

Finding Radio Transients



Alexander Andersson
St Anne's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2024

Acknowledgements

One of the things I have learnt during my DPhil is how true it is to say that research does not take place in a vacuum. I am grateful to those with whom I have had the pleasure of meeting, working and interacting over the course of my DPhil. I can't name them all but there are a few who deserve special mentions.

First and foremost to Rob and Chris for guiding my research over these years and providing excellent supervision, without which I would have been truly lost at many points. It has been a joy to discuss science with you both and I have learnt, amongst a great many other things, more than I ever thought I would about Real Tennis, Leeds United and some great music.

To my research colleagues and friends, I really can't thank you enough. Starting my DPhil during COVID lockdowns was not easy and I spent my first year wondering ever I would feel part of anything. Since then you have all welcomed me into this community of kind, funny and interesting people, for which I am so grateful. To Lauren, Stefano, Joe, Katie, Francesco, James, Jakob, Amy, Dave, Sara, Ed and Ian, without your scientific and technical discussions and advice, I would not be half the scientist I am now. Special mentions go to the foremost few, with whom memories attending conferences, meetings and workshops were some of the highlights of the DPhil experience.

Similarly, to the Zooniverse team I extend immense gratitude. Without the amazing platform and dedication you all put in, many millions of lives would be less enriched by the beautiful act of human collaboration. I am lucky to have been a very small part of an amazing story. Likewise to the rest of the ThunderKATs, it has truly been a pleasure getting to know you all.

I was also lucky enough to spend a significant amount of my DPhil living in and visiting South Africa, where I had the most brilliant time. Special thanks go to Patrick and Michelle, both of whom have looked after me and welcomed me to UCT and UWC, where I really felt I grew as a scientist and as a person. To Moses, Reika, Kira, Kelebogile, Nathan, Amo, Miranda, Okkert, Sumari, Tim, Simon, Nikki, Henco, Krishna, Senate and Martin thank you for making me feel welcome. I look forward to visiting and working with you some time soon.

To twist a great quote from a great scientist, whilst I thank my research colleagues for their helpful discussions, I thank my office comrades for their discussions. To

Mads, Curro, Casey, Jaime, Eric and James - it has been one hell of a ride and the memories I have made with you all will last a lifetime. I can't wait to make more memories with you all and look back on the times we have had together. The Tower has fallen, long live the Tower.

To my friends, past and present, our interactions have enriched my life more than I can say. Special mentions go to the physics cohort of UoM 2020 and especially Martin and Meg, who put up with my antics during our undergrad days. To my many housemates and theatrical friends, you are always in my heart. Friends in Oxford, it has been the utmost pleasure to get to know you all and I look forward to at least 3 more years visiting the White Rabbit.

It goes without saying that my family made me the person I am today. To Mum, without your never-ending support and belief, I would never have gotten this far. To Dad, for your enthusiasm for (and participation in!) my science, I really can't express how much it means to me to have you in my corner. To Geoff, your advice on academia and friendship has been a guiding light through the DPhil. JJ, your support means the world to me and I am so proud of where we both are now - who would have thought that two nerds playing Halo, Fable and WoW in their rooms together could grow up and be doing what we're doing? To the Jacksons, being part of your family has been such a blessing. Steph you are the best sibling I could have ever hoped to inherit and to Rosita and Ariella, I hope the world you grow up into is one in which we can all watch the stars together.

Finally, to Tang. Oh Tang. What a journey we have had. Getting my decision letter from Oxford whilst we were locked down together in your uni house was a strange time in our lives and we didn't know then what would come of it all. Here we are, 4 years later, with a whole life built together. Your never-ending patience, your ability to ground me and your love, support and care have made this thesis possible. To bask in the warmth of your light has been the greatest joy of all and I can't wait to see where we go from here.

Statement of Originality

I carried out the work presented in this thesis as a student at the University of Oxford between October 2020 and May 2024, supervised by Profs. Rob Fender and Chris Lintott, who have both provided guidance throughout this work and are therefore co-authors on all articles mentioned below. This work was funded by a Science Technology Facilities Council Studentship, project reference 2444981. I hereby declare that no part of this thesis has been submitted in support of another degree, diploma or other qualification at the University of Oxford or any other higher learning institute. Except where otherwise stated or where reference is made to the work of others, the work in this thesis is entirely my own.

The work in Chapter 3 is based on the peer-reviewed paper Andersson et al., 2022 (MNRAS Volume 513, Issue 3, pp.3482-3492), for which I am the lead author and contributed all of the analysis and writing. As part of this work, data reduction was done by David Williams, technical and scientific advice was given by Laura Driessen, Alexander van der Horst, Sara Motta, Rachel Osten and Lauren Rhodes. Initial *TESS* analysis and description was done by Nora Eisner, whilst SALT spectra were taken and reduced by David Buckley. Paul Groot, Steven Bloemen and Paul Vreeswijk are responsible for the development of MeerLICHT and its data products, which were used in this work, whilst Patrick Woudt and Rob Fender were the principal investigators of the ThunderKAT project, from which the radio data came.

The work in Chapter 4 is based on the peer-reviewed paper Andersson et al., 2023 (MNRAS Volume 523, Issue 2, pp.2219-2235), for which I am the lead author and contributed all of the technical work and written text. As part of this work, data calibration and reduction was done by Joe Bright, Francesco Carotenuto, Laura Driessen, Mathilde Espinasse, Kelebogile Gasealahwe, Sara Motta, Lauren Rhodes, Evangelia Tremou, David Williams and Xian Zhang. Technical and scientific advice was given by Stefano Giarratana, Ian Heywood, Alexander van der Horst, Payaswini Saikia. Paul Groot, Steven Bloemen and Paul Vreeswijk are responsible for the development of MeerLICHT and its data products, which were used in this work, whilst Patrick Woudt and Rob Fender were the principal investigators of the ThunderKAT project, from which the radio data came. Jonas Andersson, Lizzeth Ruiz Arroyo, Loïc Baert, Matthew Baumann, Wilfried Domainko, Thorsten Eschweiler, Tim Forsythe, Sauri Gaudenzi, Rachel Ann Greiner, Davide Iannone,

Karla Lahoz, Kyle Melville, Marianne De Sousa Nascimento, Leticia Navarro, Sai Parthasarathi, Pilonen, Najma Rahman, Jeffrey Smith, B. Stewart, Newton Temoke, Chloe Tworek and Isabelle Whittle were all volunteers whose contributions to the work warranted co-authorship.

The work in Chapter 5 is based on a scientific paper still in preparation, for which I am the lead author and have contributed all of the technical work and written text. Michelle Lochner provided advice and guidance throughout the development of this work, as well as writing a significant amount of the public codebases used herein. Advice on the text has been given by Jakob van den Eijnden, Assaf Horesh, Alexander van der Horst, Payaswini Saikia, Gregory Sivakoff, Evangelia Tremou and Mattia Vaccari. Again the data come from ThunderKAT, of which Patrick Woudt and Rob Fender were the principal investigators.

The copyright of this thesis rests with the author. The author asserts his moral rights.

Abstract

Modern radio telescopes are data-intensive machines, producing many TB of data every night. Amongst this deluge of data are transient and variable phenomena, whose study can shed new light on processes as varied as stellar dynamos and the accretion discs in supermassive black holes. In this work I demonstrate the applicability of different methods to the discovery of these astrophysical transients and variables coming from telescopes such as MeerKAT.

I first consider a standard approach to discovering transients by characterising their variability. By making use of even modest sampling with the high sensitivity and wide field of view of MeerKAT, I demonstrate how we are now able to uncover new transients almost by accident - if we exclude the vast amount of time spent planning, building and operating excellent telescopes, efficient pipelines and well-crafted observing proposals. In this work I found a stellar flare from a nearby M dwarf, which was then followed up and complemented by optical and X-ray photometry and spectroscopy, providing new insights on the system.

Next I built a citizen science platform in order to perform such transient searches at scale, making use of a wide range of data available in the MeerKAT archive. I detail the process of review and beta-testing that resulted in the final design of the *Bursts from Space: MeerKAT* project. Over 1000 volunteers took part, demonstrating a healthy appetite for further Zooniverse data releases. Volunteers discovered or recovered a wide range of phenomena, from flare stars and pulsars to scintillating AGN and transient OH maser emission. I was also able to use the known transients in our fields to understand some reasons why interesting sources may be missed and will fold this learning through to future iterations of the project. This is the first demonstration of volunteers finding radio transients in images.

Finally, I show how anomaly detection, an unsupervised machine learning approach, is a suitable tool for finding these variable phenomena at scale, as is required for modern astronomical surveys. I use three feature sets as applied to two anomaly detection techniques in the `ASTRONOMALY` package and analyse anomaly detection performance by comparison with citizen science labels. By using transients found by citizen scientists as a ground truth I demonstrate that anomaly detection techniques can recall over half of the radio transients within 10% of the sample dataset. I find that the choice of feature set is crucial, especially when considering

available resources for human inspection and follow-up. I find that active learning on $\sim 2\%$ of the data improves recall by up to 10%, depending on the feature-model pair. The best performing feature-model pairs result in a factor of 5 times fewer sources requiring vetting by humans. This is the first effort to apply anomaly detection techniques to finding radio transients and shows great promise for application to other datasets, a real-time transient detection system and upcoming large surveys.

Contents

List of Figures	xi
List of Tables	xviii
List of Abbreviations	xix
1 Introduction	1
1.1 Serendipity in astronomy	1
1.2 Sources of radio transients	3
1.2.1 Incoherent radio sources	6
1.2.2 Coherent emitters	11
1.3 Emission mechanisms and propagation	14
1.3.1 Incoherent radio emission	16
1.3.2 Coherent emission mechanisms	18
1.3.3 Propagation effects	23
1.4 The search for radio transients	25
2 MeerKAT and the ThunderKATs	30
2.1 MeerKAT	30
2.1.1 Data reduction	32
2.2 ThunderKAT	34
2.2.1 TraP	38
2.2.2 MeerLICHT	42
3 The detection of stellar flares with MeerKAT	44
3.1 Introduction	45
3.2 Radio observations	47
3.3 Archival multiwavelength association	53
3.3.1 Optical and near-IR photometry	56
3.3.2 X-ray observations and correlation	58
3.4 SALT spectroscopy	65
3.5 Discussion	67
3.6 Conclusions	71

4	Citizen science for radio transients	73
4.1	Introduction	74
4.1.1	Rationale	74
4.1.2	Citizen science in astronomy	76
4.2	ThunderKAT observations and pre-processing	82
4.2.1	Source extraction and subject generation	82
4.3	Citizen science platform	83
4.3.1	Beta testing and development	83
4.3.2	Final design	89
4.3.3	Data release	89
4.4	Results	93
4.4.1	Comparison to target sources	94
4.4.2	Previous commensal studies	96
4.4.3	Scintillation analysis	99
4.5	Counterparts and associations	102
4.5.1	Highlights	104
4.6	Discussion	111
4.7	Conclusions	115
5	Anomaly detection	117
5.1	Introduction	118
5.1.1	Data-intensive astronomy	118
5.1.2	Machine learning in transient astronomy	119
5.2	Observations and classifications	125
5.3	Feature extraction	128
5.3.1	Baseline features	128
5.3.2	Feets features	129
5.3.3	Wavelet features	130
5.3.4	Feature space and clustering	134
5.4	Anomaly detection	135
5.4.1	Anomaly recall and purity	139
5.5	Active learning	144
5.5.1	Active learning improvement	146
5.6	Discussion	147
5.7	Conclusions	152
6	Conclusions and Future Work	153
6.1	Stellar activity	153
6.2	Citizen science and transient searches	155
6.3	Anomaly detection with large surveys	157
6.4	Final remarks	158

Appendices

A	Transient searches and citizen science in South Africa	161
A.1	Transient Searches	161
A.2	Engagement in Carnarvon, Cape Town and online	163
B	Bursts from Space:MeerKAT results table	169
	Bibliography	175

List of Figures

1.1	The luminosity-duration parameter space for radio variables and transients, taken from Keane (2018) as originally presented by Pietka, Fender, and Keane (2015). The maximum brightness temperature of incoherent transients delimits a boundary above which all variables must emit via a coherent mechanism.	5
1.2	Left: A schematic of the processes involved in a solar flare, where magnetic reconnection in the corona is the energy source from which all subsequent processes arise. Electrons are accelerated, generating the radio burst, some of which are channelled down into the chromosphere, causing brightening in X-rays, as well as the enhancement of emission lines such as H α . Right: A 17.1 nm image from the <i>Transition Region And Coronal Explorer</i> satellite taken 2 hours after a flare. Figures both from Lang (2009).	7
1.3	An artist’s impression of a black hole X-ray binary, displaying the accretion from a companion star onto the black hole and the emission of radio jets. <i>Image Credit: John A. Paice, https://ras.ac.uk/media/344</i>	9
1.4	The first detected FRB (Lorimer et al., 2007). The dispersive sweep over 100s of ms indicates a large distance to the source and is typical of many FRBs.	13
1.5	The characteristic intensity spectra for the emission mechanisms discussed here, as shown by Dulk (1985), where δ is what I have called p . In all cases we can see increasing functions at low frequencies before reaching a peak after which the spectra flatten or decrease. .	19
1.6	Flowchart from Melrose (2017), demonstrating how plasma emission can be generated by plasma oscillations, as originally discussed by Ginzburg and Zhelezniakov (1958).	20
1.7	A schematic of a pulsar magnetosphere, from Burke and Graham-Smith (2014). The light cylinder, within which the magnetosphere co-rotates with the neutron star, is defined by radius r_c	22

2.1	Logarithmic distribution of all MeerKAT antennae in the Karoo. The dense core of antennae allows for high sensitivity to extended structures, whilst the sparse outer antennae provide longer baselines for higher resolution.	31
2.2	Map of the XRB fields (blue circles) used in this work, with names given. This is overplotted on a map of the sky taken at 408 MHz, from Remazeilles et al. (2015) and originally presented by Haslam et al. (1982).	36
2.3	Schematic of the Transients Pipeline, adapted from Swinbank et al. (2015) Data are read in and processed by the pipeline in a combination of Python scripts and SQL database operations. The final products are a set of light curves and their properties (fluxes, positions, variability statistics) which users can then query to search for transients and variables.	39
3.1	Variability parameters for all sources detected in the H1743–322 field during ThunderKAT observations in late 2018. Clear outliers, particularly in V can be seen, including the XRB H1743–322 itself, several manually vetted artifacts and an unknown source.	50
3.2	Three example false positive light curves from this dataset. These sources were all found to be around bright sources, with detections only in the first three epochs. This was likely due to a different calibrator being used for these observations.	51
3.3	One arcminute square images of MKT J174641.0–321404 and surrounding noise during all ThunderKAT observations, running chronologically from top left, across then down. The source can be clearly seen to appear and disappear multiple times across all epochs. The synthesised beam of each image can be seen in the bottom left of each panel, whilst the white central cross-hairs show the source’s mean MeerKAT position.	52
3.4	Upper: Radio light curve of MKT J174641.0–321404 across all 11 ThunderKAT observation epochs. Non-detections are defined as measurements below a 3σ threshold (dashed line) where σ is the local RMS noise in that epoch. The non-physical measurements in some epochs are caused by the forced photometry and epoch-specific noise structure, visible in Figure 3.3. The higher detection threshold (due to more noise) in the first three epochs is likely due to the use of a different phase calibrator than the later observations. Lower: Radio light curve compiled from 5×3 minute slices of the brightest 15 minute epoch. From this there is evidence for marginal variability within the epoch.	54

- 3.5 MeerLICHT q -band (440-720 nm) reference image (pixels) of the region near radio source MKT J174641.0–321404. 1 and 2σ positional uncertainties on the MeerKAT position (circles) show an overlap between the MeerLICHT object and the *Gaia* proper motion attributed to SCR 1746. 56
- 3.6 The *TESS* light curve of SCR 1746 (Tess Input Candidate 111898820) spanning 27.9 days from May to June of 2021, including the 2-minute cadence data (purple) and the same data binned down to 30 minute cadence. Flare amplitudes can be seen up to $30\times$ above non-flaring times. 59
- 3.7 **Upper:** The phase folded *TESS* light curve of SCR 1746, with the corresponding Lomb-Scargle periodogram inset. Flares tend to occur between 0.0 and 0.4 in phase, corresponding to the brighter parts of the quiescent light curve. A sinusoidal best fit to the unbinned data is plotted, showing a relative amplitude of 0.52 ± 0.01 . **Lower:** A closer view at the non-flare behaviour of the phase-folded light curve, in which the amplitude of the best fit sine curve is clearly visible. The 30 minute cadence data reduces the scatter and falls in good agreement with the sinusoid. 60
- 3.8 **Upper:** Optical data from MeerLICHT and ASAS-SN. Flaring behaviour can be seen in data from the ASAS-SN Sky Patrol portal between radio observations and mid-2021. Downwards facing triangles indicate upper limits in their respective bands. **Lower:** The radio light curve of SCR 1746 overlaid with the simultaneous optical data. No flares are simultaneous to within < 0.4 days with those in the other bands. 61
- 3.9 A stacked image of the total count rate across all Swift-XRT energies towards SCR 1746, as provided by the 2SXPS catalogue (Evans et al., 2020). The image is $\sim 6.3'$ to a side and shows a detection of a source that, whilst consistent with the MeerKAT uncertainties, does not formally overlap with the *Gaia* J2000 position of SCR 1746. The increased background counts to the right of the image are caused by H1743–322. 62
- 3.10 X-ray and radio luminosities of several types of active star (Guedel and Benz, 1993) - reproduced from <https://github.com/AstroLaura/GuedelPlot>. Limits on SCR 1746's quiescent emission show this is an intrinsically faint source and could be approximately consistent with the Güdel and Benz relation for M dwarf (dM/dMe) stars. The radio flare is also plotted, showing a brightness increase of over an order of magnitude above quiescent limits. 64

3.11	Caption with citation	66
4.1	A map of the path of totality for the 1715 eclipse, as compiled by Halley and Senex (1715) updated from his initial prediction based on observations from citizen scientists. Edited from the Institute of Astronomy Library’s copy of this map, accessible at http://www.dspspace.cam.ac.uk/handle/1810/221308	78
4.2	Examples of the four observational classes within the workflow on the Zooniverse, showing both the light curve and an image of each. From top to bottom these are: Stable - no variation in the light curve given the error bars and a point source; Extended - variations caused by changes to the PSF and a source that is larger than the beam (lower left); Transient - a clear variable light curve for a source the same size as the PSF; Artefact - a spuriously transient light curve and a faint source on the outskirts of a very bright object, with non-Gaussian noise structure. The final class, Unsure , by definition has no archetypal characteristic no figure is shown here.	86
4.3	Volunteer responses to each of the three tested workflows in the <i>BfS:MKT</i> beta test.	87
4.4	Classification workflow for <i>BfS:MKT</i> , showing the light curve and local sky figures for GX339–4.	90
4.5	Histogram of the number of classifications made per user on the <i>BfS:MKT</i> project. As might be expected, many users perform only a few classifications, whilst a few volunteers dedicate thousands of votes to the project.	91
4.6	Variability plane for the 168 sources identified by citizen scientists to be variable, along with those they find to not be so in grey. The outer histograms are the marginal distributions in η and V of volunteer identified variables (purple) and everything else (grey), normalised such that the area under each curve sums to one. The colour bar denotes the fraction of classifications as a transient/variable source. Known variable sources (e.g. XRBs) in our fields are circled. Imaging artefacts appear at low η and high V , whilst flux calibration uncertainties can produce high η , low V sources (due to lack of systematic uncertainty in equation 2.2). Most known transients are found by citizen scientists, whilst many new sources are identified and show a wide spread of values in this parameter space.	95

- 4.7 The TRAP light curves of EXO 1846–031, *Swift* J1858.6–0814 and SAX J1808.4–3658 (upper left, upper right and lower respectively), shown here exactly as they were to volunteers on *BfS:MKT*. The foremost was classified as transient by citizen scientists whilst the others were not. I believe that *Swift* J1858.6–0814 was misclassified due to the mis-scaled, automatically generated legend and the lack of apparent variability aside from the first data point. Similarly SAX J1808.4–3658 was likely misclassified due to the lack of data points on which to make a judgement. 97
- 4.8 The ratio between observed V and the variability predicted by Hancock et al. (2019)’s model for Refractive Interstellar Scintillation (RISS), for our sample of transients and variables. Most variability can be explained by this model of a scintillating, extragalactic source, apart from the known XRBs and jetted systems. 103
- 4.9 The mean optical and radio flux densities of our sample of radio variables, atop an underlying distribution of astrophysical classes (Stewart et al., 2018). Black crosses denote counterparts within the MeerLICHT database whilst grey triangles are upper limits. Diagonal lines denote a constant ratio between radio and optical flux density, whilst the A_R marker indicates the horizontal displacement caused by 5 magnitudes of optical extinction. The majority of our radio sources are likely extragalactic as they overlap in parameter space with quasars and GRBs. 105
- 4.10 **Upper:** Light curve of OH Maser 30.1–0.7 picked up by volunteers. **Lower:** Overlay of MeerKAT radio contours over Spitzer imaging from the GLIMPSE survey (Benjamin et al., 2003). Contours are spaced linearly in 0.5 mJy increments from -0.5 to 3.5 mJy. 107
- 4.11 Light curve of PSR B1845–01 with MeerKAT. 7 of 10 volunteers voted for this as a transient/variable source. The low amplitude variability observed over a timescale of 100s of days is consistent with RISS. 109
- 4.12 Light curve of source VLASS1 J181955.28+074418.7, showing smooth, near sinusoidal variations. The epochs denoted upper limits are in error and due to not filtering out low quality images in pre-processing. This light curve was described as “too perfect” (to be real) and “beautiful” by volunteers.¹⁰ 112

5.1	Example light curves of an XRB (MAXI J1820+070, top), a false positive caused by bad data (middle) and a stable light curve (bottom). My aim is for unsupervised machine learning algorithms to preferentially select light curves similar to the foremost over the latter two. The line and shaded region of each panel denote the mean and 1σ uncertainty of the computed Gaussian Process regression, detailed in section 5.3.3.	127
5.2	The first eight symlet wavelets, used in this work to build a feature set for describing light curves. Taken from the Mathworks wavelet introduction. ¹	132
5.3	The UMAP of the <i>feets</i> feature space, colour-coded by observational field. Stars indicate volunteer-verified anomalies (transients). We see that clustering is mainly a function of observational field, aligning with the heterogeneous nature of the dataset. I note that the axes are unlabelled as these numbers are arbitrary combinations of many features and have no physical meaning. The inset is a demonstration of how the data of one particular field (GX339) is distributed, as mentioned in the text.	136
5.4	The UMAP of the wavelet coefficient feature space, colour-coded and labelled as in Figure 5.3. In this instance, the observational fields with longer light curves (e.g. J1858, GX339–4) show distinct patterns, whilst most other sources are in a very dense region of parameter space, as shown by the inset zoom.	137
5.5	The (η, V) feature space, colour-coded by observational field as in previous figures. This is the same as Figure 4.6 but with a different colour-coding.	138
5.6	Upper panel: The recall curves for both of the anomaly detection algorithms as applied to three sets of features derived from the radio light curves. The dotted line in each corresponds to $y = x$ i.e. random chance. Lower panel: A combined plot of all six recall lines for better comparison between feature sets. Dashed lines make comparisons (see text) for recall rates at 10 and 30% of the dataset. The values in the legend correspond to the area under each recall curve (AUC).	141
5.7	The Astronomaly GUI, showing subjects in anomaly-order. Using this interface, I can easily inspect the data and provide labels on some subset of objects in order to perform active learning.	144
5.8	Changes in performance due to active learning for each model-feature pair, where $y = 0$ indicates no improvement. This plot has been truncated at $x = 0.4$ as, after this point, all curves vary minimally about 0. I show two figures for clarity when comparing between curves.	147

5.9	A comparison between how many citizen scientists voted for a source as a transient and its ranking in the anomaly detection work, where a rank of 1 indicates the most anomalous source. That is, the most anomalous source in my anomaly score ordered list has a ranking of 1, whilst the least anomalous source has a ranking of 0. Sources ‘not recovered’ refers to known variables/transients in the light curves that fall below the 4/10 threshold used in Chapter 4 and were therefore regarded as missed by volunteers. The background points are all remaining sources, coloured by the density of points. The anomaly list used here corresponds to the wavelet features applied to the LOF algorithm, but results are consistent for all feature-model pairs. . . .	151
A.1	The mean optical and radio flux densities of the radio stars found by the ThunderKAT team, atop an underlying distribution of astrophysical classes from (Stewart et al., 2018).	164
A.2	Detections of radio stars found in ThunderKAT observations, compared with recent work from the ASKAP telescope. The dotted, dashed and solid lines correspond to 5σ detection limits for the data from Pritchard et al. (2024), 15 minutes with MeerKAT and 1 hour with the SKA-Mid array, i.e. 1.25 mJy, 60 and 22 μ Jy respectively.	164
A.3	One epoch of the Chandra Deep Field South, taken as part of the MIGHTEE survey with MeerKAT. It is light curves of these sources that volunteers classified as part of DR2.	166
A.4	Number of classifications per day over the first two months of DR2 of <i>BfS:MKT</i> . Each uptick corresponds to the release of an article, podcast or engagement session.	168

