

## Research Article

# A long period transient search method for the Murchison Widefield Array

Csanád Horváth<sup>1</sup>, Natasha Hurley-Walker<sup>1</sup>, Samuel McSweeney<sup>1</sup>, Timothy James Galvin<sup>2</sup> and John Morgan<sup>2</sup>

<sup>1</sup>International Centre for Radio Astronomy Research, Curtin University, Bentley, WA, Australia and <sup>2</sup>Commonwealth Scientific and Industrial Research Organisation (CSIRO), Kensington, WA, Australia

### Abstract

We present an automated search method for radio transients on the minute timescale focused on the emerging long period transients (LPTs) in image-plane radio data. The method is tuned for use with the Murchison Widefield Array (MWA) and tested on archival observations from the GaLactic and Extragalactic All-Sky MWA Extended Survey (GLEAM-X) in the 70–300 MHz range. The images are formed from model-subtracted visibilities, before applying three filters to the time series of each pixel in an image, with each filter designed to be sensitive to a different transient behaviour. Due to the nature of radio interferometry and the refraction of the fluctuating ionosphere, the vast majority of candidates at this stage are artefacts which we identify and remove using a set of flagging measures. Of the 336 final candidates, 7 were genuine transients: 1 new LPT, 1 new pulsar, and 5 known pulsars. The performance of the method is analysed by injecting modelled transient pulses into a subset of the observations and applying the method to the result.

**Keywords:** catalogues; surveys; radio continuum: general; atmospheric effects; techniques: image processing

(Received 25 June 2025; revised 22 August 2025; accepted 4 September 2025)

## 1. Introduction

Transient astrophysical events typically require highly energetic environments, providing opportunities to study otherwise inaccessible physical regimes. Some transient phenomena are unique to radio or obscured by dust at other frequencies, and many radio transients feature emission from the most extreme environments in the universe such as pulsars, magnetars, rotating radio transients, fast radio bursts, gamma ray bursts, active Galactic nuclei, and long period transients (LPTs).

LPTs are a recently discovered class of object which produce beamed radio emission characterised by their extremely long minute to hour timescale periods, complex temporal structure, and wide 10-s to 20-min pulse widths (Dong et al. 2024; Lee et al. 2025). Their long periods combined with their slow spin-down rate make them difficult to reconcile with existing models as there is not enough angular momentum to power their emission (Rea et al. 2024). There are only a handful of discovered LPTs, some of which are confirmed to be white dwarf (WD) binary systems (Rodríguez 2025; de Ruiter et al. 2025), while others (Hurley-Walker et al. 2022a; Hurley-Walker et al. 2023; Dong et al. 2024; Lee et al. 2025; Wang et al. 2025) have no confirmed counterpart at other frequencies and could be WD binaries, neutron stars, or other exotic objects.

LPTs represent a new parameter space for transient searches at radio frequencies, requiring tailored survey specifications and data

processing methods. Looking for objects which appear for  $\sim 30$  s every  $\sim$ hour and often cease emission for extended periods of time requires long dwell times, well-filled snapshot ( $u,v$ )-coverage, and a wide field-of-view (FOV) to have a chance at surveying an appreciable portion of the sky. At these timescales only single-pulse searches are practicable, and so far all LPTs have been discovered through imaging searches by interferometers. It is difficult to detect such sources using single-dish telescopes due to the presence of red noise on this timescale, which is typically subtracted by gain control systems. The Murchison Widefield Array (MWA) – used in this work – is a radio telescope well suited to this task operating in the 70–300 MHz frequency range with a 30.72 MHz processed bandwidth. It has a FOV of  $\sim 300$  to  $\sim 1\,600$  square degrees in the 215 and 87 MHz frequency bands, respectively (Tingay et al. 2013b).

Two recent searches for radio transients and variables similar to this work are described by de Ruiter et al. (2024) and Wang et al. (2023) using the LOFAR Two-metre Sky Survey (LoTSS) at 144 MHz and the Australian SKA Pathfinder (ASKAP) pilot surveys at 1 GHz, respectively. In both, model visibilities were subtracted from the observations with the goal of subtracting the non-variable components, and the result was imaged. de Ruiter et al. generated 8-s snapshot images from their model-subtracted observations and performed source-finding on the results. Wang et al. generated 15-min snapshot images from their 8 to 10 h model-subtracted observations. From the image cubes, they generated three statistical maps for each observation: chi-squared, peak value, and Gaussian correlation, to which they applied source-finding. Visibility differencing was also used by Dobie et al. (2023), Caleb et al. (2024), and Fijma et al. (2024) to similar effect. The resulting images contain the variable and transient components

**Corresponding author:** C. Horváth, Email: [csanad.horvath@postgrad.curtin.edu.au](mailto:csanad.horvath@postgrad.curtin.edu.au).

**Cite this article:** Horváth C, Hurley-Walker N, McSweeney S, James Galvin T and Morgan J. (2025) A long period transient search method for the Murchison Widefield Array. *Publications of the Astronomical Society of Australia* **42**, e129, 1–12. <https://doi.org/10.1017/pasa.2025.10093>

of the continuum, but they are also contaminated by scintillating non-variable sources, synthesised beam sidelobes, aliases, and other imaging artefacts.

Ionospheric scintillation is especially challenging due to the overwhelming quantity of non-variable sources, and the fact that the fluctuating ionosphere of the Earth varies the apparent position of those sources on the sky (Waszewski, Morgan, & Jordan 2022). To combat this, de Ruiter *et al.* (2024) defined a list of mask radii for ranges of artefact flux densities around known sources. Similarly, Wang *et al.* rejected all candidates within 30 arcsec of a known source unless they varied significantly when compared with the known brightness of the associated source. Filtering out artefact candidates is a cornerstone of this work because it is crucial to improving the sensitivity of radio-transient surveys without being overwhelmed by false positives.

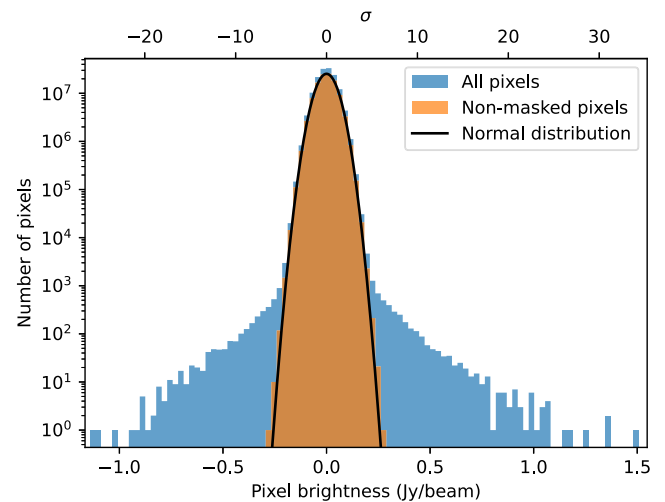
In this work, we present a method to search for radio transients on the second to minute timescales in MWA data. Here, we test the method on 7 099 archival observations from GLEAM-X data releases I and II (DRI and DRII) (Hurley-Walker *et al.* 2022b; Ross *et al.* 2024), although results of the method (Hurley-Walker *et al.* 2023; McSweeney *et al.* 2025b) from other Galactic plane MWA data (Hurley-Walker *et al.* in preparation) already exist. The method design is focused on the properties of known LPTs (although limited by the observation lengths), but other transient and variable sources are also detected. These observations are not concentrated on the Galactic plane where the highest number density of LPTs is expected, but they served as an accessible and extensive dataset for development purposes. Despite this, we found a new LPT GLEAM-X J0704–37 (Section 4.1) (Hurley-Walker *et al.* 2024) and a new pulsar PSR J0031–57 (Section 4.3) (McSweeney *et al.* 2025a). The transient search pipeline is described in Section 3, composed of model subtraction, artefact flagging, self-cross-matching, and visual inspection. Resulting transient candidates are presented in Section 4, results, and the effectiveness of the search method is discussed in Section 5.

## 2. Data

The data used in this work, GLEAM-X Data Releases I and II, described fully by Hurley-Walker *et al.* (2022b) and Ross *et al.* (2024), respectively, is a low-frequency continuum survey by the MWA covering the sky south of Dec +30, broken into seven drift scans. The frequency coverage is composed of five bands between 72 and 231 MHz, avoiding the Orbcomm satellite band by excluding the band around 137 MHz. The observing was executed as a series of week-long campaigns between 2018 and 2020, where each night produced 120-s scans for each of the five frequency settings for a single declination. The GLEAM source catalogue (Hurley-Walker *et al.* 2017) was used in the cross-matching step (Section 3.4) because it is sensitive enough to find all the sources in the 4-s integrations of GLEAM-X (Section 3.1).

