


## Research Article

# A census of compact sources at 162 MHz: First data release from the MWA Phase II IPS Survey

J. S. Morgan<sup>1</sup>, R. Chhetri<sup>1,2</sup> and R. Ekers<sup>3</sup>

<sup>1</sup>International Center for Radio Astronomy Research, Curtin University, GPO Box U1987, Perth, WA 6845, Australia, <sup>2</sup>CSIRO Space and Astronomy, P.O. Box 1130, Bentley, WA 6102, Australia and <sup>3</sup>CSIRO Space and Astronomy, P.O. Box 76, Epping, NSW 1710, Australia

### Abstract

We present a catalogue of over 7000 sources from the GLEAM survey which have significant structure on sub-arcsecond scales at 162 MHz. The compact nature of these sources was detected and quantified via their Interplanetary Scintillation (IPS) signature, measured in interferometric images from the Murchison Widefield Array. The advantage of this approach is that all sufficiently compact sources across the survey area are included down to a well-defined flux density limit. The survey is based on  $\sim 250 \times 10$ -min observations, and the area covered is somewhat irregular, but the area within  $1 \text{ h} < \text{RA} < 11 \text{ h}$ ;  $-10^\circ < \text{Decl.} < +20^\circ$  is covered entirely, and over 85% of this area has a detection limit for compact structure below 0.2 Jy. 7839 sources clearly showing IPS were detected ( $>5\sigma$  confidence), with a further 5550 tentative ( $>2\sigma$  confidence) detections. Normalised Scintillation Indices (NSI; a measure of the fraction of flux density coming from a compact component) are reported for these sources. Robust and informative upper limits on the NSI are reported for a further 31081 sources. This represents the largest survey of compact sources at radio frequencies ever undertaken.

**Keywords:** radio continuum: radio continuum: galaxies – scattering – Sun: heliosphere – techniques: interferometric – techniques: high angular resolution – Radio continuum: ISM

(Received 24 June 2022; revised 19 August 2022; accepted 11 October 2022)

## 1. Introduction

The presence or absence of radio source structure at sub-arcsecond scales allows the separation of more compact structure indicative of present or recent activity (such as cores and hotspots) from the large radio lobes that dominate the source population at low radio frequencies. Nonetheless, large, unbiased surveys of compact structure have, until recently, been lacking in the literature, due to the technical difficulties of surveying large areas of the sky at very high resolution (i.e. with long baselines), particularly at low frequencies. As a result, with a handful of exceptions (e.g. Porcas et al. 2004; Deller & Middelberg 2014) most observing campaigns aimed at identifying compact sources have focused on relatively small fields using the widefield VLBI approach (Garrett et al. 2001; Middelberg et al. 2011; Morgan et al. 2013; Radcliffe et al. 2018) and/or have selectively observed sources likely to be compact, based on their flat spectra. (e.g. Beasley et al. 2002; Moldón et al. 2015; Jackson et al. 2016).

The discovery of Interplanetary Scintillation (IPS) by Clarke (1964) led Hewish, Scott, & Wills (1964) to propose IPS as an alternative method for determining which radio sources have a compact component. Since IPS arises from interference between radio waves that traverse the solar wind several hundred kilometres apart (for metre wavelengths), coherence is destroyed for sources much larger than  $1''$ , and the scintillation signal will be

suppressed. This led eventually to several catalogues of IPS sources (Purvis et al. 1987; Balasubramanian et al. 1993).<sup>a</sup>

In previous work, IPS was reintroduced as a viable method for identifying compact sources (Morgan et al. 2018) using the most recent generation of low-frequency radio interferometers, in particular the Murchison Widefield Array (MWA; Tingay et al. 2013). With a handful of proof-of-concept observations, new compact sources were identified and their properties explored (Chhetri et al. 2018a,b; Sadler et al. 2019). Jaiswal et al. (2022) extended this work to show that compactness as measured by MWA IPS observations was correlated with structure measured with GHz VLBI, which probes scales an order of magnitude more compact at much higher frequencies. A key point is that almost all the sources that we observed were also present in the GLEAM survey (Wayth et al. 2015; Hurley-Walker et al. 2017), which provides detailed flux density measurements in the range 72–231 MHz. We also explored the feasibility of an all-sky IPS Survey (Morgan et al. 2019) using the MWA, demonstrating that it would be possible to detect many thousands of sources within  $30^\circ$  of the ecliptic.

Here we present the culmination of several years of effort to generate a catalogue of compact sources over a large area using the IPS technique. We use data from extended configuration of the Phase II MWA (Wayth et al. 2018), which, as discussed by Beardsley et al. (2019), is superior to the Phase I MWA for IPS observations due to its more uniform ( $u, v$ ) coverage. Although this IPS survey has already been used in a number of publications

**Corresponding author:** J. S. Morgan, Email: [john.morgan@icrar.org](mailto:john.morgan@icrar.org)

**Cite this article:** Morgan JS, Chhetri R and Ekers R. (2022) A census of compact sources at 162 MHz: First data release from the MWA Phase II IPS Survey. *Publications of the Astronomical Society of Australia* 39, e063, 1–16. <https://doi.org/10.1017/pasa.2022.56>

<sup>a</sup>The Ooty catalogue itself was never published, but catalogues from several IPS observatories are available on request. See Mejia-Ambriz et al. (2010).

(Drouart *et al.* 2021; Jackson *et al.* 2022; Broderick *et al.* 2022), using preliminary versions of the catalogue presented here, this is the first release of IPS data to the community from this observing campaign, and the first comprehensive description of the scheduling, observations, data reduction, and synthesis of the catalogue.

An important recent development is the demonstration of the International LOFAR Telescope (van Haarlem *et al.* 2013) to be a viable instrument for conducting widefield surveys, even while using the international (100–1000 km) baselines, and thus probing very similar spatial scales to IPS (Morabito *et al.* 2022); indeed Sweijen *et al.* (2022) detected a staggering 2483 sources with a compact component in a synthesis image with a resolution of  $0.38'' \times 0.3''$ , 6.6 square degrees in extent. Our survey is highly complementary, since it covers a far wider field of view to a much shallower depth. In addition, we are observing from the Southern Hemisphere, from the site of the future Low-frequency Aperture Array of the Square Kilometre Array (SKA\_LOW). While we are most sensitive when surveying along the ecliptic, much of this area is too far south to be easily observed by LOFAR (though there is, by design, some overlap in this first data release). Our approach therefore allows us to survey a very large fraction of the sky down to a flux density limit that encompasses a large fraction of GLEAM sources, while avoiding difficulties of low-frequency VLBI imposed by the ionosphere, and the enormous computational cost required to generate very large interferometric images.

This first data release has at least three uses. Primarily we intend it to be an astrophysical dataset which is useful in its own right, while being highly synergistic with the GLEAM survey which we use as our reference catalogue. Secondly, the catalogue will be useful in providing compact calibrators for other low-frequency, high-resolution instruments such as LOFAR and the future SKA\_LOW (including the use of the latter as the core of a VLBI array). Finally, our catalogue will be invaluable for space weather studies using IPS, providing a network of IPS sources (with their baseline scintillation indices) which is unprecedented in its sky density.

The paper is organised as follows. In Section 2 we describe comprehensively the process of scheduling observations, selecting the first data release, calibrating and imaging our observations, making image-based measurements of source brightness both in continuum and variability. In Section 3 we describe how multiple measurements of each source are synthesised into a single catalogue entry per source. In Section 4 we discuss the sensitivity of our survey, and discuss issues that have the potential to impact on reliability and completeness. In Section 5 we discuss future work.

## 2. Methods 1: Scheduling, calibration, imaging and source-finding

The basic approach of using interferometric imaging to make IPS measurements in a single observation is described by Morgan *et al.* (2018). Morgan *et al.* (2019) expand this to performing a survey with multiple observations, and provide some prospective ideas on data reduction and processing, with a focus on MWA IPS data taken in late 2015 through to mid 2016.

Since the start of 2019, we have taken a great deal more IPS data, with the MWA in its Phase II extended configuration (Wayth *et al.* 2018), which has longer baselines (up to 5.3 km) and more uniform ( $u$ ,  $v$ ) coverage than the Phase I MWA. The potential advantages of the Phase II MWA for IPS studies are set out in

detail by Beardsley *et al.* (2019) Section 6.2.1. In light of these advantages, and the more uniform coverage (mostly due to the automatic scheduling algorithm described in Section 2.1), we have chosen to use only MWA Phase II data for this data release.

Below we describe our full methodology from scheduling the observations to generating the final catalogue. Much of this methodology has been previously presented by Morgan *et al.* (2018) and (2019). Where this is the case we have provided a summary and a reference to the relevant section of the other paper.

### 2.1. Scheduling

In order to provide good coverage for both astrophysical and space weather studies, we made daily 10-min observations of a number of target fields at an elongation from the Sun of approximately  $30^\circ$ , at a number of different orientations relative to ecliptic North. As in previous work, we split the available 30.72 MHz of instantaneous bandwidth into two equal bands centred on approximately 80 and 162 MHz. Only the upper band has been processed so far. The principal daily targets were due East and West along the ecliptic (i.e.  $90^\circ$  and  $270^\circ$  relative to Ecliptic North) as well as  $60^\circ$ ,  $120^\circ$ ,  $240^\circ$ , and  $300^\circ$ . At times (depending on the Declination of the Sun, and how much time was available) we added target fields at position angles of  $30^\circ$  and  $330^\circ$ ; or at  $150^\circ$ ,  $180^\circ$  and  $210^\circ$ , all at  $30^\circ$  solar elongation. These pointings overlap at the half-power point or closer, therefore providing uniform sensitivity.

As well as maximising the target sensitivity it is also necessary to minimise the response of the instrument at the location of the Sun. To facilitate automated, optimal scheduling of these observations we used the model of Sokolowski *et al.* (2017) to pre-calculate beams for all 197 ‘sweetpot’ pointings of the MWA at 162 MHz (the reader is referred to Morgan *et al.* 2019, Section 3.1 for a much more detailed description of MWA beams).

For each day of observations, an optimised pointing and observing time for each target was chosen as follows. All possible pointings and solar hour angles were exhaustively searched (the latter with a resolution of  $1^\circ$ ) to find which best match our criteria: the highest target sensitivity with at least 20 dB of suppression at the location of the Sun.

Next, a higher-resolution search was carried out to determine the precise 10-min observing interval over which the Sun is best nulled. Again, this was an exhaustive search, this time with a resolution of 8 s, since all MWA observations must start and stop at a time when the seconds since midnight is divisible by 8.