List of Tables

1.1	Summary of how the emission mechanisms discussed relate to the observations of different variable and transient astrophysical sources as discussed in section 1.2.	15
2.1	Properties of the 11 ThunderKAT datasets used in this work. Each field’s approximate Galactic latitude is given for relevance in section 4.4.3 and Figure 4.8. The number of sources and central RMS values are calculated by the TRAP (see sections 2.2.1 and 4.2.1).	37
3.1	Select magnitudes of red dwarf SCR 1746–3214 across the visible and near-IR spectrum. Bands are listed with their effective wavelength λ_{eff}	57
5.1	Features used in this work as calculated by the <i>feets</i> package (Cabral et al., 2018b). Each light curve consists of $N = \sum i$ measurements of flux density $F_i \pm \sigma_i$ taken at time t_i . A similar version, for the original FATS code is described by Webb et al. (2020).	131
5.2	Specific recall (R) and purity (P) values, rounded to two decimal places, for each model-feature set pair, measured after ten and thirty percent of the most anomalous sources have been ‘seen’. Each pair of values listed are the default, active iteration of that feature-model instance.	142
B.1	Table of transients and variables found by volunteers during BfS:MKT, with their associated positions, median flux densities (F_{med}) and variability parameters η and V . The observed date given here is that in which the source was detected at highest S/N by the TRAP, whilst the distance recorded is how far an entry is from the pointing centre of that observation. The Transient Fraction (TF) is the fraction of classifications as a transient/variable, from 10 volunteers. Note that their Zooniverse subject ID is a unique identifier from which one can find the image and light curve online and freely accessible at https://www.zooniverse.org/projects/alex-andersson/bursts-from-space-meerkat/talk/subjects/ID , replacing ID with the numeric value below. This table will be available in machine-readable form online.	170

List of Abbreviations

1-D, 2-D	One- or two-dimensional, referring in this thesis to spatial dimensions in an image.
AAVSO	A merican A ssociation of V ariable S tars O bservers
AGB	A symptotic G iant branch
AGN	A ctive G alactic N ucleus/ N uclei
APEC	A strophysical P lasma E mission C ode
ASAS-SN	A ll S ky A utomated S urvey for S upernovae
ASKAP	A ustralian S quare K ilometre A rray P athfinder
ATA	A llen T elescope A rray
ATCA	A ustralian T elescope C ompact A rray
AU	A stronomical U nit
BfS:MKT	B ursts from S pace: M eer K AT
BL	B reakthrough L isten
BLUSE	B reakthrough L isten U ser S upplied E quipment
CHIME	C anadian H ydrogen I ntensity M apping E xperiment
CSM	C ircumstellar M edium / M edia
CNN	C onvolutional N eural N etwork
CV	C ataclysmic V ariable
DISS	D iffractive I nterstellar S cintillation
DM	D ispersion M easure
DR	D ata R elease
ECMI	E lectron C yclotron M aser I nstability
EW	E quivalent W idth
FFT	F ast F ourier T ransform
FIRST	F aint I mages of the R adio S ky at T wenty C entimetres [

FITS	Flexible Image Transport System
FoV	Field of View
FRB	Fast Radio Burst
GOTO	Gravitational-wave Optical Transient Observer
GRB	Gamma-Ray Burst
GUI	Graphical User Interface
H-R	Hertzsprung-Russell (diagram)
HRS	High Resolution Spectrograph
ICRS	International Celestial Reference System
IF	Isolation Forest
IR	Infrared
ISM	Interstellar Medium
LADUMA	Looking At the Distant Universe with MeerKAT
LOF	Local Outlier Factor
LOFAR	Low Frequency Array
LSP	Large Survey Project
LSST	Legacy Survey of Space and Time
Maser	Microwave Amplification by Stimulated Emission of Radiation
MeerKAT	The Meer ('more') Karoo Array Telescope
MeerTRAP	More TRAnsients and Pulsars
MIGHTEE	MeerKAT International GigaHertz Tiered Extragalactic Exploration
MJD	Modified Julian Date
MWA	Murchison Widefield Array
NASA	National Aeronautics and Space Administration
NN	Neural Network
NRAO	National Radio Astronomy Observatory
NVSS	NRAO VLA Sky Survey
PCA	Principal Component Analysis
PLAsTiCC	Photometric LSST Astronomical Time-Series Classification Challenge
PSF	Point Spread Function

RA	R ight A scension
RF	R andom F orest
RFI	R adio F requency I nterference
RISS	R efractive I nterstellar S cintillation
RNN	R ecurrent N eural N etwork
SALT	S outhern A frican L arge T elescope
SARAO	S outh A frican R adio A stronomy O bservatory
SCR	S uper C OSMOS- R ECONS
SDSS	S loan D igital S ky S urvey
SETI	S earch for E xtraterrestrial I ntelligence
SIMBAD	T he S et of I dentifications, M easurements and B ibliography for A stronomical D ata
SKA	S quare K ilometre A rray
SKAO	S KA O bservatory
SN(e)	S upernova(e)
SQL	S tructured Q uery L anguage
TDE	T idal D isruption E vent
TESS	T ransiting E xoplanet S urvey S atellite
ThunderKAT	T he H unt for D ynamic and E xplosive R adio T ransients with M eer K AT
TraP	L OFAR T ransients P ipeline
UHF	U ltra H igh F requency
UT	U niversal T ime
UV	U ltraviolet
VLA	V ery L arge A rray
VLA SS	V LA S ky S urvey
VAST	V ariables A nd S low T ransients
XRB	X - R ay B inary
ZTF	Z wicky T ransient F acility

*And you may ask yourself,
“Well, how did I get here?”*

— David Byrne, from Talking Heads’
Once in a Lifetime

1

Introduction

Contents

1.1 Serendipity in astronomy	1
1.2 Sources of radio transients	3
1.2.1 Incoherent radio sources	6
1.2.2 Coherent emitters	11
1.3 Emission mechanisms and propagation	14
1.3.1 Incoherent radio emission	16
1.3.2 Coherent emission mechanisms	18
1.3.3 Propagation effects	23
1.4 The search for radio transients	25

1.1 Serendipity in astronomy

The very existence of radio astronomy as a field is owed to an accidental discovery. Whilst determining interference levels in radio communications as an engineer at Bell Laboratories, Karl Jansky discovered a sidereal source of radio waves from outside the Solar System (Jansky, [1933](#)), emanating from somewhere within the constellation Sagittarius. We now know the source of this “hiss type atmospheric” to be the Galactic Centre and the supermassive black hole at its core, Sagittarius A*. In the intervening ~ 90 years, the Galactic Centre has remained a region of active study due to its extreme physical conditions (e.g. Heywood et al., [2022b](#)).

This serendipitous finding demonstrates one of the two ways discoveries in astrophysics can be made. Conventionally, measurements can be made that confirm a hypothesis, therefore discovering the existence of a phenomenon. The detection of gravitational waves (Abbott et al., 2016) is an excellent example of this, testing General Relativity a century after initial predictions were made (Einstein, 1916; Einstein, 1918). However, in observation-based sciences, discoveries can also be made through a process of exploration and “accidental” discovery, with no specific hypothesis being tested (Harwit, 2003).

Throughout the 20th century, further serendipitous discoveries with radio telescopes have advanced our understanding of the Universe immensely. I list just three relevant examples here but note that many more cases exist (see Fender et al., 2015; Norris, 2017). Observational evidence for a Big Bang was found by accident (again by members of Bell Laboratories) in the form of a temperature reading ~ 3 K higher than expected from a horn-reflector antenna (Penzias and Wilson, 1965). This temperature was, in the same journal issue, identified as the remnant blackbody radiation from the Universe’s hot and dense initial conditions (Dicke et al., 1965). Only a few years after this, then PhD student Jocelyn Bell investigated a strange-looking signal on a chart recorder connected to a telescope designed for studying interplanetary scintillation and ended up discovering the very first pulsar (Hewish et al., 1968). Finally, in 2007, whilst searching through archival data from a survey of the Small Magellanic Cloud, a one-off, ms-duration radio burst was observed, likely coming from a distance much further than any Milky Way satellite galaxy (Lorimer et al., 2007). The discovery of each of these; the cosmic microwave background, pulsars and fast radio bursts (FRBs), has led to the development of whole sub-fields within radio astronomy and each has been recognised with awards; the 1974 Nobel, the 2018 Special Breakthrough and the 2023 Shaw prizes respectively. These discoveries highlight the idea that untargeted searches with radio telescopes are a fruitful endeavour and can shed light on a wide range of phenomena.

In modern astronomy, the volume of observations from telescopes is ever-increasing, with more data than ever before in which to find interesting, rare

and even new phenomena. The Square Kilometre Array (SKA) telescope will be the world's largest radio observatory, consisting of two separate facilities - a low frequency component (50 – 350 MHz) in Australia and a mid-frequency component (0.35 – 15 GHz) in South Africa (Dewdney et al., 2009). Some of the key science aims of the SKA include the detection of gravitational waves through observations of pulsars, constraining the origin of cosmic magnetic fields and looking for tracers of life from other planets (Braun et al., 2015). However, the expected data rates are unprecedented in astronomy, with cumulative data products reaching 300 PB/year (Scaife, 2020).

The aim of this thesis is to find and test methods for finding sources that vary in brightness as seen by radio telescopes, such as the pulsars and FRBs already mentioned. Some of these variables will vary to such an extent that for the majority of the time they cannot be seen by telescopes at all when not in a bright phase i.e. they are transient. I will use the terms variables and transients interchangeably. Previous studies have shown that $< 5\%$ of sources in the radio sky are variable when observing around 1 GHz (Ofek et al., 2011; Mooley et al., 2016; Driessen et al., 2022), making searches for transients and variables an exercise in searching for anomalies. This, combined with the large data rates from radio observations, creates a needle-in-a-haystack problem of astronomical proportions. As such, in this thesis I will describe the techniques I have used to address this problem and present the results of searches for variables and transients in radio observations, after having given an overview of the relevant theory and observations.

1.2 Sources of radio transients

At radio frequencies the spectral luminosity L_ν observed at frequency ν from a spherical blackbody of radius R and temperature T_B can be approximated by

$$L_\nu = \frac{2\nu^2 k_B T_B}{c^2} 4\pi R^2, \quad (1.1)$$

where k_B is the Boltzmann constant and c is the speed of light in a vacuum. Rearranging this Rayleigh-Jeans approximation we then get

$$T_B = \frac{L_\nu}{8\pi\nu^2 k_B} \frac{c^2}{R^2}. \quad (1.2)$$

For a transient source of radio waves that lasts for time Δt , the emission region cannot be larger than $R = c\Delta t$ by light travel time arguments and therefore we estimate a lower limit to a source's brightness temperature as

$$T_{B,min} = \frac{\Delta L_\nu}{8\pi\nu^2 k_B \Delta t^2}. \quad (1.3)$$

This brightness temperature therefore encodes two of the most important qualities of a transient radio source: its duration and its luminosity. For this reason, different radio transient phenomena are typically classified according to their brightness temperature. Specifically, sources with $T_B > 10^{12}$ K must be coherent, where the responsible particles (usually electrons) are emitting with some phase relationship, whilst those below this threshold are incoherent emitters. This boundary is based on the finding that incoherent sources whose temperature is above this threshold will rapidly undergo energy losses and return to a lower T_B (Kellermann and Pauliny-Toth, 1969).

A broad and useful representation of these ideas is shown in Figure 1.1, where the timescales and luminosities of various classes of astrophysical transients are plotted. In this figure it can be seen that there exist populations of transients either side of the divide, with coherent emitters generally being shorter in duration. This difference in timescales dictates the observational techniques used to find each kind of emitter; incoherent transients are usually observed by taking repeated images with a radio telescope, whilst coherent sources are observed through beamforming techniques. Many of the time-variable phenomena present in Figure 1.1 such as pulsars and decametric emission from Jupiter were first observed using beamforming techniques (Hewish et al., 1968; Burke and Franklin, 1955), whereby the voltage signals from a set of radio dishes or antennae are squared and then summed. Modern beamforming datastreams typically employ sub-second time cadence, at the expense




Figure 1.1: The luminosity-duration parameter space for radio variables and transients, taken from Keane (2018) as originally presented by Pietka, Fender, and Keane (2015). The maximum brightness temperature of incoherent transients delimits a boundary above which all variables must emit via a coherent mechanism.

of spatial resolution in order to allow for data storage and transient searches within computational limits. By contrast, searches for slower variability have been done by forming radio images, the core operations of which involve multiplications between each pair of antennae in an array and Fourier transforms, making image plane searches computationally expensive. However, recent increases in computing power and improved baseline coverage (see Section 2.1 for more on this) have allowed for untargetted transient searches with imaging surveys, observing large fields of the sky whilst simultaneously providing precise localisations (Fender et al., 2015). These advances have allowed for image plane searches to push to higher time resolution and more sensitive radio images i.e. filling in the parameter space outlined in Figure 1.1. Below I give a summary of some of the most common transient and

variable radio phenomena observed in the Universe, focusing on observations taken at frequencies around 1 – 4 GHz unless otherwise specified.

1.2.1 Incoherent radio sources

1.2.1.1 Solar and stellar radio flares

The Sun is the brightest radio source in the sky, emitting transient radio bursts with flux densities up to at least 10^9 Jy (Castelli, Aarons, and Michael, 1968). These bursts are related to the optical flares and plasma ejections that have been seen to occur on the Sun for over 150 years (Carrington, 1859). The exact mechanisms that drive flares and radio bursts are still not precisely understood (e.g. Benz, 2017), but there is a consensus that the motion of magnetic fields within the corona leads to a build-up of energy as the field lines become twisted (Yamada, Kulsrud, and Ji, 2010). Eventually, magnetic reconnection occurs, where two field lines approach one another and restructure into a lower energy configuration, leading to an impulsive release of energy as plasma is freed and accelerated (Yamada, Kulsrud, and Ji, 2010). This release of energy is thought to power the incoherent radio bursts seen on the Sun and other stars, along with visible and X-ray emission. A schematic of this process can be seen in Figure 1.2. Importantly, type IV solar bursts, which often occur during and after large flares, are correlated with space weather events such as coronal mass ejections (Robinson, 1986; Salas-Matamoros and Klein, 2020). These coronal mass ejections, where magnetised plasma is expelled from the corona (Webb and Howard, 2012), have the potential to cause damage to spacecraft, personnel in space and electricity networks on Earth (Serra et al., 2020).

Moving further away from the Earth, a variety of magnetically active stars have been seen at radio frequencies. The first such radio flare was observed over 50 years ago (Lovell, Whipple, and Solomon, 1963; Lovell, 1969) from YZ Canis Minoris. Many of these active stars are M-dwarfs, with masses $0.07 - 0.6 M_{\odot}$ (Rajpurohit et al., 2013), including the frequently observed, nearby and bright systems UV Ceti, Proxima Centauri and AD Leonis (Lacy, Moffett, and Evans, 1976; Zic et al., 2019; Bastian, 1990; Slee, Willes, and Robinson, 2003; Villadsen and Hallinan, 2019).

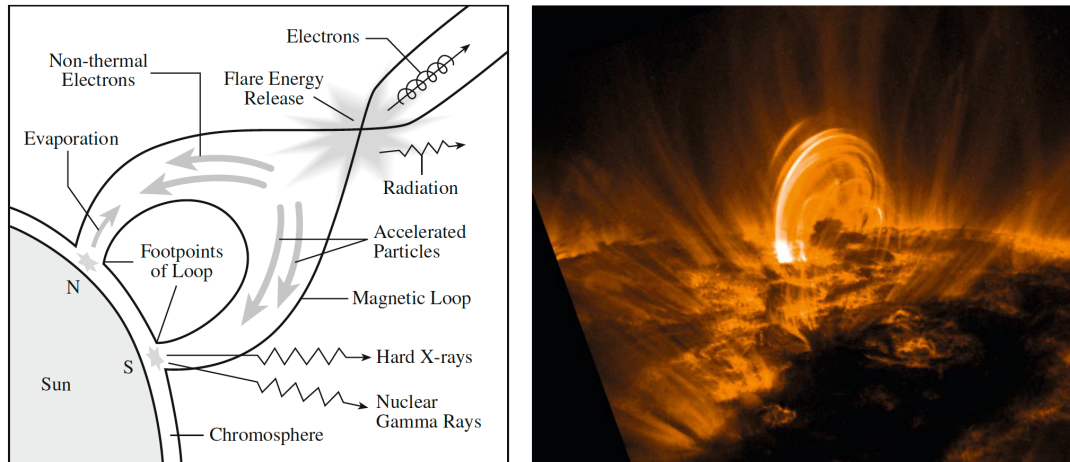


Figure 1.2: **Left:** A schematic of the processes involved in a solar flare, where magnetic reconnection in the corona is the energy source from which all subsequent processes arise. Electrons are accelerated, generating the radio burst, some of which are channelled down into the chromosphere, causing brightening in X-rays, as well as the enhancement of emission lines such as $H\alpha$. **Right:** A 17.1 nm image from the *Transition Region And Coronal Explorer* satellite taken 2 hours after a flare. Figures both from Lang (2009).

As with the Sun, the radio bursts seen on M dwarfs are indicative of their space weather environments, relevant to the habitability of any hosted exoplanets (Zic et al., 2020). In addition to the incoherent emission described above, some stars have been observed to emit coherent, circularly polarised flares (Osten and Bastian, 2006; Pineda and Villadsen, 2023). These coherent flares can be seen on the same stars as the incoherent flares and vary on timescales from milliseconds to hours, in contrast to the incoherent emission occurring over tens of minutes to even days (Osten, 2007).

In addition to these low mass single stars, binary systems can produce transient radio emission. RS CVn binaries, which consist of a main sequence star and a chromospherically active giant or subgiant (Hall, 1976), are known sources of radio emission (García-Sánchez, Paredes, and Ribó, 2003). The subgiant components of these systems have large outer convective envelopes and rotate quasisynchronously with the orbital motion with a period of 1 – 15 days, generating magnetic fields over 2 orders of magnitude larger than that of the Sun (Lanza, 2005). Similarly, Algol binaries, where one evolved star fills its Roche lobe (Budding et al., 2004), are known radio emitters, displaying flares with similar timescales to RS CVn flares (Richards et al., 2003).

1.2.1.2 Jets and outflows on all scales

Of the incoherent emitters, many are caused by astrophysical jets. These jetted systems are observed across a wide range of energy and mass scales, both in our Milky Way and at extragalactic distances (Beall, 2016). These jets are highly collimated outflows, typically emitting some form of synchrotron radiation (see section 1.3.1.1). Historically, many of these sources were first observed at higher energies (e.g. X-rays and gamma-rays), before the multiwavelength association was made with high angular resolution radio observations. This has resulted in the perhaps unfortunate phenomenological nomenclature of referring to objects that emit radiation at all wavelengths as e.g. X-ray binaries or gamma-ray bursts.

In our Galaxy we see radio jets emitted from binary systems that consist of a compact object and a star. Systems whose compact object is a white dwarf are known as cataclysmic variables (CVs) due to their eruptive nature - regular outbursts have been seen on CVs dating back to the beginning of modern astronomy (Pickering and Fleming, 1896). In these systems the white dwarf accretes matter from the companion via Roche lobe overflow, where the accumulation of material on the white dwarf can ignite a thermonuclear runaway explosion, known as a classical nova (Starrfield, 1989). The ionised gas ejected by this explosion can produce radio emission for months to years following the initial nova (Hjellming et al., 1979). Recurrent novae are similar, though smaller, outbursts, occurring every few decades, with detectable radio emission lasting for months to years (de Ruiter et al., 2023; Nyamai et al., 2023). These nearby systems that enter outbursts regularly provide an excellent local laboratory to test outstanding questions in jet physics (Coppejans and Knigge, 2020).

If the compact object in a binary is a black hole or a neutron star then the system is referred to as an X-ray binary, in reference to the X-ray emission observed as matter from a companion is accreted onto the compact object (Fender, 2006). Again these compact object systems launch jets, with a clear association seen between changes in accretion state (probed by X-ray emission) and (radio) jet properties (Fender, Belloni, and Gallo, 2004), including the launch of relativistic




Figure 1.3: An artist's impression of a black hole X-ray binary, displaying the accretion from a companion star onto the black hole and the emission of radio jets. *Image Credit:* John A. Paice, <https://ras.ac.uk/media/344>

blobs of plasma that have resolved proper motions (Mirabel and Rodríguez, 1994; Carotenuto et al., 2021). An illustration of a black hole XRB is shown in Figure 1.3, with similar geometries believed to hold for other compact object binaries, with the accretion disk perpendicular to the jet outflow.

Supernovae are the explosive deaths of massive stars and are categorised broadly into type I and type II, based on the detection or non-detection of hydrogen emission lines (Turatto, 2003). Supernovae are seen to emit radio emission as fast-moving ejecta collides with circumstellar media (CSM). Radio observations of such supernovae both probe the density of the CSM and trace the fastest moving ejecta in the explosion (Chevalier, 1982; Chevalier, 1998; Weiler et al., 2002; Chevalier and Fransson, 2006). Similarly, the deaths of some massive stars can cause bursts of gamma-rays (gamma-ray bursts; GRBs) to be emitted from relativistic ejecta (Granot and van der Horst, 2014). As the jet from a GRB shocks its CSM it is possible to observe an afterglow of emission from X-rays to radio waves. There are also GRBs seen that are thought to originate from the merger of two binary neutron stars, which also produce radio-loud afterglows (e.g. Schroeder

et al., 2023). Radio emission is not extinguished by dust in the same way that optical light is, nor affected as much by the relativistic beaming of a GRB’s jet as the eponymous gamma rays are, meaning that widefield radio surveys provide a way to test the rates of GRBs and supernovae via orphan afterglow candidates (Levinson et al., 2002; Gal-Yam et al., 2006).

A tidal disruption event (TDE) occurs when a star passes close to a supermassive black hole and is pulled apart by tidal forces (Rees, 1988). The resulting stellar debris forms an accretion disc and again radio emission is seen from outflows and jets. As discussed in the radio-focused review of Alexander et al. (2020), only a few percent of TDEs are radio loud, but these exhibit emission detectable for years post-disruption. The observed outflows from radio-loud TDEs can be sub-relativistic (e.g. Bright et al., 2018) or relativistic (e.g. Rhodes et al., 2023b). Both observational strategy and relativistic beaming induce selection biases in our understanding of TDEs, in addition to the rates of sub-relativistic TDEs being poorly constrained (Alexander et al., 2020). Therefore, large area radio surveys might play an important role in studying the TDE population, with a recent example indicating that the radio emission from TDEs may brighten decades after any higher energy counterpart has faded (Ravi et al., 2022).

Stepping up in terms of scale one last time, many jets are seen originating from active galactic nuclei (AGN). These jets are expelled from accretion disks associated with supermassive black holes at the cores of galaxies, with masses from $10^5 - 10^9 M_{\odot}$ (Padovani et al., 2017). Many subclasses of AGN are defined by the viewing angle with respect to their jet, whilst some large and nearby AGN are some of the brightest radio sources in the low-frequency sky, sometimes referred to as the ‘A Team’ galaxies - e.g. Cyg A and Vir A (de Bruyn et al., 2009). AGN can vary on both year-long timescales (Condon et al., 1979; Taylor and Gregory, 1983) and display short outbursts over months (Hovatta, Lehto, and Tornikoski, 2008), thought to be caused by the propagation of shocks within the jets (Hughes, Aller, and Aller, 1989; Hovatta et al., 2008; Lainela, 1994). These AGN jets can be highly relativistic, where the emitting particles are moving at velocities

approaching the speed of light and so can experience a preferential beaming of their radiation. This boosts their observed flux, causing these sources to cross the incoherent-coherent boundary in Figure 1.1.

1.2.2 Coherent emitters

1.2.2.1 Pulsars and rotating transients

The most common coherent radio-variable sources in the sky are pulsars - neutron stars with masses $\sim 1.4 M_{\odot}$, magnetic fields of strength $> 10^{11}$ G and rotation periods of typically $\lesssim 1$ s (Lorimer and Kramer, 2004). Their rapid rotation and the misalignment of their rotational and magnetic axes produces a lighthouse-like effect where telescopes on Earth receive a brief pulse for every rotation of the neutron star. Studying pulsars allows for investigations of extremely dense matter (Link, Epstein, and Lattimer, 1999), tests of general relativity (Taylor and Weisberg, 1982; Colom i Bernadich et al., 2023) and studies into exotic emission mechanisms (Ruderman and Sutherland, 1975). Furthermore, the propagation of their pulsed light gives many lines of sight through the Galaxy to help us understand the interstellar medium (ISM), whilst arrays of ms-period pulsars are allowing for searches for a stochastic gravitational wave background (Agazie et al., 2023; EPTA Collaboration et al., 2023; Reardon et al., 2023; Xu et al., 2023). Some pulsars emit giant pulses that are up to 1000 times brighter than the mean pulse (Hankins et al., 2003). To date, 11 such giant pulse emitters have been detected, out of over 3000 pulsars total (Kuzmin, 2007; Manchester et al., 2005) Finally, there exist some pulsars that display sharp changes in behaviour, with in some cases the radio emission switching off for more than a month at a time (Kramer et al., 2006).

A related class of rotating neutron stars are called magnetars, whose magnetic field strengths are over 100 times stronger than that of a typical pulsar (Lyne and Graham-Smith, 2012). Of the 30 magnetars known, only 6 are radio-bright, with the majority being detected via their X-ray emission (Olausen and Kaspi, 2014)¹. There has been recent interest surrounding radio-loud magnetars due to their potential

¹<http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>, accessed 25-04-2024

links to fast radio bursts (FRBs), including the observation of a bright radio flare that is consistent with FRB observations (Bochenek et al., 2020; CHIME/FRB Collaboration et al., 2020, see below), as well as the recently observed conversion between linear and circularly polarised emission, indicative of a similar environment to that of potential FRB progenitors (Lower et al., 2024).

