.
- Platts, E., Weltman, A., Walters, A., Tendulkar, S. P., Gordin, J. E. B., & Kandhai, S. 2019, *PhRp*, 821, 1–27. <https://doi.org/10.1016/j.physrep.2019.06.003>. arXiv: 1810.05836 [astro-ph.HE].
- Sadler, E. M. 1999, in *Looking Deep in the Southern Sky*, Vol. 103, ed. R. Morganti, & W. J. Couch.
- Scott, D. R., et al. 2023, *A&C*, 44, 100724. <https://doi.org/10.1016/j.ascom.2023.100724>. arXiv: 2301.13484 [astro-ph.IM].
- Shannon, R. M., et al. 2024, arXiv e-prints (August): arXiv:2408.02083. <https://doi.org/10.48550/arXiv.2408.02083>. arXiv: 2408.02083 [astro-ph.HE].
- Virtanen, P., et al. 2020, *NM*, 17, 261. <https://doi.org/10.1038/s41592-019-0686-2>.
- Wang, Z., et al. 2024, arXiv: 2409.10316 [astro-ph.HE].
- Wright, E. L., et al. 2010, *AJ*, 140, 1868. <https://doi.org/10.1088/0004-6256/140/6/1868>. arXiv: 1008.0031 [astro-ph.IM].