To avoid clashes between observations, the targets for the day’s observing were scheduled in strict order (those closest to the ecliptic were typically scheduled first). Code for automated scheduling is available on github.<sup>b</sup>

Overall, this resulted in 1448 observations scheduled almost every day between 2019 February 4 and 2019 August 18.


### 2.2. Array calibration

Each observation was downloaded as a measurement set,<sup>c</sup> and calibrated (Offringa *et al.* 2015) against a sky model based on the GLEAM survey (Hurley-Walker *et al.* 2017) using publicly available code (Hurley-Walker *et al.* 2022a),<sup>d</sup> with a few brighter

<sup>b</sup>[github.com/johnsmorgan/ips\\_plan](https://github.com/johnsmorgan/ips_plan).

<sup>c</sup>[asvo.mwatelescope.org](https://asvo.mwatelescope.org).

<sup>d</sup>[github.com/nhurleywalker/GLEAM-X-pipeline](https://github.com/nhurleywalker/GLEAM-X-pipeline).



**Figure 1.** Pointing centres for 263 observations selected for data release 1 are marked with blue points. The 7839 sources appearing in the final catalogue are plotted in grey to indicate the final coverage (an apparent lack of sources associated with the westernmost pointings is due to the criteria that a source be detected in 5 observations—see Section 2.8). Very bright ‘A-team’ sources are also shown.

sources from outside the GLEAM survey area characterised on the basis of other radio surveys.<sup>e</sup> Only baselines between 130 and 2600 m were used for calibration. The higher cut-off was due to the maximum baseline of (MWA phase I) GLEAM. The lower cut-off was to avoid contributions from the Sun (which was not included in the sky model due to its variable nature) which is mostly resolved out on baselines  $\geq 65\lambda$  (see Section 2.4 for further information). These calibration solutions consist of complex gains for each spectral channel for all 4 correlations products (XX, XY, YX, YY), for all 128 tiles. We then used these calibration solutions to triage our data. Two metrics were then used to determine the goodness of the calibration. The first was the fraction of the calibration solution for which a solution was obtained, the 2nd was the residual of a linear fit of the calibration solution phase as a function of frequency. These two metrics were then used to select observations for further analysis. Broadly speaking, around 88% of observations met our threshold for further analysis. Most observations that did not reach this threshold failed due to known issues with the array at that time. Some useful information can likely be gleaned from these observations in the future with more careful calibration and flagging.

### 2.3. Selection of data release 1

Observations were chosen to have good overlap with deep infrared surveys and high-resolution radio surveys, to allow a continuation of the work begun by Sadler et al. (2019). We have also focused initially on Northern-hemisphere observations, to facilitate the comparison we have made with International LOFAR (Jackson et al. 2022). These were then supplemented with observations covering the Galactic Plane and Orion region for pulsar searches (Chhetri et al., in preparation) and measurements of Galactic scattering (Morgan et al., in preparation). For the Easternmost part of the survey area we processed all observations taken over 49 d for two pointings, resulting in a very dense oversampling of the sky. For the rest of the survey we imaged only a subset of observations, resulting in around half the density (as can be seen in Figure 1).

In all, 263 observations were selected, all observed between 2019 February 21 and 2019 August 18 (close to solar minimum).


### 2.4. Imaging

For imaging purposes, measurement sets were generated with the native time resolution of 0.5 s, and a spectral resolution of 160 kHz. The outer 160 kHz channels of each 1.28-MHz coarse channel were flagged. This level of spectral averaging will cause some degree of bandwidth smearing towards the edge of the beam (Ord et al. 2015); but since we are mostly interested in the scintillation index of our sources, anything that affects the numerator (variability) and the denominator (mean brightness) in proportion will cancel out (except for a modest reduction in signal-to-noise due to slight decoherence on the longer baselines). CHGCENTRE, a companion tool to WSCLEAN (Offringa et al. 2014), was used to adjust the phase centre to the true primary beam maximum, and then to rotate the phases to the minimum- $w$  direction (i.e. close to the zenith). The latter dramatically reduces the computation required for imaging by reducing the number of  $w$ -layers required<sup>f</sup>, as well as ensuring that the PSF is as uniform as possible (in pixel space) across the resulting image.

Each observation was imaged twice: first a continuum image using all 10 min of data with a fairly deep clean was generated (hereafter called the ‘standard image’). This is to provide a means to measure the (mean) flux density of each source of interest (the divisor in the scintillation index). Secondly, each 0.5 s timestep of the observation is imaged separately, with only a shallow image-based clean (hereafter called the ‘snapshot images’). In order to ensure that precisely the same data were used for the dividend and divisor of the scintillation index, we made the decision to use precisely the same imaging image size, pixel size and visibility averaging and weighting for both imaging runs. This approach is feasible because the extended Phase II MWA has excellent ( $u, v$ ) coverage, and weighting schemes close to natural produce very good images. An innovation we have introduced since previous work is to subtract the model (that results from the deep

<sup>e</sup>[github.com/johnsmorgan/marco](https://github.com/johnsmorgan/marco).

<sup>f</sup>See, for example, Perley (1999) for a description of the  $w$ -term problem, and Offringa et al. (2014) for the approach that WSCLEAN uses to solve it.



**Figure 2.** The solid blue line shows the weighting scheme used, in arbitrary units. This scheme consists of zero weight for all baselines  $< 50\lambda$  (where  $\lambda$  is taken to be 1.85 m); a Tukey taper from  $50\lambda - 100\lambda$ , and a Gaussian taper equivalent to  $2'$  FWHM in the image plane. The vertical dashed lines delimit the range of baseline lengths used for calibration (see Section 2.2). The dash-dotted line is the Nyquist limit imposed by the  $1'$  pixel size in the image plane. The grey bars indicate the density of baselines in each annulus of the  $(u, v)$  plane.

cleaning of the standard image) from the visibilities before imaging the snapshots. This ensures that the cleaning done on the snapshot images is directed towards cleaning (and reducing the sidelobes of) varying sources. This reduces some artefacts due to very bright continuum sources.

WSCLEAN (Offringa *et al.* 2014; Offringa & Smirnov 2017), and associated tools have been used throughout for calibration, imaging and analysis. We have found WSCLEAN and associated tools for MWA data reduction (Offringa *et al.* 2015) to be highly performant, reliable, and flexible. In particular, the ease of imaging with a model subtracted, and the ability to fine-tune the visibility weighting scheme were extremely valuable for this project. The latter allowed us to find a weighting scheme which balanced a number of competing factors. Firstly, our IPS measurements are limited by the system noise (rather than confusion which is often the case e.g. Wayth *et al.* 2015), which favours the use of a weighting scheme that weights all baselines equally (i.e. natural weighting). Secondly, the Sun remains a contaminant in spite of strong suppression by the beam response of the instrument. The quiet Sun's power is concentrated in a relatively small number of short baselines, and while this large-scale structure is relatively static and therefore does not change on IPS timescales, it does cause strong ripples in a standard image.

We also wish to minimise the size of our images, both to reduce the computation required for imaging, and the storage required for the snapshot images. Inclusion of the longest MWA baselines would necessitate a small pixel size and therefore much larger images than used in our MWA phase I pilot studies.

Our compromise weighting scheme is summarised in Figure 2. Uniform weighting is applied to the data by WSCLEAN by default. Baselines  $< 50\lambda$  were then discarded completely. A Tukey taper (Harris 1978) was applied to baselines of length  $50\lambda < B \leq 100\lambda$ .

Finally, a Gaussian taper was applied to the data, equivalent to an image-plane FWHM of  $2'$ .

As shown in Figure 2, the Gaussian taper follows the baseline density as a function of baseline length, so its main effect is to reverse the uniform weighting and restore a more natural weighting scheme, albeit with strongly reduced weighting on the longest baselines in line with the resolution in line with the large pixel size. The large pixel size also excludes approximately 7% of baselines from being gridded. Nonetheless, tests showed that with these parameters the image noise (as measured by subtracting two consecutive snapshot images and measuring the standard deviation) was only approximately 10% higher than natural weighting.

The  $2'$  Gaussian taper also has the advantage that our resolution is well-matched to GLEAM (Hurley-Walker *et al.* 2017) at the appropriate frequency, making it easy to perform absolute flux density corrections.

For the snapshot images only, we also limited the number of  $w$ -layers to 24 (the number of CPU cores on the machine we were using). This reduced the computation time further, and in spite of the WSCLEAN 'suggested' number of  $w$ -layers being up to an order of magnitude higher, this did not appear to cause any problematic effects. All snapshots excluding the first 4 s and the last 16 s were imaged (c.f. Tian *et al.* 2022), leaving 1152 timesteps per 10-min observation.

Both XX and YY correlation products were imaged separately, and the resulting snapshot images (along with XX and YY standard images and beams) were stored in a single HDF5 precisely as described by Morgan *et al.* (2018), Appendix I. Calibration scaling factors as described below were subsequently stored in the same file, and these are applied on-the-fly when any relevant data are extracted for subsequent analysis.



## 2.5. Post-imaging calibration

In order to both combine both polarisations into ‘pseudo Stokes  $I$ ’ ( $I_l$ ) with optimal weightings, and apply an absolute calibration factor, we use Sault, Staveley-Smith, & Brouw (1996), Equation (1), in a very slightly modified form:

$$I_l(l) = \frac{\sum_p A(p)B(l,p)I(l,p)/\sigma^2(p)}{\sum_p A^2(p)B^2(l,p)/\sigma^2(p)}. \quad (1)$$

Here the sum is over the two polarisations  $p$  (XX and YY).<sup>g</sup> We separate the beam into an absolute correction factor  $A$  per polarisation (which, like the variance  $\sigma^2$  is assumed to be direction-independent), and a beam,  $B(l)$  which is direction-dependent (denoted by it being a function of  $l$ ), and generated for each pixel using the model of Sokolowski et al. (2017). For these, for efficiency, we pre-calculated all-sky beams and interpolated these as required using software available on github<sup>h</sup> and archived on Zenodo (Morgan & Galvin 2021).

After imaging, AEGEAN 2.2.0 (Hancock et al. 2012; Hancock, Trott, & Hurley-Walker 2018) was used to generate a source catalogue for the standard image for XX and YY separately. After applying primary beam corrections, selecting a subset of bright, unresolved sources, and comparing with the relevant measurements from GLEAM (Hurley-Walker et al. 2017) these could be used to determine the absolute calibration factor  $A$  for each polarisation. A small variability image (Morgan et al. 2018) for the central quarter of the full image was then generated from the snapshots for each polarisation for the purpose of measuring  $\sigma$ , which was taken to be the median of this image. Note that the direction-independent  $\sigma(p)$  values are only used for combining the two polarisations. When the noise is measured for a particular source, either the standard or variability images, as described below, a local measurement of Root-Mean-Square (RMS) is used.