In the past two years there have been recent discoveries of pulsing sources with periods beyond those typical of pulsars. For example, Caleb et al. (2022) report the discovery of a pulsar with a 76 s period, which challenges both theoretical models of pulsar emission as well as traditional observational searches for neutron stars. Similarly, Hurley-Walker et al. (2022) and Hurley-Walker et al. (2023) have identified sources that emit strongly polarized coherent radio pulses on periods of tens of minutes. These systems show some similarities to magnetars and bear a resemblance to as-yet unidentified transients found in the direction of the galactic centre (Hyman et al., 2005; Hyman et al., 2009; Wang et al., 2021b; Wang et al., 2022). However, simulations and theoretical models struggle to reconcile the production of bright coherent radio emission with either a population of dipolar radio emitting neutron stars or white dwarfs (Rea et al., 2024). Further searches for these long period systems may give more clues as to their behaviour and origins.

1.2.2.2 Fast radio bursts

Fast radio bursts are ms-duration signals, first seen by Lorimer et al. (2007). The dispersive sweep seen in Figure 1.4 is caused by the dispersion of radio waves through free electrons along the line of sight to the burst's origin (see section 1.3.3.1). This large dispersion seen in FRBs is indicative of their very distant origin i.e. that they traverse through more free electrons than are in just the Milky Way. FRBs of cosmological origins with confirmed host galaxies have therefore been used to probe the baryon content of the Universe (Macquart et al., 2020). As of yet the only FRB observed with a confirmed progenitor has been the Galactic magnetar already discussed (Bochenek et al., 2020; CHIME/FRB Collaboration et al., 2020). There are many (> 50) theoretical models proposing what progenitors cause these




Figure 1.4: The first detected FRB (Lorimer et al., 2007). The dispersive sweep over 100s of ms indicates a large distance to the source and is typical of many FRBs.

phenomena, with almost all invoking the extreme densities and magnetic fields found around neutron stars (Platts et al., 2019)².

Many FRBs are detected as one-off signals, though some are seen to repeat, with both populations displaying diverse morphologies (Pleunis et al., 2021). Repeating FRBs have been particularly useful in the discussion of their progenitors, as they immediately rule out cataclysmic origins for reoccurring bursts and allow for localisation and follow-up. For example, the first repeating fast radio burst, FRB121102 (Spitler et al., 2016), was localised to a host galaxy (Chatterjee et al., 2017), allowing for very long baseline interferometry to identify a coincident persistent radio source, constrained to a size of $\lesssim 0.7$ pc (Marcote et al., 2017), which itself shows moderate variability (Rhodes et al., 2023a).

²https://frbtheorycat.org/index.php/Main_Page accessed 06-05-2024

1.2.2.3 Masers around giant stars

Stars with initial masses in the range $0.8 - 8 M_{\odot}$ lose up to 80% of their mass during the asymptotic giant branch (AGB) phase of their lives, as stellar pulsations exert radiation pressure on dust formed in their extended, cool atmospheres, generating a wind of mass loss rate $\lesssim 10^{-5} M_{\odot} \text{yr}^{-1}$ (Höfner and Olofsson, 2018). The illumination of this circumstellar gas by the inner star and collisions in this dense environment can cause population inversions in the energy levels of different molecular species (i.e. there are more molecules in a higher energy state than a lower energy state), causing masers to occur (Burke and Graham-Smith, 2014). Masers - microwave amplification by the stimulated emission of radiation - are seen in different species moving radially out from an AGB star: at ~ 4 AU ($2 - 3$ stellar radii) SiO masers are seen at 42.8 and 43.1 GHz, whilst water masers form at 15 AU and are seen at 22 GHz, with OH masers occurring at 1612, 1665 and 1667 MHz at a distance of 40 AU (Reid, 2007). SiO and water masers form at different distances due to requiring a high enough temperature and density but low enough opacity, whilst OH masers form when ambient ultraviolet radiation dissociates water molecules (Herman and Habing, 1985). Masers are expected to be variable, for example varying with an offset phase with respect to the pulsations of the AGB star over years (Harvey et al., 1974), as well as displaying flares that last hundreds of days, interpreted as the propagation of a shock wave driven out by stellar pulsations (Esipov et al., 1999). Long term monitoring of some stellar masers has revealed a range of variability timescales, from decade-long slow changes to the monthly lifetimes of individual cloud conditions, with at least some variations being due to varying excitation at particular locations (Winnberg et al., 2008; Winnberg, Brand, and Engels, 2024).

1.3 Emission mechanisms and propagation

This section provides an overview of some of the important radio emission mechanisms and propagation effects seen in this thesis. Many of these results come

Table 1.1: Summary of how the emission mechanisms discussed relate to the observations of different variable and transient astrophysical sources as discussed in section 1.2.

Mechanism	Sources	Comment
Cyclotron	The Sun Flaring dwarfs Active binaries	$\propto \nu^2$ at low ν , fall-off at high ν
Synchrotron	CVs, XRBs, SNe TDEs, AGN	$\propto \nu^{2.5}$ at low ν , $\propto \nu^{\frac{1-p}{2}}$ at high ν Can be used to find p
Blackbody	?	Rarely seen at such low frequencies
Free-free	Planets, CVs	Low velocity outflows
Plasma	Active stars	Probes n_e , equation 1.13
ECMI	Active stars Jupiter-Io Star-Planet?	Probes B , equation 1.7 In aurorae, at 10s of MHz
Pulsar-like	Pulsars Magnetars ULPs and FRBs?	Steep spectra, $\alpha \lesssim -2$ Unknown progenitors

from the essential books of Rybicki and Lightman (1986) and Longair (2011), whose derivations I follow here.

I start by considering the radiative transfer of emitted radio waves. In general the radiative transfer equation for a light ray of specific intensity I_ν travelling path s is given by

$$\frac{dI_\nu}{ds} = -\kappa_\nu I_\nu + j_\nu \quad (1.4)$$

where j_ν and κ_ν are the emission and absorption coefficients respectively. The solution to the radiative transfer equation, when no emission is present before entering a cloud of emitting plasma, is given by

$$I_\nu(\tau_\nu) = S_\nu(1 - e^{-\tau_\nu}) \quad (1.5)$$

where constant source function S_ν is the ratio of emission to absorption coefficients (j_ν/κ_ν) and τ is the optical depth ($d\tau_\nu = \kappa_\nu ds$). It is common to consider the optically thin ($\tau \ll 1$) and optically thick ($\tau \gg 1$) regimes of radiative transfer.

A summary of the emission mechanisms I will describe and how they relate to the observed sources discussed in section 1.2 is presented in Table 1.1.

1.3.1 Incoherent radio emission

The observed emission mechanisms can provide important constraints on the physical conditions of the source. It is therefore often useful to compare the observed spectrum of a source to the known spectra of different emission mechanisms. This is often done by calculating the spectral index α for $I_\nu \propto \nu^\alpha$. Below I outline the spectral shape of emission that arises from electrons spiralling in magnetic fields, in free-free interactions with ions and when emitted as a black body.

1.3.1.1 Gyroemission

The motion of an electron of mass m_e and charge e in a uniform magnetic field \mathbf{B} is described by the equation

$$\frac{dv_\perp}{dt} = \frac{e}{m_e} v_\perp \times \mathbf{B} \quad (1.6)$$

where v_\perp is the component of velocity perpendicular to the B-field. The solution to this equation is uniform circular motion in the plane perpendicular to the magnetic field, with rotational frequency

$$\nu_c = \frac{eB}{2\pi m_e}. \quad (1.7)$$

In the non-relativistic limit collisions with ambient plasma are important and the velocities of electrons are well described by a Maxwellian distribution (Dulk, 1985). Therefore the emission seen is concentrated at the local electron cyclotron frequency ν_c and its low harmonics ($n < 10$). This is referred to as cyclotron emission.

At the other extreme, when considering highly relativistic electrons, the radiation will be highly beamed, such that light seen by an observer will be confined to a small angle ($\sim 1/\gamma$) over a time period smaller than the gyration period, where $\gamma = 1/\sqrt{1-\beta^2}$ and $\beta = v_\perp/c$. This results in an increase in gyrofrequency to $\nu_g \approx \gamma^2 \nu_c$. This is referred to as synchrotron emission.

When observing a synchrotron-emitting source, there is of course not a single electron but an entire population. Often electrons are accelerated to relativistic

velocities via Fermi acceleration (Fermi, 1949; Bell, 1978), with the population of electrons $N(E)$ at energies E accelerated into a power law distribution

$$N(E)dE = N_0 E^{-p} dE. \quad (1.8)$$

with spectral index p . Following Pacholczyk (1970), the solution to the transfer equation for this power law of electrons can be written as

$$I_\nu = S(\nu_1) z^{5/2} (1 - \exp(-z^{-(p+4)/2})) \quad (1.9)$$

where $S(\nu_1)$ depends on the magnetic field strength and pitch angle of emitting electrons, $z \equiv \nu/\nu_1$ and ν_1 is defined by the condition that $\tau(\nu_1) = 1$. We can immediately see the similarities between this solution and equation 1.5. Therefore in the optically thin regime, when $\tau \ll 1$ we see

$$I_\nu = S_\nu \tau_\nu \propto z^{5/2} \cdot z^{-(p+4)/2} \propto \nu^{\frac{1-p}{2}}. \quad (1.10)$$

This is one of the core results of synchrotron radiation - that at high frequencies it is optically thin and the slope of its spectrum is a direct measure of the power law index of the accelerated electrons. By contrast, when propagating through an optically thick region $\tau \gg 1$, therefore

$$I_\nu \approx S_\nu \propto \nu^{5/2}. \quad (1.11)$$

This is referred to as synchrotron self absorption, where at low frequencies an observer ‘sees’ only the surface of the emitting region and the slope of the spectrum has no dependence on p .

In between these two limits (cyclotron and synchrotron) falls gyrosynchrotron emission, where radiation is generated by electrons at the high-energy tail of a thermal velocity distribution, or with a power law distribution in energy. Emission at harmonics $10 < n < 100$ are important for gyrosynchrotron emission. However, analytic computation of emission and absorption is almost intractable, with some numerical approximations presented by Dulk and Marsh (1982) and Robinson and Melrose (1984). These approximations have shown that the absorption and

emission coefficients depend on global plasma properties, such as the electron density, the effective temperature and magnetic field strength. This emission is therefore relevant for determining the plasma properties in M dwarfs and RS CVn binaries as demonstrated by Mutel et al. (1987), Osten et al. (2006), and Zic et al. (2019).

1.3.1.2 Blackbody and free-free radiation

There are at least two important thermal processes to consider, beginning with black body radiation. I have already shown in section 1.2 how a black body at low frequencies has a spectrum $\propto \nu^2$. At high frequencies, the Wien approximation holds and Planck's law becomes

$$I_\nu = B_\nu(T) = \frac{2h\nu^3/c^2}{\exp(h\nu/k_B T) - 1} \approx \frac{2h\nu^3}{c^2} e^{-h\nu/k_B T} \quad (1.12)$$

where h is the Planck constant. So at high frequencies, blackbody radiation falls off sharply, which was historically important for avoiding the ultraviolet catastrophe (Rybicki and Lightman, 1986).

Bremsstrahlung, also known as free-free radiation, refers to the emission of photons due to the deceleration of a free electron in the electric field of a heavier ion i.e. it is a thermal emission mechanism occurring in charged plasmas. For example, free-free radiation is commonly seen from HII regions. A full derivation of bremsstrahlung requires a quantum treatment, with consideration of photon energies comparable to the energies of emitting particles (Rybicki and Lightman, 1986). At high frequencies, the bremsstrahlung spectrum is approximately flat (i.e. $\alpha \approx 0$), with a turnover at lower frequencies, becoming optically thick and emitting as a black body i.e. $\alpha = 2$. A comparison of the spectra discussed in this section can be seen in Figure 1.5.

1.3.2 Coherent emission mechanisms

I next summarise some of the key ideas around coherent emission in the cases of stellar and neutron star magnetospheres. In general, these mechanisms have high (non-physical) brightness temperatures, are highly polarised and exhibit complex




Figure 1.5: The characteristic intensity spectra for the emission mechanisms discussed here, as shown by Dulk (1985), where δ is what I have called p . In all cases we can see increasing functions at low frequencies before reaching a peak after which the spectra flatten or decrease.

structure. The majority of these results come from Dulk (1985), Melrose (2009), and Melrose (2017).

1.3.2.1 Plasma emission

Plasma emission is thought to be responsible for the majority of low frequency radio bursts from the Sun, generated by electrons accelerated in shocks or high velocity streams. It has also been invoked to explain some coherent stellar radio bursts (Osten and Bastian, 2006). Longitudinal plasma oscillations, also known as Langmuir waves, generated by high velocity streams of electrons (also known as electron beams) are the cause of plasma emission. As the high velocity beam streams through the plasma, the ambient plasma develops a bump in the high

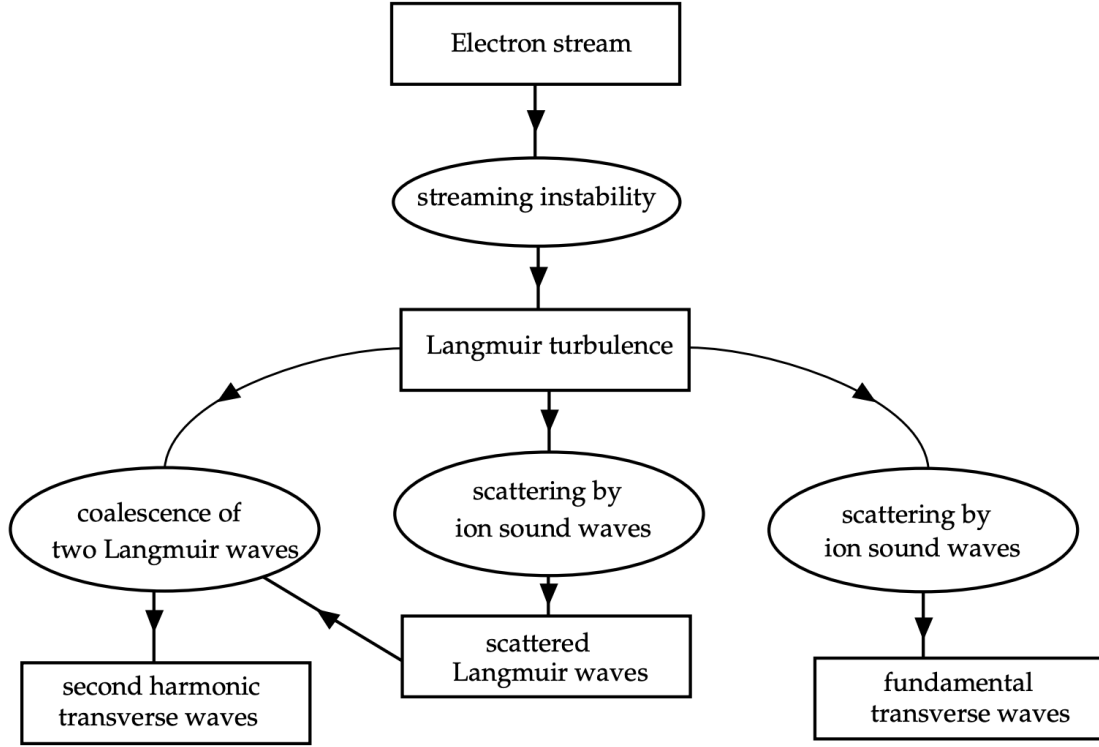


Figure 1.6: Flowchart from Melrose (2017), demonstrating how plasma emission can be generated by plasma oscillations, as originally discussed by Ginzburg and Zhelezniakov (1958).

velocity tail of the parallel component of its velocity distribution. This bump leads to instabilities, causing exponential growth of Langmuir waves, which extract free energy from the velocity distribution. Subsequent interactions, outlined by Figure 1.6 then convert plasma oscillations into radiation at plasma frequency

$$\nu_p = \left(\frac{e^2 n_e}{\pi m_e} \right)^{1/2} \quad (1.13)$$

and its second harmonic (Melrose, 2017), where n_e is the number density of electrons. Observationally, plasma emission is seen as high brightness temperature emission that is highly polarised and provides a probe of the plasma density.

1.3.2.2 Electron cyclotron maser instability

The electron cyclotron maser instability (ECMI) is another mechanism for generating emission in the magnetospheres of stars and planets. This mechanism is seen in the magnetospheres and aurorae of planetary bodies such as the Jupiter-Io

system (Zarka, 1998), whilst interactions between host stars and their planets are thought to produce similar emission (Zarka, 2007; Vedantham et al., 2020). As with plasma emission, an anisotropy in the electron velocity distribution drives this emission, leading to the growth of radiation. In plasma emission I mentioned that the anisotropy was parallel to the magnetic field (i.e. in v_{\parallel}), however for ECMI this anisotropy is in the perpendicular component, leading to an effective negative absorption.

Much of ECMI theory is based on a loss-cone distribution, where there are very few electrons with a small v_{\perp} i.e. they have small pitch angles $\theta = \arctan(v_{\perp}/v_{\parallel})$. To begin, electrons high in magnetic loops are heated via magnetic reconnection (see Figure 1.2). The electrons stream down to the footprints of a magnetic loop and as they do so, the magnetic field strength increases. Correspondingly, v_{\perp} also increases as the electrons orbit on their gyroradius. Conservation of energy implies that v_{\parallel} must therefore decrease. If the pitch angle is larger than some critical value $\theta^* = \arcsin(B_{top}/B_{foot})$ then the parallel velocity will go to zero and reverse, resulting in a magnetic mirror (Dulk, 1985). Particles with pitch angles lower than θ^* pass through the mirror, collide with the atmosphere below the loop and are lost. The resulting velocity distribution is therefore depleted of particles with small pitch angles. This now-asymmetric distribution is unstable to the growth of electromagnetic waves i.e. there are more high energy electrons able to donate energy to a wave than there are low energy electrons that can absorb energy from the same wave, leading to rapid growth (Dulk, 1985). The wave growth rate is fastest for the cyclotron frequency that has already been introduced in equation 1.7.

1.3.2.3 Pulsar emission

The characteristic radio emission of pulsars is a result of their strong magnetic fields ($B > 10^{11}$ G), which take an approximately dipolar shape. This dipolar field co-rotates with the neutron star out to a radius defined by the distance at which the angular velocity reaches c (Goldreich and Julian, 1969). A schematic of this is shown in Figure 1.7. The exact mechanism by which radio emission




Figure 1.7: A schematic of a pulsar magnetosphere, from Burke and Graham-Smith (2014). The light cylinder, within which the magnetosphere co-rotates with the neutron star, is defined by radius r_c .

arises in pulsars is not fully understood (Philippov and Kramer, 2022), though it almost certainly relates to the acceleration of particles in the magnetosphere, with many models invoking the pair production of photons in this regime (Chen and Ruderman, 1993). Whatever the original cause of radiation, the net result is that the pulsar loses energy over time, slowing its rotation. From measuring a pulsar's period and its first derivative, it is possible to estimate both the age and the magnetic field strength of the neutron star, assuming an entirely dipolar field (Lyne and Graham-Smith, 2012). The observed spectrum of pulsars is typically steep,

with $\alpha < -3$ in extreme cases, often with a low frequency turnover (Sieber, 1973). Finally, pulsar emission is typically linearly polarised at the 20% level or more, with most pulse profiles exhibiting a small amount of circular polarisation (Lorimer and Kramer, 2004), though their polarisation properties are both frequency and phase-dependent (Oswald et al., 2023).

1.3.3 Propagation effects

Propagation effects are extrinsic mechanisms, caused by the intervening material between an observer and a source of radio waves. I briefly outline two of the most important propagation effects that are relevant to this thesis.

1.3.3.1 Dispersion

Free electrons in the interstellar and intergalactic medium result in a frequency-dependent refractive index. In turn, the light travel speed through the ISM depends on this index. As a result, radio waves can be dispersed, so that lower frequency signals arrive later than those at higher frequencies. This dispersion can be described by a dispersion measure (DM). As given by Lorimer and Kramer (2004), the time delay Δt induced by dispersion can be written

$$\Delta t \propto \text{DM} \nu^{-2}. \quad (1.14)$$

The value of the DM, for a source at distance d , is given by the integrated column density of electrons along the line of sight l towards a source:

$$\text{DM} = \int_0^d n_e(l) dl. \quad (1.15)$$

Therefore, if dispersion is observed for a source, it can be used to measure the intervening electron distribution if the distance is known (as in Macquart et al., 2020). Similarly, if the underlying electron distribution is known, or at least can be modelled, then dispersion provides a rough distance estimate to a source. This is particularly relevant for pulsars in the Milky Way, where two models of Galactic free electrons are used to determine the distances to these sources referred to as

NE2001 and YMW16 (Cordes and Lazio, 2002; Yao, Manchester, and Wang, 2017). The dispersion measure therefore provides a way to distinguish if a signal originates from within the Milky Way or has travelled across extragalactic distances, giving immediate constraints on the potential source.

1.3.3.2 Scintillation

Unresolved radio sources are seen to change in brightness as a function of time and frequency, caused by inhomogeneities in the intervening media. Ionospheric and interplanetary scintillation were the first of these effects to be seen by radio astronomers (Hewish, Scott, and Wills, 1964; Rickett, 2001), though these have stronger effects at lower frequencies and are not relevant to this thesis. Inhomogeneities in the ISM can cause significant variations at GHz frequencies, known as refractive and diffractive interstellar scintillation (RISS and DISS), with their effects often modelled by a thin screen of plasma between the observer and the source (Rickett, 1990). This screen effectively disturbs the phase of a set of waves, producing interference and therefore an amplitude modulation. As described in Walker (1998), the amplitude m and timescale t of the modulation of radio waves from an extragalactic source depends on the observing frequency, with scalings of the form

$$t_r \propto \frac{\nu_0^{11/5}}{\nu}, \quad t_d \propto \frac{\nu_0^{6/5}}{\nu} \quad (1.16)$$

and

$$m_r \propto \frac{\nu_0^{17/30}}{\nu}, \quad m_d = 1 \quad (1.17)$$

where subscripts r and d refer to refractive and diffractive regimes respectively and ν_0 is the transition frequency between strong and weak scattering. A more quantitative description of the predictions of RISS-induced variability is discussed in more detail in section 4.4.3 when used in relation to sources of interest. In summary, DISS occurs on rapid timescales (seconds to minutes), at high amplitude and over small fractional bandwidths. This becomes relevant for very compact sources such as pulsars, but does not affect sources of larger angular sizes or those that have multiple emission regions (Rickett, 2007). RISS occurs at lower amplitudes and

over longer timescales and is caused by inhomogeneities in the ISM on large scales. RISS variability is more pronounced at GHz frequencies, with its peak depending on the line of sight through the Milky Way (Walker, 1998). Therefore RISS can be the cause of variability on timescales greater than hours for unresolved sources (e.g. pulsars and AGN). When observed, this scintillation can be used as a probe of the structure of the Milky Way's ISM, as demonstrated by Wang et al. (2021a), who identify a nearby (4 pc) plasma filament by observing the hour-timescale variability of five sources oriented in a line on the sky spanning almost 2 degrees.

1.4 The search for radio transients

Despite the broad range of interesting systems already described, radio telescopes have until recently been primarily of use for variable and transient searches as follow-up instruments, typically informed by detections at higher energies. Nevertheless, many searches have been conducted at radio frequencies from tens of MHz to a few GHz. I outline here some of these searches, their techniques and their findings before describing some of the current transient hunting efforts.

A number of transients have been discovered around 1 GHz with the Karl G. Jansky Very Large Array (VLA), though many have yet to be associated with multiwavelength counterparts or definite progenitor systems. Some of these searches made use of frequent observations of a calibration field or those taken as part of an extragalactic deep field. For example, Bower et al. (2007) made use of weekly images of a calibrator field for over 22 years, manually inspecting the light curves of sources found in each epoch to search for transients. 8 transients with no multiwavelength counterparts were found that were each present in only one epoch, along with two that were consistent with radio supernovae and GRB afterglows. Similarly, Mooley et al. (2013) and Mooley et al. (2016) made use of VLA observations of the Chandra Deep Field South and the Sloan Digital Sky Survey to search for transients on month-timescales with good multiwavelength coverage. By computing two variability statistics they found two stellar flaring sources (one dKe star and one RS CVn system) and concluded that scintillation explains the variability seen

in $\sim 1\%$ of all detected radio sources. Levinson et al. (2002) searched for radio transients in two large VLA sky surveys - Faint Images of the Radio Sky at Twenty Centimeters (FIRST; Becker, White, and Helfand, 1995) and the National Radio Astronomy Observatory (NRAO) VLA Sky Survey (NVSS; Condon et al., 1998). They compared sources between these surveys and searched for radio afterglows of GRBs, finding 9 candidate transients, of which only one was confirmed by Gal-Yam et al. (2006) and Ofek et al. (2010) in dedicated follow-up observations. These follow-up studies concluded that two of the observed candidates were variable (an AGN and a pulsar), while the others were all caused by side lobes, the mis-match in resolution between the two surveys and a processing bug in the FIRST survey. Ofek et al. (2010) also provided a re-analysis of VLA searches, concluding that the rotational energies provided by old, isolated neutron stars, in addition to accretion from the ISM, provide enough energy on sufficient timescales to explain the one-off transients found in surveys such as Bower et al. (2007). Ofek et al. (2011) also conducted a transient search in the Galactic plane with the VLA, targeting 141 fields with 12-16 50s observations per field. They specifically chose these observations to be far from bright sources and supernova remnants in NVSS and conclude that on long timescales much of the variability seen is caused by scintillation, finding a single transient and concluding that the areal density of such transients is consistent with a neutron star origin. More recently, Sarbadhicary et al. (2021) utilised the VLA's COSMOS Hi Legacy Survey to create a Variable and Explosive Radio Dynamic Evolution Survey (CHILES VERDES), which made use of the multiwavelength data of the COSMOS field and 5.5 years worth of data reaching RMS sensitivities of $\sim 10 \mu\text{Jy beam}^{-1}$ per epoch over moderate FoV of radius $22.5'$. They found 58 AGN-type sources, based on the multiwavelength COSMOS data from Smolčić et al. (2017), that were primarily varying due to scintillation.