## 3. Search method

The transient search pipeline looks for transient candidates by applying filters to image cubes. The image cubes consist of two spatial dimensions and one temporal dimension. The sky model is subtracted from the cube to remove non-variable sources as described in Section 3.1. Three filters, described in Section 3.2, are applied to each temporal column to generate 2D filter images.



**Figure 1.** Pixel brightness distribution of a representative 154 MHz model-subtracted data cube, as well as the same data with a 4′ mask applied around the sources in the GLEAM catalogue. The thermal noise dominating the data is approximately normally distributed with  $\sigma = \text{RMS}$ . The deviation from normal above  $\sim 5\sigma$  is primarily due ionospheric scintillation of real radio sources.

They are intended to probe different timescales and pulse morphologies. Candidate transients are found by looking for islands of large values in the filter maps, described in Section 3.3. Candidates are cross-matched with the GLEAM catalogue in Section 3.4 and those likely to be artefacts rather than genuine transients are flagged according to the criteria in Section 3.5. The catalogue of remaining un-flagged candidates is cross-matched with itself, and a S/N cutoff is applied which decreases for candidates with multiple detections in Section 3.6. The pipeline is written in Python and makes use of the Astropy library<sup>a</sup> and parts of the GLEAM-X imaging pipeline.

### 3.1 Imaging

The formation of the image data is described by Hurley-Walker *et al.* (2022b); to summarise, deep images and visibility models are formed by deconvolution of each two-minute snapshot; these model visibilities are then subtracted and every 4-s interval of a given snapshot observation is imaged. The data are stored as HDF5 cubes which are largely thermal-noise-dominated as shown in Figure 1, storing only the differences between each time step and the continuum average. Note that the cubes are not cleaned or primary beam corrected.

Spatially, the cubes are  $2\,400 \times 2\,400$  pixels covering  $1\,742\text{ deg}^2$  to  $290\text{ deg}^2$  with point spread function (PSF) full width at half maximum (FWHM) semi-major axes of  $156''$  to  $62''$  for the 88–216 MHz bands respectively. The PSF is slightly under-sampled with 2.5 pixels across the semi-major axis to reduce computational cost. For more details see Table B1. These cubes form the basis of the searches described in this paper.

RFI sometimes overwhelms a few time steps in an observation, so time steps whose root-mean-squared (RMS) was greater than  $1.5 \times$  the cube RMS were dropped before applying the transient filters.

<sup>a</sup><https://www.astropy.org/>.

**Table 1.** Transient search filter settings. The lower threshold is the threshold for island area flood-filling. The upper threshold is the minimum peak pixel brightness for a candidate to be considered. “Cube RMS” refers to the RMS of the transient cube.  $A = \{1.5 \text{ for } 87 \text{ MHz}; 1.25 \text{ for } 118 \text{ MHz}; 1 \text{ otherwise.}\}$ .

Filter name	Lower threshold	Upper threshold	Parameters
TCG	$5.5 \times \text{Cube RMS}$	$7 \times A \times \text{Cube RMS}$	$\sigma_{\text{tcg}} = 4\text{s}$ $W = 100\text{s}$
Spike	5.5	8.5	$\sigma_{\text{clip}} = 4$
RMS	$2 \times \text{Cube RMS}$	$2.5 \times A \times \text{Cube RMS}$	

### 3.2 Transient filters

The following transient filters take as input a model-subtracted observation cube and output a 2D image whose values aim to roughly scale with how variable the brightness of a pixel is.

#### 3.2.1 Spike

The Spike filter is intended to look for unresolved bursts narrower than the bin width of 4 s. A sigma-clip  $\sigma_{\text{clip}} = 4$  is applied to the time series of each pixel to estimate the noise level (Equation 1), and the maximum number of standard deviations above the clipped pixel time series is recorded as the Spike value at that pixel (Equation 2):

$$I_{\text{noise}} = \left\{ J \in I \mid \frac{J - \bar{I}}{\text{std}(I)} \leq \sigma_{\text{clip}} \right\} \quad (1)$$

$$\text{spike}(I) = \max \left( \frac{I - \bar{I}_{\text{noise}}}{\text{std}(I_{\text{noise}})} \right) \quad (2)$$

where  $I$  is the time series of the pixels at a coordinate.

#### 3.2.2 Time-Correlated Gaussian (TCG)

The time-correlated Gaussian (TCG) filter looks for wider pulses on the order of 10-s. It is a matched filter that correlates the time series of a cube at each pixel with a Gaussian kernel with standard deviation  $\sigma_{\text{tcg}}$  to find approximately Gaussian shaped pulses:

$$K_i = \begin{cases} \exp\left(-\frac{t_i^2}{2\sigma_{\text{tcg}}^2}\right) & \text{for } -\frac{W}{2} < t_i < \frac{W}{2} \\ 0 & \text{elsewhere} \end{cases} \quad (3)$$

$$\text{tcg}(I) = \left( \frac{K - \bar{K}}{\sum_{i=1}^n (K_i - \bar{K})^2} \right) * (I - \bar{I}) \quad (4)$$

where  $K = \{K_1, \dots, K_n\}$  is the kernel,  $t$  is time in seconds since the start of the observation,  $W$  is the width of the kernel,  $i$  is the index of the time step, and  $n$  is the number of time steps. The  $*$  represents the convolution operation. The choice of parameter values are shown in Table 1 and discussed in Section 5.

#### 3.2.3 Root mean square (RMS)

The RMS filter simply calculates the RMS through time of each pixel as:

$$\text{rms}(I) = \sqrt{\frac{1}{n} \sum_i (I_i - \bar{I})^2} \quad (5)$$

It is intended to detect highly variable, stochastic sources without a clear pulse shape. The island thresholds for the RMS filter are proportional to the cube RMS as per Table 1.

### 3.3 Island detection

A 2D boolean map is generated which is true where any of the three filters are above the lower threshold (Table 1). Islands are then found using a 3-by-3 connectivity matrix of all ones. The choice of lower threshold is a balance between flood-filling the entire image and having enough pixels per island for a defined shape.

Islands whose maximum TCG, Spike, and/or RMS filter values are greater than the upper threshold are assigned the valid\_tcg, valid\_spike, and/or valid\_rms flags, respectively. Islands which do not have either of the aforementioned flags are immediately discarded. The upper threshold is a balance between the true and false positive rates, but it was chosen purposefully low such that the bottom candidates are entirely unconvincing so that stricter selection criteria can be chosen in later steps without redoing the island detections. Each island is fitted with the smallest area ellipse which encloses all of the island pixels, defined by the ellipse centre, major and minor radii, and rotation angle.

Several artefacts which would ordinarily be removed by standard imaging techniques are still present in our image cubes because the 4-s intervals have not been individually cleaned due to the computational expense (see Section 3.1). We therefore have to ensure that the selection criteria described in subsequent sections are able to identify the following broad classes of artefacts which dominate the candidates at this stage:

1. **Ionospheric scintillation:** The ionosphere is constantly in motion with waves of electron density on the scale of tens of metres to tens of kilometres. The varying electron density both diffracts and refracts radio waves causing amplitude and phase scintillation which increases at longer wavelengths. The neighbourhoods of non-variable sources are frequently falsely detected as their apparent positions shift in and out of neighbouring pixels (Jordan et al. 2017; Hurley-Walker & Hancock 2018). Additionally, amplitude scintillation from both the ionosphere and the interplanetary medium can impose flux variability on otherwise non-variable sources (Waszewski et al. 2022).
2. **Synthesised beam sidelobes:** In radio imaging, the sky is convolved with the synthesised beam which, if the calibration is imperfect, results in subtracted sources being surrounded by extended artefacts. Due to scintillation of the parent source and the rotation of the sky relative to the primary beam, the variation of the sidelobe artefacts (which exhibit structure on spatial scales comparable to the PSF), can be erroneously detected as transients.
3. **Aliases:** The primary beam of the MWA tiles has sidelobes with sensitivity  $\sim 10\%$  the main beam response (Sokolowski et al. 2017), so bright ( $> 100\text{-Jy}$ ) sources outside of the FOV can make significant contributions to the visibilities (e.g. Briggs, Schwab, & Sramek 1999, section 3.2); such sources are aliased into the image and often appear to move. It is prohibitively expensive to image the whole sky, but the peeling algorithm (Hurley-Walker et al. 2017) removes most aliases by subtracting a model of bright sources from the visibilities.