Once the factors  $A(p)$  and  $\sigma(p)$  were calculated, it was possible to combine both polarisations of the standard image together using Equation (1). A variability image was then constructed using the timeseries data, with polarisations combined in the same manner, exactly as described by Morgan et al. (2018), Section 2.3. Essentially, the variability image is produced by taking the timeseries corresponding to each pixel of the image, applying a filter with a bandpass of 0.1–0.4 Hz, and then computing the RMS. The filter emphasises the frequency range where the IPS signal is strongest, while filtering out almost all ionospheric scintillation (Waszewski, Morgan, & Jordan 2022). Additionally, for this survey, we found that by applying a Tukey window (Harris 1978) to the timeseries before filtering we were able to vastly reduce the number of spurious detections, probably due to instrumental effects near the start and end of observations.

## 2.6. Standard image and variability image characteristics

The result of the process at this point was a single standard image and a single variability image for each observation; these were the

<sup>g</sup>NB: X and Y here are the instrumental polarisations corresponding to ground-based dipoles oriented East-West and North-South respectively. This does not accord with the IAU/IEEE definition of X and Y (Hamaker & Bregman 1996, and references therein) and the two are not necessarily orthogonal. However, this is unimportant for our purpose of combining these measurements while minimising noise.

<sup>h</sup>[github.com/johnsmorgan/mwa\\_pb\\_lookup](https://github.com/johnsmorgan/mwa_pb_lookup).

only data products that were carried forward for further analysis. BANE (part of the AEGEANTOOLS suite; Hancock et al. 2018) was used to produce ‘background’ and ‘rms’ images. MWA continuum images are confusion-limited at a level well above the thermal noise, and the noise tends to be higher around bright sources. The noise in variability images are generally thermal noise dominated, since the number of variable sources is lower than the total number of sources (and the scintillating flux density is less than the mean flux density). Sidelobes are only seen around the very brightest sources, and then only occasionally.

Furthermore, since the variability image measures standard deviation, all pixels have a positive value. The ‘background’ value, is the thermal noise level which is present whether a pixel corresponds to variable source or not.

The ‘rms’ image measures the spatial deviation of the background about its mean. If the lightcurves consisted of Gaussian random noise, the RMS would be  $\chi$  distributed, with 564 degrees of freedom (d.o.f.; after our application of a low-pass filter with a timescale of 1 s and Tukey window with the taper covering 4% of the timeseries). In practice, we observe a ratio of 32.25 between the background and RMS of the variability images, and this is remarkably consistent across all observations for all points in the image. This ratio implies a  $\chi$  distribution with  $\sim 520$  d.o.f. The  $\chi$  distribution converges towards a normal distribution with increasing degrees of freedom (more rapidly than the better known  $\chi^2$  distribution), justifying our assumption of Gaussian noise in the variability image (though we assume a  $\chi$  distribution with 520 d.o.f. when analysing our false detection rate; see Section 4.2).

## 2.7. Source finding and characterisation

Most MWA continuum surveys (e.g. MWACS, GLEAM; Hurley-Walker et al. 2014, 2017) mosaic together multiple observations in order to increase ( $u, v$ ) coverage and sensitivity before source-finding (see Carroll et al. 2016, for a counterexample). This allows the detection of sources that would fall below the level of significance in a single observation. We take a different approach here, and source-find separately for each observation. This is primarily due to the detection limit in a variability image only falling with the fourth root of time, meaning that the sensitivity increase from combining observations is fairly modest (40  $2\text{-}\sigma$  observations would be required to match the significance of a single  $5\text{-}\sigma$  detection).

For us, measuring at least a subset of sources in each observation independently is vital in order to determine how the scintillation index depends on the solar latitude of the piercepoint (i.e. the latitude of the point on the Sun beneath the point of closest approach of the line of sight to the Sun; see Section 3.1 below and also Morgan et al. 2018, Section 2.4), as well as checking for drastic space weather variations in a particular part of the sky in a particular observation. We know exactly to what extent (if at all) each individual observation contributes to a source’s catalogued properties, making it easier to exclude particular problematic measurements on a case-by-case basis.

We therefore, used AEGEAN to catalogue all sources detected at  $5\text{-}\sigma$  in all variability and standard images. To capture information on sources scintillating below  $5\text{-}\sigma$ , we also made a measurement in the variability image at the location of each detection in the standard image (there were typically an order of magnitude more  $5\text{-}\sigma$  detections in the standard image than the variability image). As will be explained in Section 3.2, this is vital in order to make an

unbiased estimate of the compactness of our sources. These measurements were derived from a simple 2D cubic interpolation of the 3×3 nearest pixels to the continuum position in the variability image, and its corresponding ‘background’ and ‘rms’ images.

Correction of source positions for ionospheric refractive shifts was carried out following Morgan *et al.* (2018), Section 3.2.1: that is the offsets of a subset of bright, isolated sources was used to construct a vector field using a radial basis function technique, and this was used to calculate corrected source coordinates for each detection. As before, we estimate that the typical error in position is well below 1′ after this correction.

### 2.8. Consolidation of measurements

The analysis so far has treated each observation entirely independently. We now organised our data so that measurements are grouped by source. At this stage we also rejected any measurements that do not meet the following criteria for inclusion. We used GLEAM as our reference catalogue, and only sources in GLEAM are in our catalogue (this excludes very few IPS sources, see Section 4.2 for further details). For each observation we first identified all GLEAM sources which lay within the quarter power beamwidth of the primary beam. For each of these sources we then determined whether a continuum or variability detection had indeed been made for that source (within a match radius of 1′). This resulted in approximately 2.5 million detections. Any variability detections which did not match with a GLEAM source (regardless of whether they had a counterpart in the continuum image) were stored for later analysis of false detections (see Section 4.2).

Measurements outside the elongation range 20°–40° were discarded and we further specified that there should be 5 continuum detections of our sources, since we found that with fewer measurements, a discrepant measurement (due to, e.g. a space weather fluctuation) could dominate and cause biased results. We also required at least one measurement within the half-power point of the primary beam. After these cuts there remained 748321 detections of 42838 individual GLEAM sources, each with between 5 and 57 measurements.

Note the two possible ways that variability image measurement may be made: either by a direct detection by source-finding in the variability image, or indirectly by measurement at the location of the continuum detection. Strong variability detections will have both, the remainder will only have the latter. For those with both, the direct detection value was used; otherwise the indirect detection was used.

## 3. Methods 2: Synthesis of the catalogue

Following Chhetri *et al.* (2018a), we quantify the IPS of each of our sources with the Normalised Scintillation Index (NSI). The normalisation is with respect to the Sun–source geometry: the scintillation index decreases with increasing distance from the Sun as the solar wind expands into 3D space. Additionally, the polar solar wind tends to be more diffuse than that emanating from the Sun’s equatorial regions. We follow Manoharan (1993) in assuming an elliptical form for the contour of equal scintillation index around the Sun, with the minor axis at the poles of the Sun and the major axis at the equator of the Sun (see Section 3.1). The result of this normalisation is that the NSI is zero for a source that shows no IPS, and unity for a source that scintillates like a point source.

In this section, we describe how we determine the NSI of each source given our data, how we classify our sources as detections or upper limits, how we determine the error on the NSI, and how we determine the sensitivity at the location of each of our sources.

While the mapping of an observed scintillation index to the NSI is straightforward, the process of deriving the scintillation index from the observed fractional variance while avoiding bias and quantifying the error is necessarily complex. The complication arises from the fact that we are measuring variance due to scintillation in the presence of variance due to thermal noise. Both types of variance are stochastic processes (albeit with different timescales).

The concept of debiasing the variability index to remove the effects of random noise is not new (Barvainis *et al.* 2005); and in previous work we noted that this leads to asymmetric errors (Morgan *et al.* 2018). However, in the present work, the use of multiple observations of each source, each with different sensitivities, moves us into an unusual domain where the variance of the signal we are looking for is often dwarfed by the variance due to thermal noise.

### 3.1. Normalisation of the scintillation index

The scintillation index of a point source is given by an empirical function of the form

$$m \propto \lambda(e \sin \epsilon)^{-b}, \quad (2)$$

(Manoharan 1993; see also Morgan *et al.* 2019 Equations (6) & (7)) where  $\lambda$  is the observing wavelength,  $e$  is the elliptical term, which depends on the solar latitude below the piercepoint and the ellipticity, and  $\epsilon$  is the solar elongation. While the constant of proportionality, the ellipticity and  $b$  may be expected to vary during the solar cycle, and from one solar cycle to the next, we find no strong evidence to vary these parameters from their respective standard values of 0.06, 1.5, and 1.6. Jackson *et al.* (2022) found that IPS estimates of compact flux density were around 20% lower than LOFAR estimates, which might justify applying a uniform amplitude scaling. However a new comparison with the most recent version of the catalogue, suggests that the discrepancy may be closer to 10%. The sources in common with LOFAR are the Northernmost of our range and very far South for LOFAR. For now, we do not apply any overall correction, though we may revisit this in a future data release.

We found a major:minor axis ratio of the contour of constant scintillation index of 1.5 to be a good fit to our data, which is a typical value for solar minimum as found by Manoharan (1993).

### 3.2. Determination of NSI

For the case of a single observation, the process required to infer the Normalised Scintillation Index for each source above a certain detection threshold is described fully by Morgan *et al.* (2018). Briefly we first determine  $\Delta S$ , the scintillating flux density, from the variability image as follows

$$\Delta S = \sqrt{P^2 - \mu^2} \quad (3)$$

(Morgan *et al.* 2018, Equation (4)) where  $P$  is the pixel value at the location of the source (as determined by fitting in the image plane), and  $\mu$  is the background level of the variability image (i.e. the variability due to thermal noise). Note that neither  $\Delta S$  nor  $S$  (the flux density measurement in the continuum image) are absolutely flux density calibrated. Any errors (due to e.g. a primary beam model

error) will cancel in the scintillation index  $\Delta S/S$ . Consequently, the current analysis of our data is insensitive to, and largely unaffected by longer-term variability such as intrinsic variability or refractive interstellar scintillation, which in any case is relatively rare at low frequencies (Bell et al. 2019). Where an absolute flux density measurement is required, such for determining sensitivity (Section 4), we assume that  $S$  is equal to the GLEAM flux density at the appropriate frequency (hereafter  $S_{162}$ ).