Similar transient searches have been conducted at around 1 GHz with other facilities. The Allen Telescope Array (ATA; Welch et al., 2009) conducted a 12 epoch survey at 1.4 GHz, comparing source catalogues with NVSS, manually inspecting candidates and finding no transients (Croft et al., 2010). Croft et al.

(2011) also conducted a search for transients occurring within epochs from the same ATA survey, finding no transients on minute-timescales above a limit of 40 mJy. Bannister et al. (2011) searched through archival 843 MHz Molonglo Observatory Synthesis Telescope data, taken at 843 MHz in order to find transients. By characterising light curves with two variability statistics, two low mass XRBs and SN 1987A were recovered, along with highly variable, scintillating AGN and a sample of 13 transients with unknown classifications.

At lower frequencies (~ 300 MHz), Hyman et al. (2005) used minute-timescale images of the Galactic Centre region with the VLA to uncover a single radio transient with a brightness temperature of $\sim 10^{16}$ K, concluding that the emission from the magnetosphere of a magnetar could provide enough energy to power this source. Similarly, Jaeger et al. (2012) and Polisensky et al. (2016) made use of visits to extragalactic deep fields observed with the VLA, sampling timescales from days to months, finding a single radio transient with no counterpart.

At lower frequencies still, Stewart et al. (2016) made use of 400 hours of observations of the celestial North pole taken at 60 MHz with the LOw Frequency ARray (LOFAR; Van Haarlem et al., 2013) to search for transients. Their 11 minute snapshots identify a single transient with no counterpart at any other wavelength. Murphy et al. (2017) combined survey catalogs from the Murchison Widefield Array (MWA; Tingay et al., 2012) and the Giant Metre-Wave Radio Telescope (Swarup et al., 1991) to search for low frequency transients and found one transient with no multiwavelength counterpart. Bell et al. (2014) similarly make use of the MWA to search at minute- and year-timescales, finding likely scintillating variables and no transients. Finally, Varghese et al. (2019) discovered a transient at 34 MHz using over 10,000 hours of observations taken by the Long Wavelength Array with a DM-inferred redshift of $z \approx 0.58$ indicating a cosmological origin.

In summary, these searches have found a sample of radio transients, a few of which have clear progenitors and counterparts, but many with no associations at other wavelengths. These searches have also shown how the variable population is dominated by AGN, primarily through scintillation and some possible intrinsic

variation, with $\sim 1 - 5\%$ of the radio sky appearing variable. The majority of these studies conducted their searches by manually inspecting light curves and images, with some calculating variability statistics on the former. From this literature, one of the common conclusions arrived at is that, in order to systematically find radio transients, high sensitivity, regular cadence and large FoVs are required, in addition to good multiwavelength coverage. The advent of the current generation of wide field radio interferometers, including SKA pathfinder instruments like MeerKAT (Jonas and MeerKAT Team, 2016) and the Australian SKA Pathfinder (ASKAP; Johnston et al., 2007) are now allowing the astronomical community to sample the radio sky regularly, with improved sensitivity and, crucially, larger fields of view (FoV).

Ongoing image plane searches for radio transients are being conducted with ASKAP and its Variables and Slow Transients survey (VAST; Murphy et al., 2013; Murphy et al., 2021). In general, the ASKAP software pipeline is set up to produce reduced images and catalogues from the appropriate data archive, from which VAST can run a purpose-built transient finding pipeline (Pintaldi et al., 2022), including commensally on other surveys (e.g. the Rapid ASKAP Continuum Survey; McConnell et al., 2020). This Python3 based pipeline takes the source catalogues produced as a standard data product from the ASKAP processing pipeline and forms light curves for detected sources, on which variability parameters are calculated. The logic behind the source association and the variability parameters used come from the LOFAR Transients Pipeline (Swinbank et al., 2015), which is used extensively in this thesis (see section 2.2.1), however the Vast Pipeline implements these procedures using Python routines rather than in SQL and provides an extensive web app and GUI (Pintaldi et al., 2022). The VAST survey and its pilot have already uncovered a wide range of interesting phenomena, including radio eclipses from a candidate pulsar binary (Zic et al., 2024), a sample of highly variable, highly polarised sources towards the Galactic centre (Wang et al., 2022), radio counterparts to classical novae (Gulati et al., 2023) and minute-timescale variables caused by both plasma filaments and stellar processes (Wang et al., 2021a; Wang et al., 2023). The use of circular polarisation images in many of these studies is a good way to eliminate many false

positives, leaving only the relevant transients, whilst the fast timescale (e.g. minute-long) imaging used is exploring a previously unstudied area of parameter space.

Similarly, large field surveys such as the VLA Sky Survey (VLASS; Lacy et al., 2019), provide coverage of the entire observable sky, giving important constraints and serendipitous detections of transients. For example, VLASS data has been used to identify luminous late-time emission from supernovae (Stroh et al., 2021), detections of two unusual cataclysmic variables (Ridder et al., 2023) and a transient source consistent with a merger-triggered supernova with no known optical detection (Dong et al., 2021).

The remainder of this thesis describes how I have made use of observations with the MeerKAT telescope to contribute to the search for radio transients. In Chapter 2 I give the relevant details of MeerKAT and the data products used in this thesis, primarily from the ThunderKAT survey (Fender et al., 2016). In Chapter 3 I use the discovery of the first flaring M dwarf with MeerKAT as a demonstration of how one might search for transients in these data. Chapter 4 then outlines the development of a citizen science platform in order to scale transient searches to larger data rates. Chapter 5 describes how I have combined volunteer classifications with anomaly detection algorithms to explore the application of unsupervised machine learning to transient finding. Finally, in Chapter 6 I draw together a summary and look forward to the future of radio transient surveys.

Somewhere beneath the stars

Beyond just more of our space

Nothing is yours, nothing is mine

We are decaying over time

— Thundercat, from *Interstellar Love*

2

MeerKAT and the ThunderKATs

Contents

2.1 MeerKAT	30
2.1.1 Data reduction	32
2.2 ThunderKAT	34
2.2.1 TraP	38
2.2.2 MeerLICHT	42

This chapter presents the MeerKAT radio telescope, which has been the primary instrument of use during my studies. I then go on to describe the primary data products that have been used in my work, from the ThunderKAT large survey project, which ran on MeerKAT from 2018 to 2023.

2.1 MeerKAT

Much of the relevant information on radio astronomy observations, telescopes and imaging comes from Thompson, Moran, and Swenson (2017) whose formulae and methods are followed here.

MeerKAT - the ‘more’ Karoo Array Telescope (the Afrikaans word for more is ‘meer’, in reference to the original KAT-7 array; Carignan et al., 2013; Jonas and MeerKAT Team, 2016; Camilo et al., 2018) - is a radio interferometer in the




Figure 2.1: Logarithmic distribution of all MeerKAT antennae in the Karoo. The dense core of antennae allows for high sensitivity to extended structures, whilst the sparse outer antennae provide longer baselines for higher resolution.

Karoo region of the Northern Cape, in South Africa. It consists of 64 13.5 m dishes, spaced as seen in Figure 2.1, with baselines varying between ~ 20 m and ~ 7.7 km and is equipped with receivers at UHF- (Ultra high frequency; 544 – 1088 MHz), L- (856 – 1712 MHz) and S-bands (1750 – 3500 MHz). These frequency bands can be split into 32768 channels, though the data used in this thesis were stored with 4096 channels.

The primary beam of a MeerKAT observation is set by the individual dish size (D) and the observing frequency - or equivalently wavelength λ - with the angular extent in radians given by

$$\theta \approx \frac{1.22\lambda}{D} = \frac{0.366}{\left(\frac{\nu}{\text{GHz}}\right) D}. \quad (2.1)$$

Therefore when observing at L-band, with central frequency 1.284 GHz, MeerKAT is sensitive to a region of the sky with a diameter of ~ 1.1 degrees.

When MeerKAT observes, the fundamental measurement made is a voltage signal at each antenna. The voltages of all antennae are correlated (multiplied) every 2- or 8-seconds in a pair-wise fashion, resulting in $N(N - 1)/2 = 2016$ voltage measurements across all baselines. Each baseline is sensitive to emission on scales equal to and smaller than as given by equation 2.1, replacing dish size with the baseline length between the two antennae. For example, observing again at the L-band central frequency, using the longest baseline of ~ 7.7 km produces a maximum resolution of $\sim 7.5''$ - this sets the size of the smallest resolvable sources in the sky i.e. the PSF.

The voltage output from each baseline can be described by a set of sine and cosine terms, which are combined to form a quantity known as the complex visibility. The amplitude of this complex visibility is related to a source's flux density, whilst its phase relates to the source position projected along a baseline. The Fourier transform of the complex visibility is the underlying sky brightness distribution. The combination of all visibilities then produces many overlapping fringe patterns that form the response of the entire interferometer. Different baselines sample the visibility plane at different scales, whilst the rotation of the Earth over an observation will change the baseline vectors over time, changing the sampling of the underlying visibility function. This process is referred to as sampling the u, v plane where u and v are the coordinates in 2-D of the visibility function, with units of inverse length. Ultimately, a well-sampled visibility is then Fourier transformed to produce an image of the sky.

2.1.1 Data reduction

MeerKAT picks up not just radio waves from space, but also a range of terrestrial and orbiting radio frequency interference (RFI). RFI is accounted for by a process called *flagging*, where measurements affected are ignored by software for the remainder of the data analysis. Some flags are known *a priori* e.g. if a particular part of the observing bandwidth is used for Wi-Fi communication, which will overpower astrophysical signatures by orders of magnitudes (Hugo et al., 2022). Similarly,

unexpected RFI can be easy to spot, as it produces large amplitudes compared to the target data, which can then be excised.

Once data have been flagged, they still require calibration. Calibration primarily consists of correcting instrumental and atmospheric factors. Instrumental factors can include correcting slight changes in antenna positions and electronic delays, as well as calibrating the receiver's response across its bandpass. Atmospheric factors can include short timescale variability due to ionospheric effects as well as attenuation as a function of zenith angle.

Calibration is typically done by observing a series of bright, well understood point sources and transferring their calibration solutions onto the target data. Before MeerKAT observations are taken for a user some *a priori* calibrations are applied to the data, primarily solving for antenna-based delays by observing a suitable calibrator. Then, at the start of an observation, it is typical to first observe a flux (or bandpass) calibrator, allowing for the flux scaling and bandpass to be correctly modelled. The recommended MeerKAT calibrator is J1939-6342, though J0408-6545 is also usable. It is recommended that the flux calibrator is observed for 10 minutes every three hours, providing stability to within 3% over the observation¹. Following the flux calibrator, a secondary gain calibrator is observed. This source is used to correct for ionospheric conditions and therefore must be relatively near a target observation on the sky - the closer the calibrator and the target field, the more accurate the calibration. For MeerKAT it is recommended to observe the gain calibrator for 2 minutes every half-hour. When reducing the target data, the calibration solutions are applied to the target data.

Once the data are calibrated, they need to be imaged. Imaging is the name given to the procedure of going from a set of (calibrated) visibilities to a science-ready image. Firstly, the observed u, v sampling is mapped to a regular grid, allowing for a fast Fourier transform (FFT) of the data. During the FFT, different baselines can be weighted between their natural (physical) distribution, which creates images of higher sensitivity at the expense of a lower resolution, to uniform, where the longer

¹<https://skaafrica.atlassian.net/wiki/spaces/ESDKB/pages/292257835/Calibration+strategies>, accessed 02-05-2024

baselines are upweighted, producing higher resolution images with a sensitivity penalty. From here, the gridded u, v plane is iteratively deconvolved through a process known as *cleaning* (Högbom, 1974). The basic process of cleaning involves making an initial ‘dirty’ image by taking an FFT of the gridded u, v plane, and then finding the brightest point in the image. A point source is then subtracted from the image at this location and stored in a separate sky model. This is repeated, removing and storing the positions and fluxes of point sources in a sky model, until some stopping criterion is met based on e.g. that the residual flux density level in the image is below some threshold. Then, a clean image is made by convolving the sky model with an elliptical Gaussian fit to the main lobe of the PSF (the clean beam) and adding these points back into the image.

2.2 ThunderKAT

ThunderKAT² (Fender et al., 2016) was the MeerKAT Large Survey Project (LSP) dedicated to image plane radio transients, running for 5 years from September 2018 to 2023. This project used MeerKAT’s 64 dishes and FoV > 1 square degree to take targeted, weekly observations of reported transients such as XRBs, CVs and GRBs, whilst also operating commensal searches of internal and cross-LSP data. ThunderKAT scheduled observations approximately every weekend over the 5 year lifetime of the project, with each observation of a source being 15 minutes long. Regular observations with a wide field of view naturally make these observations well suited for uncovering any transients caught in the field of view.

I have made use of 11 ThunderKAT XRB datasets in order to perform commensal searches for transients and variables. In Chapter 3 I made use of 11 observations of the field around H1743–322, as these were available for early testing. Following this, I used 11 datasets (including that of H1743–322) for the citizen science work described in Chapter 4, all taken between mid-2018 and late 2021. This selection of datasets was based on those available at the time of the public launch of the Zooniverse project and contained enough epochs to form light curves with greater

²The HUNt for Dynamic and Explosive Radio transients with MeerKAT

than 3 points. Chapter 5 makes use of the resulting citizen science classifications and so used the same 11 datasets.

A table of the 11 datasets used in this study, with the number of observations in each, is given in Table 2.1, with a map of these sources seen in Figure 2.2. The varying number of epochs between datasets is a direct result of the radio activity of the central XRB i.e. if a source is seen to fade below detection it is removed from the weekly scheduling block. One exception to this is the GX339–4 field, which was observed every week since ThunderKAT observations began, in contrast to the one or two outburst cycles followed for all other datasets. The central root mean square flux density (RMS) varies due to baseline coverage and the presence of diffuse structures, where the values presented in Table 2.1 are the median values of the RMS calculated across all images of a given dataset, evaluated in the central 8th of an image. The number of sources and this RMS are calculated by the Transients Pipeline used throughout this work and described in more detail in section 2.2.1. It is important to remember that the commensal nature of this work constrains us to whatever observational cadence was used for monitoring the XRB.

Every dataset used here was reduced by the primary users as listed in Table 2.1, following the MeerKAT recommendations discussed earlier. These data were typically reduced using OXKAT (Heywood, 2020), a semi-automatic set of scripts that perform flagging, calibration and imaging of MeerKAT data. OXKAT makes use of several existing radio astronomy packages including CASA (McMullin et al., 2007) for tasks such as gain and bandpass calibration, self calibration and flagging, CUBICAL (Kenyon et al., 2018) for further self calibration procedures, TRICOLOUR for further flagging (Hugo et al., 2022) and WSCLEAN (Offringa et al., 2014) for imaging. These steps are broken into 1st generation calibration (1GC; direction independent effects), flagging and self-calibration (2GC), with optional 3GC steps to account for direction dependent effects. In practice none of the observations here required 3GC steps. Some of the earlier observations were reduced prior to the release of OXKAT, however these still follow the same basic reduction of flagging using AOFLAGGER (Offringa, 2010), bandpass, phase calibration and flux scaling in




Figure 2.2: Map of the XRB fields (blue circles) used in this work, with names given. This is overlotted on a map of the sky taken at 408 MHz, from Remazeilles et al. (2015) and originally presented by Haslam et al. (1982).

Table 2.1: Properties of the 11 ThunderKAT datasets used in this work. Each field’s approximate Galactic latitude is given for relevance in section 4.4.3 and Figure 4.8. The number of sources and central RMS values are calculated by the TRAP (see sections 2.2.1 and 4.2.1).

Dataset/Central XRB (Number of Epochs)	Galactic Latitude(°)	Duration	Number of TRAP Sources	Average Central RMS (μ Jy)	XRB Paper
GX339-4 (167)	-04.33	2018-04-14 – 2021-10-31	714	35	Tremou et al. (2020); Tremou et al. in prep
MAXI J1820+070 (77)	+10.16	2018-09-28 – 2020-11-22	1838	26	Bright et al. (2020)
GRS 1915+105 (60)	-00.22	2018-12-08 – 2021-04-10	510	136	Motta et al. (2021)
MAXI J1348-630 (51)	-01.10	2019-01-29 – 2020-09-26	533	45	Carotenuto et al. (2021)
MAXI J1848-015 (35)	-00.10	2021-02-28 – 2021-11-19	271	290	Tremou et al. (2021); Bahramian et al. (2023)
MAXI J1803-298 (28)	-03.84	2021-05-04 – 2021-11-19	1093	22	Espinasse et al. (2021)
EXO1846-031 (26)	-00.92	2019-08-04 – 2020-04-10	366	88	Williams et al. (2022)
Swift J1858.6-0814 (25)	-05.32	2018-11-10 – 2020-03-02	1512	22	Rhodes et al. (2022)
4U1543-47 (21)	+05.42	2021-06-19 – 2021-11-14	904	27	Zhang et al. (2021)
H1743-322 (11)	-01.83	2018-09-05 – 2018-11-10	379	52	Zhang et al. in prep Williams et al. (2020)
SAX J1808.4-3658 (6)	-08.15	2019-07-31 – 2019-08-31	754	25	Gasealahwe et al. (2022)

CASA, and imaging with WSCLEAN, DDFACET (Tasse et al., 2018) or CASA. The commensal nature of this work means that, due to different science requirements and observational conditions for each dataset, the resultant images are heterogeneous in their properties, although mostly homogeneous within a particular field. For specific details on how each set of XRB data is reduced I refer the reader to the relevant paper as listed in Table 2.1. Some of these differences could of course have been mitigated if I were to re-reduce each observation in a homogeneous fashion, however the compute time for the 507 epochs used here is prohibitively high and beyond the scope of a single project. Furthermore, in keeping with the commensal attitude taken, I am interested in finding out what science can be done with in-hand observations, with minimal extra effort required in processing the observations. Based on the results of this work and other commensal findings, there are plans to re-reduce a large part of the ThunderKAT data in a near-homogeneous fashion.

2.2.1 TraP

The LOFAR Transients Pipeline (TraP; Swinbank et al., 2015) aims to detect serendipitous varying and transient sources in the large field of view observations. This pipeline is used throughout the remainder of this thesis and so I detail its working below, before its use in Chapters 3 and 4. The TraP searches for transients and variables in a given dataset by finding sources above a user-defined noise threshold, which then are associated between prior and subsequent images. An overview of the workflow of this pipeline is given in Figure 2.3. All user inputs are defined by the `job_params.cfg` input file - an example of which is given in Listing 2.1 - before the pipeline is run. Images or cubes (multi-frequency images) are fed into the pipeline and it adopts the following procedures, split between Python routines and SQL database operations:

1. The data, in FITS format (Wells, Greisen, and Harten, 1981), are accessed by the software and stored in memory.

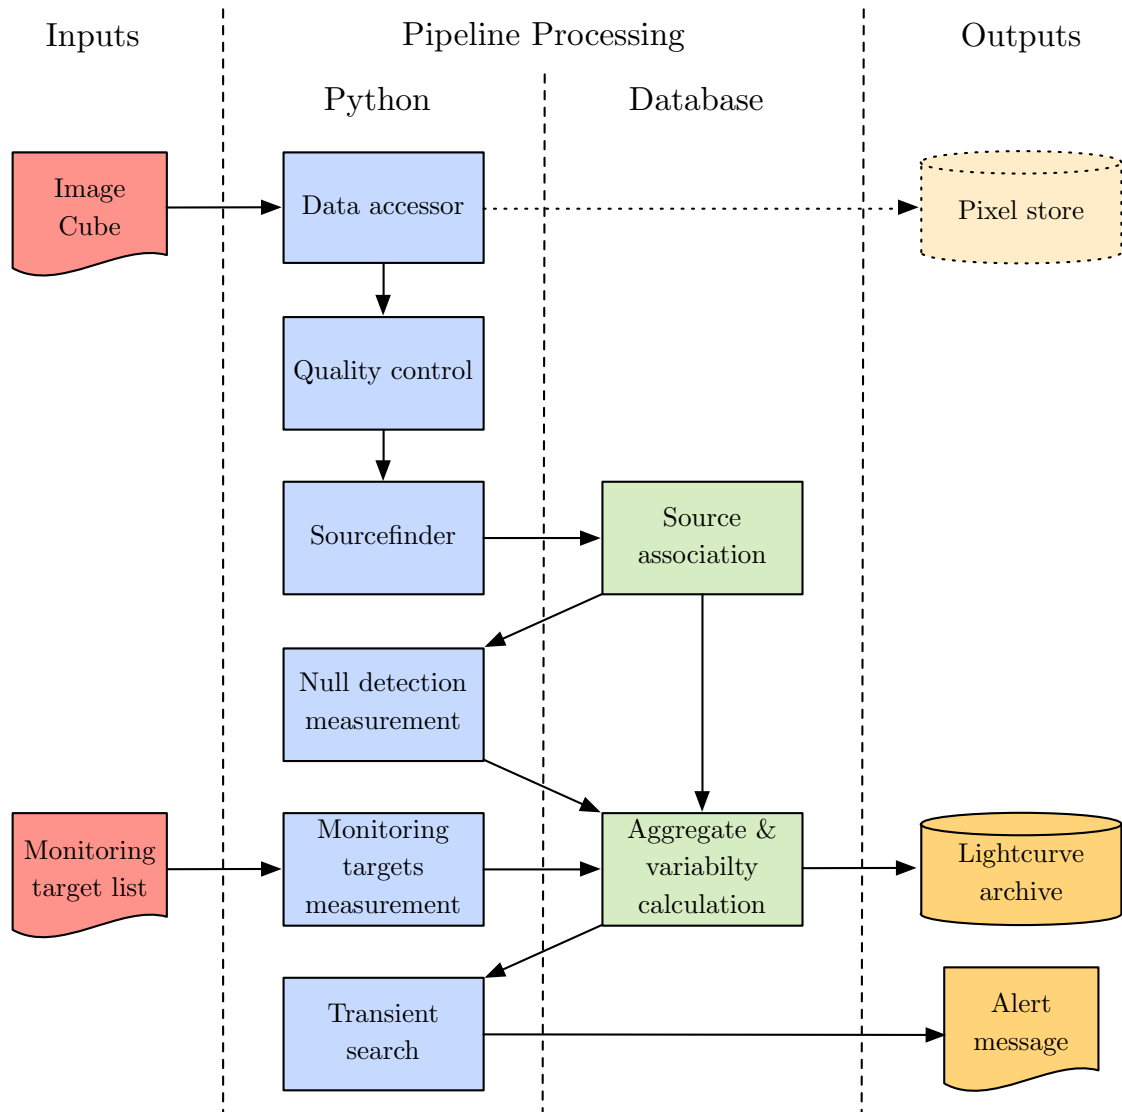


Figure 2.3: Schematic of the Transients Pipeline, adapted from Swinbank et al. (2015) Data are read in and processed by the pipeline in a combination of Python scripts and SQL database operations. The final products are a set of light curves and their properties (fluxes, positions, variability statistics) which users can then query to search for transients and variables.

2. Quality control is performed, accounting for high RFI and poor calibration. To do this, a central subsection (typically $1/8^{\text{th}}$) of the image is sigma-clipped to remove all sources, after which the central RMS flux is calculated. If this RMS falls outside globally acceptable maximum and minimum values (also set by the user) then the image is rejected. An appropriate PSF shape is also checked, ensuring that its semi-major and semi-minor axis are not infinite, sufficiently sampled (between 2 and 30 pixels) and not excessively elliptical (a ratio of > 2). Finally, images taken nearby ($< 10^\circ$) to bright radio sources are removed. These sources include the Sun, Jupiter and the ‘A Team’ galaxies - bright radio galaxies such as Cyg A and Vir A (de Bruyn et al., 2009).
3. Sourcefinding is performed, based on SExtractor (Bertin and Arnouts, 1996). The sourcefinding routine first splits an image into a grid of 50 by 50 pixel squares, in which all sources are clipped out by removing independent pixels that do not follow a Gaussian distribution. The median and RMS pixel values in each cell are calculated and then interpolated across the image to form a background and noise map respectively. Note that positions less than 10 pixels (set by `margin`) from the edge of an image are excluded from this analysis, due to the incomplete interpolation across partial grid cells. The background map is then subtracted from the data. After this, sources are defined as all groups of pixels with peaks above a user-defined `detection_threshold` (e.g. 8σ where σ is the interpolated RMS at that position), flood-filled down to an `analysis_threshold` of typically 3σ . These source pixel groups are optionally deblended (`deblend_nthresh` > 0). Each source is then fit with an elliptical Gaussian which can, crucially, be forced to be a point source (`force_beam = True`). The results of this fit - its position, peak and integrated fluxes, lengths of major and minor axes, position angle and signal-to-noise ratio - are stored to a table in the database.
4. A source association operation is conducted, in which sources extracted in new images are either identified as new or as updated points in the light curves of

previously known sources. Sources are associated if their positions are closer than n semi-major axes of the restoring beam and the de Ruiter radius - the angular distance between two sources divided by their positional errors - is less than a (typically fixed) threshold. Typically $n = 1$ beam semi-major axes is sufficient for association, though up to 3 might be considered if there are considerable systematic pointing errors and/or significant proper motion, to the detriment of database performance at run time.