Assuming an accurate sky model, this method is effective, but the computational expense prohibits the perfect peeling of all sources. Notwithstanding efforts to suppress aliasing, sources from outside the FOV can be detected.

4. Radio frequency interference (RFI): Interference from terrestrial sources of radiation such as satellites and aeroplanes interfere with the antennas, usually in short bursts. Generally, RFI is easily identified in the image plane as diagonal stripes as one baseline, and therefore spatial frequency, overwhelms the rest.
5. Thermal noise: The background of the images is full of approximately Gaussian thermal noise (see Figure 1). Along the time axis, MWA noise is remarkably clean of systematic effects. The thermal noise is not correlated between time-steps, tending to result in sparse low (S/N) Spike candidates.

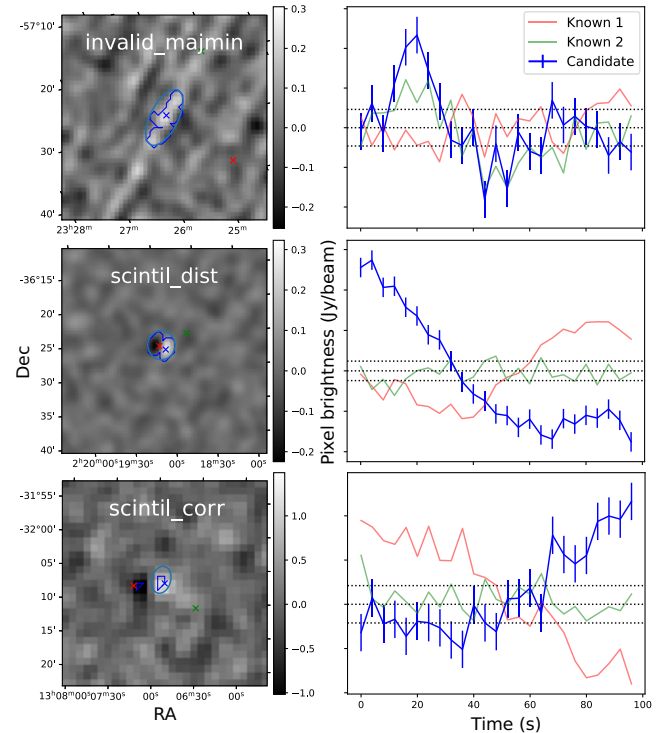
### 3.4 Catalogue cross-matching

In the interest of flagging scintillating known sources, the catalogue of islands is cross-matched with the GLEAM catalogue. It is sometimes ambiguous which source an island is associated with in its neighbourhood, so the following two possible matches are recorded for each candidate: the nearest source to the candidate, and the brightest source within the greater of  $4'$  (the maximum scale at which we observe position shifts due to ionospheric scintillation) and the island's radius. Throughout the rest of the paper, 'nearest known source' refers to either of the aforementioned, and conditions involving the 'nearest known source' are applied to both.

### 3.5 Flagging of artefacts

Candidates are assigned flags by the below criteria. Candidates for which any of the following flags are true are rejected. The decision-making process which lead to the design of the flags and their specific parameters is discussed in Section 5.

1. *invalid\_beam*: Is the primary beam at the candidate coordinates less than 50% of the maximum value within the observation? The edges of observations have too low sensitivity to be useful.
2. *invalid\_majmin*: Is the ratio between the major and minor radii of an ellipse fitted to contain the island greater than 2, AND does the island include more than 3 pixels? This flag is intended to catch extended imaging artefacts which appear either radially around bright stationary sources, or as a result of a moving source. Moving sources can be aliases from outside the FOV, or real sources like airplanes or satellites (Tingay *et al.* 2013a). For an example of a candidate with this flag, see Figure 2. Limiting this flag to sources with at least 3 pixels prevents the rejection of candidates with islands too small to have a meaningful shape.
3. *invalid\_area*: Does the candidate island have more than 20 pixels? Aliases often appear as distorted patches rather than point sources.
4. *invalid\_mean*: Is the mean brightness of the candidate greater than the RMS of the observation? Sources or aliases which are not completely removed in the model subtraction step tend to have flat, positive lightcurves.



**Figure 2.** Examples of candidates which were excluded due to the *invalid\_majmin*, *scintil\_dist*, or *scintil\_corr* flags alone. At left are  $0.5^\circ$  cutouts of the time step of the transient cubes corresponding with the maximum brightness of the candidates. The blue cross marks the coordinates of the peak, the red and green crosses mark the two nearest known sources. At right are the lightcurves at the marked coordinates. The dark blue contour marks the flood-filled island, to which the light blue ellipse is fitted. The horizontal dotted lines mark the mean and RMS of the transient cube.

5. *scintil\_dist*: Is the separation between the candidate and any of the nearest known sources less than the greater of  $4'$  and the major island radius, AND
  - is it an 87 MHz observation AND is the candidate peak beam-corrected flux density less than  $1.5 \times$  the nearest known source's flux density, OR
  - is it a 118 MHz observation AND is the candidate peak beam-corrected flux density less than  $1.25 \times$  the nearest known source's flux density, OR
  - is the candidate peak beam-corrected flux density less than the nearest known source's flux density?

Here, we assume that candidates are unresolved, subtending at most one synthesised beam and thus implicitly convert from pixel brightness (Jy/beam) to flux density (Jy). We then correct for the primary beam to allow comparison with the catalogue.

In the low-frequency regime, ionospheric scintillation often boosts the apparent brightness of sources. A  $4'$  circular mask was deemed as a good compromise between loss of sky area and reduction of false positives due to scintillation of non-variable sources. Candidates are not masked if they are brighter than their masking known source, as that is unlikely to be scintillation. For an example of a candidate with this flag, see Figure 2.

6. *scintil\_corr*: Is the separation between the candidate and any of the nearest known sources less than the greater of  $8'$  and twice the major island radius, AND is the known source



brighter than the candidate peak beam-corrected brightness, AND is the lightcurve at the known source coordinate negatively correlated with the candidate lightcurve (Pearson coefficient  $> 0.85$ )? If ionospheric scintillation shifts the apparent position of a source by more than  $4'$ , then the known source is usually in the negative lobe of the artefact (red in row 3 of Figure 2), which is negatively correlated with the positive lobe (blue in row 3 of Figure 2).

7. `close_to_ateam`: Is the candidate within  $5^\circ$  of any of the following sources? These bright sources produce imaging artefacts too severe to analyse in a more careful way than a simple circular mask.

- Centaurus A (13:25:28,  $-43:01:09$ )
- Hydra A (09:18:06,  $-12:05:44$ )
- Pictor A (05:19:50,  $-45:46:44$ )
- Hercules A (16:51:08,  $+04:59:33$ )
- Virgo A (12:30:49,  $+12:23:28$ )
- Crab (05:34:32,  $+22:00:52$ )
- Cygnus A (19:59:28,  $+40:44:02$ )
- Cassiopeia A (23:23:28,  $+58:48:42$ )

8. `close_to_bright`: Is the candidate within  $1^\circ$  of a source brighter than 10 Jy in the GLEAM source catalogue?

9. `is_moon`: Is the candidate within  $4'$  of the moon? The moon reflects Earth's RFI, which is sometimes detected as a transient candidate, particularly in the 118 MHz and 154 MHz observing bands (McKinley et al. 2013).

10. `invalid_freq`: Is the candidate from an observation centred on 87 MHz? We decided to exclude the bottom frequency band due to the prevalence of extreme ionospheric scintillation making it difficult to draw useful conclusions. Further justification is given in Section 5.

### 3.6 Candidate grouping

The goal of the preceding section is not to build a perfect classifier, but rather to minimise the number of candidates that need to be visually inspected while also minimising the risk of discarding a candidate which would be interesting enough to warrant further investigation if it were to be visually inspected. At this stage, we have eliminated the most obvious false positives, but we have not yet incorporated multi-observation information into the automated part of the decision making. A low-S/N candidate which only appears once is unlikely to be real, but if an equally dim candidate appears multiple times then it may be worth investigating. In practice, this means that the S/N requirements for repeating candidates can be lowered.