We reiterate that in the absence of scintillation,  $P - \mu$  has a Gaussian distribution with an expected value of zero, and a variance of  $\sigma^2$  (where  $\sigma$  is the spatial RMS in the variability image). In contrast, while there is a monotonic relationship between  $\Delta S$  and  $P$ , it is non-linear. The errors on  $\Delta S$  are non-Gaussian, and  $\Delta S$  is not real unless  $P > \mu$ .

We now wish to determine the NSI on the basis of multiple measurements. If we were to restrict ourselves to measurements where  $P - \mu > 5\sigma$  as in previous work, this would introduce a positive bias. Even with the  $5\sigma$  threshold relaxed,  $P - \mu$  can be negative in the presence of noise, and these values do not map to a real-valued  $\Delta S$  or NSI.

We have devised a scheme to determine the NSI that best fits all our measurements of variability without any need to remove negative or low-S/N measurements. First, we define the variability S/N  $\rho_{\text{var}}$ :

$$\rho_{\text{var}} = \frac{P - \mu}{\sigma}. \tag{4}$$

This is a ‘standardised’ statistic: in the absence of a scintillation signal it has zero mean and unit variance. For a particular NSI, the scintillating flux density that a source *would* have is simply

$$\Delta S(\text{NSI}) = S \cdot m_{\text{pt}} \cdot \text{NSI}, \tag{5}$$

where  $m_{\text{pt}}$  is the scintillation index a point source would have (a function of the Sun/source geometry; see Section 3.1). We can then determine  $\rho(\text{NSI})$ , the S/N the variability detection *would* have for a given NSI:

$$\rho(\text{NSI}) = \frac{\sqrt{\Delta S(\text{NSI})^2 + (32.25\sigma)^2} - (32.25\sigma)}{\sigma}. \tag{6}$$

32.25 is the ratio between  $\mu$  and  $\sigma$  in the variability image (see Section 2.6),

We then determine (by iterative fitting)  $\text{NSI}_{\text{fit}}$ , the NSI that minimises the residual sum of squares:

$$\text{RSS}(\text{NSI}) = \frac{\sum_i w_i (\rho(\text{NSI})_i - \rho_{\text{var},i})^2}{\sum w_i}, \tag{7}$$

where the sum is taken over all observations and  $w_i$  is the weight of the  $i$ th observation.

For the weights  $w_i$  we use the S/N using the continuum flux density  $S$  as the numerator and the noise in the *variability* image as the denominator. This weights our measurements by inverse thermal noise.<sup>1</sup> This fitting process is summarised in Figure 3, which

<sup>1</sup>This is a compromise between uniform weighting and inverse variance weighting. In addition, an arguable choice would be to further scale the weights by  $m_{\text{pt}}$ . This might marginally increase sensitivity by up-weighting measurements where the scintillating flux density is expected to be higher. However, observations close to the Sun also move closer to the strong scintillation regime where the expected relationship between solar elongation and  $m_{\text{pt}}$  may not hold. On the other hand, our chosen scheme weights observations close to the centre of the primary beam highest, moving the effective elongation of our NSIs towards 30°, making for a more uniform dataset.

illustrates the relationship between scintillation index, NSI,  $\rho_{\text{var}}$ , and  $w$ , as well as showing RSS as a function of NSI.

For purely practical purposes we allow negative NSIs (by giving  $\rho(\text{NSI})$  the sign of the NSI), which are non-physical, but provide a best fit for sources where  $\sum \rho_{\text{var}}$  is negative. The existence of sources with negative NSIs may be indicative of the reliability of the catalogue, so we report their number in Section 4.4. However, for the final catalogue (and in particular for the detection statistics presented in the next section),  $\text{NSI}_{\text{fit}}$  is set to zero for all sources with a best-fit NSI < 0.

### 3.3. Classification of detection/non-detection

Next we use  $\text{RSS}(\text{NSI})$  to generate simple statistics to determine the significance of the detection of IPS for each source. IPS is deemed to have been detected if

$$l_0 = \text{RSS}(0) - \text{RSS}(\text{NSI}_{\text{fit}}) > 25. \tag{8}$$

$l_0$  is twice the log-likelihood ratio between an NSI of  $\text{NSI}_{\text{fit}}$  and an NSI of zero. In the signal-absent case (and in the asymptotic limit of a large number of datapoints),  $l_0$  is  $\chi^2$  distributed with a single degree of freedom (since we have only one free parameter in our fit). As is well known, the range over which  $\text{RSS}(\text{NSI})$  is within 1 of its minimum value is the 68% (i.e.  $1\sigma$ ) confidence interval. Due to the quadratic form of a Gaussian log-likelihood,  $l_0 = 25$  is roughly equivalent to a  $5\sigma$  detection (James 2006), 25 being the  $l_0$  of a single  $5\sigma$  measurement.

We use a similar statistic,  $l_1$ , with a slightly less stringent threshold, to determine if we can place an informative upper limit on a source’s NSI (for those sources with  $\text{NSI}_{\text{fit}} < 1$ ):

$$l_1 = \text{RSS}(1) - \text{RSS}(\text{NSI}_{\text{fit}}) > 9. \tag{9}$$

Finally, for sensitivity calculations, it is useful to estimate, for each source, the NSI at which the source would only just be detected. We estimate this by finding the  $\text{NSI}_{\text{lim}}$  which satisfies the following equation

$$\frac{\sum w_i \rho(\text{NSI}_{\text{lim}})_i^2}{\sum w_i} = 25, \tag{10}$$

where  $\rho(\text{NSI})$  is as given in Equation (6). This is the NSI at which the source would reach our detection threshold in the absence of scatter.

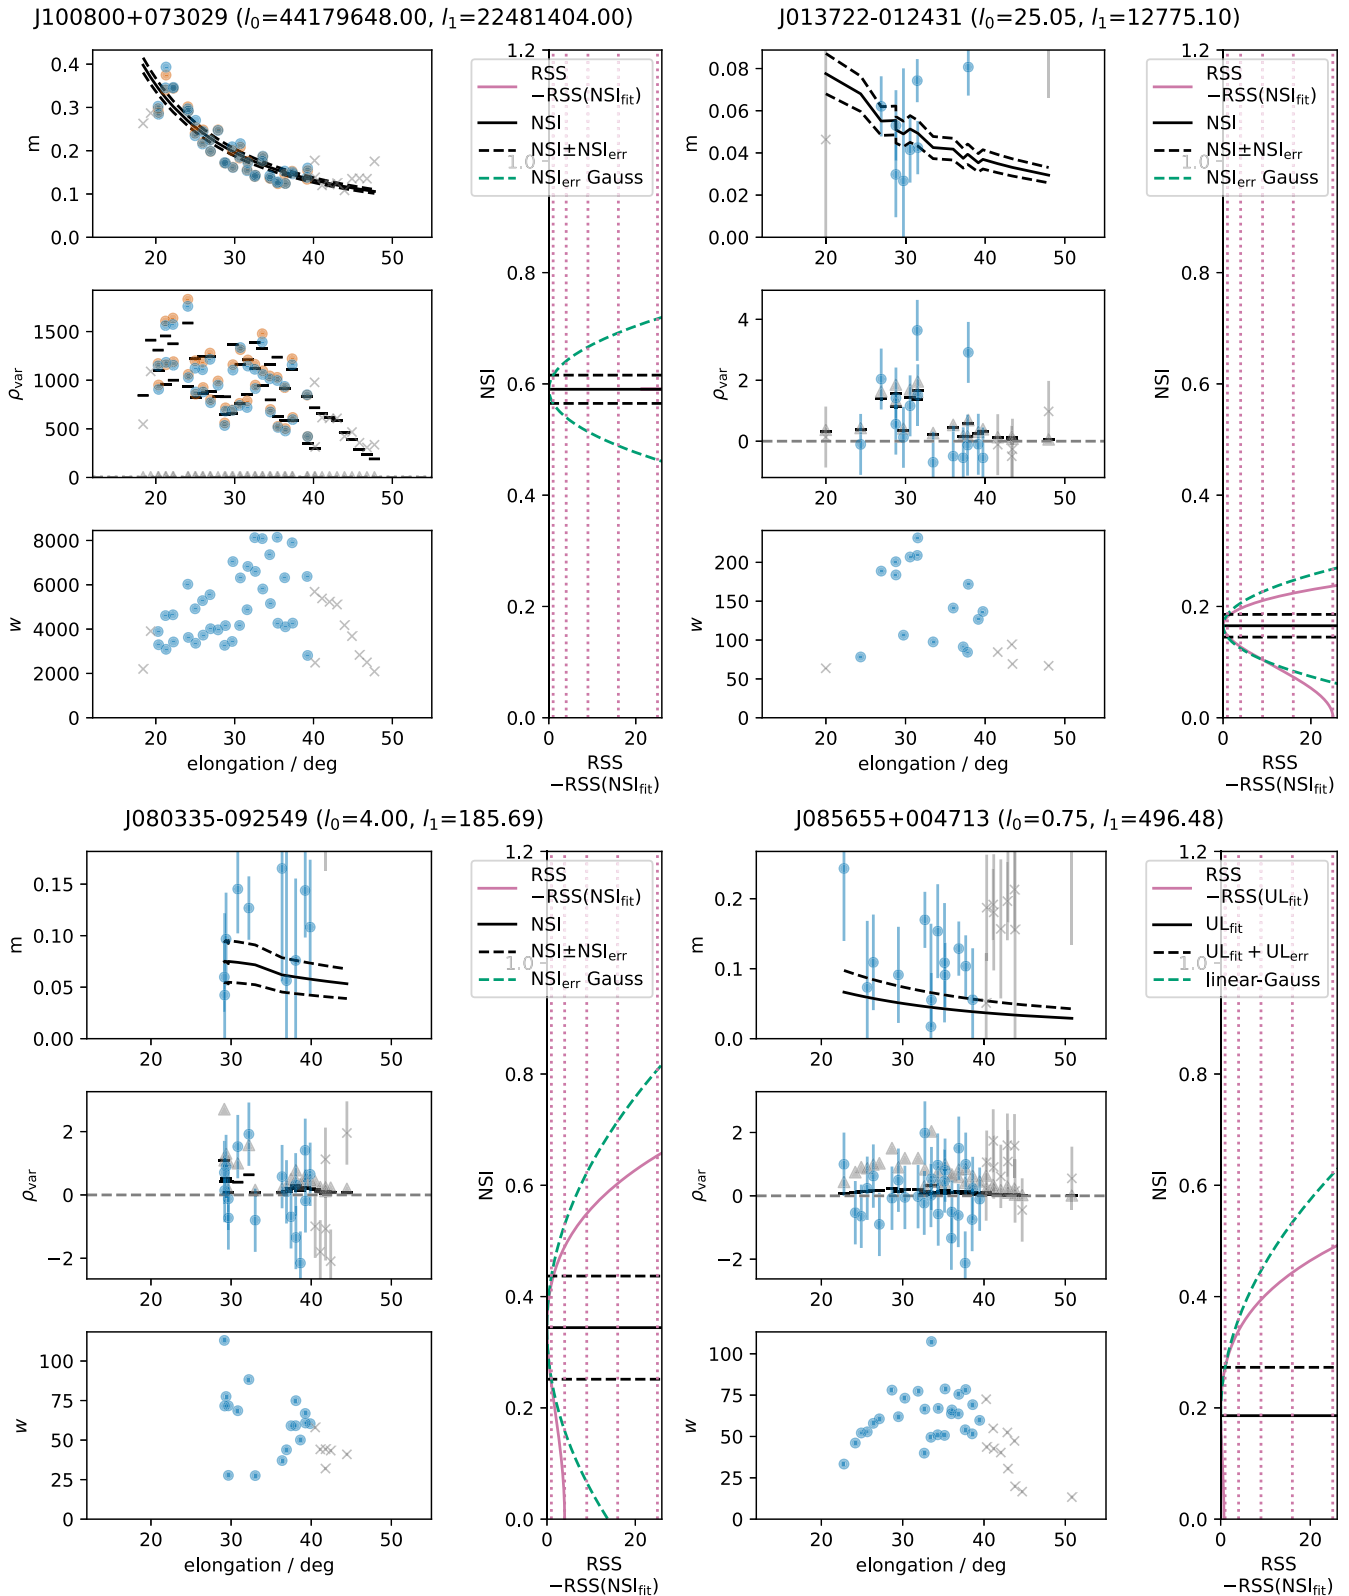
Figure 4 plots  $\sqrt{l_0}$  (essentially the S/N of  $\text{NSI}_{\text{fit}}$ ) against the ratio  $\text{NSI}_{\text{fit}}/\text{NSI}_{\text{lim}}$ . The scatter at higher S/N is the effect of space weather (see Section 3.4).