5. The above steps are repeated for all images in a dataset. The pipeline will, in addition to blind sourcefinding, check the positions of all sources seen in previous images. Notably, the reverse is not true so checking pre-discovery images requires an additional pipeline run. Users can specify for how many images to check the positions of previously found sources with the `expiration` parameter, though throughout this thesis this value was always set to be larger than the number of epochs in a given dataset, ensuring the most complete light curves possible. This same process of source finding and association can also be performed for a user-specified list of positions to monitor and ‘force fit’.
6. A series of aggregate properties are calculated for the updated light curve of each source. These include the weighted mean position, its average integrated and peak flux densities and two variability statistics η and V . The former, defined for N flux density measurements $F_i \pm \sigma_i$ of weighted mean \overline{F}_ν^* as

$$\eta \equiv \chi_{N-1}^2 = \frac{1}{N-1} \sum_{i=1}^N \frac{(F_{i,\nu} - \overline{F}_\nu^*)^2}{\sigma_i^2} \quad (2.2)$$

describes their reduced chi-squared value when compared to a stable source i.e. constant flux density sources are expected to have $\eta \approx 1$ whilst variable sources produce values greater than unity.

V is sometimes known as the modulation parameter or co-efficient of variability and is the ratio of the sample standard deviation s to the mean of its flux measurements i.e.

$$V \equiv \frac{s_\nu}{\overline{F_\nu}} = \frac{1}{\overline{F_\nu}} \sqrt{\frac{N}{N-1} (\overline{F_\nu^2} - \overline{F_\nu}^2)}. \quad (2.3)$$

These aggregate properties can then be queried from the SQL database in order to find variable and transient sources, for example by selecting sources with large values of η and V .

The resulting light curves and their statistics can then be used to search for varying and transient sources. If it is assumed there are very few truly transient sources in our fields then the majority of light curves should produce values of $\eta \sim 1$ and $V \ll 1$. By plotting the distributions of these two statistics one can see if any sources produce anomalously high values of either, which would be likely real astrophysical transients or highly variable sources. For example, Driessen et al. (2020) demonstrate the first use of the TraP as applied to ThunderKAT data of the GX339–4 field and, by simply plotting the distributions of η and V found a likely active binary system. Similarly, Sarbadhicary et al. (2021) use observations of the COSMOS field (Scoville et al., 2007) taken by the VLA to search for variables in 209 epochs taken between 2013 and 2019. They identify 58 variable sources and, due to the excellent multiwavelength coverage of this field, are able to identify them as AGN whose variability is primarily due to scintillation.

2.2.2 MeerLICHT

Complementary to the ThunderKAT project is MeerLICHT (‘more light’; Bloemen et al., 2016), a 65cm primary mirror optical telescope with a 2.7 square degree FoV, located in Sutherland, South Africa. One of MeerLICHT’s primary goals is to co-point with MeerKAT, providing optical data commensurate with radio observations. MeerLICHT observations consist of 1 minute exposures of a given field, alternating between q -band (440 – 720 nm) and each of 5 Sloan bands u , g , r , i and z , reaching approximately 22nd magnitude. Therefore over the course of a 15 minute ThunderKAT observation, MeerLICHT is able to characterise the MeerKAT field of view across the entire optical spectrum. I have used MeerLICHT to search for optical counterparts to radio transients as found in this thesis.

```

1 [persistence]
2 description = "H1743 dataset"
3 dataset_id = -1
4 rms_est_sigma = 4           #; Sigma value used for iterative
   clipping in RMS estimation
5 rms_est_fraction = 8       #; Determines size of image
   subsection used for RMS estimation
6 rms_est_history = 100      #; how many images used for
   calculating rms histogram
7 rms_est_max = 100          #; global maximum acceptable rms
8 rms_est_min = 0.0          #; global minimum acceptable rms
9 bandwidth_max = 0.0        #; if non zero override bandwidth of
   image, determines which images fall in same band
10
11 [quality_lofar]
12 low_bound = 1              #; multiplied with noise to define
   lower threshold
13 high_bound = 80           #; multiplied with noise to define
   upper threshold
14 oversampled_x = 30         #; threshold for oversampled check
15 elliptical_x = 2.0         #; threshold for elliptical check
16 min_separation = 10        #; minimum distance to a bright
   source (in degrees)
17
18 [source_extraction]
19 detection_threshold = 8     #; extraction threshold (S/N)
20 analysis_threshold = 3
21 back_size_x = 50
22 back_size_y = 50
23 margin = 100
24 deblend_nthresh = 0        #; Number of subthresholds for
   deblending#; 0 disables
25 extraction_radius_pix = 1500 #; for 1.2lmbda/D *0.5 (radius) = 31
   arcmin, this is ~1.5x primary beam size.
26 force_beam = True
27 box_in_beampix = 10
28 ew_sys_err = 10            #; Systematic errors on ra & decl (
   units in arcsec)
29 ns_sys_err = 10
30 expiration = 11           #; number of forced fits performed
   after a blind fit
31
32 [association]
33 deruiter_radius = 5.68
34 beamwidths_limit = 3.0
35
36 [transient_search]
37 new_source_sigma_margin = 1

```

Listing 2.1: Configuration file for TraP as run on the H1743 field. The majority of these parameters are kept to their default values, with exceptions described in the text. The same listing holds for the other datasets used in this thesis.

*Let us go to the shore;
there the waves will kiss our feet.
With mysterious sadness
the stars will shine down on us.*

— Aleksey Plescheyev, from
Tchaikovsky's *June: Barcarolle*

3

The detection of stellar flares with MeerKAT

Contents

3.1	Introduction	45
3.2	Radio observations	47
3.3	Archival multiwavelength association	53
3.3.1	Optical and near-IR photometry	56
3.3.2	X-ray observations and correlation	58
3.4	SALT spectroscopy	65
3.5	Discussion	67
3.6	Conclusions	71

This chapter presents the detection and follow-up of the radio transient MKT-[J174641.0–321404](#), which is detailed in the journal article [Serendipitous discovery of radio flaring behaviour from a nearby M dwarf with MeerKAT \(Andersson et al., 2022\)](#). This is one of the first examples of serendipitous transients found with MeerKAT, providing excellent grounds for further searches with this instrument, in addition to the inherent interest in the atmospheres of our stellar neighbours.

3.1 Introduction

Radio emission has been seen originating from stars for over 60 years (Lovell, Whipple, and Solomon, 1963). Since this time, stellar radio detections have been seen across the Hertzsprung-Russell diagram, from the luminous O and B stars, whose emission is related to mass loss via ionised winds, through to ultra-cool M dwarf systems (Güdel, 2002). The emission mechanism underlying detections from the cooler half of the H-R diagram is described by Benz and Guedel (1994), whereby magnetic reconnection events accelerate electrons to produce non-thermal emission (see section 1.3). As these electrons lose energy they also heat the stellar corona, producing thermal X-rays. As a result, the conditions of stellar coronae can be probed by radio observations of stars. Furthermore, the emission from these low- to intermediate-mass systems can be both quiescent in nature, or display rapid flaring on timescales of minutes to days (Osten, 2007). These flares can be incoherent in nature, attributed to gyrosynchrotron emission from mildly relativistic electrons, or else caused by coherent mechanisms with high degrees of circular polarisation and brightness temperatures $> 10^{12}\text{K}$ (Dulk, 1985) - see for example the polarised detections of UV Ceti by both the MWA and ASKAP (Lynch et al., 2017; Zic et al., 2019). These coherent bursts are attributed to either plasma- or electron cyclotron maser-emission, which probe the local electron density and magnetic field strength respectively.

Of particular interest amongst these radio emitters are the ‘dMe’ stars - these are M-dwarfs, stars of surface temperatures between 3700 K (M0) and 2500K (M9.5), corresponding to masses $0.07 - 0.6 M_{\odot}$ (Rajpurohit et al., 2013). The ‘e’ in ‘dMe’ refers to the near ubiquitous presence of hydrogen (and calcium) emission lines seen when observing these stars, indicative of heating in their chromospheres (Cram and Mullan, 1979; Cram and Giampapa, 1987). Due to their low mass, the internal structure of M dwarfs is mainly (or sometimes entirely) convective in nature, which, in combination with stellar rotation, leads to strong surface magnetic fields through a dynamo effect (see e.g. Newton et al., 2017). M dwarfs make up approximately 70% of the stars in the Milky Way (Henry, 2007), making them the most common

place to understand stellar coronae and magnetic activity. Furthermore, the M dwarf population has been found to harbour many potential exoplanets, with an estimated occurrence rate of ~ 1.2 planets per star, increasing at later spectral types (Hardegree-Ullman et al., 2019). Therefore, host star variability is a crucial consideration in understanding planet habitability. The flares produced by M dwarfs can be orders of magnitude larger than solar flares (Lacy, Moffett, and Evans, 1976) and, along with associated space weather events such as coronal mass ejections (CMEs), are likely to erode the atmospheres of planets (Lammer et al., 2007), affecting UV surface dosage, greenhouse warming efficiency and surface water retention (Airapetian et al., 2017). This is thought to most affect planets close enough to their host to sustain liquid water, as at this distance from an M dwarf they are expected to be tidally locked to their host, resulting in small magnetic moments and a weak magnetic shield to protect against CME plasma (Khodachenko et al., 2007). Optical flares and superflares have been recorded from known planet or planet-candidate hosting M dwarf systems (Günther et al., 2020; Lin et al., 2021) including the 7 planets detected around TRAPPIST-1, a nearby (12 pc) active M8 dwarf (Gillon et al., 2017; Luger et al., 2017). Flaring has even been seen from our planet hosting, stellar nearest-neighbour Proxima Centauri (Anglada-Escudé et al., 2016; Damasso et al., 2020), where Zic et al. (2020) identified the analogue of a solar type IV burst, whose occurrence is strongly associated with space weather events.

Detections of radio emission and flares from M dwarfs have focused on a few nearby sources such as UV Ceti, Proxima Centauri, AD Leonis and YZ Canis Minoris (Lacy, Moffett, and Evans, 1976; Bastian, 1990; Slee, Willes, and Robinson, 2003; Villadsen and Hallinan, 2019). However, the recent advent of interferometers that can regularly sample the radio sky at high sensitivity with wide fields of view has allowed for new discoveries in this space. One method of detecting radio stars, demonstrated by Driessen et al. (2021), is by cross-matching commensal radio observations with the positions of known flaring stars detected at other wavelengths. In this work, the authors found radio emission coming from the previously known X-ray flare star EXO 040830-7134.7, whose location was observed commensally in

MeerKAT observations. However, the lack of polarisation measurements and only three observations make it unclear whether the observed emission is incoherent or coherent, or indeed whether the source was flaring or being observed in quiescence. The other way to look for radio stars is simply to search wide fields of the sky and, upon making radio detections (transient or otherwise), cross-match against optical sky survey data (e.g. *Gaia* and *TESS*; Ricker et al., 2015; Bailer-Jones et al., 2018) for nearby stellar systems. This method is demonstrated by Lynch et al. (2016), who made use of circular polarisation and multi-frequency images taken with ATCA to conduct a volume-limited survey of ultracool dwarfs in the southern hemisphere, detecting quiescent, gyrosynchrotron emission from 3 of 15 sources observed. More recently, Pritchard et al. (2021) and Pritchard et al. (2024) utilised the wide field of view of ASKAP (Hotan et al., 2021), with multi-epoch sampling and use of Stokes V imaging, to detect dozens of radio-loud stars from across the H-R diagram with no previous radio counterparts. Similarly, LOFAR (Van Haarlem et al., 2013), due to its eponymous low frequencies, has a correspondingly large field of view, which has allowed for untargeted, flux-limited surveys to observe coherent emission from 19 M dwarfs not before detected at radio wavelengths (Callingham et al., 2021).

This chapter describes the search, detection and follow-up of a stellar flaring source with MeerKAT, one of the first discoveries made as part of commensal efforts of the ThunderKAT team. This served as a precursor to more systematic searches detailed later in the thesis.

3.2 Radio observations

XR B H1743–322 was observed as part of the ThunderKAT monitoring program at a weekly cadence for 11 epochs, from 9th September until the 10th of November 2018. The observational schedule for this, and indeed most, ThunderKAT target followed reports of new outburst behaviour from satellites sensitive to high energies (Williams et al., 2020). Due to the commensal nature of this work, the reduction of all observations was performed as part of Williams et al. (2020), the basic procedure for which was outlined in Chapter 2 and which I summarise here.

Each observation was taken using the L-band (900–1670 MHz) receiver, with a central frequency of 1284 MHz, bandwidth of 856 MHz and 4096 frequency channels. The first observation consisted of four 10 minute scans, with every subsequent observation consisting of 15 minutes on source. The on source observations were interleaved by 2 minutes on the phase calibration source J1712–281 for the first three observations and J1830–3602 for the remainder of the epochs. This choice of secondary calibrator seems to have had an effect on the systematic properties of the resulting images, as discussed later. The primary calibrator J1939–6342 was observed for 10 minutes before beginning each observation block to set the flux scale and band pass solutions. The reduction of these radio visibilities followed the standard protocol of flagging, calibration and imaging. Some initial manual flagging was done to remove RFI before the extremal ~ 209 kHz bandwidths were removed (due to band edge effects), along with zero amplitude visibilities. Further flagging was performed with `AOFlogger` (Offringa, Van De Gronde, and Roerdink, 2012), in total leaving around $\sim 80\%$ of the initial data remaining. Calibration was performed with `CASA` (McMullin et al., 2007) and imaging was done with `WSClean` (Offringa et al., 2014, Offringa and Smirnov, 2017), producing square images of 204 arcminutes in size, with a pixel scale of $1.5''$. A Briggs weighting of -0.85 was applied to clean the images, with resulting images having a synthesised beam of size $\sim 5''$. The epoch in which the resulting stellar transient was found to be brightest was also split into 5×3 minute images, in order to search for intra-epoch variability (the lower panel of Figure 3.4).

These commensal images form the basis of this work and were searched through the Transients Pipeline for new transients. I ran the `TraP` on the 11 images of the H1743 field, with input parameters to the pipeline shown in Listing 2.1. One of the altered parameters is the `extraction_radius_pix`, which simply identifies out to how many pixels in radius from the centre sourcefinding should be carried out. I chose to analyse images out to a radius approximately $1.5 \times$ MeerKAT’s primary beam radius of $\sim 45'$, as has been done in previous work - Sarbadhicary et al. (2021) have shown that this is a typical radius beyond which source densities fall

significantly. During fitting, each source is assumed to be the size of the synthesised beam (`force_beam=True`) as here I am only interested in the unresolved point source variability. Using the default 8σ detection threshold above the local noise, 377 individual sources were detected across all epochs. This threshold is a trade-off between being inundated with many false positives and missing genuine transients. The resultant TraP variability parameter distributions are detailed in Figure 3.1, showing outliers to typical η and V parameters (i.e. $V \gtrsim 0.75$ and $\eta \gtrsim 20$). I note that 50% of sources lie below $\eta = 1$, whilst those above are either bright sources with very little variability (η is approximately proportional to signal-to-noise ratio squared), or extended sources whose variation is artificially caused by the shape of the synthesised beam changing between epochs.

As seen in Figure 3.1, there is a grouping of sources in the same area of parameter space, all with $V \sim 1.6$ and $\eta \sim 10$. These were all inspected by eye and found to be imaging or processing artefacts near bright sources. I have plotted three example light curves from this group in Figure 3.2. As mentioned earlier, the change from J1712–281 to J1830–3602 as the phase calibration source after the first three epochs has systematically affected the observations, as seen by these sources only showing variability in the first three epochs. This systematic effect can also be seen in the noise levels shown in Figures 3.3 and 3.4. I was able to quickly identify these false positives due to their small number and the well documented change in calibration procedure. However, for datasets of much greater size, this kind of systematic would be much harder to identify. This is one line of reasoning that led me to develop a citizen science project in Chapter 4. After removing these false positives, two highly significant astrophysical sources were identified with values of $\eta > 20$ and $V > 1$. One, reassuringly, was the target of this field, the candidate black hole XRB H1743–322 (Markwardt and Swank, 2003) and the other was MKT J174641.0–321404, a hitherto unknown radio source located $\sim 5'$ East of the aforementioned XRB.

This radio source was detected serendipitously twice during the final four epochs of observation, at a maximum of $590 \pm 60 \mu\text{Jy}$ and can be seen to appear and




Figure 3.1: Variability parameters for all sources detected in the H1743–322 field during ThunderKAT observations in late 2018. Clear outliers, particularly in V can be seen, including the XRB H1743–322 itself, several manually vetted artifacts and an unknown source.

disappear in the radio images of Figure 3.3. As described in the previous section, it possible to give the TraP a set of coordinates to monitor during every time step. I ran the TraP with instructions to monitor the position of the radio source, in order to better understand the emission from MKT J174641.0–321404 and provide better constraints on its variability. The resulting light curve from such forced photometry of the source over all epochs can be seen in Figure 3.4’s upper panel, where non-detections (defined as flux measurements detected at $< 3\sigma$ above local noise) are indicated with faint points. The source can be seen to be detected during the first epoch of observation, at $\sim 3.5\sigma$. The brightest detection of MKT J174641.0–321404 was also split into 5 3-minute integrations, as seen in Figure 3.4’s lower panel. The flux densities seen in this intra-epoch imaging show variation from marginal detections and non-detections to no excess brightness at all. This provides marginal evidence for variability which I use later to constrain the brightness

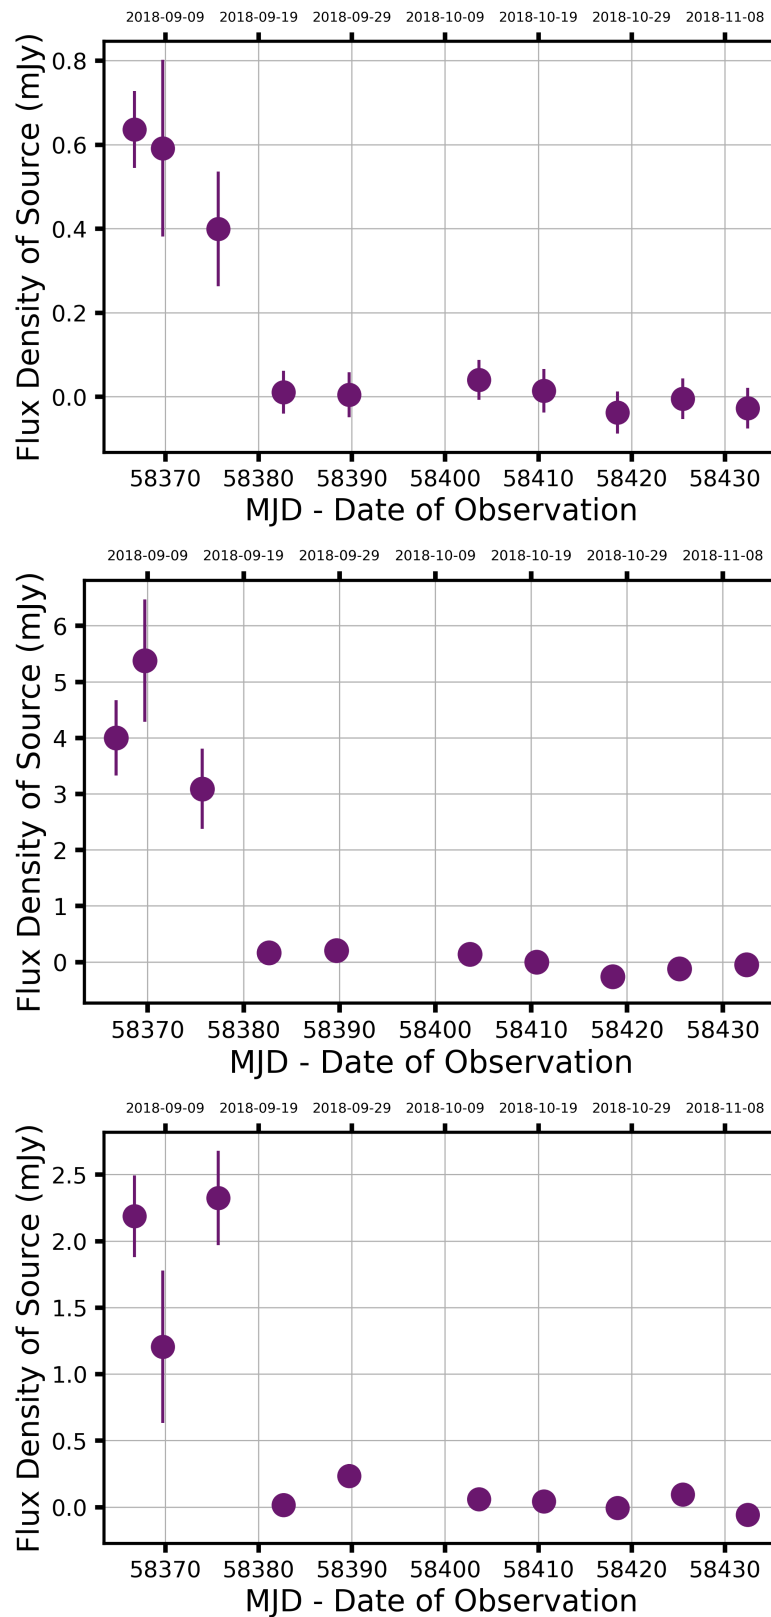


Figure 3.2: Three example false positive light curves from this dataset. These sources were all found to be around bright sources, with detections only in the first three epochs. This was likely due to a different calibrator being used for these observations.




Figure 3.3: One arcminute square images of MKT J174641.0–321404 and surrounding noise during all ThunderKAT observations, running chronologically from top left, across then down. The source can be clearly seen to appear and disappear multiple times across all epochs. The synthesised beam of each image can be seen in the bottom left of each panel, whilst the white central cross-hairs show the source’s mean MeerKAT position.

temperature of the source. As the source is $\sim 5'$ removed from the phase-centre of the radio images, the flux calibrations are imperfect. However, the true values will not be so dissimilar as to invalidate these findings, especially as I am primarily interested in relative variations, and so no further observations have been made nor primary beam corrections applied to correct the light curve values. Using measurements of the shape of MeerKAT’s L-band primary beam from de Villiers and Cotton (2022) I estimate that the intrinsic flux densities measured are $< 5\%$

different from those reported here.

3.3 Archival multiwavelength association

The radio position of the source is taken to be the weighted mean of the three detection epochs using the Python Blob Detector and Source Finder (PyBDSF; Mohan and Rafferty, 2015). This differs slightly from TraP as the former includes the first epoch’s detection, whilst the blind TraP detection does not. Cross-matching this mean position of $17^{\text{h}}46^{\text{m}}41.0^{\text{s}}, -32^{\circ}14'04''$ ($266.6709^{\circ}, -32.2343^{\circ}$) with the SIMBAD database¹ (Wenger et al., 2000) returns one object within $10''$, the optical source SCR 1746–3214 (henceforth SCR 1746, detection detailed in Boyd et al., 2011). SCR 1746 is an M dwarf and has an ICRS J2000 position of $17^{\text{h}}46^{\text{m}}40.659^{\text{s}} - 32^{\circ}14'04.50''$ with uncertainties of 0.10 and 0.8 mas respectively, calculated from *Gaia* Data Release 2 (Gaia Collaboration et al., 2016; Gaia Collaboration et al., 2018), placing it at an on-sky distance from the radio coordinates of $4.5''$. SCR 1746 has a *Gaia*-measured distance from the Sun of $12.06_{-0.2}^{+0.1}$ pc (Bailer-Jones et al., 2018). Due to this relative proximity, the star has a high proper motion μ of $(205.4 \pm 0.2, 103.0 \pm 0.1)$ mas yr⁻¹ in Right Ascension and Declination respectively. Accounting for this motion, SCR1746 is $1.3''$ from the radio source. This is within the 2σ positional uncertainty of the radio position where $\sigma = 1.2''$ and so we can associate the radio transient events of MKT J174641.0–321404 with SCR 1746. The next nearest optical source in the *g*-band MeerLICHT image (see below and Figure 3.5) is more than $20''$ away. As a final point on positions, Driessen et al. (2022) investigated the absolute astrometry of MeerKAT pointings in other ThunderKAT data and found an offset of, on average, $0.9''$ with respect to bright point sources seen by the Australian Telescope Compact Array (ATCA). I have made no attempt to correct for a systematic offset in this work but note that even if this moved our radio position further from the *Gaia* system, the sources would still be within 3σ positional uncertainties.

¹Set of Identifications, Measurements, and Bibliographies for Astronomical Data, accessed via <http://simbad.u-strasbg.fr/simbad/>.