The candidate database was cross-matched with itself with a cross-match radius of  $4'$  to create candidate 'groups'. For each group, the mean peak pixel brightness S/N and the mean pixel fluence S/N was calculated. Fluence was chosen in addition to the peak pixel brightness to bias pulses which may not be so bright but have a longer duration (and so are equally as convincing as a shorter bright pulse).

We calculate the peak pixel brightness S/N  $P$  as

$$P = \frac{I_k}{\sigma_I} = \frac{N_k + R_k}{\sigma_I} \quad (6)$$

where  $N$  and  $R$  are the noise and real contributions to the pixel brightness, respectively, and  $k$  is the time-step index at which

the pixel brightness is at maximum. We assume  $N$  is normally distributed about 0 with standard deviation  $\sigma_I \simeq$  cube RMS. The mean peak pixel brightness S/N of a group of  $m$  candidates is

$$\bar{P} = \frac{1}{m} \sum_{j=1}^m P_j = \frac{1}{m} \sum_{j=1}^m ((N_k)_j + (R_k)_j) = \bar{P}_N + \bar{P}_R \quad (7)$$

where  $j$  is the candidate index in the group, and  $\bar{P}_N$  and  $\bar{P}_R$  are the noise and real contributions to  $\bar{P}$ . The standard deviation of  $\bar{P}_N$  across many groups of size  $m$  will be  $\sigma_{\bar{P}} = \frac{1}{\sqrt{m}}$  because  $\sigma_P = 1$  by definition. We choose a cutoff  $C_P$  and select candidate groups whose  $\bar{P}$  is more than  $C_P$  times the noise level according to

$$\bar{P} > C_P \sigma_{\bar{P}} = \frac{C_P}{\sqrt{m}} \quad (8)$$

For a group size of 1, this reduces to a simple S/N cutoff.

The cutoff for the group mean pixel fluence S/N is derived similarly. We calculate the pixel fluence  $F$  as:

$$F = \Delta t \sum_{i=1}^n I_i = \Delta t \sum_{i=1}^n (N_i + R_i) = F_N + F_R \quad (9)$$

where  $n$  is the number of time steps in the observation,  $F_N$  and  $F_R$  are the noise and real contributions to the pixel fluence respectively, and  $\Delta t$  is the time step.  $F_N$  will have a standard deviation across many sources of  $\sigma_F = \sigma_I \Delta t \sqrt{n}$ . Hence, we define the pixel fluence S/N  $f$  as:

$$f = \frac{F}{\sigma_I \Delta t \sqrt{n}} \quad (10)$$

and the group mean pixel fluence S/N as:

$$\bar{f} = \frac{1}{m} \sum_{j=1}^m f_j = \frac{1}{m} \sum_{j=1}^m (f_{Nj} + f_{Rj}) = \bar{f}_N + \bar{f}_R \quad (11)$$

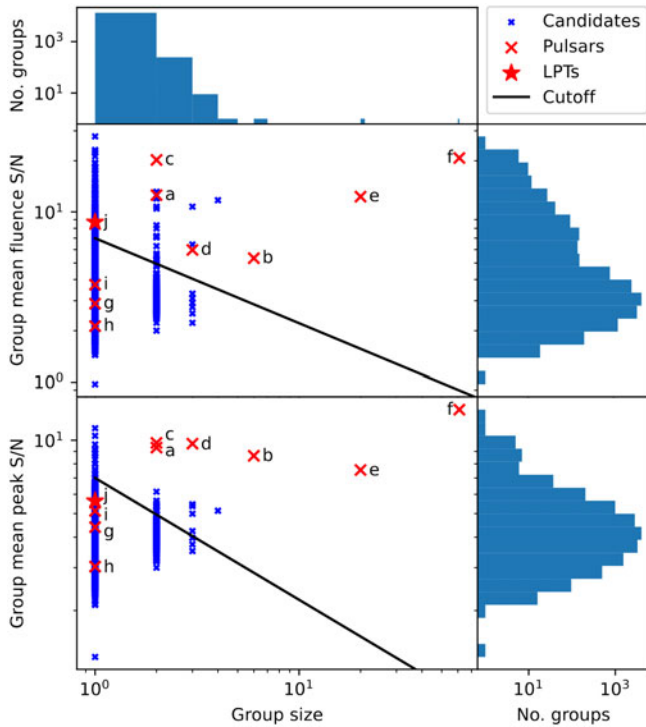
where  $f_N$  and  $f_R$  are the noise and real contributions to  $f$ , respectively. The standard deviation of  $\bar{f}_N$  across many groups of size  $m$  of will be  $\sigma_{\bar{f}} = \frac{1}{\sqrt{m}}$  because  $\sigma_f = 1$  by definition. We choose a cutoff  $C_F$  and select candidate groups whose  $\bar{f}$  is more than  $C_F$  times the noise level according to

$$\bar{f} > C_F \sigma_{\bar{f}} = \frac{C_F}{\sqrt{m}} \quad (12)$$

This is a simplistic approach, but for the purpose of creating a reasonable cutoff curve for candidate believability, it is adequate. A value of  $C_P = C_F = 7$  was chosen, and candidates which fall below both cutoffs are rejected. The cutoff curves are plotted in Figure 3. The value of 7 aligns the most closely with the visual believability of candidates and can be interpreted as a more advanced  $7\sigma$  signal requirement.

### 3.7 Classification interface

All candidates which make it through the above cuts need to be visually inspected to determine suitability for further investigation. A candidate classification web interface was developed by Astronomy Data and Computing Services (ADACS). The interface presents the user with a diagnostic plot, such as the one for GLEAM-X J0704-37 in Figure 4, as well as a GIF cutout, cutouts from the three filter maps, significance of the candidate's brightness relative to the rest of the observation, further diagnostic information about the candidate and the observation, options to



**Figure 3.** Distribution of mean pixel fluence and peak pixel brightness S/N statistics against group size. The cutoff curve is drawn on both. The groups marked in red are identified as the following real sources: (a) PSR J0630–2834; (b) PSR J0031–57; (c) PSR J0437–4715; (d) PSR J0410–31; (e) PSR J2048–1616; (f) PSR J0034–0721; (g) PSR J2241–5236; (h) PSR J0502–6617; (i) PSR J1244–1812; (j) GLEAM-X J0704–37.

categorise and rate the candidate, other nearby candidates, and nearby objects from the SIMBAD and ATNF catalogues. This was an invaluable resource both to generate the final shortlist of candidates, and during development to tune the parameters.

#### 4. Results

Of the 13 844 candidates found in the dataset, 610 formed groups of 2 or more, making 13 493 total groups. Of those, there are 330 candidate groups above the fluence S/N cutoff in Figure 3, and 33 above the peak pixel brightness S/N cutoff. Only 6 of the peak pixel brightness S/N groups are not among the 330 pixel fluence S/N groups, and none of them passed visual inspection. This makes the final number of candidate groups 336, 7 of which were found to be genuine transients. One source was the real LPT GLEAM-X J0704–37 (see Section 4.1), detected by TCG. The following 6 were found to be pulsars:

- |                   |                     |
|-------------------|---------------------|
| a. PSR J0630–2834 | (TCG & RMS)         |
| b. *PSR J0031–57  | (TCG & Spike)       |
| c. PSR J0437–4715 | (TCG & RMS)         |
| d. *PSR J0410–31  | (TCG & RMS & Spike) |
| e. PSR J2048–1616 | (TCG & RMS)         |
| f. PSR J0034–0721 | (TCG & RMS)         |

The periods of pulsars are much shorter than our time resolution – in fact, PSR J0437–4715 and PSR J2241–5236 are millisecond pulsars (Johnston et al. 1993; Keith et al. 2011) – so it was either minute-timescale brightness fluctuation that was detected,

unresolved giant pulses, or simply the scintillation of their time-averaged brightness (Dai et al. 2016), rather than individual pulses. At least PSR J0034–0721 is known to null on timescales of minutes (Huguenin, Taylor, & Troland 1970) and well studied with the MWA (e.g. McSweeney et al. 2018), so it is possible that its variability due to nulling was what we detected. The starred items are discussed in subsequent subsections.

Of the candidates which fell below the cutoff in Figure 3, three of their coordinates coincided with known pulsars, but after visual inspection they were indistinguishable from the overwhelming majority of dim candidates. It is difficult to determine whether these are chance coincidences, but since they do not pass visual inspection they do not represent false negatives. They would have had to be detected at least 5–15 times rather than only once to be convincing given the criteria in Section 3.6 and how faint they are.