Note that in addition to  $\text{NSI}_{\text{lim}}$  we also calculate  $\text{NSI}_{5\text{lim}}$ . This is identical to  $\text{NSI}_{\text{lim}}$  except that only the 5 highest weighted measurements are used. The significance of 5 is that  $5 \times 5\sigma$  continuum measurements are required in order for a source to be included in a catalogue, so when multiplied by the flux density of the source to give  $S_{5\text{lim}}$ , this statistic measures the flux density of the weakest detectable NSI=1 source. See Section 4.1.1 for further details.

### 3.4. Error on the NSI for detected sources


Recall that  $\text{NSI}_{\text{fit}}$  is our estimate of the NSI of each source. We take the ( $1\sigma$ ) error due to thermal noise to be  $(\text{NSI}_{\text{upper}} - \text{NSI}_{\text{lower}}) / 2$ , where

$$\begin{aligned} \text{RSS}(\text{NSI}_{\text{lower}}) - \text{RSS}(\text{NSI}_{\text{fit}}) &= 1; \text{NSI}_{\text{lower}} < \text{NSI}_{\text{fit}} \\ \text{RSS}(\text{NSI}_{\text{upper}}) - \text{RSS}(\text{NSI}_{\text{fit}}) &= 1; \text{NSI}_{\text{upper}} > \text{NSI}_{\text{fit}} \end{aligned} \tag{11}$$



**Figure 3.** Figures summarising detections and NSI for a high-S/N source (upper left); a 5-sigma source (upper right); a 2-sigma marginal detection (bottom left) and a source for which there is a robust upper limit (bottom right). In each case, top panel shows scintillation index  $m$  (Equation (2)) as a function of elongation, 2nd panels from the top show  $\rho_{\text{var}}$  (Equation (4)) and bottom panel shows the weights (Equation (4)). In the top 2 panels, orange points indicate direct detections, blue indicates indirect detections as defined in Section 3.2. In the top 2 panels, black lines indicate the expected value for the NSI determined by the fit. In the middle panel the grey triangles show the expected value for  $\text{NSI}_{\text{lim}}$  (Equation (10)). The vertical panels on the right of each figure show the RSS statistic as a function of NSI (Equation (7)). Also plotted is the RSS that would be expected for Gaussian errors given by  $\text{NSI}_{\text{err}}$  (Section 3.4). For the upper-limit source, Equation (13) is used (Section 3.5).





**Figure 4.**  $\sqrt{I_0}$  (a S/N-like quantity) against the ratio  $NSI_{fit}/NSI_{5\sigma}$ . Blue points are detections, green points are marginal detections, pink points are sources with upper limits only. Two trend lines are shown to illustrate that in the weak signal limit, the S/N increases only with the square root of the NSI, and only for the strong detections is there a linear relationship.

(i.e. half the range in NSI over which RSS increases by  $< 1$ ).  $NSI_{lower}$  and  $NSI_{upper}$  are calculated automatically using the MINOS routine of IMINUIT (James & Roos 1975; Venzon & Moolgavkar 1988; Dembinski et al. 2022), which is used throughout for iterative fitting.

There is an additional error due to stochastic changes in the solar wind, as well as a much smaller uncertainty arising from the stochastic nature of scintillation itself (see Morgan et al. 2018, Equation (7)). To quantify the combined effect of these, we chose a sample of 120000 IPS measurements with sufficient S/N that thermal noise is negligible. The ‘g’-factor—the observed scintillation index compared to expected (from  $NSI_{fit}$ )—was then calculated for each. The distribution of g values had a standard deviation very close to 25%.

This variance will be amplified by the non-uniform weighting of measurements used when determining  $NSI_{fit}$ . To mitigate this, we determine the error due to space weather by dividing 25% by the square root of the ‘effective sample size’ (Kish 1965) given by


$$N_{eff} = \frac{(\sum w_i)^2}{\sum (w_i)^2}. \tag{12}$$

These two sources of error are plotted separately in Figure 5. They are combined in quadrature as  $NSI_{err}$ , leading to a fractional error distribution which is reasonably uniform in the range 4.8%–11.4% (which encompasses 80% of detected sources).

Consistent with previous studies (e.g. Tappin 1986), we found an excess of extreme values compared to a Gaussian distribution.

### 3.5. Upper limits

Upper limits on the NSI are useful, since setting robust limits on the fraction of the radio source that is sufficiently compact to scintillate is a useful astrophysical tool when interpreting radio source



**Figure 5.** The two main sources of error on the NSI (plotted as a fractional error) as a function of S/N. Dashed lines show logarithmically-spaced density contours for the Space Weather fractional errors.

morphology. In part as a result of the non-linearity between the strength of an IPS detection, and the implied scintillating flux density (Equation (3)) we have a very large number of sources for which scintillation is not unambiguously detected, but where we can rule out an NSI of 1. This is illustrated clearly in the bottom-right panel of Figure 3, which shows the residual sum of squares for such a source as a function of NSI.  $RSS(NSI)$  is almost flat below  $NSI = 0.25$  before increasing more steeply than quadratically for higher values of NSI.

We characterise the upper limits via two numbers:  $UL_{fit}$  and  $UL_{err}$ . The former is the maximum-likelihood NSI using a least squares fit (or zero if the fitted  $NSI < 0$ ).  $UL_{err}$  is upper error on  $UL_{fit}$  (allowing for asymmetric errors), determined as being the NSI at which the sum of squares increases by 1 (c.f. Equation (11)).

While  $UL_{fit}$  and  $UL_{err}$  are derived in a very similar way to  $NSI_{fit}$  and  $NSI_{err}$ , it is critical to emphasise that  $UL_{fit}$  is not a useful estimate of the NSI, and any NSI less than  $UL_{fit}$  is consistent with our data. For this reason we give this column a name that cannot be confused with an NSI by a casual user of our catalogue.

For sources for which upper limits are given, an approximate (un-normalised) log probability distribution function for the NSI given our data is given by

$$\log(p(NSI)) = \begin{cases} -\left(\frac{NSI - UL_{fit}}{UL_{err}}\right)^2 & ; NSI > UL_{fit} \\ 0 & ; \text{otherwise} \end{cases} \tag{13}$$

that is, Gaussian in likelihood above the maximum likelihood NSI, uniform below.

### 3.6. Catalogue contents

We now split all of our sample of sources (as defined in Section 2.8) into the following categories: ‘detected’ sources are those as defined above. We further define ‘marginal’ detections as those for which  $4 < I_0 \leq 25$  (i.e. 2-5 $\sigma$  detections). ‘upper-limit’ sources

are sources that are not ‘detected’ or ‘marginal’, but satisfy the criterion for an informative upper limit ( $NSI_{fit} < 1$ ;  $l_1 > 9$ ). Any remaining sources are not listed in the final catalogue.  $NSI_{fit}$  and  $NSI_{err}$  are provided only for ‘detected’ and ‘marginal’ sources;  $UL_{fit}$  and  $UL_{err}$  are only provided for ‘upper-limit’ sources.

The low threshold of  $2\sigma$  for the most ‘marginal’ detections is chosen since these sources have roughly symmetric errors within  $\pm 1\sigma$  (see Figure 3) and so are best described as a ‘marginal’ detection with symmetric errors rather than an ‘upper limit’ with asymmetric errors. Clearly some will be false detections (even if the majority are true). In any case,  $l_0$  and  $l_1$  are provided in the catalogue (allowing, e.g.  $< 3\sigma$  detections to be excluded or treated as upper limits where appropriate).

The GLEAM ID and the flux density of the GLEAM source at the relevant frequency is also provided. Our observing bandwidth is aligned with GLEAM, and the GLEAM flux density we report for each source is an (inverse variance weighted) average derived from the relevant GLEAM measurements and their errors. We also list  $S_{slim}$ : the product of  $NSI_{slim}$  (see Section 3.3) and the GLEAM flux density. This provides an estimate of the detection limit of the scintillating component of a source at that location.  $S_{cont\ lim}$ , the detection limit in the standard images (see Section 4.1.1) is also provided.