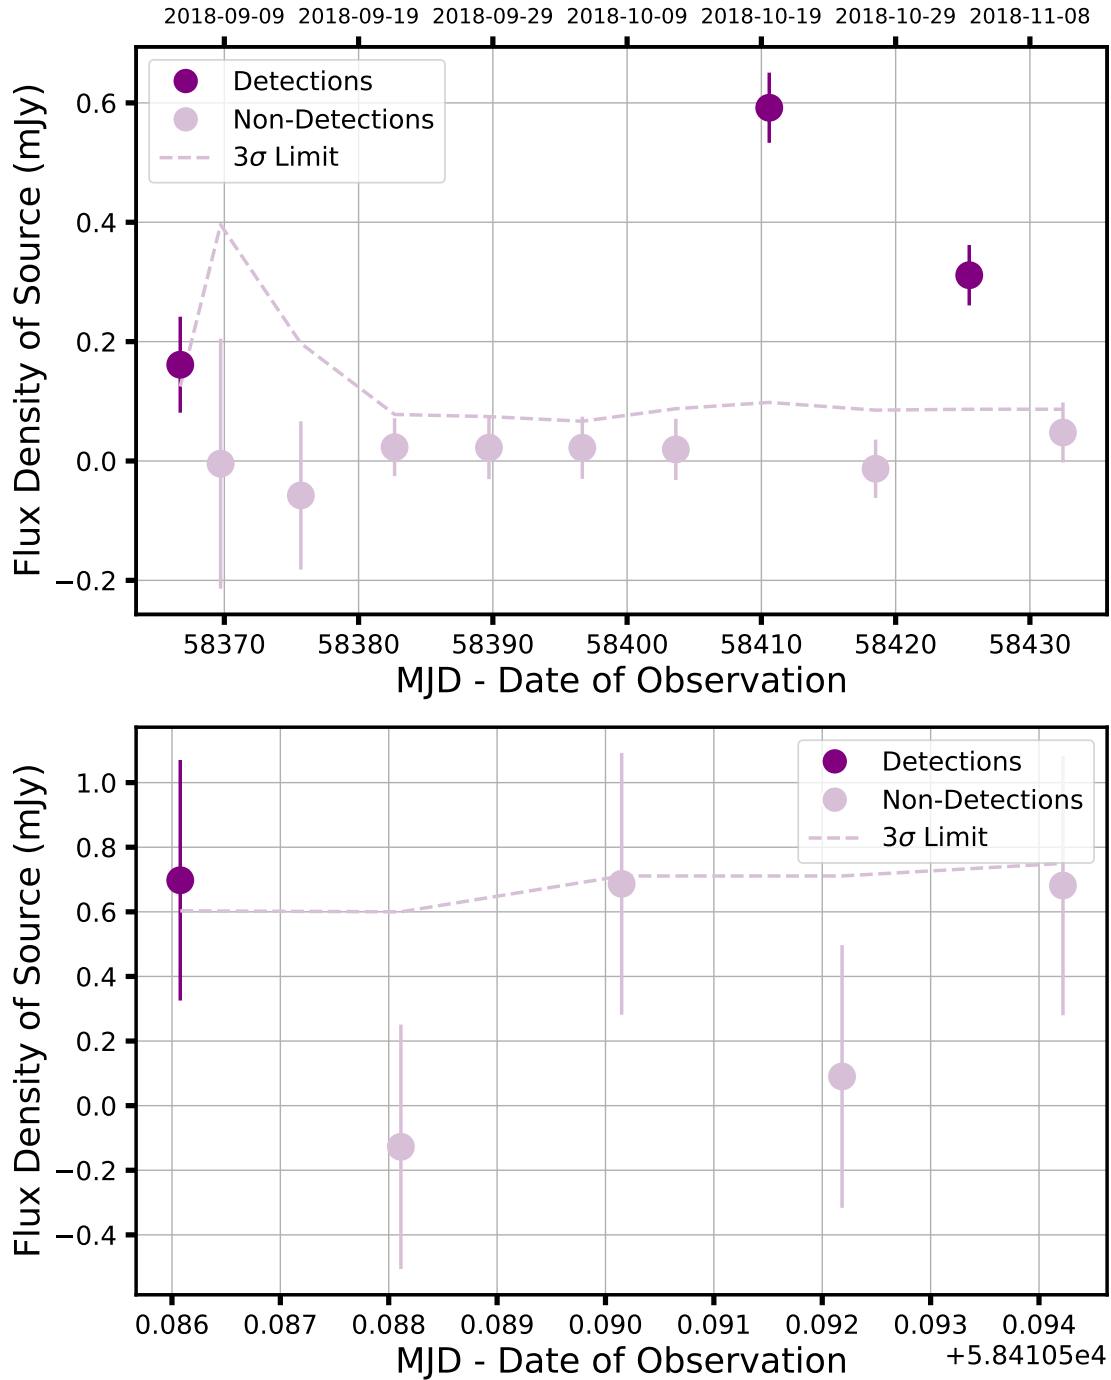


Figure 3.4: Upper: Radio light curve of MKT J174641.0–321404 across all 11 ThunderKAT observation epochs. Non-detections are defined as measurements below a 3σ threshold (dashed line) where σ is the local RMS noise in that epoch. The non-physical measurements in some epochs are caused by the forced photometry and epoch-specific noise structure, visible in Figure 3.3. The higher detection threshold (due to more noise) in the first three epochs is likely due to the use of a different phase calibrator than the later observations. **Lower:** Radio light curve compiled from 5×3 minute slices of the brightest 15 minute epoch. From this there is evidence for marginal variability within the epoch.

Multiwavelength coverage from MeerLICHT (Bloemen et al., 2016) confirms this association. A stack of q -band (440–720 nm) exposures taken on 7th June, 9th July and 16th July 2019 and 28th July 2020 form the MeerLICHT reference image for this field, where the small time period between the radio and optical observations renders proper motion effects negligible. The region surrounding MKT J174641.0–321404 can be seen in Figure 3.5, showing the proper motion of the *Gaia* source to within positional uncertainties of both the MeerKAT and MeerLICHT objects. We therefore can be confident that the MeerLICHT source is SCR 1746 and moreover that MKT J174641.0–321404 represents radio flaring behaviour from this likely M dwarf. I note that there are no observations simultaneous with the radio data as the strict MeerKAT-MeerLICHT coupling was not fully in-place during 2018, early in the life of the projects. In contemporary observations MeerLICHT follows MeerKAT pointings when possible, during clear observations taken at night.

The archival associations and MeerLICHT data allow for the star’s spectral type to be estimated from its photometry. The latest version of the *TESS* Input Catalog (TIC v8; Stassun et al., 2019) estimates SCR 1746 to have an effective temperature of 2870 ± 160 K, a mass of $0.12 \pm 0.02 M_{\odot}$ and radius of $0.146 \pm 0.004 R_{\odot}$. Sebastian et al. (2021) compile a target list of ultracool dwarf systems in preparation for *TESS*, where an empirical relationship for $T_{\text{eff}}(M_H)$ from Filippazzo et al. (2015) is inverted to show that the target has a spectral type of M5.1. Their calculated temperature, mass and radius are all slightly different than, but consistent with, those from Stassun et al. (2019). The effective temperature scale of M dwarfs given by Rajpurohit et al. (2013) also indicates that a temperature of ~ 2900 K is in agreement with a mid-M spectral typing. Finally, MeerLICHT observations provide quasi-simultaneous colours $(r - i) = 2.22 \pm 0.02$ and $(i - z) = 1.09 \pm 0.02$, in agreement with typical values for M6 stars (West, Walkowicz, and Hawley, 2005; Douglas et al., 2014).

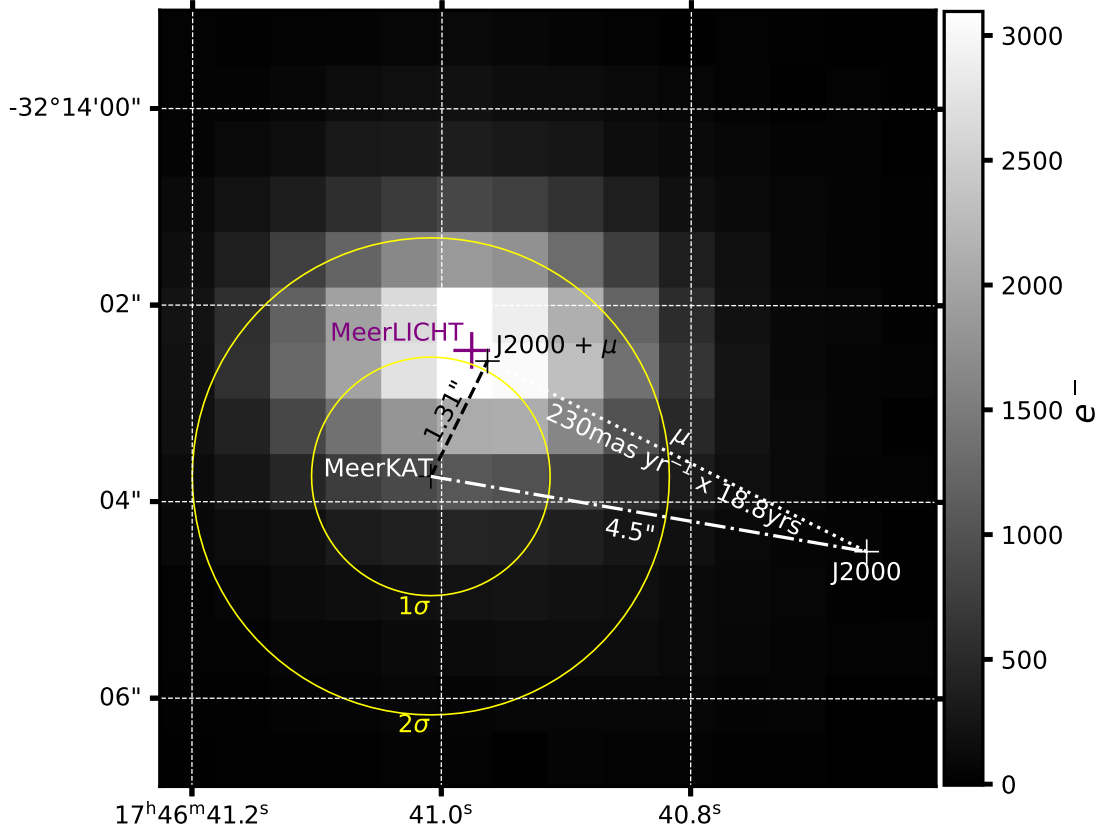


Figure 3.5: MeerLICHT q -band (440–720 nm) reference image (pixels) of the region near radio source MKT J174641.0–321404. 1 and 2σ positional uncertainties on the MeerKAT position (circles) show an overlap between the MeerLICHT object and the *Gaia* proper motion attributed to SCR 1746.

3.3.1 Optical and near-IR photometry

There is substantial photometric sky survey coverage of SCR 1746. Catalogued measurements of this source include the Two-Micron All-Sky Survey (2MASS; Cutri et al., 2003; Skrutskie et al., 2006) the AllWISE data release (Cutri et al., 2013), US Naval Observatory B catalogue (USNO-B; Monet et al., 2003), GLIMPSE Source Catalog (Spitzer Science Center, 2009) and the optical SuperCOSMOS-RECONS (SCR) southern sky proper motion searches (Boyd et al., 2011). A non-exhaustive summary of some of the recorded magnitudes of this source can be found in Table 3.1, intended to illustrate the source’s photometric properties of being relatively bright and red.

SCR 1746 (Tess Input Candidate 111898820; Stassun et al., 2019) was also

Table 3.1: Select magnitudes of red dwarf SCR 1746–3214 across the visible and near-IR spectrum. Bands are listed with their effective wavelength λ_{eff} .

Band (λ_{eff})	Magnitude	Uncertainty	Reference
μm	mag	mag	
B (0.45)	17.97	0.04	Page et al., 2012
V (0.55)	15.6	0.2	Stassun et al., 2019
G (0.62)	13.9734	0.0006	Gaia Collaboration et al., 2018
I (0.88)	12.98	0.30	Sebastian et al., 2021
J (1.25)	10.35	0.02	Cutri et al., 2003
H (1.65)	9.74	0.03	"
K (2.16)	9.38	0.02	"
W1 (3.35)	9.18	0.03	Cutri et al., 2013
W2 (4.6)	9.07	0.3	"
W3 (11.6)	8.82	0.10	"
W4 (22.1)	8.1	0.4	"

observed by NASA’s Transiting Exoplanet Survey Satellite (*TESS*; Ricker et al., 2015) during Sector 39 of its extended mission. The data, which span 27.9 days between 27 May and 24 June 2021, consist of observations obtained every 2-seconds. These images were combined into 2 minute cadence data products on board of the spacecraft prior to being processed and reduced by the Science Processing Operations Center (SPOC; Jenkins et al., 2016). The full *TESS* dataset for this star, binned to 2 minute and 30 minute cadence, can be seen in Figure 3.6, showing clear flares that reach up to 30 times quiescent flux levels². The data gap seen in the centre of the full light curve corresponds to the time taken (~ 1 day) for the spacecraft to send the data to Earth and re-orient itself.

The *TESS* data were searched using a Lomb-Scargle periodogram (Lomb, 1976; Scargle, 1982) to identify periodic signals in the light curve, revealing a significant periodicity with a period of 0.2292 ± 0.0025 days. When phase folding the *TESS* light curve at this period, a sinusoidal fit to the data produces a best fit relative amplitude of 0.52 ± 0.01 (bottom panel of Figure 3.7). There are some background sources on the same *TESS* pixel, but at G magnitudes of 18.9 or fainter (c.f. 13.95 for SCR 1746) I conclude that these are all too faint to be causing the observed

²I note that the *TESS* Full Frame Image timescale is also 30 minutes but that this matches the chosen binning only by coincidence.

modulation. As can be seen in the top panel of Figure 3.7, more flares tend to occur between phase 0 and 0.5, corresponding to the brightest parts of the phase curve. This could be due to the photometric modulation being caused by magnetically active spots and associated bright plages rotating into or out of view. If it is assumed that the star rotates as a rigid body with radius R then the photometric period can be expressed as a tangential velocity $v_t = 2\pi R/P = 32.2 \pm 0.9 \text{ km s}^{-1}$, demonstrating that SCR 1746 is a fast rotator, comparable to other low mass systems (c.f. Gizis et al., 2017 who find values of $v_t \sim 45 \text{ km s}^{-1}$ for 2M0335+23 and the solar value of $\sim 2 \text{ km s}^{-1}$). Even without this assumption on radius, a rotational period of ~ 0.23 days is much faster than the Sun (~ 25 days) and very much in line with the findings of Pass et al. (2023), who show that approximately 3 in every 4 active M dwarfs have rotation periods shorter than 2 days.

Additional optical monitoring of the source is provided by both MeerLICHT and the All-Sky Automated Survey for Supernovae (ASAS-SN, a global network of 24 40-cm telescopes; Shappee et al., 2014; Kochanek et al., 2017). The optical data in Figure 3.8 spans almost 4 years of observations in Johnson’s V, Sloan *ugriz* and *g*-band filters. MeerLICHT rotates between filters every 2 minutes, allowing for quasi-simultaneous colour determination. From the MeerLICHT data we can clearly see the star is much brighter at redder wavelengths. Aperture photometry at a 5σ detection threshold from the ASAS-SN Sky Patrol portal shows that optical flares are likely to have been observed from this source in the *g*-band. Despite this, none of the optical or radio flares are simultaneous to within < 0.4 days. Given that flares can be as short as minutes, this is not unexpected behaviour. Similarly, the photometric period observed by *TESS* is not observed in these data due to infrequent sampling - the average number of ASAS-SN *g*-band detections per period is less than two i.e. the underlying periodicity is not being Nyquist sampled.

3.3.2 X-ray observations and correlation

As demonstrated by Driessen et al. (2021), radio-detected M dwarfs can be X-ray sources. Therefore when searching for multiwavelength counterparts, I also queried




Figure 3.6: The *TESS* light curve of SCR 1746 (Tess Input Candidate 111898820) spanning 27.9 days from May to June of 2021, including the 2-minute cadence data (purple) and the same data binned down to 30 minute cadence. Flare amplitudes can be seen up to $30\times$ above non-flaring times.

X-ray catalogues. X-ray counts from the 7 year Swift-XRT point source catalogue (1SWXRT; D’Elia et al., 2013), taken over a total exposure time of 4631.1s on the H1743–322 field were found to be coincident within positional uncertainties of SCR 1746 in the epoch J2000. The uncertainty in the X-ray position is $6.6''$, larger than the uncertainties in either the radio or optical data used here. This catalogue converts the raw count-rate of the detector into a flux value by assuming a model spectrum of an absorbed power-law. I note that this model is appropriate for non-thermal electrons but not for thermal coronal X-ray emission, as mentioned by Stelzer et al. (2013). The associated X-ray source has a 0.3-10keV flux upper limit of $1.68 \times 10^{-13} \text{ mW m}^{-2}$, measured at $S/N = 2.8$.

1SWXRT has been superseded by the Swift-XRT Point Source catalogues 1 and 2 (1SXPS; Evans et al., 2014, 2SXPS; Evans et al., 2020), both of which contain sources close to the radio transient (separation in 1SXPS = $7.5''$, 2SXPS = $2.31''$) that do not formally lie close enough to the *Gaia* position for association within a 90 per cent confidence interval region. They do, however, overlap with the

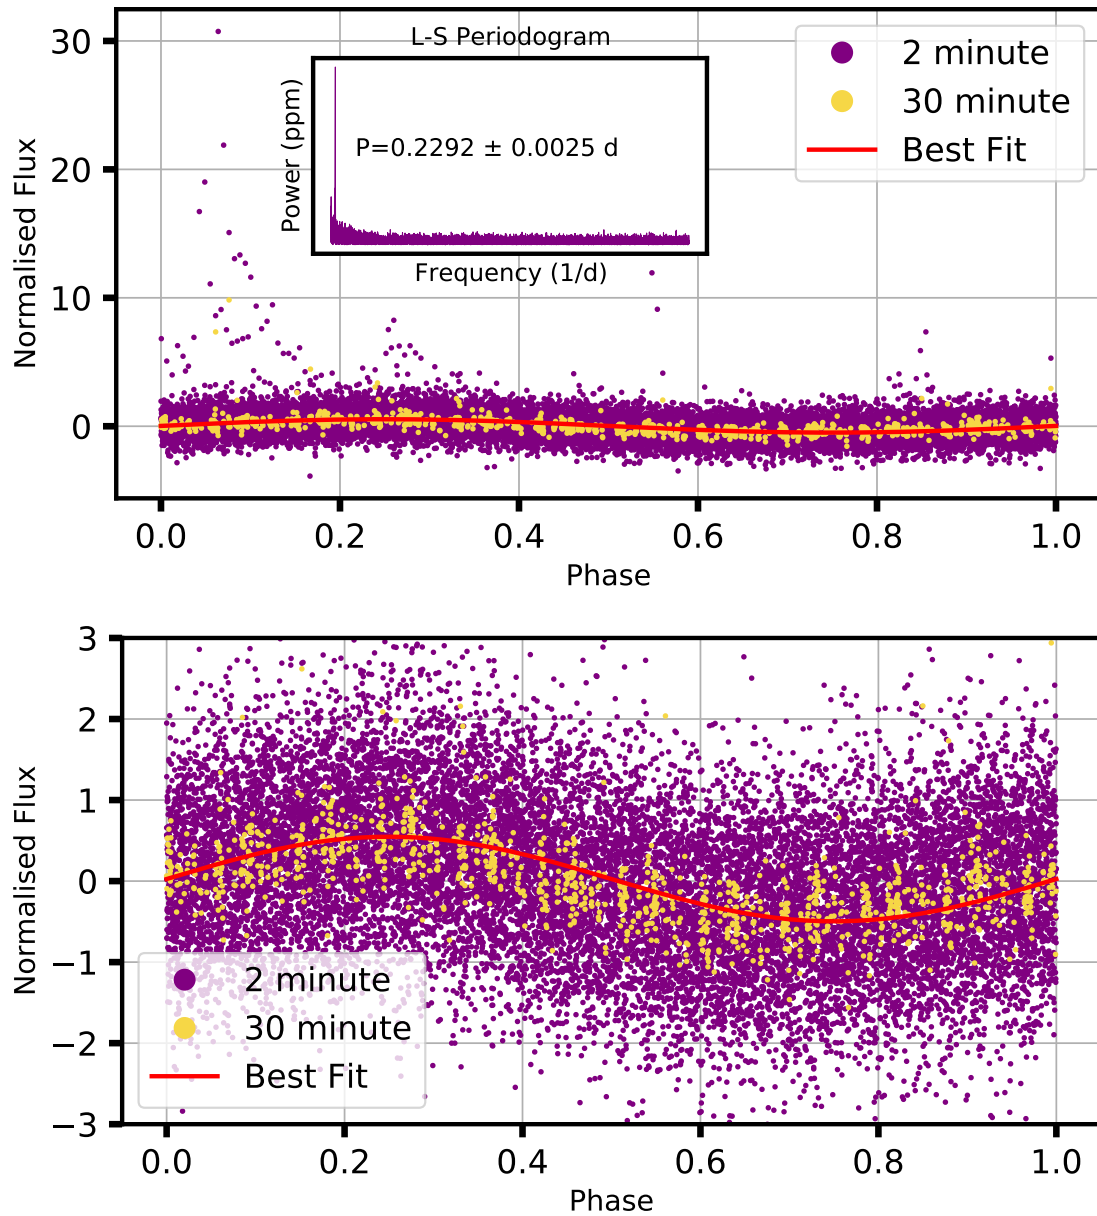


Figure 3.7: Upper: The phase folded *TESS* light curve of SCR 1746, with the corresponding Lomb-Scargle periodogram inset. Flares tend to occur between 0.0 and 0.4 in phase, corresponding to the brighter parts of the quiescent light curve. A sinusoidal best fit to the unbinned data is plotted, showing a relative amplitude of 0.52 ± 0.01 . **Lower:** A closer view at the non-flare behaviour of the phase-folded light curve, in which the amplitude of the best fit sine curve is clearly visible. The 30 minute cadence data reduces the scatter and falls in good agreement with the sinusoid.




Figure 3.8: Upper: Optical data from MeerLICHT and ASAS-SN. Flaring behaviour can be seen in data from the ASAS-SN Sky Patrol portal between radio observations and mid-2021. Downwards facing triangles indicate upper limits in their respective bands. **Lower:** The radio light curve of SCR 1746 overlaid with the simultaneous optical data. No flares are simultaneous to within < 0.4 days with those in the other bands.




Figure 3.9: A stacked image of the total count rate across all Swift-XRT energies towards SCR 1746, as provided by the 2SXPS catalogue (Evans et al., 2020). The image is $\sim 6.3'$ to a side and shows a detection of a source that, whilst consistent with the MeerKAT uncertainties, does not formally overlap with the *Gaia* J2000 position of SCR 1746. The increased background counts to the right of the image are caused by H1743–322.

relatively large MeerKAT uncertainties. The 2SXPS entry, seen in 46 observations of this field and whose counts can be seen in Figure 3.9, includes more detail. Specifically, the counts from observations in this field have been converted into flux values using a model from the Astrophysical Plasma Emission Code (APEC; Smith et al., 2001), which is specifically intended for hot plasmas such as those found in the coronae of stars. The calculated APEC model has a mean flux of $7.1(\pm 0.6) \times 10^{-14} \text{ mW m}^{-2}$, which is a factor of $\sim 0.42\times$ the 1SWXRT flux. I decided to use the analysis from 1SWXRT but note that the same overall conclusions made here would still apply if the 2SXPS entry was used.

As discussed in Section 3.1, there is a known relation between the quiescent radio (L_R) and X-ray (L_X) luminosities of several types of active star, indicating a connection between the nonthermal, energetic electrons causing the radio emission

and the bulk coronal plasma responsible for thermal X-rays:

$$\log(L_X) \lesssim \log(L_R) + 15.5 \quad (3.1)$$

(Guedel and Benz, 1993; Benz and Guedel, 1994). The standard interpretation of this Güdel and Benz relation is that magnetic reconnection in the corona accelerates a population of non-thermal electrons that emit at radio frequencies via gyrosynchrotron emission. These electrons also heat the chromosphere of the star, producing thermal X-ray emission. Coherent emission processes are known to violate this relationship, as shown by Callingham et al. (2021), making agreement or otherwise with the Güdel and Benz relation an effective diagnostic of emission mechanism (see Section 3.5).

Figure 3.10 shows the quiescent limits and flare magnitude of SCR 1746 in comparison to other M dwarfs and other types of active star. This figure shows that M- and K-type dwarfs, some suffixed ‘e’ to denote emission lines (dM/dMe and dKe, pink triangles and yellow diamonds respectively) tend to have lower radio and X-ray luminosities than the RS CVn binaries denoted by black circles. The quiescent radio limit for SCR 1746 is $3\times$ the lowest RMS noise floor across the individual MeerKAT epochs, measured locally using PyBDSF to be $\sim 22\mu\text{Jy}$. This corresponds to a specific luminosity upper limit of $1\times 10^{13}\text{erg s}^{-1}\text{Hz}^{-1}$, an order of magnitude below the brightest observed flare at $(1.0\pm 0.1)\times 10^{14}\text{ erg s}^{-1}\text{Hz}^{-1}$. These are only approximate results for SCR 1746 as shown in the figure, as the quiescent radio and X-ray emission measurements should ideally be taken simultaneously, which has not been the case here. Nevertheless, that measurements of SCR 1746 are consistent with the Güdel and Benz relation provides evidence for the system undergoing magnetic reconnection and displaying gyrosynchrotron emission in the radio and thermal X-rays. As mentioned earlier, the 2SXPS catalogue entry does not lie close enough to the optical position to allow for formal association. However, it does overlap with the MeerKAT position and if it were associated then the lower catalogued flux and definitive detection of X-rays from SCR 1746 would place the star more clearly in line with the Güdel and Benz relation, strengthening the case for this being a flaring M dwarf undergoing magnetic reconnection.