##### 4.1 GLEAM-X J0704–37

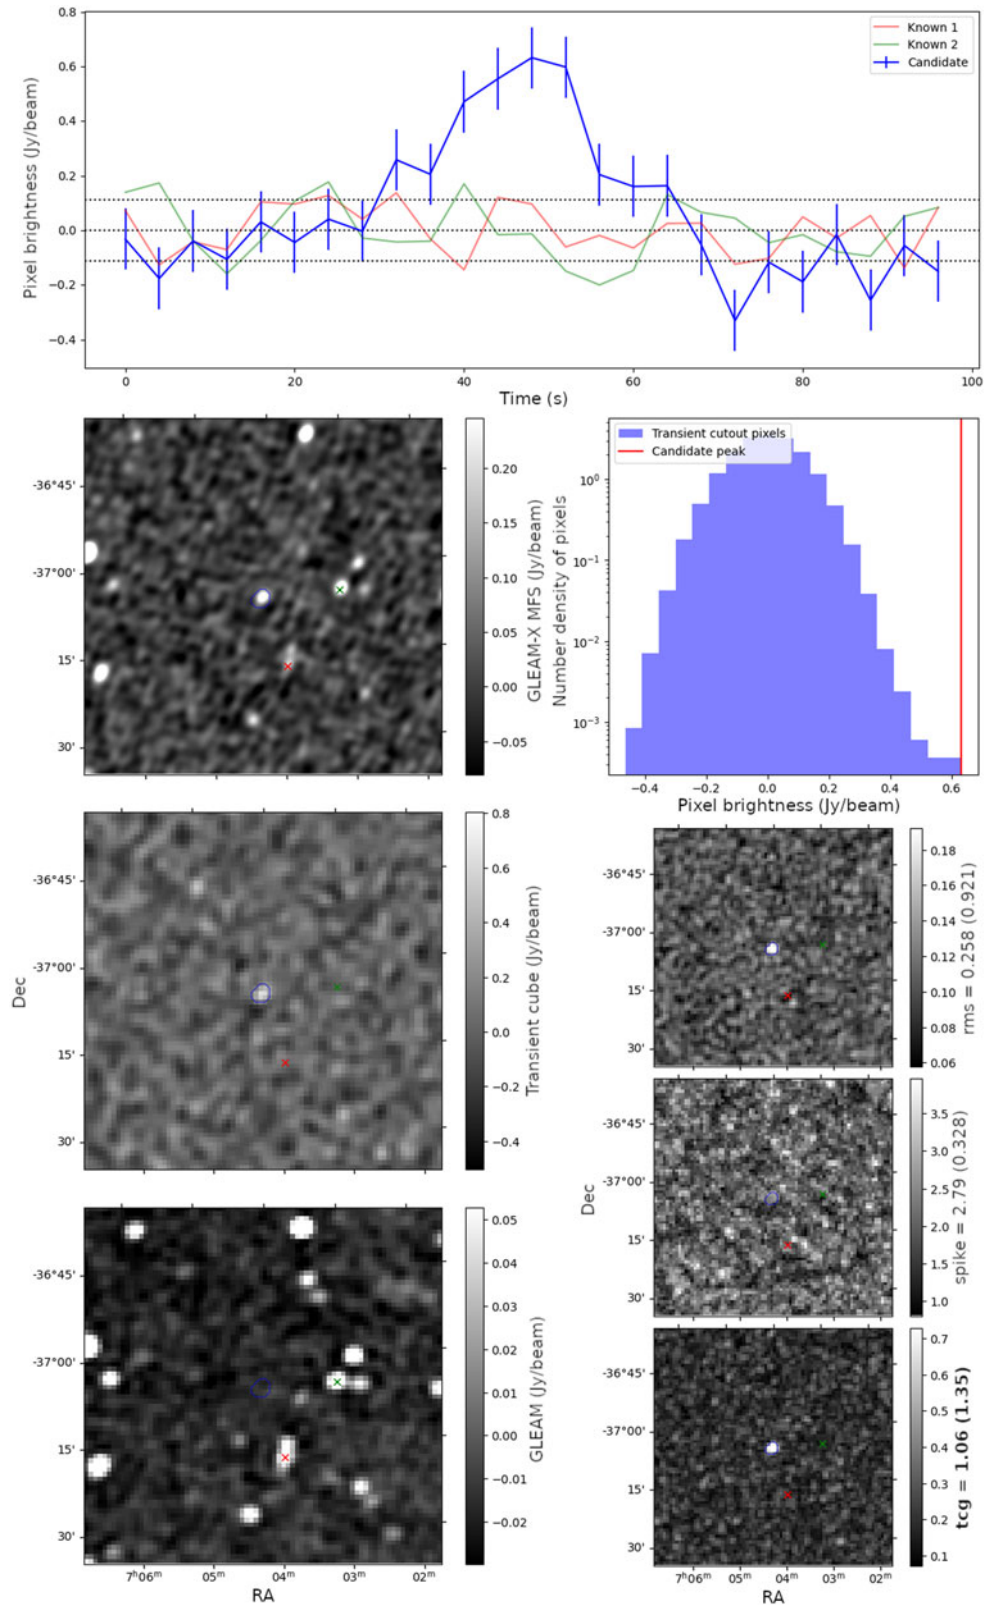
GLEAM-X J0704–37 was discovered using an earlier iteration of the method described here in an archival observation centred on 118 MHz of GLEAM-X DRI and DRII. Figure 4 is the diagnostic candidate plot for the detection of GLEAM-X J0704–37, which was at the very edge of the searched area (see Figure B1).

In follow-up observations and archival observations published by Hurley-Walker et al. (2024), it was found to be a long-period radio transient, with a periodicity of  $P = 2.9$  h and a DM of  $36.54(1)$  pc cm<sup>−3</sup>. An optical counterpart was found, consistent with a cool main-sequence star of spectral type M3; further follow-up by Rodríguez (2025) identified a WD companion, characterising this system as a polar M-dwarf/WD in a tight, tidally locked  $P = 2.9$  h orbit.

##### 4.2 PSR J0410–31

PSR J0410–31 was discovered in a single-pulse search of the High Time Resolution Universe (HTRU) survey with the Parkes ‘Murriyang’ radio telescope at 1.352 GHz (Burke-Spolaor et al. 2011). It was found to have a periodicity  $P = 1.8785(2)$  s and a dispersion measure (DM) of  $9.2(3)$  pc cm<sup>−3</sup>. Confirmed by Rane et al. (2016), also with Murriyang, this source is an intermittent Rotating Radio Transient (RRAT), capable of producing high-fluence bursts and sometimes invisible (presumably nulling) in standard periodicity searches.

PSR J0410–31 was detected in our analysis as a series of  $\sim$ Jy-bright unresolved pulses, consistent with its RRAT classification. It was observed to lie  $\sim 10'$  away from its catalogued location, likely due to the large positional error of Parkes, and the low accuracy of timing solutions generated for RRATs. To further improve the localisation, we selected the ObsID 1201778784, the highest-frequency (170–200 MHz) observation with the brightest (1.8-Jy; S/N $\sim$ 45) pulse. We imaged the two 0.5-s time steps in which the brightest pulse was detected, performed source-finding with AEGER (Hancock, Trott, & Hurley-Walker 2018) and compared the nearby radio galaxy positions to the Sydney University Molonglo Sky Survey (SUMSS; Mauch et al. 2003) and the NRAO VLA Sky Survey (NVSS; Condon et al. 1998). For the 30 S/N > 10 sources in a  $4^\circ \times 4^\circ$  region around the pulsar, bulk shifts were observed, likely due to ionospheric refraction (Hurley-Walker & Hancock 2018). Against SUMSS,  $\Delta RA = -13''.8 \pm 7''.2$ , and  $\Delta Dec = -0''.7 \pm 3''.5$ ; against NVSS,  $\Delta RA = -14''.7 \pm 7''.6$ , and  $\Delta Dec = -1''.8 \pm 4''.3$ . We therefore applied corrections of  $14''$  and  $1''$  to our detected position to account for these shifts, and expect



**Figure 4.** Diagnostic plot for the detection of GLEAM-X J0704-37. The top panel is the lightcurve of the detected island (blue), and the lightcurves of the two nearest known sources (red and green). The dashed horizontal lines are the observation mean pixel brightness (centre line) and the positive and negative RMS. In the left column are 1° cutouts from the multi-frequency synthesis (MFS) image formed during the routine GLEAM-X imaging (top), peak time step from the model-subtracted data cube (middle), and GLEAM image cutout (bottom). The blue contour marks the candidate island, and the crosses are the nearest known sources. At the top of the right column is a histogram of the pixel values in the model-subtracted cube with the candidate peak marked in red. Below are 1° cutouts from the RMS, Spike, and TCG filter maps, with the filter that triggered the detection in bold (in this case TCG). The numbers following the filter name in the vertical labels are the peak filter value followed by the filter value divided by the upper threshold from Table 1 in parentheses.

a residual astrometric error of  $\sim 3''$ . Our updated position of PSR J0410–31 is therefore  $04^{\text{h}}09^{\text{m}}55.5^{\text{s}} - 31^{\circ}10'40''5$ , which is consistent with the original position reported by Burke-Spolaor et al. (2011). Use of these updated coordinates should slightly improve the S/N of future pulsation and single-pulse searches towards this source.