These are shown in Table 1. A description of all columns in the catalogue is given in Table 2. The full table is available online<sup>j</sup>

#### 4. Sensitivity, reliability and completeness

The result of the methodology outlined in the previous two sections is a table of 7839 detections, 5550 marginal detections, and 26367 sources with upper limits. 74% of these lie within an area of 4875 square degrees bounded by  $1\text{ h} < RA < 11\text{ h}$ ;  $-10 < Decl. < +20$ . For that region, Figure 6 shows the sources for each of these categories in the context of all the GLEAM sources in the same region. Almost all sources brighter than 160 mJy at our observing frequency appear in our catalogue. In the remainder of this section we quantify more clearly the sensitivity as a function of survey area, and describe various complications which may affect the reliability and completeness of our survey.

##### 4.1. Sensitivity

In contrast to most astronomical surveys, our sensitivity at a given location on the sky depends on two factors. This is because detections must fulfil both the standard image detection criteria, and the variability detection criteria. While the detection limits imposed by these two sets of criteria are correlated, they are quite distinct: with the former dominated by sidelobe confusion, and the latter dominated by thermal noise (though Sun–source geometry also plays an important role). Over much of the survey area, the continuum detection threshold is significantly lower than the variability detection threshold, and so, with certain caveats discussed below, the sensitivity to variability alone determines whether a source is detected or not, and the detection limit for a particular compact flux density is well-defined.

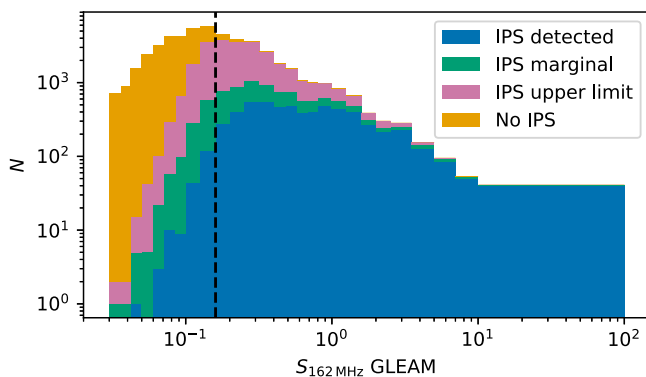
<sup>j</sup><https://doi.org/10.5281/zenodo.7256913>.

**Table 1.** First 11 lines of catalogue. Columns are fully described in Table 2.

GLEAM (1)	RAJ2000 (2)	DEJ2000 (3)	s_162 (4)	elongation (5)	class (6)	nsi_fit (7)	nsi_err (8)	ul_fit (9)	ul_err (10)	l_0 (11)	l_1 (12)	s_cont_lim (13)	s_slm (14)	n_fit (15)	n_eff (16)
J000311+093114	0.79639	+9.52069	0.263	28.3	upper_limit	-	-	0.146	0.139	0.1	217.3	0.245	0.146	5	4.35
J000315+111538	0.81368	+11.26067	0.535	31.3	upper_limit	-	-	0.000	0.191	0.1	310.5	0.269	0.276	5	4.65
J000327+105553	0.86297	+10.93145	0.312	30.3	upper_limit	-	-	0.313	0.082	2.9	198.2	0.177	0.174	6	5.34
J000501+095054	1.25791	+9.84843	0.265	31.2	upper_limit	-	-	0.451	0.119	2.9	39.5	0.262	0.207	5	4.35
J000515+102442	1.31362	+10.41180	0.249	31.7	upper_limit	-	-	0.448	0.111	3.4	46.5	0.222	0.186	5	4.42
J000556+093701	1.48703	+9.61714	0.304	30.2	detected	0.802	0.101	-	-	78.1	20.1	0.194	0.186	6	5.14
J000608+105209	1.53332	+10.86923	0.352	30.3	upper_limit	-	-	0.119	0.177	0.0	155.2	0.200	0.223	6	5.53
J000634+084808	1.64536	+8.80223	0.262	30.3	upper_limit	-	-	0.131	0.172	0.1	147.4	0.195	0.167	6	5.21
J000649+082014	1.70647	+8.33732	0.756	30.6	detected	0.212	0.033	-	-	26.2	4926.3	0.166	0.158	5	4.67
J000707+093004	1.78035	+9.50130	0.274	30.2	upper_limit	-	-	0.303	0.093	2.0	164.9	0.158	0.162	6	5.24
J000712+105434	1.80150	+10.90963	0.315	30.5	marginal	0.596	0.098	-	-	17.4	49.2	0.182	0.212	6	5.47

**Table 2.** Description of all columns in catalogue (Table 1).

Number	Name	ASCII name	Units	Description
1	–	GLEAM	–	GLEAM ID from Hurley-Walker et al. (2017) (searchable in NED if prepended with GLEAM)
2	–	RAJ2000	deg	GLEAM RA
3	–	DEJ2000	deg	GLEAM Decl.
1	$S_{162}$	s_162	jansky	GLEAM flux density at 162 MHz (see Section 3.6)
5	$\epsilon$	elongation	deg	Weighted mean solar elongation of observations
6	–	class	–	Class: ‘detected’, ‘marginal’ or ‘upper_limit’ (see Section 3.6)
7	$NSI_{fit}$	nsi_fit	–	Normalised Scintillation Index (NSI; see Equation (7))
8	$NSI_{err}$	nsi_err	–	Error on NSI (see Section 3.4)
9	$UL_{fit}$	ul_fit	–	Upper limit on NSI (see Section 3.5)
10	$UL_{err}$	ul_err	–	Error on Upper limit on NSI (see Section 3.5)
11	$l_0$	l_0	–	Measure of likelihood that NSI is non-zero (see Equation (8))
12	$l_1$	l_1	–	Measure of likelihood that true NSI is non-unity (see Equation (9))
13	$S_{cont\ lim}$	s_cont_lim	jansky	Continuum limit on detection (see Sections 3.3 & 3.6)
14	$S_{slim}$	s_slim	jansky	Limit on variability detection (see Sections 3.3 & 3.6)
15	–	n_fit	–	Number of observations used in fit
16	$N_{eff}$	n_eff	–	Effective number of observations used in fit taking weights into account (see Equation (12))



**Figure 6.** Stacked histogram of all GLEAM sources within 1 hr  $< RA < 11$  h;  $-10^\circ < Decl. < +20^\circ$ , with colour showing status in our catalogue: detected, marginal, upper limit, not in catalogue. Bins following Franzen et al. (2016). Vertical dashed line indicates 0.16 Jy, the lower edge of the bin in which 82.5% of sources are in the IPS catalogue.

In limited sky areas, our variability measurements are more sensitive than the continuum measurements. This might mean, for example, that a 100 mJy compact component can be detected if it is a component of a 1 Jy continuum source (i.e. NSI of 10%) but not if it is an isolated compact source (i.e. NSI of 100%).

Another less fundamental issue is the difficulty of assessing our sensitivity for an arbitrary point on the sky. We sidestep this problem by calculating our sensitivity metrics only at the locations of GLEAM sources which fulfil our selection criteria. The reader can estimate the sensitivity for an arbitrary location by assuming the sensitivity is the same as it is at the location of nearest catalogued source. The area of the celestial sphere which is closest to a given source than any other (i.e. the Voronoi cell of each source) has been calculated (Caroli et al. 2010) and these areas are used to determine sensitivity metrics as a function of survey area presented below.

#### 4.1.1. Sensitivity in continuum and variability


The key criterion for continuum detection is 5 detections at  $5\sigma$ , so the detection threshold for a given source location is simply  $5 \times$  the RMS noise in the 5th highest S/N continuum detection  $S_{cont\ lim}$  (see Section 3.6). The detection threshold for detection of variability (i.e. compact structure) is complicated, as it depends on the precise data, but it can be estimated by  $S_{slim}$ . Figure 7 shows the spatial distribution of sensitivity to compact structure. The detection limit varies from  $\sim 0.1 - 1$  Jy with good sensitivity being associated with low ecliptic latitude and high Galactic latitude.

The cumulative distribution of sensitivity to compact structure as a function of area is presented in Figure 8. To make this more clear, we have spatially smoothed the sensitivity at the location of each source, by making it a weighted average of all sources that lie within a radius of  $5^\circ$ , where, following (Epanechnikov 1969), the relative weight is given by


$$w(r) = \begin{cases} 1 - \left(\frac{r}{5}\right)^2 & ; r < 5^\circ \\ 0 & ; \text{otherwise.} \end{cases} \quad (14)$$

This shows more clearly that the majority of the survey area has a detection limit below 0.2 Jy with more than 2000 square degrees with a sensitivity  $< 0.15$  Jy.

Figure 8 also shows that for most locations, the continuum sensitivity limit is below the variability limit. This is explored further in Figure 9 which highlights the rare areas where the inverse is the case. These areas are associated with very bright continuum sources, since the sidelobes of these sources will cause the noise in the continuum images to be increased. The edges of the survey area, particularly the Western edge, are also affected since these areas are more sparsely sampled (see Figure 1). In many cases, only the bright source itself (and perhaps a small area around it) are affected. As noted above, in these areas it is possible that a compact source whose IPS variability might have been detected will not make it into the catalogue as it will not have fulfilled the criteria for continuum detection. However, these areas are quite



**Figure 7.** Map of all catalogued sources. Colour bar is the compact flux detection limit in jansky.



**Figure 8.** Cumulative distribution of (variability) sensitivity as a function of survey area. To construct the black line, the sensitivity measurements at the location of each source (shown in Figure 7) are spatially smoothed, and the sources are put in order of smoothed sensitivity. Each point in the survey area is associated with its nearest source (Voronoi tiling) and thus each source has an area associated with it. The  $x$ -axis is then the cumulative area represented by these sources. Thus, approximately 5000 square degrees of survey area have a detection limit (to compact structure)  $\leq 0.2$ Jy. Blue points are the variability sensitivity measurements without spatial smoothing. Pink points are the (unsmoothed) continuum sensitivity for the same sources.

limited in extent, and only a small fraction of IPS sources are likely to be affected. For example, if the ratio between variability sensitivity and continuum sensitivity is 0.9, then only sources with an NSI  $> 0.9$  within 10% of the variability detection limit will be excluded.


There is another mechanism by which source with a lower NSI might be detected while one with a higher NSI (but the same compact flux density) might not. Sources with a higher NSI may be detected in more continuum images, and therefore will have more variability measurements. A number of factors combine to make this effect negligible. First, most sources in our catalogue are detected in continuum in every observation where the source is within the quarter power point. Secondly, the impact of additional continuum detections on the IPS sensitivity is modest, both because they will have low weight, and because IPS sensitivity only increases with the fourth power of the number of observations for sources near the detection limit.