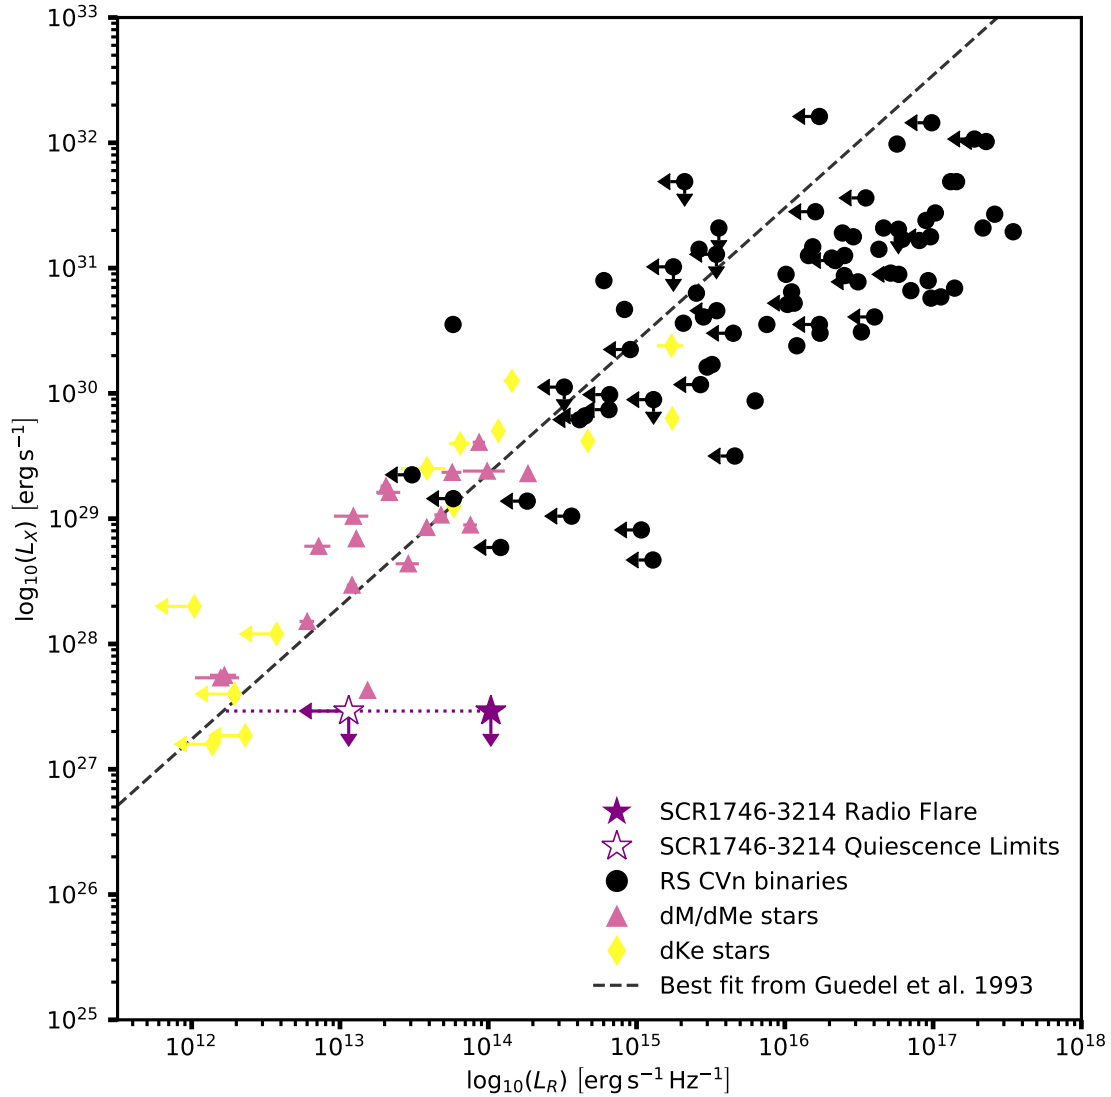


Figure 3.10: X-ray and radio luminosities of several types of active star (Guedel and Benz, 1993) - reproduced from <https://github.com/AstroLaura/GuedelPlot>. Limits on SCR 1746's quiescent emission show this is an intrinsically faint source and could be approximately consistent with the Güdel and Benz relation for M dwarf (dM/dMe) stars. The radio flare is also plotted, showing a brightness increase of over an order of magnitude above quiescent limits.

3.4 SALT spectroscopy

There exists a survey programme on the 11-m Southern African Large Telescope (SALT; Buckley, Swart, and Meiring, 2006) to follow-up and provide spectra on transients detected as part of the ThunderKAT project. As such, following the radio detection and some preliminary optical analysis, a spectrum was requested from SALT in order to provide confirmation of spectral type and investigate the magnetic activity of the star. Two consecutive 1000s exposures were taken using SALT’s High Resolution Spectrograph (HRS; Bramall et al., 2012; Crause et al., 2014) starting at UT 23:49:05 on the 11th of July 2021 at a seeing of $\sim 2.0''$. The HRS is a high dispersion échelle spectrograph and was operated in low-resolution mode at a spectral resolving power $\frac{\lambda}{\Delta\lambda} \sim 16500$. Wavelength calibration for HRS is performed relative to Thorium-Argon arc spectra on a weekly basis, in order to account for instrumental shifts (e.g. changes in optical path lengths due to expansion or contraction). The data were reduced with PySALT³ (Crawford et al., 2016), accounting for trimming, bias subtraction, gain correction, cosmic ray cleaning and flat fielding. Whilst HRS is dual beam, providing both blue (370-555nm) and red (555-890nm) spectra, as archival photometry indicates SCR 1746 is highly red (Table 3.1), I opt to only analyse the latter.

The two red spectra taken of SCR 1746 can be seen in Figure 3.11. The extensive structure at wavelengths $>7000 \text{ \AA}$ is typical for late M dwarfs and is caused by absorption due to molecules such as TiO, whilst $H\alpha$ can also be seen in emission at $\sim 6560 \text{ \AA}$. This $H\alpha$ emission is a key diagnostic of magnetic activity on the star, caused by heating in the chromosphere. One can quantify this emission by calculating the equivalent width (EW) of the line. The EW of a spectral line is the width of the adjacent continuum that has the same area as is taken up by the feature itself. It is therefore essentially a measure of strength of the line, defined for total spectrum $F(\lambda)$ relative to a continuum level F_c as

$$EW = \int \frac{F_c - F(\lambda)}{F_c} d\lambda = \int \left(1 - \frac{F(\lambda)}{F_c}\right) d\lambda, \quad (3.2)$$

³<http://pysalt.salt.ac.za/>




Figure 3.11: Two reduced spectra of SCR 1746–3214, taken by SALT. Some notable spectral features have been indicated, including clear H α emission and molecular lines, estimated from Kirkpatrick, Henry, and McCarthy (1991). Inset is a view of the H α emission line, which shows rotational broadening (see Section 3.5).

which has units of wavelength and for emission lines, whose $F(\lambda) > F_c$ by definition, is a negative value. Following the prescription of Newton et al. (2017), the EW of the H α emission line integrated between 6558.8 Å and 6566.8 Å were calculated relative to a continuum level either side of the feature, between 6500–6550 Å and 6575–6625 Å. The mean EW and its standard deviation of the two spectra are -5.5 ± 0.2 Å placing SCR 1746 clearly in the population of magnetically active stars ($\text{EW} \leq -1$ Å; see Newton et al., 2017).

In order to estimate a spectroscopic classification and metallicity for SCR 1746, data were compared to co-added templates of individual stellar spectra from the Sloan Digital Sky Survey’s Baryon Oscillation Spectroscopic Survey (SDSS BOSS; Dawson et al., 2013). BOSS is primarily a redshift survey designed to measure the scale of baryon acoustic oscillations, but its ample coverage of stellar spectra make it useful for ancilliary purposes. Comparisons to template spectra were made using the spectral typing code PyHammer (Kesseli et al., 2017; Roulston, Green, and

Kesseli, 2020)⁴. 34 spectral indices between 3650 and 10200 Å, calculated from co-added spectral types O5 through L3 with metallicities $-2.0 < [\text{Fe}/\text{H}] < +1.0$ dex, form the model templates for each stellar class. The weighted mean and variance of these 34 features in an input spectrum are compared to the template values via chi-squared minimisation. The model with the smallest chi-squared is the closest to the input spectra and visual inspection allows for confirmation of this. Using this analysis, the best fitting empirical template for both SALT spectra is of an M8 star with metallicity $Z = -0.5$. I take the standard deviations of the PyHammer results when tested on their own SDSS BOSS spectra i.e. those used to build the empirical templates, as pessimistic estimates for uncertainties in classification and metallicity Z , which come to ± 1.5 spectral subtypes and ± 0.4 dex respectively (see the Appendix of Kesseli et al., 2017 for more on this).

3.5 Discussion

Commensal analysis of MeerKAT images of the sky surrounding H1743–322 identified a new radio transient, MKT J174641.0–321404, coincident with the high proper motion star SCR 1746–3214. The radio flaring seen in late 2018 was detected three times over 11 epochs of data from MeerKAT. Due to its position only a few arcminutes from the XRB H1743–322, more commensal observations of the flare star are likely if the XRB enters an outburst. H1743–322 has yet to enter any outburst in the 6 years since 2018, despite showing outbursts every 1-3 years between 2003 and 2018 (see e.g. Coriat et al., 2011; Shui et al., 2023). Nevertheless further commensal observations of SCR 1746 will provide further constraints on, for example, how often the star emits radio-loud flares.

MKT J174641.0–321404 is the second serendipitous, galactic transient discovered by the ThunderKAT team after Driessen et al. (2020), both of which are associated with kinds of active star. SCR 1746 is also the second M dwarf detected by MeerKAT (see likely quiescent emission from targeted searches of ThunderKAT commensal

⁴<https://github.com/BU-hammerTeam/PyHammer>

data in Driessen et al., 2021). These recent studies highlight the strength of highly sensitive, wide FoV observations taken regularly with instruments like MeerKAT.

As discussed in Section 1.2, the brightness temperature T_B of a source is a useful diagnostic of emission mechanism. Burgasser and Putman (2005) have converted equation 1.2 into a scaling form for radio emission of flux density F_ν observed at distance d such that

$$T_B = 2 \times 10^9 \left(\frac{F_\nu}{\text{mJy}} \right) \left(\frac{\nu}{\text{GHz}} \right)^{-2} \left(\frac{d}{\text{pc}} \right)^2 \left(\frac{L}{R_{\text{Jup}}} \right)^{-2} \text{ K} \quad (3.3)$$

where R_{Jup} is the radius of Jupiter. I have taken length scale L to be 2 stellar radii (radius $\sim 0.146R_\odot$ from the TIC; Stassun et al., 2019), the approximate size of M dwarf corona (Benz, Alef, and Guedel, 1995). For the brightest radio detection this produces a brightness temperature of $\sim 10^{10}$ K. As demonstrated by Pietka, Fender, and Keane (2015) one can relax the assumption of a length scale by taking the variation in a light curve as being indicative of the maximum size of the emitting region. That is, a source seen to vary on timescale Δt must have a maximum emitting size of $R = c\Delta t$ - this is a maximum as the radiating electrons will be moving at speeds less than c . Substituting this variability timescale into Equation 1.2 to calculate a value of T_B with fewer assumptions. From the lower panel of Figure 3.4, SCR 1746 varies conservatively by 0.1 mJy over a timescale $\Delta t = 3$ minutes, producing a minimum brightness temperature of 9×10^7 K. Incoherent, synchrotron emission is limited to at most $T_B \sim 10^{12}$ K due to rapid inverse Compton losses (known as the inverse Compton catastrophe, see e.g. Kellermann and Pauliny-Toth, 1969; Readhead, 1994), however no such limit exists for coherent mechanisms, which can have arbitrarily high T_B . The limits derived are all much less than 10^{12} K and so the mechanism could be either incoherent or coherent in nature. The clearest evidence for a coherent mechanism would be the presence of high degrees of circular polarisation, a measurement not currently available as the appropriate calibration for short-integration MeerKAT data not at phase-centre are still being developed. A factor of $\gtrsim 10$ brightness increase compared to the measured upper limits, lasting for timescales of at least minutes is also in keeping with incoherent radio bursts

(Osten, 2007). Furthermore, the magnetic activity seen by the measured $H\alpha$ EW and the approximate position of SCR 1746 on the Güdel and Benz relation (Figure 3.10) suggest that gyrosynchrotron emission is responsible for the observed emission. Further radio observations such as polarimetry, broader frequency coverage to constrain the spectral index, or simply a longer observation campaign will help determine the exact nature and frequency of the radio flares.

Optical photometry clearly indicates SCR 1746 is bright and red, with photometric relations suggesting a mid-M spectral type. This is in mild disagreement with the findings from the SALT spectrograph when run through PyHammer, which indicates a later spectral type of M8 (see Section 3.4 and Figure 3.11 therein), thus suggesting it falls under the ultracool dwarf designation. The properties of ultracool dwarfs beyond spectral type M7 depart from expectations set by earlier types (Berger, 2002; Burgasser and Putman, 2005; Berger et al., 2010), in the most extreme cases deviating from the Güdel and Benz relation (plotted in Figure 3.10) by four orders of magnitude (e.g. Berger et al., 2001). These ultracool dwarfs show a marked decrease in magnetic activity (seen in reduced relative $H\alpha$ flux) and relative X-ray flux, but an almost constant level of radio emission, i.e. the relevant coupling between wavebands appears to no longer hold in ultracool dwarfs. For example, Callingham et al. (2021) demonstrate that radio emission of 19 M dwarfs detected at 144 MHz deviate from the Güdel and Benz relation and are caused by coherent processes i.e. different to that detected for late-type stars at gigahertz frequencies and not related to their chromospheric activity. To provide a relative $H\alpha$ to bolometric flux density I employ the χ method of Walkowicz, Hawley, and West (2004). Calculating χ from SCR 1746's $(i - J)$ colour (Douglas et al., 2014), then $L_{H\alpha}/L_{bol} = EW_{H\alpha} \times \chi = (8.1 \pm 2.5) \times 10^{-5}$. This relatively strong $H\alpha$ flux is again in agreement with a spectral type M5-M6, before the breakdown of the radio-X-ray- $H\alpha$ couplings (see Figure 5 of Berger et al., 2010). Furthermore, the spectral indices of late M dwarfs are not well characterised - Kesseli et al. (2017) calculate their spectral indices using 7 M8 dwarfs to calculate, compared to a sample of 184 M5 stars - so it is perhaps not surprising that the empirical spectroscopic

analysis is in disagreement with other evidence. As such I conclude that SCR 1746 is of spectral type at least M5 but note that a later spectral-type may be possible.

The *TESS* monitoring of SCR 1746 (Figures 3.6 and 3.7) shows clear optical flares, along with a photometric modulation of relative amplitude 0.52 ± 0.01 at a period of 0.2292 ± 0.0025 days. I suggest that the photometric period observed is the rotational period of the star (as in e.g. Gizis et al., 2017) and can check this by comparing the inferred tangential velocity to that calculated by the broadening of spectral lines in the SALT data (Figure 3.11, inset). For example, Lane et al. (2007) reconcile the photometric period of a radio-detected M dwarf with its spectroscopically inferred rotational velocity, invoking magnetically induced spots as the cause of periodic variability. Spectral lines are expected to be broadened due to rotation and from their width I can estimate $v \sin(i)$, not accounting for the star's inclination i with respect to our line of sight. I can quantify the width of our line by measuring the full width at half maximum (FWHM) of the $H\alpha$ emission. I assume a simple relationship for rotational velocity to be

$$v \sin(i) = \frac{\Delta \lambda c}{\lambda} = \frac{\text{FWHM}_{H\alpha} c}{2 \lambda_{H\alpha}} \quad (3.4)$$

where the factor of 2 accounts for the effect of both blue- and red-shifted radial motion. Using the SALT data $v \sin(i) = 33.0 \pm 1.0 \text{ km s}^{-1}$, in excellent agreement with the photometric inferred velocity of $32.2 \pm 0.9 \text{ km s}^{-1}$ (from Section 3.3). This therefore agrees with the idea that the modulation is rotational in nature, is consistent with previous findings that fast rotators are magnetically active (e.g. Newton et al., 2017) and even implies a high inclination of the star with respect to our line of sight. Additionally, the flares in the *TESS* data tend to occur when the source is brighter or increasing in brightness (Figure 3.7), also in agreement with the light curve modulation being caused by large magnetically active regions rotating into view. A more thorough analysis of the spectroscopic broadening could be done by cross-correlating our spectrum with a set of synthetic spectra with known rotational broadening. However, that these two methods agree is post-hoc justification that a simple measure on a clear emission describes the physics adequately for this work.

The optical broad-band behaviour of SCR 1746 according to ASAS-SN (Figure 3.8) appears to show clear and regular optical flares, which again is to be expected from a mid-late M dwarf (Paudel et al., 2018). During the times of radio observations, no flares in either waveband are simultaneous to within 0.4 days of observations in the other. This demonstrates the value of simultaneous radio-optical observations to further constrain source properties, something now provided by the MeerKAT-MeerLICHT coupling. Furthermore, polarimetry, high resolution IR spectra and precise radial velocities may be of interest to further constrain the properties of SCR 1746, all of which could be provided by a southern hemisphere spectropolarimeter such as SALT or CRIRES+⁵. In depth multiwavelength monitoring is beyond the scope of this thesis but would be of particular note if any detections of exoplanets were made in this system.

3.6 Conclusions

I have presented the serendipitous detection of MKT J174641.0–321404, a radio flaring transient source discovered in commensal searches of images from the H1743–322 radio field taken by MeerKAT. Weekly images from the ThunderKAT survey have allowed us to detect and characterise this radio emission over 11 epochs and then perform multiwavelength searches. This was the first dataset that I used during my studies, allowing me to gain familiarity with tools such as the TraP) and develop investigative procedures for finding transients.

Archival data from ASAS-SN, MeerLICHT and *Gaia* show that MKT J174641.0–321404 is coincident with star SCR 1746–3214, located at a distance of ~ 12 pc. Archival photometry shows that the star is an M dwarf, showing red colours consistent with a spectral type M5. X-ray measurements have been combined with our radio data to place this M dwarf on the Güdel and Benz relation, indicating that the observed emission is likely gyrosynchrotron in nature.

Dedicated follow-up spectroscopy with the SALT instrument broadly confirms the star’s spectral type, near the transition between mid M dwarfs and ultracool

⁵https://www.eso.org/sci/facilities/develop/instruments/criages_up.html

systems. The observed spectra also showed clear $H\alpha$ emission lines, indicative of heating in the chromosphere of the star, as expected for magnetically active systems. Spectroscopy also showed that the emission lines of SCR 1746 were broadened due to the star's rotation, again an indicator of activity.

TESS observations of SCR 1746–3214 show clear stellar flares, whilst the observed periodic variability is rotational in nature and likely caused by the presence of magnetically active regions. The rotation rate seen in this photometry matches within uncertainties the broadening seen in spectral lines.

This is the second Galactic serendipitous transient found with MeerKAT, showcasing the unique capabilities of the current generation of radio telescopes for commensal science. Radio detections of stars such as this are a direct measure of particle acceleration in the atmospheres of stars, in contrast to the proxy measures of heating seen in, for example, X-rays and $H\alpha$. Furthermore, these kinds of detections can be useful probes of magnetic fields and electron densities around other stars, which are difficult to ascertain by other means. Untargeted searches with sensitive, wide-field observations such as with MeerKAT and ASKAP are therefore not biased in the same way as, for example, follow-up of known systems or multiwavelength searches. Of course some combination of studies at all wavelengths, focusing on in-depth study of known systems as well as the search for further examples, is the optimal way forward. These observations emphasise the necessity for multiwavelength association in understanding the nature of the transient sky, which has been one of the core drivers of the MeerLICHT project. This need for multimodal data will only become more pressing as our transient detection capabilities increase along with our data volumes.

I'm scanning all my radars

Well she said she's from a quasar

*Forty thousand million light years away**

— Jay Kay, from Jamiroquai's *Cosmic Girl*

4

Citizen science for radio transients

Contents

4.1	Introduction	74
4.1.1	Rationale	74
4.1.2	Citizen science in astronomy	76
4.2	ThunderKAT observations and pre-processing	82
4.2.1	Source extraction and subject generation	82
4.3	Citizen science platform	83
4.3.1	Beta testing and development	83
4.3.2	Final design	89
4.3.3	Data release	89
4.4	Results	93
4.4.1	Comparison to target sources	94
4.4.2	Previous commensal studies	96
4.4.3	Scintillation analysis	99
4.5	Counterparts and associations	102
4.5.1	Highlights	104
4.6	Discussion	111
4.7	Conclusions	115

*Kay does not specify if the distance described is the luminosity distance of said quasar, the co-moving radial distance, or indeed the proper (physical) distance. If this is the co-moving distance then this 40 Gly measurement corresponds to a redshift of $z \sim 54$ for this quasar, where I have assumed a flat cosmology of $H_0 = 67.37 \text{ km s}^{-1}/\text{Mpc}$, $\Omega_M = 0.3147$ and $\Omega_\Lambda = 0.6853$ (Planck Collaboration et al., 2020). This redshift is well beyond humanity's telescopic reach and indeed is likely before we expect quasars to have formed. We can therefore assume, given that neither the eponymous Cosmic Girl nor Jay Kay have (yet) disclosed information pertaining to the most powerful telescope in the world, nor any novel theory on quasar formation in the early universe, that the quoted distance is that which was scanned by “all [his] radars” i.e. the luminosity distance observed by a radio telescope. Under the same cosmological assumptions as

This chapter presents the development of and results from the first citizen science project dedicated to radio transients. This work was published in Bursts from Space: MeerKAT - the first citizen science project dedicated to commensal radio transients (BfS:MKT; Andersson et al., 2023) and demonstrates both how well citizen scientists can find transients and how keen they are to do so. Fitting in this piece of the transient-finding puzzle holds great promise for further study and broader/deeper observations. Further data releases have been launched on the Zooniverse website, of which analysis is ongoing.

4.1 Introduction

4.1.1 Rationale

In Chapter 3 I demonstrated how MeerKAT can be used to find interesting and rare transients. Crucially it is the combination of sensitivity, regular cadence and wide field of view (FoV) that allow for commensal searches for these known transient phenomena and new classes of objects as yet undiscovered.

As discussed in section 1.4, previous investigations from both MeerKAT and other instruments have found that fewer than 5% of point sources are transient or variable at 1.4 GHz (Ofek et al., 2011, and references therein), with source classes spanning a wide range of timescales and physical processes. The majority of radio variables found are active galactic nuclei (AGN; see e.g. Thyagarajan et al., 2011), whose variations can be attributed primarily to refractive scintillation (Rickett, 1990). While these AGN dominate samples of variables, active or flaring stars have been found in untargeted radio surveys (Chapter 3, Mooley et al., 2016; Driessen et al., 2020; Andersson et al., 2022), as have supernovae and GRB orphan afterglow candidates (Levinson et al., 2002; Gal-Yam et al., 2006). Pulsars can vary intrinsically in the image plane - indeed some of the slowest known pulsars

before, this corresponds to a redshift of $z \sim 1.61$ and so I conclude that the titular cosmic girl must have travelled from a galaxy observable by modern quasar surveys. Nevertheless, that such deep radio observations were available “on a Saturday in 1999” is of great interest, and perhaps concern, to organisations such as the SKAO. Finally, how the song’s subject arrived on Earth from such a distance is never clarified, with only a brief mention of a teleportation-ready “transporter” in the source material - I defer this topic to those more familiar with quantum entanglement.

are discovered in imaging data (Tan et al., 2018; Caleb et al., 2022). Furthermore, scintillation through the interstellar medium can cause both short timescale, large amplitude variations and lower amplitude variability occurring on timescales of hours to years for point-like sources (see section 1.3.3.2 and Rickett, 2001; Hancock et al., 2019). There are also numerous accounts of radio transients being discovered without clear progenitor systems or multiwavelength counterparts (Bower et al., 2007; Stewart et al., 2016; Murphy et al., 2017). These include the elusive sources near the Galactic centre (Davies et al., 1976; Zhao et al., 1992; Hyman et al., 2005; Chiti et al., 2016; Wang et al., 2022). The serendipitous discoveries, elusive nature and broad physics at play in this zoo of radio transients all point towards the need for new searches and the development of novel methods to maximise the science yield of our observations.