### 4.3 PSR J0031–57

Similar to PSR J0410–31, PSR J0031–57 was also detected as a series of unresolved bursts with a small but discernible DM of  $\sim 6 \text{ pc cm}^{-3}$ , consistent with a Galactic pulsar or RRAT. However, no source at that location was found in existing pulsar or RRAT catalogues. Archival MWA data were then searched, yielding detections of PSR J0031–57 in high-time resolution MWA data in both 2018 and 2020. These observations confirmed that PSR J0031–57 is a pulsar with spin period  $P = 1.57 \text{ s}$  and  $\text{DM} = 6.8 \text{ pc cm}^{-3}$ . Despite being discovered via its relatively rare, unusually bright pulses, PSR J0031–57 also exhibits much dimmer pulses that appear to be persistent, casting doubt on whether it should be classified as an RRAT. A full description of these results is given by McSweeney et al. (2025a).


## 5. Evaluation of performance

The filter parameters in Table 1 were optimised to maximise detection probability while minimising the false positive rate in a sample of 100 observations from the dataset, 20 from each frequency band, which did not contain any genuine transients. The sample was searched to estimate the false positive rate, the results of which are presented in Section 5.1. Into a copy of the same 100 observations, modelled transients were injected. The copy was searched to estimate the true positive rate, the results of which are summarised in Section 5.2.

Standard statistics commonly used in the context of automatic classifiers such as precision, recall, and accuracy are difficult to use in this context. The recall estimate (fraction of real transients detected) depends on the distribution of injected transients in terms of brightness and duration, as well as the density of other sources on a particular region of the sky and the wavelengths observed. Figure 5 attempts to disambiguate these dependencies, except for the spatial one. We see that at 87 MHz the curve of recall against peak pixel brightness begins to plateau below 50% which, combined with the large number of candidates it generates, was the justification for its exclusion by the `invalid_freq` flag.

Regarding TCG, the choice of  $\sigma_{\text{tcg}} = 4 \text{ s}$  corresponds to the kernel representing a pulse template  $\sim 6\sigma_{\text{tcg}} = 24 \text{ s}$  wide, about a quarter of the observation duration. In Figure 5, we see that the TCG recall biases longer pulses as desired, and is successful even for injected transients with pulse widths  $\sigma_{\text{time}} \leq 5 \times \sigma_{\text{tcg}}$ . This kernel width was found to be a good compromise to maximise sensitivity between the shorter Spike sensitivity timescale and the longer RMS sensitivity timescale (although RMS was not initially conceived for this purpose).

The precision (fraction of candidates which are real transients) and accuracy (fraction of correctly classified samples) lose their meaning entirely because (a) we do not have a good idea of the density of real transient sources relative to artefacts (the density of injected transients is arbitrary), (b) we expect a very low precision in the first place for a blind search like this (and this is not



**Figure 5.** Histograms of injected modelled transients which were recovered after rejecting flagged candidates, as a percentage of the total number injected to estimate the recall. The injected pulse profiles have a Gaussian shape with standard deviation  $\sigma_{\text{time}}$  and peak pixel brightness  $I_{\text{peak}}$ .

a problem), and (c) the distinction between classified samples is ambiguous because our method treats the sky as a continuous area.

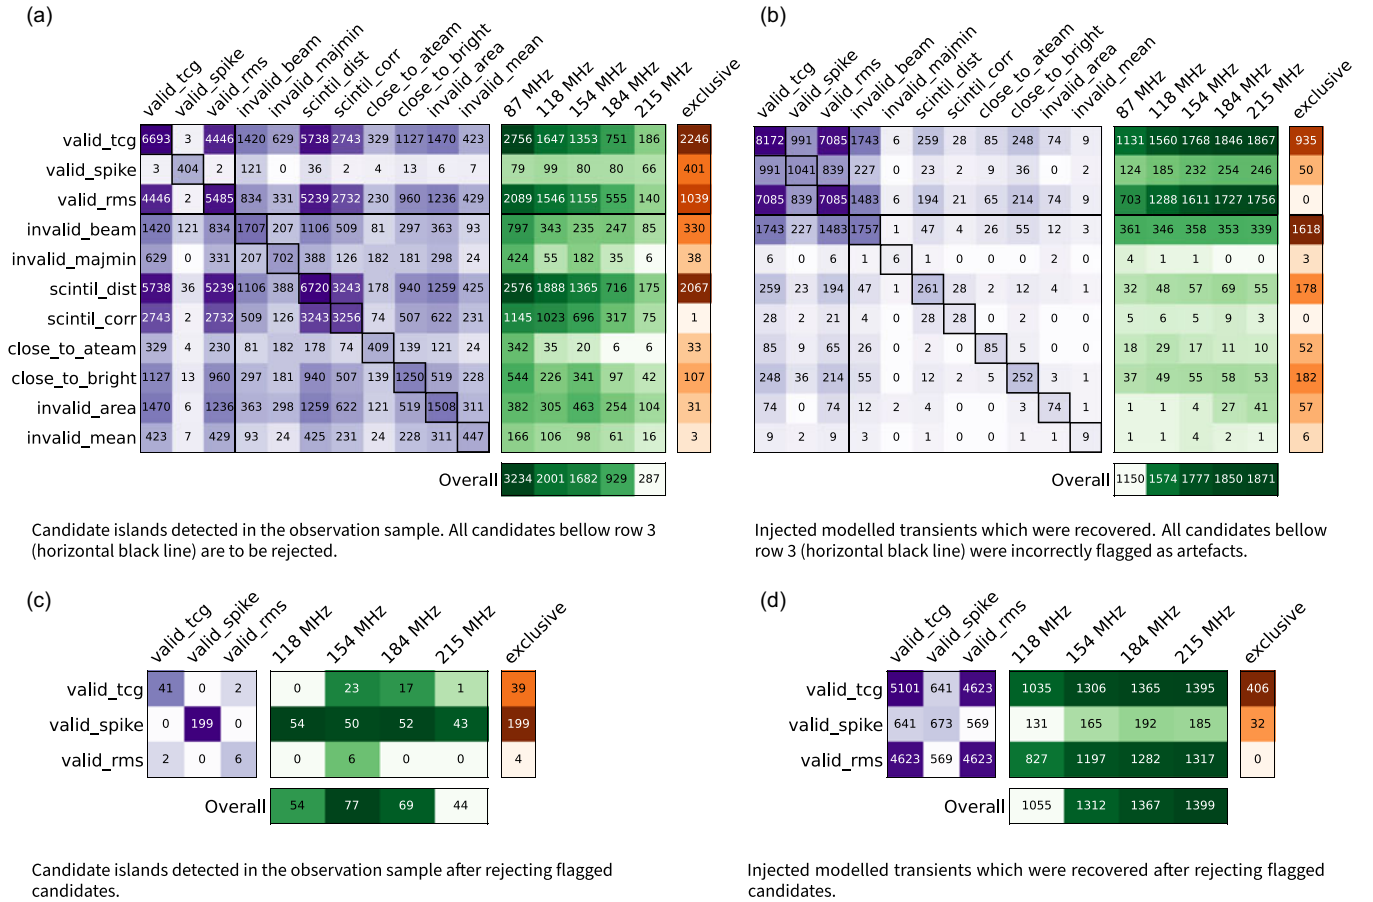
### 5.1 False positives

The transient search steps were applied to the 100 sample observations. Figure 6a is a confusion matrix of the Section 3.5 flags assigned to the candidate islands found in the sample observations, and Figure 6c is a confusion matrix after all flagged candidates have been removed. All of these candidates are assumed not to be true transients, and the down-selected candidate set was visually inspected to verify this.