#### 4.2. False detections

In addition to the 2.5 million  $5\sigma$  detections within  $1'$  of a GLEAM source, there are an additional 854 unassociated  $5\sigma$  detections of variability, within the GLEAM coverage area (and fulfilling the criteria of being within the quarter power point of the primary beam). Most of these can be explained: the Crab Nebula and Hydra A are deliberately excluded from GLEAM and account for 45 detections, as well as 68 detections which are deemed to be their sidelobes. 10 sources appear to be genuine IPS detections from the Western hotspot of PKS0945+07, a wide double which is characterised in GLEAM (J094746+072509) as an ellipse not fully encompassing the source. Similarly 3 detections appear to be associated with the Northern component<sup>k</sup> of GLEAM

<sup>k</sup>NB White *et al.* (2020) consider the two components of GLEAM J101051-020137 to be unassociated.





**Figure 9.** Comparison of variability sensitivity to continuum sensitivity. Grey points (the majority of the map) are those sources where the continuum sensitivity exceeds the variability. Coloured points are those where the variability sensitivity exceeds the continuum sensitivity and the colour bar indicates the ratio. Bright GLEAM sources are shown in magenta, with the size of the points being proportional to the flux density at 162 MHz.

J101051-020137. 17 detections appear to be genuine detections of IPS from TGSSADR J093908.4+020059, a 1 Jy source (not in GLEAM, but situated approximately  $2'$  from GLEAM J093918+015948, a 4 Jy source). The pulsar PSR 0953+10, which is not in GLEAM, is also detected 3 times, the only source below the sensitivity limit of GLEAM from which we have unambiguously detected variability. The Moon is detected twice (probably due to reflected RFI; McKinley et al. 2013), while Jupiter (which might be expected to give rise to continuum and variability detections) is outside our survey area.

A notable source of false detections is the Pulsar B0950+08 (GLEAM J095309+075539), which is an extremely bright pulsar with a period of  $\sim 0.25$  s and has one of the lowest dispersion measures known (Pilkington et al. 1968; Manchester et al. 2005). This pulsar produces strong, sidelobe-like artefacts within a radius of approximately 7 degrees, accounting for 116 false detections. These artefacts are much more pronounced, and more symmetric, than the sidelobes we occasionally see around IPS sources, and therefore are unlikely to be caused by time variability alone. The highly symmetric nature of these artefacts suggests an amplitude error, and we suggest that it arises due the pulses only occupying a narrow line in the dynamic spectrum, meaning that the pulsar cannot be correctly deconvolved using the PSF for the full bandwidth.

The remaining 500 (presumed) false detections are not evenly distributed. Just 62 observations account for 301 of the false detections, while 41 observations have no false detections and 83 have just one. A couple of observations have visible artefacts due to solar bursts, and we presume that most of the remaining observations with 5 or more false detections have similar issues at a much lower level (i.e. either solar activity or RFI). The number of observations with 0–4 false detections is consistent with a Poisson distribution with a mean of approximately 1.75. With approximately 737000 resolution units within the quarter power point of the primary beam for a typical observation, the appropriate  $\chi$  distribution (see Section 2.6) would predict only 0.34 detections/observation more than 5 standard deviations from the mean. On the other hand, the rate 1.75 detections/observation is reached at only a slightly lower significance:  $4.65\sigma$ . This underlines the exceptionally clean

nature of our data, and is testament to the extremely quiet observing conditions at the Murchison Radio Observatory. It bodes extremely well for the use of the MWA to detect further transients similar to GLEAM-X J162759.5-523504.3 (Hurley-Walker et al. 2022b) as well as other radio transients that might be expected on similar timescales to IPS (e.g. prompt emission from gamma-ray bursts: Tian et al. 2022).

#### 4.3. Effect of sub-arcminute structure

As described above, we have used GLEAM as our reference catalogue, and it matches extremely well in resolution with our continuum images. In order to allow measurements of scintillating flux density in cases where the scintillation signature is not detectable at a high level of significance, we also make point measurements of the variability measurements at the location of the continuum source (see Section 2.8). This approach may lead to an underestimate of the scintillating flux density if the centroid of the source at GLEAM resolution does not match with the centroid of the scintillating emission.

We can measure this effect for sources that have sufficiently high S/N in individual observations that we can determine the offset between continuum and variability centroid. 3920 sources have at least 3 direct detections in the variability image. Although this sample of ‘detected’ sources may be expected to be biased towards more compact objects, the full range of NSIs is represented, and only 1% of sources show an offset between continuum and variability centroid of more than  $35''$ . Even an offset of this magnitude would only incur a 21% reduction in the measured scintillating flux density. However, it is nonetheless possible that some scintillating sources embedded in asymmetrically in extended structure (e.g. single hotspots) may have underestimated NSIs due to this effect.

#### 4.4. Reliability and completeness

These investigations demonstrate that, for the most part,  $S_{5\text{lim}}$  (the proxy for sensitivity, given in the catalogue for the location of each source) gives a good indication of the sensitivity.

Examining the distribution of scatter about the trend-line in Figure 4 in the range  $2 < \sqrt{l_0} \leq 10$  we estimate that our ‘detected’ sample is 10% complete at  $0.61S_{5\text{lim}}$ , 50% complete at  $0.80S_{5\text{lim}}$ , 90% complete at  $0.96S_{5\text{lim}}$ , and 99% complete at  $1.08S_{5\text{lim}}$ . The reason that the 50% completeness is not reached closer to  $0.5S_{5\text{lim}}$  is due mainly to the difference between  $S_{\text{lim}}$  and  $S_{5\text{lim}}$  (see Section 3.3).

The extremely small number of false detections of variability give us confidence that our reliability is extremely high, since it would require the chance coincidence of a noise spike with a known GLEAM source on more than one occasion. The only indication that we have that the reliability is less than 100% is a small number of sources for which the sum of variability S/N ratios ( $\rho_{\text{var}}$ ; see Section 3.2) is less than 0. This indicates that the variance at the location of the source is less than would be expected due to thermal noise at the location of the source, which is unlikely to be physical. A very small number of sources appear to have uniformly negative noise to an extent which is unlikely to be due to chance, and we suggest that this may be due to a poor estimate of the thermal noise level ( $\mu$ ; see Equation (3)). 240 such sources have a level of significance that would place them in the ‘marginal category’ were the sign of the variability reversed. None would reach the threshold for the ‘detected’ category.

## 5. Discussion

### 5.1. Comparison with previous work and predictions

As shown in Figure 9, we achieve a detection limit for a compact component  $\leq 0.2$  Jy over almost 5000 square degrees. This is at least a factor of two improvement in sensitivity over our previous work, and is in line with the predicted effect of switching to a more natural visibility weighting scheme (Morgan *et al.* 2019, Section 3.4). The use of data from multiple observations (whether a particular source is detected at  $5\sigma$  in variability or not) has also led to an increase in sensitivity, and it is likely that even greater sensitivity will be realised in the RA range  $21 \text{ h} < \text{RA} < 24 \text{ h}$ , where the Galactic latitude is high, but the ecliptic is further south, and therefore closer to the zenith at the MWA.

A fundamental property of our IPS measurements near the detection limit is the asymmetric probability distribution function for the scintillating flux density and parameters derived from it. This arises since it is necessary to subtract the variance due to scintillation from that due to thermal noise, expressed in Equation (3) (c.f. the ‘Debiased Variability Index’ Barvainis *et al.* 2005, Equation (1)). Since we have 1152 time samples at 0.5 s resolution (or approximately 500 scintillation timescales) per observation, we are able to detect and measure a scintillation signature even when the variance due to scintillation is only a fraction of that due to thermal noise. This means we can robustly detect sources with only approximately 50 mJy of scintillating flux (corresponding to a compact component of approx. 160 mJy at  $30^\circ$  solar elongation).

The unavoidable consequence is that for sources weaker than this, the scintillating flux density cannot be derived from our data, which is consistent with a range of values for the scintillating flux including zero. However, we derive upper limits on any scintillating flux density for a very large number of sources, ensuring that we extract the maximum amount of information from our data.

### 5.2. Future work

In this work we have set out in detail the methodology that we will use to conduct IPS surveys. While there remains some scope for optimisation, we expect that the fundamental approach that we have taken here to combine measurements from multiple observations while keeping track of errors and sensitivity is one that should serve us well into the future as we increase our survey area, and as we move to new instruments, such as ASKAP (Chhetri *et al.* 2022), and, in time, SKA\_LOW.


The current data release contains more than an order of magnitude more sources than our previous catalogues. We also probe down to fainter sources and determine the NSI with greater precision. Therefore, in future work, we can revisit previous topics with greater statistical power (Chhetri *et al.* 2018a,b; Sadler *et al.* 2019). We expect that as we are now pushing towards lower flux densities we will be able to provide compactness constraints on sources that are being searched for HI absorption (Sadler *et al.* 2020).

We now have IPS measurements for over 250 sources classified by Callingham *et al.* (2017) as having peaked or otherwise curved spectra, over 90% of which we class as ‘detected’. These sources have NSIs ranging from 0.5–1.0 (strongly weighted towards the high end) providing at least some information on angular size, which has been so important in understanding the nature of GHz Peaked Spectrum sources (O’Dea 1998).

The current catalogue also allows us to use the MWA for Space Weather research, by providing a dense network of IPS sources with known scintillation indices. By comparing the scintillation index observed in an individual MWA observation with this baseline level for each source (the ‘g-level’), we are able to map out structures in the solar wind such as CMEs (Morgan *et al.*, in preparation). In Section 3.4 we conducted a preliminary analysis of ‘g-levels’ with the purpose of determining the error that this imposes on our IPS measurements. By plotting these ‘g-levels’ spatially for each observation, we have already discovered a variety of structures in the solar wind which, due to the space density of our sources, we are able to map in unprecedented detail (Waszewski *et al.*, in preparation).

Another aspect of our data which remains largely unexplored is the wealth of information contained within the timeseries, which in the present work we have boiled down to a single standard deviation measurement. IPS power spectra are routinely used to characterise the solar wind (e.g. Tokumaru *et al.* 2021, and references therein). However, weak scintillation power spectra also encode the visibility amplitude of the radio source on scales close to the Fresnel scale (Macquart & de Bruyn 2007). Sources with an NSI very close to unity are unresolved to IPS, and the power spectrum may not contain source structure information even for IPS-resolved sources (i.e.  $\text{NSI} < 1$ ) if the features responsible lie on scales much larger than the Fresnel scale (but sufficiently small that they are unresolved at the  $2'$  interferometric resolution of the MWA; see Section 2.6 of Morgan *et al.* 2019). Preliminary research indicates that data on intermediate scales is responsible for many sources being IPS-resolved, but that substantial structure is also seen close to the Fresnel scale for many sources (Hedge *et al.*, in preparation).

In Section 4.3 we discuss the potential offset of GLEAM source positions from the location of the IPS source (see also Morgan *et al.* 2018, Figure 3 and accompanying text). The relative locations of the continuum and compact component centroids provides a further quantum of information on the structure of our sources on



**Figure 10.** Galactic plane (near anti-centre) and Orion region H intensity from all-sky image of Finkbeiner (2003) overlaid with our IPS detections (circles) and tentative detections (crosses). Note the very distinct areas where there is a much lower number density of IPS sources, even towards the centre of the survey area. These areas are often associated with H $\alpha$  emission.

compact scales. We defer a more detailed analysis to future work; however we note that approximately half of ‘detected’ sources have at least 3 variability detections at  $5\text{-}\sigma$ , and many of these align well with a particular source component visible in images from TGSS ADRI (Intema et al. 2017).

Also evident in the current data is that there are areas of the sky with a notable lack of IPS sources. We interpret this as the effect of scatter broadening due to turbulence in the ionised Interstellar Medium (ISM), which has previously been invoked to explain a well-known secular decrease in the number density of IPS sources towards the Galactic plane (Readhead & Hewish 1972; Rao & Ananthakrishnan 1984; Hajivassiliou 1992). Now, however, for the very first time, we have sufficient source density to demonstrate not just a trend with galactic latitude, but a very clumpy distribution of ionised turbulence; correlated strongly with H $\alpha$  emission (see Figure 10). This new science will be explored in a planned publication (Morgan et al., in preparation).

**Acknowledgements.** We dedicate this paper to the memory of our friend and colleague Jean-Pierre Macquart. We thank Natasha Hurley-Walker for making her sky model publicly available. We thank Natasha Hurley-Walker and Phil Edwards for their careful reading of the manuscript. This scientific work makes use of the Murchison Radio-astronomy Observatory, operated by CSIRO. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the operation of the MWA is provided by the Australian Government (NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited. We acknowledge the Pawsey Supercomputing Centre which is supported by the Western Australian and Australian Governments. We acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA), part of the High Energy Astrophysics Science Archive Center (HEASARC). HEASARC/LAMBDA is a service of the Astrophysics Science Division at the NASA Goddard Space Flight Center.

## References

- Balasubramanian, V., Janardhan, P., Ananthakrishnan, S., & Manoharan, P. K. 1993, *BASI*, **21**, 469
- Barvainis, R., Lehar, J., Birkinshaw, M., Falcke, H., & Blundell, K. M. 2005, *ApJ*, **618**, 108
- Beardsley, A. P., et al. 2019, *PASA*, **36**, e050
- Beasley, A. J., Gordon, D., Peck, A. B., Petrov, L., MacMillan, D. S., Fomalont, E. B., & Ma, C. 2002, *ApJS*, **141**, 13
- Bell, M. E., et al. 2019, *MNRAS*, **482**, 2484
- Broderick, J. W., et al. 2022, arXiv e-prints, p. [arXiv:2204.08490](https://arxiv.org/abs/2204.08490)
- Callingham, J. R., et al. 2017, *ApJ*, **836**, 174
- Caroli, M., de Castro, P. M. M., Lorient, S., Rouiller, O., Teillaud, M., & Wormser, C. 2010, in *Experimental Algorithms*, ed. P. Festa (Berlin, Heidelberg: Springer Berlin Heidelberg), 462, doi:[10.1007/978-3-642-13193-6\\_39](https://doi.org/10.1007/978-3-642-13193-6_39)
- Carroll, P. A., et al. 2016, *MNRAS*, **461**, 4151
- Chhetri, R., Ekers, R. D., Morgan, J., Macquart, J. P., & Franzen, T. M. O. 2018b, *MNRAS*, **479**, 2318
- Chhetri, R., Morgan, J., Ekers, R. D., Macquart, J. P., Sadler, E. M., Giroletti, M., Callingham, J. R., & Tingay, S. J. 2018a, *MNRAS*, **474**, 4937
- Chhetri, R., et al. 2022, arXiv e-prints, p. [arXiv:2208.04981](https://arxiv.org/abs/2208.04981)
- Clarke, M. 1964, PhD thesis, Cambridge University
- Deller, A. T., & Middelberg, E. 2014, *AJ*, **147**, 14
- Dembinski, H., et al. 2022, scikit-hep/iminuit: v2.11.2, Zenodo, doi:[10.5281/zenodo.3949207](https://doi.org/10.5281/zenodo.3949207)
- Drouart, G., et al. 2021, *PASA*, **38**, e049
- Epanechnikov, V. A. 1969, *TP&A*, **14**, 153
- Finkbeiner, D. P. 2003, *ApJS*, **146**, 407
- Franzen, T. M. O., et al. 2016, *MNRAS*, **459**, 3314
- Garrett, M. A., et al. 2001, *A&A*, **366**, L5
- Hajivassiliou, C. A. 1992, *Nature*, **355**, 232
- Hamaker, J. P., & Bregman, J. D. 1996, *A&AS*, **117**, 161
- Hancock, P. J., Murphy, T., Gaensler, B. M., Hopkins, A., & Curran, J. R. 2012, *MNRAS*, **422**, 1812



- Hancock, P. J., Trott, C. M., & Hurley-Walker, N. 2018, *PASA*, **35**, e011
- Harris, F. 1978, *PROC*, **66**, 51
- Hewish, A., Scott, P. F., & Wills, D. 1964, *Natur*, **203**, 1214
- Hurley-Walker, N., et al. 2014, *PASA*, **31**, e045
- Hurley-Walker, N., et al. 2017, *MNRAS*, **464**, 1146
- Hurley-Walker, N., et al. 2022a, *PASA*, **39**, e035
- Hurley-Walker, N., et al. 2022b, *Natur*, **601**, 526
- Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, *A&A*, **598**, A78
- Jackson, N., et al. 2016, *A&A*, **595**, A86
- Jackson, N., et al. 2022, *A&A*, **658**, A2
- Jaiswal, S., An, T., Wang, A., & Tingay, S. 2022, *MNRAS*, **509**, 2122
- James, F. 2006, *Statistical Methods in Experimental Physics* (2nd edn.; Hackensack, NJ: World Scientific)
- James, F., & Roos, M. 1975, *CoPhC*, **10**, 343
- Kish, L. 1965, *Survey Sampling* (New York: Wiley & Sons)
- Macquart, J. P., & de Bruyn, A. G. 2007, *MNRAS*, **380**, L20
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, **129**, 1993
- Manoharan, P. K. 1993, *SoPh*, **148**, 153
- McKinley, B., et al. 2013, *AJ*, **145**, 23
- Mejia-Ambriz, J. C., Villanueva-Hernandez, P., Gonzalez-Esparza, J. A., Aguilar-Rodriguez, E., & Jeyakumar, S. 2010, *SoPh*, **265**, 309
- Middelberg, E., et al. 2011, *A&A*, **526**, A74
- Moldón, J., et al. 2015, *A&A*, **574**, A73
- Morabito, L. K., et al. 2022, *A&A*, **658**, A1
- Morgan, J., & Galvin, T. 2021, *johnsmorgan/mwa\_pb\_lookup*: Version 1.0, Zenodo, doi:10.5281/zenodo.5083990
- Morgan, J. S., Argo, M. K., Trott, C. M., Macquart, J.-P., Deller, A., Middelberg, E., Miller-Jones, J., & Tingay, S. J. 2013, *ApJ*, **768**, 12
- Morgan, J. S., et al. 2018, *MNRAS*, **473**, 2965
- Morgan, J. S., Macquart, J. P., Chhetri, R., Ekers, R. D., Tingay, S. J., & Sadler, E. M. 2019, *PASA*, **36**, e002
- O'Dea, C. P. 1998, *PASP*, **110**, 493
- Offringa, A. R., & Smirnov, O. 2017, *MNRAS*, **471**, 301
- Offringa, A. R., et al. 2014, *MNRAS*, **444**, 606
- Offringa, A. R., et al. 2015, *PASA*, **32**, e008
- Ord, S. M., et al. 2015, *PASA*, **32**, e006
- Perley, R. A. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 180, *Synthesis Imaging in Radio Astronomy II*, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley, 383
- Pilkington, J. D. H., Hewish, A., Bell, S. J., & Cole, T. W. 1968, *Natur*, **218**, 126
- Porcas, R. W., Alef, W., Ghosh, T., Salter, C. J., & Garrington, S. T. 2004, in *European VLBI Network on New Developments in VLBI Science and Technology*, 31 (arXiv:astro-ph/0412651)
- Purvis, A., Tappin, S. J., Rees, W. G., Hewish, A., & Duffett-Smith, P. J. 1987, *MNRAS*, **229**, 589
- Radcliffe, J. F., et al. 2018, *A&A*, **619**, A48
- Rao, A. P., & Ananthakrishnan, S. 1984, *Natur*, **312**, 707
- Readhead, A. C. S., & Hewish, A. 1972, *Natur*, **236**, 440
- Sadler, E. M., Chhetri, R., Morgan, J., Mahony, E. K., Jarrett, T. H., & Tingay, S. 2019, *MNRAS*, **483**, 1354
- Sadler, E. M., et al. 2020, *MNRAS*, **499**, 4293
- Sault, R. J., Staveley-Smith, L., & Brouw, W. N. 1996, *A&AS*, **120**, 375
- Sokolowski, M., et al. 2017, *PASA*, **34**, e062
- Sweijen, F., et al. 2022, *NatAs*, **6**, 350
- Tappin, S. J. 1986, *P&SS*, **34**, 93
- Tian, J., et al. 2022, *PASA*, **39**, e003
- Tingay, S. J., et al. 2013, *PASA*, **30**, e007
- Tokumaru, M., Fujiki, K., Kojima, M., & Iwai, K. 2021, *ApJ*, **922**, 73
- van Haarlem, M. P., et al. 2013, *A&A*, **556**, A2
- Venzon, D. J., & Moolgavkar, S. H. 1988, *JRSSC (AS)*, **37**, 87
- Waszewski, A., Morgan, J., & Jordan, C. H. 2022, *PASA*, **39**, e036
- Wayth, R. B., et al. 2015, *PASA*, **32**, e025
- Wayth, R. B., et al. 2018, *PASA*, **35**, e033
- White, S. V., et al. 2020, *PASA*, **37**, e018