There were 6 693, 404, and 5 485 islands detected by TCG, Spike, and RMS respectively, of which 41, 199, and 6 made it through the flagging. There was effectively no overlap between islands detected by Spike and the other filters owing to the much shorter timescale of sensitivity, while TCG and RMS overlapped by over 80%. While more than 10 times as many islands were detected by TCG and RMS than Spike, Spike had the most candidates after flagging. Spike candidates are often the result of random fluctuations in the background thermal noise, but even a low S/N detection could be real if it reoccurs multiple times. For this reason, rather than raising the Spike cutoff threshold, we allowed the grouping step in Section 3.6 to reject the majority of low S/N spike detections.

The greatest number of islands were rejected by the `scintil_dist` flag for being too close to a known source, and most of those were from the TCG or RMS filters. Unfortunately, the TCG filter is especially sensitive to variations at exactly the timescale of ionospheric scintillation. These effects are the worst at the lowest frequencies due to the square relationship with wavelength, which is the reason for scaling the cutoff at low frequencies for `scintil_dist` in Section 3.5. This relationship between false positive rate and frequency is visible in the gradient across the frequency columns in the top three rows (or lack thereof for Spike) of Figure 6a. The Spike false positive rate does not correlate with observing frequency because it does not tend to pick up scintillation due to its short timescale sensitivity. The sensitivity of RMS and TCG to scintillation is evident in the large values at the intersection of rows `scintil_dist` and `scintil_corr`, and columns `valid_tcg` and `valid_rms` in the confusion matrix. Their sensitivity to the long synthesised





**Figure 6.** Confusion matrices of candidate flags detected in a sample of 100 observations (20 at each of the 5 frequencies). Each cell counts the number of candidates in both the row category AND the column category. The left square matrices count the flags assigned to candidates, and the values along the diagonal are the total with that flag. The middle matrices count the flags in observations in different frequency channels, with the bottom ‘Overall’ row counting the total number of candidates in each channel. The ‘Exclusive’ column at right counts the number of candidates with only a particular flag (i.e. 38 candidates were rejected purely because of invalid\_majmin in (a)). Note that is\_moon wasn’t included because the moon wasn’t in this subset of observations, and invalid\_freq wasn’t included because it is equivalent to the 87 MHz column.

beam sidelobes around sources is evident in their intersection with invalid\_majmin.

## 5.2 Modelled transients

Transients were modelled as a Gaussian pulse through time defined by the peak pixel brightness  $I_{\text{peak}}$  and the pulse profile standard deviation  $\sigma_{\text{time}}$ , multiplied by a spatial 2D circular Gaussian with  $\sigma_{\text{space}} = 1$  pixel. When adding a transient to a data cube, the mean through time was subtracted to simulate the model subtraction step. For  $N_{\text{inj}}$  transients at time  $t_i$  and pixel coordinates  $(x_i, y_i)$  injected into an observation  $I_{\text{obs}}$ , the resulting data cube  $I_{\text{inj}}$  is described by:

$$I_{\text{trans}}^i(t, x, y) = I_{\text{peak}} \exp \left( -\frac{(t - t_i)^2}{2\sigma_{\text{time}}^2} - \frac{(x - x_i)^2 + (y - y_i)^2}{2\sigma_{\text{space}}^2} \right) \quad (13)$$

$$I_{\text{inj}}(t, x, y) = I_{\text{obs}}(t, x, y) + \sum_{i=1}^{N_{\text{inj}}} (I_{\text{trans}}^i(t, x, y) - \langle I_{\text{trans}}^i(t', x, y) \rangle_{t'}) \quad (14)$$

A circular Gaussian is a good approximation because we do not expect transients bright enough to produce significant sidelobes

with the MWA’s  $(u, v)$ -coverage, and  $\sigma_{\text{space}} = 1$  pixel corresponds to a FWHM of 2.355 pixels which is approximately the PSF FWHM. Although an even distribution of modelled source brightnesses in an image is not realistic,  $I_{\text{peak}}$  was not multiplied by the primary beam because statistics relative to the background noise level (which does not vary with the primary beam) are more useful in this analysis.

In each observation,  $N_{\text{inj}} = 100$  modelled transients were injected with uniform distributions of position,  $I_{\text{peak}} \in [0.1, 3.0]$  Jy, and  $\sigma_{\text{time}} \in [0.4, 20]$  s. Note that  $\sigma_{\text{time}} = 20$  s corresponds to a pulse width of  $\sim 120$  s if we consider the a Gaussian pulse to extend to  $\sim 3\sigma_{\text{time}}$ , which is the observation length. Pulses on a longer timescale would not be imaged anyway due to the model subtraction (Section 3.1). The same filtering and flagging steps as before were applied to the data cubes with injected transients.

The number of islands detected which corresponded with injected transients with each flag are shown in Figure 6b in the same format as Figure 6a, and the number recovered after rejecting flagged islands is shown in Figure 6d. The majority of false negatives were due to the invalid\_beam flag. This is unavoidable because the edges of the observations where sensitivity is low are too noisy to trust detections. The next most common were scintil\_dist and close\_to\_bright due to injected transients falling on

real sources by chance. After candidate selection, 64% of modelled transients were recovered. For the 118 MHz band, only 53% were recovered, while for 154 MHz and above, 69% were recovered. TCG, Spike, and RMS detected 64%, 8.4%, and 59% of injected transients respectively, but RMS detected only 1 which was not also detected by either Spike or TCG. The right panel of Figure 5 shows that spike is most sensitive on short timescales, while TCG and RMS are sensitive on long timescales, especially when it comes to dim sources.

## 6. Conclusion

In this work, we have developed and tested a new method for searching for radio transients on the minute timescale. Testing the method on the GLEAM-X DRI and DRII we discovered a new pulsar PSR J0031–57 and a new LPT GLEAM-X J0704–37 which, due to its high Galactic latitude out of the busy Galactic plane, could for the first time be associated with an optical counterpart. Due to the high latitudes of most of the searched data (see Figure B2) we might expect a relatively low number density of real transients, making these discoveries all the more interesting. Nevertheless, it served as a diverse and extensive dataset of false positives to develop and tune this algorithm. The MWA Galactic plane monitoring programme is a more targeted search, with observations at 185–215 MHz covering  $|b| < 15^\circ$  and  $73^\circ < l < 284^\circ$  (Hurley-Walker *et al.* in prep) in 2022 (with a second phase in 2024/25). It discovered the  $P \sim 22$ -min LPT GPM J1839–10 (Hurley-Walker *et al.* 2023) and confirmed ASKAP J1755–25 (McSweeney *et al.* 2025b) as a  $P \sim 1.16$  h LPT using iterations of the method described in this paper.

The number of hand-tuned parameters and decision layers in this search method prompts the consideration of machine learning (ML). While beyond the scope of this work, an ML model – perhaps image-based, decision-tree-based, or otherwise – could be trained on a dataset consisting of synthetic transients and real artefacts generated by our method. A pre-selection of candidates would still be necessary in such a method because convolving an ML model over the entire sky area would be prohibitively expensive.

The most effective methods for separating the true positives from the false positives were the `scintil_dist` flag (see the ‘exclusive’ column of Figure 6a and b) and the candidate grouping (see Figure 3). The latter highlights that the best way to be sure a repeating source is real is to see it multiple times, but whether or not a given source is detected multiple times is a complicated function of observation length and observing cadence. Generally, however, the chances of re-detection go down for longer period sources (see Appendix A). The small number of candidates with multiple detections did not justify the implementation of any automated periodicity detection in this study, but this may be effective in future larger surveys.

**Acknowledgement.** This scientific work uses data obtained from Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory. 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. ASVO has received funding from the Australian Commonwealth Government through the National eResearch Collaboration Tools and Resources (NeCTAR) Project, the Australian National Data Service (ANDS), and the National Collaborative Research Infrastructure Strategy.

This work was supported by resources provided by the Pawsey Supercomputing Centre, which is also used by the MWA. Establishment of Inyarrimanha Ilgari Bundara, and the Pawsey Supercomputing Centre are initiatives of the Australian Government, with support from the Government of Western Australia and the Science and Industry Endowment Fund. We acknowledge the Wajarri Yamaji People as the Traditional Owners and Native Title Holders of the observatory site.

We acknowledge ADACS and Dr Paul Hancock for the development of the transient classification web interface.

**Data availability statement.** A FITS table of candidates will be made available on the PASA Dastore. Further data used in this work is available upon request from the authors.

**Funding statement.** This research is supported by an Australian Government Research Training Program (RTP) Scholarship doi.org/10.82133/C42F-K220. N.H.-W. is the recipient of an Australian Research Council Future Fellowship (project number FT190100231).

**Competing interests.** The authors declare no competing interests.

**Code availability statement.** The code used in this work is available at <https://github.com/CsanadHorvath/Transients.git>.

## References

- Anumalapudi, A., *et al.* 2025, *MNRAS*, **542**, 1208
- Blout, S., *et al.* 2025, *A&A*, **699**, A341
- Briggs, D. S., Schwab, F. R., & Sramek, 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, 127
- Burke-Spolaor, S., *et al.* 2011, *MNRAS*, **416**, 2465
- Caleb, M., *et al.* 2024, *NatAs*, **8**, 1159
- Condon, J. J., *et al.* 1998, *AJ*, **115**, 1693
- Dai, S., *et al.* 2016, *MNRAS*, **462**, 3115
- de Ruiter, I., *et al.* 2024, *MNRAS*, **531**, 4805
- de Ruiter, I., *et al.* 2025, *NatAs*, **9**, 672
- Dobie, D., *et al.* 2023, *MNRAS*, **519**, 4684
- Dobie, D., *et al.* 2024, *MNRAS*, **535**, 909
- Dong, F. A., *et al.* 2024, arXiv e-prints, arXiv:2407.07480
- Dong, F. A., *et al.* 2025, *ApJ*, **988**, L29
- Fijma, S., *et al.* 2024, *MNRAS*, **528**, 6985
- Hancock, P. J., Trott, C. M., & Hurley-Walker, N. 2018, *PASA*, **35**, e011
- Huguenin, G. R., Taylor, J. H., & Troland, T. H. 1970, *ApJ*, **162**, 727
- Hurley-Walker, N., & Hancock, P. J. 2018, *A&C*, **25**, 94
- Hurley-Walker, N., *et al.* 2017, *MNRAS*, **464**, 1146
- Hurley-Walker, N., *et al.* 2022a, *Natur*, **601**, 526
- Hurley-Walker, N., *et al.* 2022b, *PASA*, **39**, e035
- Hurley-Walker, N., *et al.* 2023, *Natur*, **619**, 487
- Hurley-Walker, N., *et al.* 2024, *ApJ*, **976**, L21
- Hyman, S. D., *et al.* 2005, *Natur*, **434**, 50
- Johnston, S., *et al.* 1993, *Natur*, **361**, 613
- Jordan, C. H., *et al.* 2017, *MNRAS*, **471**, 3974
- Keith, M. J., *et al.* 2011, *MNRAS*, **414**, 1292
- Lee, Y. W. J., *et al.* 2025, *NatAs*, **9**, 393
- Li, D., *et al.* 2024, arXiv e-prints, arXiv:2411.15739
- Mauch, T., *et al.* 2003, *MNRAS*, **342**, 1117
- McKinley, B., *et al.* 2013, *AJ*, **145**, 23
- McSweeney, S. J., Bhat, N. D. R., Tremblay, S. E., & Deshpande, A. A. 2018, *PIAU*, **337**, 376
- McSweeney, S. J., *et al.* 2025a, *ApJ*, **981**, 143
- McSweeney, S. J., *et al.* 2025b, *MNRAS*, arXiv:2507.14448
- Rane, A., *et al.* 2016, *MNRAS*, **455**, 2207
- Rea, N., *et al.* 2024, *ApJ*, **961**, 214
- Rodríguez, A. C. 2025, *A&A*, **695**, L8

Ross, K., et al. 2024, *PASA*, **41**, e054  
Sokolowski, M., et al. 2017, *PASA*, **34**, e062  
Tingay, S. J., et al. 2013a, *AJ*, **146**, 103  
Tingay, S. J., et al. 2013b, *PASA*, **30**, e007  
Wang, Y., et al. 2023, *MNRAS*, **523**, 5661  
Wang, Z., et al. 2025, *Natur*, **642**, 583  
Waszewski, A., Morgan, J., & Jordan, C. H. 2022, *PASA*, **39**, e036

Appendix A. Periodic Transient Detectability

To get a sense of the likelihood of a periodic source being in the field of view while it is active at all, we ran a simple simulation as follows. For an even distribution of coordinates in the dataset coverage shown in Figure B2, we counted the number times a periodic source at that coordinate would have both been on and in the field. In Figure A1, we show the fraction of those hypothetical sources which would have been seen at least some number of times. We see that a ~10 min source has a ~95% chance of being in the field at least 5 times, whereas a ~200 min source like GLEAM-X J0704–37

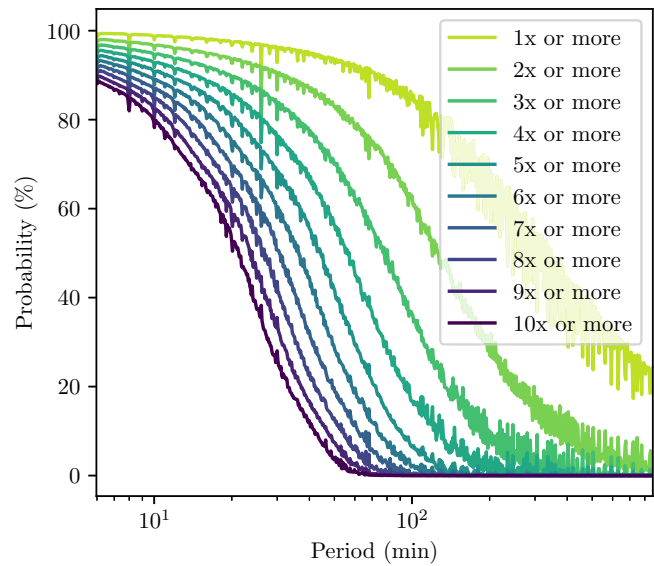


Figure A1. Probability that a periodic transient is pointed towards Earth and in the MWA field x number of times in GLEAM-X DRI and DRII.

has a ~60% chance of being seen just once. The vertically aligned spikes in the curves are caused by harmonics between the periods and observing times. Critically, this ignores the actual detectability of sources; LPTs are difficult to find even with a perfect telescope. The problem is further confounded by the fact that many LPTs have very short (sometimes repeating) windows of activity.


Appendix B. Extended Figures

The LPTs shown in Figures B1 and B2 are:


- CHIME J0630+25 (Dong et al. 2024)
- GLEAM-X J0704–37 (Hurley-Walker et al. 2024)
- ILT J1101+55 (de Ruiter et al. 2025)
- ASKAP J1448–68 (Anumalapudi et al. 2025)
- GLEAM-X J1627–52 (Hurley-Walker et al. 2022a)
- CHIME/ILT J1634+44 (Dong et al. 2025; Bloat et al. 2025)
- GCRT J1745–30 (Hyman et al. 2005)
- ASKAP J1755–25 (Dobie et al. 2024; McSweeney et al. 2025b)
- ASKAP/DART J1832–09 (Li et al. 2024; Wang et al. 2025)
- GPM J1839–10 (Hurley-Walker et al. 2023)
- ASKAP J1839–07 (Lee et al. 2025)
- ASKAP J1935+21 (Caleb et al. 2024)

Table B1. Properties of the model-subtracted image cubes towards zenith for each frequency band. All images are 2 400 × 2 400 pixels.

	88 MHz	118 MHz	154 MHz	185 MHz	216 MHz
PSF FWHM	156"	115"	87"	72"	62"
semi-major axis	2.5 pix	2.5 pix	2.5 pix	2.5 pix	2.5 pix
PSF FWHM	121"	89"	68"	56"	48"
semi-minor axis	1.9 pix	1.9 pix	1.9 pix	1.9 pix	1.9 pix
Pixel size	63"	47"	36"	30"	26"
Image sky area	1742 deg <sup>2</sup>	959 deg <sup>2</sup>	567 deg <sup>2</sup>	394 deg <sup>2</sup>	290 deg <sup>2</sup>



**Figure B1.** Sky map of transient candidates from GLEAM-X DRI and DR11. The blue points are the candidate groups after artefact flagging in Section 3.5. The red points are the final candidates after applying candidate grouping in Section 3.6. The bright sources are those listed under `close_to_ateam` in Section 3.5.



**Figure B2.** Sky map of GLEAM-X DRI and DR11 observations. The red points are the final candidates after applying candidate grouping in Section 3.6. The colour map is the total (non-continuous) minutes of data at a coordinate. The bright sources are those listed under `close_to_ateam` in Section 3.5.