Despite these methods of searching for radio transients bearing fruit (see Chapter 3), they are not optimal. Firstly, the volume of data to analyse is far greater than any one person can achieve by eye on reasonable timescales. In the 5 years of its operation, ThunderKAT observed over 50 XRBs in total at weekly cadence, typically following each source for over a month. The resulting archive consists of over 400 observing blocks, each of which followed 1 – 6 XRBs in any given week, totalling over 100 TB of raw data to reduce and analyse. In this chapter I looked at only the available data as of December 2021, consisting of over 500 final images. Each image then contains of order several hundred sources, from just a single 15 minute observation. ThunderKAT also had memoranda of understanding with many of the other Large Survey Projects on MeerKAT, such as LADUMA (Blyth et al., 2016), MIGHTEE (Jarvis et al., 2016) and MHONGOOSE (de Blok et al., 2016), to use their data commensally. As a result there are thousands of observations in the growing archive, in which transients may reside, probing right down to 1σ sensitivity limits $\sim 1\mu\text{Jy}$ (Heywood et al., 2022a). These data overload issues are only exacerbated when imaging on shorter timescales, including down to the 8s integration time of MeerKAT, as has been explored somewhat within ThunderKAT (e.g. Chastain et al., 2023; Fijma et al., 2024). Searches with telescopes such as

the MWA (Lonsdale et al., 2009) and MeerKAT have now started to find variable and transient sources active on timescales between the limits of pulsar searches and image-plane transients searches, with periodicities of a few minutes (Caleb et al., 2022; Hurley-Walker et al., 2023), demonstrating that when searches are conducted on new timescales there are interesting things to be found.

Furthermore, a common approach used to search for radio transients is calculating the two variability statistics η and V (see section 2.2.1 for more detail). V is the standard deviation of a set of flux densities relative to their mean, while η quantifies statistical significance of variability as a reduced chi-squared relative to a constant source. These statistics themselves are correlated and anti-correlated with observed brightness respectively. The chi-squared is proportional to the sum of fluxes divided by their statistical uncertainties - i.e. the signal-to-noise ratio - squared. For V it is only possible to detect smaller deviations clearly for brighter sources and so this statistic is anti-correlated with observed brightness. As a result, these two statistics are not ‘silver bullets’ for finding all the transients in a given dataset. Similarly, Driessen et al. (2022) have shown how the sampling rate of underlying variability can alter the calculated η and V . Furthermore, radio observations are not free from false positives. Two causes of these false positives are the non-Gaussian artefacts that typically occur around bright sources in radio images, and the changes in the point spread function (PSF, or restoring beam) caused by differing elevations over a set of observations, which induces non-intrinsic variability in the measurements of resolved objects. As these issues might plague only one observation in a dataset, they can lead to measurements easily confused for bona fide transients by automated methods.

In order to overcome these limitations and biases in searches, I have explored the application of citizen science to this problem, as is detailed below.

4.1.2 Citizen science in astronomy

The involvement of amateur (unpaid) astronomers with the professional community has been longstanding, dating back at least 300 years to when Edmund Halley called for observations of the 1715 total eclipse over central England (Halley, 1716).

The map seen in Figure 4.1 shows the path of totality of the eclipse, which was updated based on timing measurements made by at least 25 curious volunteers. It is only because these volunteers of 1715 were numerous, distributed and observant that such a map could be created, all of which are qualities which persist and relate to the role of citizens in modern citizen science for astronomy. For a review of citizen science and how volunteers contribute to astronomy, Marshall, Lintott, and Fletcher (2015) provides a good overview as of the mid-2010s.

One of the largest contributions that volunteers make to astronomy is in observations which are made at high cadence, across the globe, and with minimal delay. One of several examples is the American Association of Variable Star Observers (AAVSO) network, whose global distribution of volunteers monitor variable stars and their transient episodes, including outbursts from recurrent novae stretching back over 60 years (de Ruiter et al., 2023; Nyamai et al., 2023). Similarly, the discovery of supernovae has long been supported by citizens, with significant contributions being made for brighter, nearby systems. For example, SN 2023ixf was first reported by prolific supernova hunter Kōichi Itakagi in May 2023, approximately 12 hours before any professional observing team (Itakagi, 2023). 2023ixf has gone on to be one of the most interesting supernovae of the past few years as, due to its proximity in M101, sensitive and high resolution data have allowed for both investigation of the progenitor system that existed pre-explosion (e.g. Kilpatrick et al., 2023) and the resultant post-explosion interactions with the system's circumstellar material (Jacobson-Galán et al., 2023). It is also worth noting that even infrequent observations can result in significant contributions to astronomy, one good example of which is Thomas Bopp's accidental co-discovery of comet Hale-Bopp when looking at globular clusters through a friend's telescope (Bopp, 1997).

Volunteers are, of course, capable of making contributions to astronomy beyond taking data. I have already mentioned how Karl Jansky's discovery of astronomical radio waves has contributed to the development of radio astronomy as a field. The advent and spread of the world wide web has been transformational in this regard, allowing for the distribution, classification and analysis of astronomical data



Figure 4.1: A map of the path of totality for the 1715 eclipse, as compiled by Halley and Senex (1715) updated from his initial prediction based on observations from citizen scientists. Edited from the Institute of Astronomy Library's copy of this map, accessible at <http://www.dspace.cam.ac.uk/handle/1810/221308>

amongst thousands of people. Some of the earliest examples of this are the use of idle compute time to search through archives for signals in, for example, SETI searches and gravitational wave data (Abbott et al., 2009; Werthimer et al., 2001). Beyond this passive compute distribution, citizens started making active contributions over the internet with the *Stardust@home* project, in which volunteers would scan through sample return images obtained from Comet Wild-2 as part of the Stardust mission in an effort to identify interstellar material (Westphal et al., 2006; Méndez, 2008). This work directly inspired *Galaxy Zoo* (Lintott et al., 2008) which has, in turn, spawned a generation of people-powered research since the start of the 2010s.

Galaxy Zoo exemplifies many of the online citizen science efforts of the past 15 years. The initial problem is that the amount of data collected by a science team - in this case from the Sloan Digital Sky Survey (SDSS) - is much greater than individuals or small groups of professionals can hope to classify on reasonable timescales. For the original *Galaxy Zoo* and its subsequent incarnations, the classification task relates to the morphology of galaxies (e.g. as spirals or ellipticals as done by Hubble, 1936), but for other projects this might consist of, for example, identifying the animal species present in a picture. Distributing this classification task to many ($> 10^5$) volunteers, who have been primed with a tutorial and the necessary introductory material, allows them to make contributions to the project. Typically several volunteers observe each subject¹, with the final classification being some average over individual results. In some cases, this aggregation is weighted by some scheme, though defining exactly what constitutes a ‘good’ classification is inherently subjective and practitioners must be careful to honour the effort put in by many volunteers to their science. One reasonable strategy is to up-weight based on consensus (Lintott et al., 2008; Schwamb et al., 2012) where volunteers whose votes agree with the majority of voters are given more weight, though such consensus might not agree with expert opinion. Another option is to introduce training or simulated data for which the ground truth is known and feed these into a live classification scheme. This allows for both assessment of and also feedback on an

¹For astronomy a subject is typically one astronomical source such as a galaxy.

individual's classifications - the current Zooniverse platform allows for volunteers to be notified if their choice and the training label do not agree, which has the potential to increase that user's understanding and performance. Once the final votes have been aggregated, the project scientists can then go on to use classifications for further science and, increasingly for more recent projects, train machine learning models.

As a final note, individuals looking at subjects is useful for more than aggregating classifications. In some cases, anomalous subjects have been identified that create whole new avenues of research. Sticking with *Galaxy Zoo*, one example of this is 'Hanny's Voorwerp', an extended region of gas ionised by the now-faded AGN of IC 2497 (Lintott et al., 2009). This was spotted by schoolteacher Hanny van Arkel and generated further research into emission-line clouds around low-redshift galaxies (Keel et al., 2012). A second serendipitous find from *Galaxy Zoo* is the discovery of 'Green Peas', a class of compact galaxies with extremely high star formation rates (Cardamone et al., 2009). These Green Peas, named due to their strong [OIII] emission lines, are thought to be low redshift analogues of Ly α emitters and thus constitute local laboratories for understanding extreme star formation processes. Since this first incarnation of *Galaxy Zoo*, further datasets have been presented to volunteers, using data from e.g. the Dark Energy Camera, Hyper Suprime-Cam and the James Webb Space Telescope. These have formed over 70 *Galaxy Zoo* papers, with a total of \sim 450 articles mentioning the project in their abstracts, demonstrating the impact such a sustained citizen science effort can have.

Since *Galaxy Zoo*, a whole ecosystem of crowd-sourcing citizen science projects has been developed on the Zooniverse platform². The scope of these projects goes well beyond astronomy, with projects in ecology, history and cell biology to name but a few. To date, over 2.5 million individual volunteers have made annotations, classifications and further contributions to citizen science via the website. There are over 450 publications listed on the platform as of May 2024, including databases and Astronomers Telegrams³. The impact of the Zooniverse goes much further than these publications - an example of which comes from the *Penguin Watch*

²<https://www.zooniverse.org>

³<https://www.zooniverse.org/about/publications>

project, which collects data on penguin colonies from time lapse photography. The data collected from this project have directly impacted government policy in South Georgia and the South Sandwich Islands, for example with the enhancement of a no-take fishing zone around the foraging sites of chinstrap penguins as part of a wider Marine Protected Environment (Clubbe et al., 2018). Finally, Zooniverse projects are fantastic tools for engaging with communities about science (Masters et al., 2015). In Appendix A I describe some of the people I have had the privilege of working and sharing science with during the course of my studies.

In this chapter I describe the first citizen science project dedicated to finding radio transients in images. This investigation was motivated by the traditional difficulties in performing such widefield radio transient searches, as well as the successes of previous transient-related Zooniverse projects. For example, Wright et al. (2017)'s *Supernova Hunters* combined a neural network with human classifications to outperform either classifier alone and is still discovering supernovae, over six years since launch⁴. Similarly, *Citizen ASAS-SN* users have discovered >10,000 new variable sources that are not present in the existing star catalogs of the southern hemisphere (Christy et al., 2022). Finally, the original *Bursts from Space* project is dedicated to the discovery of low-signal-to-noise FRBs from the CHIME telescope. CHIME only saves FRBs above a S/N of 8.5, due to data storage constraints. *Bursts from Space* presented FRBs below this cut to volunteers and has uncovered sources with dispersion measures up to twice as far as those published on the Transient Name Server as of early 2021 (Walmsley, 2021). The rest of this chapter details the development and first analysis of the *Bursts from Space: MeerKAT* (BfS:MKT) project, as described by Andersson et al. (2023). The aims of this project are to discover new transients, eliminate spurious false positives and provide complementary analysis to other commensal search methods, as well as allowing us to assess the viability of further citizen science work.

⁴see <https://www.wis-tns.org/object/2022aeee>

4.2 ThunderKAT observations and pre-processing

The observations used as part of the first data release consist of a subset of ThunderKAT XRB images, as presented in Chapter 2. These were based on which datasets were available at the time of the public launch of the project which contained enough epochs to form reasonable light curves - more than 3 observations. The observations used in this work were taken between mid-2018 and late 2021. Generally, the observing strategy is determined by reports from X-ray facilities of activity from an XRB, which is then observed at weekly cadence by ThunderKAT in 15 minute images. It is important to remember that the commensal nature of this work constrains us to whatever observational cadence was used for monitoring the XRB.

ThunderKAT data were typically reduced using OxKAT (Heywood, 2020), a semi-automatic set of scripts that perform calibration, flagging and imaging of MeerKAT data. The commensal nature of this work means that, due to different science requirements and observational conditions for each dataset, the resultant images are heterogeneous in their properties, although mostly homogeneous within a particular field. For specific details on how each set of XRB data is reduced I refer the reader to the relevant paper as listed in Table 2.1.

4.2.1 Source extraction and subject generation

Each set of images was processed using the TraP (Swinbank et al., 2015). As a case study of using the TRAP was presented in section 2.2.1, I will summarise only the main points here. The TRAP finds sources above a user-defined threshold in a set of astronomical images, creating light curves of each unique source and calculating statistics for each source. This is done by fitting a Gaussian component to each source in every epoch and associating it with those found at that position in all previous images, updating the database as new observations are added. Most TRAP parameters are kept at their default values (see Listing 2.1). The `detection_threshold` was fixed at 8 throughout. The `expiration` i.e. the number of force fits to a position where a source was found in a previous epoch, was always kept at greater than the total number of observations in a dataset, meaning wherever

a source had been found, a light curve with data points for all remaining time steps was created. I am interested in unresolved, point sources and their variability, so I set all source fits to be fixed at the size of the PSF via `force_beam = TRUE`. For extended regions of emission, the change in size and position angle of the PSF between observations can lead to non-intrinsic variability measurements, as discussed in section 4.3. To allow for deblending I set `deblend_nthresh = 10`, which accounts for overlaps between nearby sources such as double lobed radio galaxies. Finally, the `extraction_radius_pix`, describing how far ‘out’ from phase centre in the image to search for sources, was always kept to approximately $1.5\times$ the main lobe of the primary beam, in line with Sarbadhicary et al. (2021).

Once all observations had been processed by the TRAP, I generated figures and uploaded them to the *BfS:MKT* page on the Zooniverse. For each of the 8874 unique entries to the TRAP database (every `runningcatalog` source), two types of figures were made - a light curve and postage stamps of the sky around a source. The local sky figure is a square arcminute image centred on the weighted mean Right Ascension and Declination (RA,Dec) of each source. These figures were then uploaded to the Zooniverse platform along with some basic metadata (RA,Dec, median flux density and the time stamp of the highest S/N observation), creating the subjects for citizen scientists to classify. For beta-testing, all three of these workflows were tested. Following feedback from volunteers discussed in section 4.3.1 below, the final workflow consisted of just LC + Best Detection.

4.3 Citizen science platform

4.3.1 Beta testing and development

The process of launching a citizen science project on the Zooniverse involves first building the appropriate website material using the online building tools, which is then subject to an internal review by the Zooniverse team. Once feedback has been given and acted on, the project then undergoes beta-testing with selected volunteers. I tested three workflows, consisting of different figures being presented to volunteers. These were:

1. Light Curve and one representative sky figure, using the image where the source was detected at the highest S/N. I called this the **LC + Best Detection** workflow.
2. Displaying sky figure cut outs for all observations, for volunteers to step through and observe variations (or lack thereof). This was labelled the **Image Plane Only** workflow.
3. Both the light curve and all cut outs, referred to as the **All Data workflow**.

Internal review from the Zooniverse team was completed by 2nd August 2021, with a few changes needing to be made. These included adding examples of interesting transients already found by MeerKAT to demonstrate the kinds of sources that might be found. I also added a clarification section explaining, for example, that the singular image used in the LC + Best Detection workflow is the highest S/N detection of the source, or what to do for certain edge cases such as light curves with very few data points. The question tree in these initial stages consisted of

- Q1: Real or Artefact?
- Q2: If Real, is the subject a stable point source, transient/variable, an extended blob or are you unsure?

There are therefore five classes to which a source can ultimately be assigned, examples of which can be seen in Figure 4.2.

- **Stable** sources are unresolved, point sources whose light curves are judged to be (within uncertainties) consistent with flat.
- The **Extended Blob** classification is intended to catch the resolved, extended sources, regardless of variability. The changing size and angle of the PSF between epochs introduces non-intrinsic variability, particularly in the case of extended sources and so volunteers were instructed to classify subjects they deem to be resolved as Extended Blobs - making use of the local sky figure, comparing the source size to that of the PSF in the lower left corner.

- The **Artefact** classification was implemented to account for any spurious, non-astrophysical sources that may be present in the images.
- **Transient/Variable** classifications are those I am searching for, which are point sources with variable light curves.
- Finally, if volunteers are uncertain if a subject fits into any of these classes - either due to visualisation issues, their own interpretation or anything other reason - they can say they are **Unsure**. I included Unsure to assess the confidence of volunteers - if a subject does not clearly fit into one class this will be seen quantitatively (not just in e.g. the Talk board). Also, without an 'unsure' option, volunteers may have settled for classifying as either stable or transient, leading to an under- or over-prediction of interesting sources.

Beta-testing was launched on 7th September 2021 and lasted until early October, with respondents providing classification on the subjects and, perhaps more importantly, giving feedback on the project. The dataset used came from the J1858 field (see Table 2.1) simply because that was data I had already processed at time of beta-testing and I knew this field had examples of all relevant classes. I received feedback from 17 volunteers on the three workflows and project design, though note that not all responses were mandatory and not all volunteers tested all workflows, resulting in heterogeneous response numbers. From this feedback, some of the most important things to consider are how easy/difficult volunteers found the different workflows and how satisfied they were with each workflow. It is easy to see how projects featuring tasks that are easy but not satisfying would be less likely to gain traction with volunteers. Similarly, even if tasks are satisfying, if they are long and difficult, that also increases the barrier of engagement with the project and results in a fewer contributions per unit time.

Firstly, there were responses to questions about the project as a whole. In general, the project was well received, with 13 of 17 respondents stating that the project goals are clear, 15 felt the questions made sense and 13 said that they would take part in the project were it to launch. Most volunteers considered questions

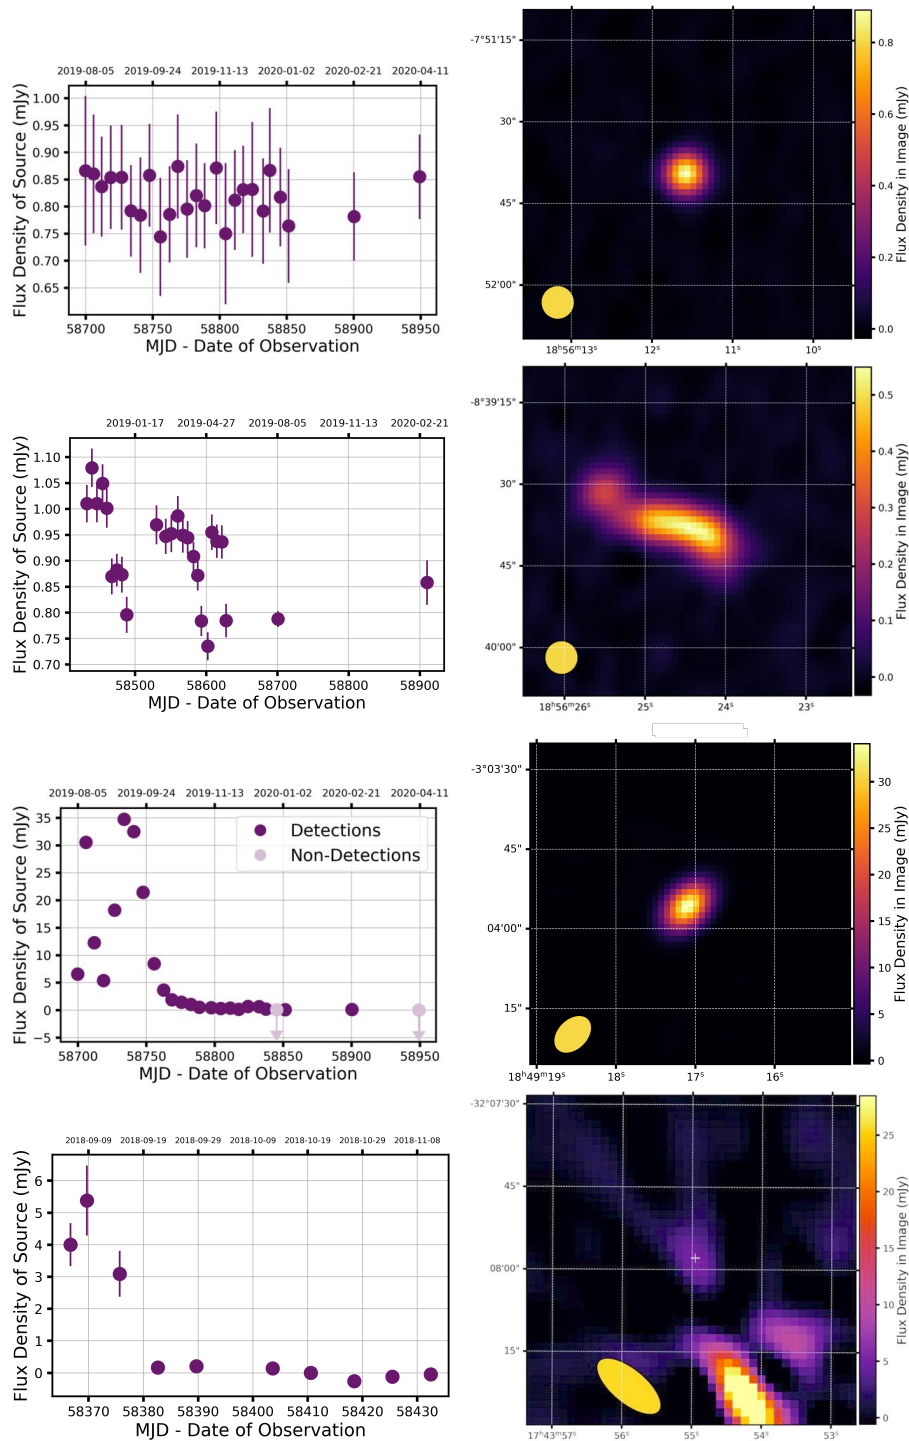


Figure 4.2: Examples of the four observational classes within the workflow on the Zooniverse, showing both the light curve and an image of each. From top to bottom these are: **Stable** - no variation in the light curve given the error bars and a point source; **Extended** - variations caused by changes to the PSF and a source that is larger than the beam (lower left); **Transient** - a clear variable light curve for a source the same size as the PSF; **Artefact** - a spuriously transient light curve and a faint source on the outskirts of a very bright object, with non-Gaussian noise structure. The final class, **Unsure**, by definition has no archetypal characteristic no figure is shown here.

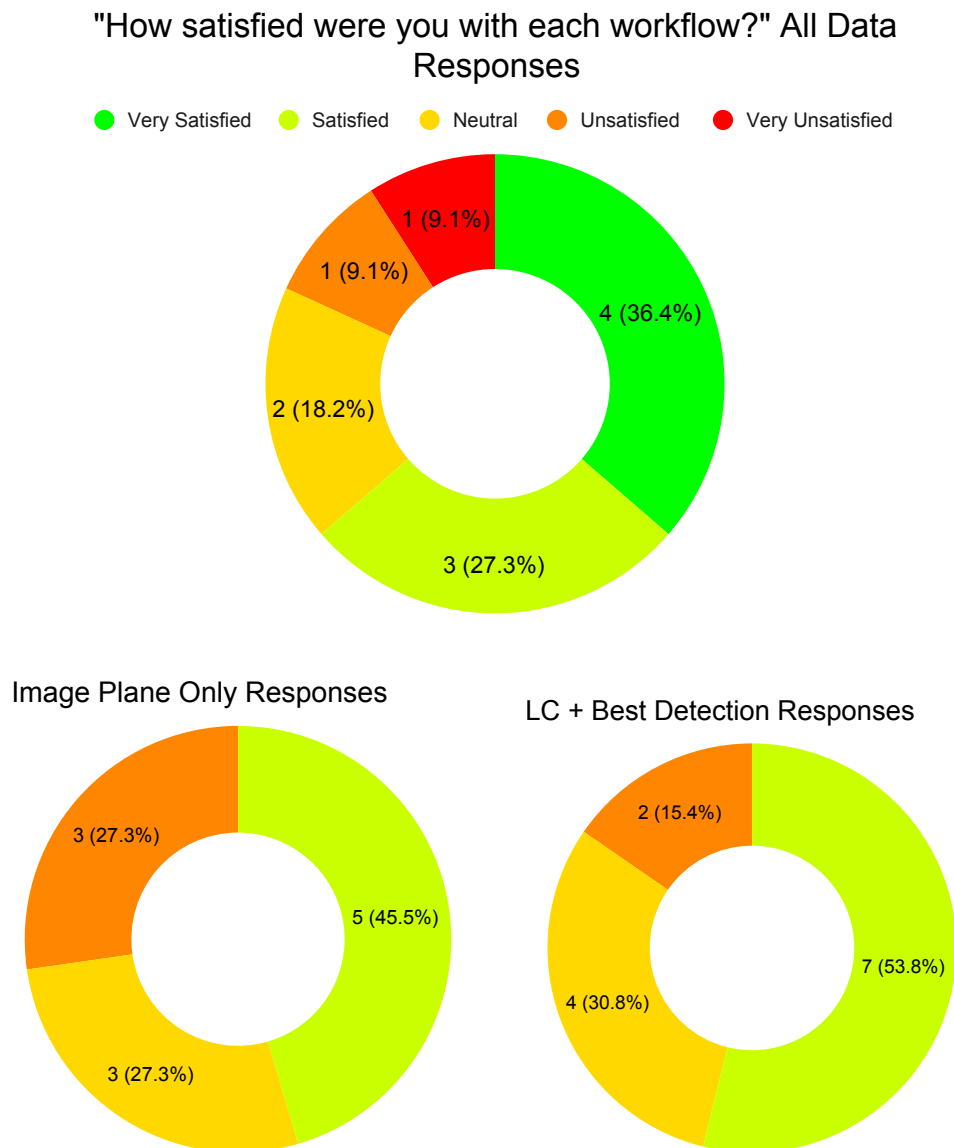


Figure 4.3: Volunteer responses to each of the three tested workflows in the *BfS:MKT* beta test.

to be moderately easy or somewhat hard (i.e. minimal responses saying ‘very easy or very hard’). Those who found the questions moderately hard felt they were generally too long or complicated. The solution to this was, as suggested by a volunteer, to collapse to a one-step workflow in which we simply ask what class the subject belongs to, as seen in Figure 4.4. All other general improvements were clarifications and tweaks to wording.

Turning now to workflow-specific responses, we can directly compare how

satisfying volunteers found each of the three layouts in Figure 4.3. Firstly, the All Data workflow had the most respondents (11), with 7 saying the project was either satisfying or very satisfying. For the Image Plane Only workflow, only 3 of 11 volunteers found this setup satisfying, with as many saying they were unsatisfied. The simplest workflow, LC+Best Detection, received 7 satisfied responses out of 13 volunteers. We can see immediately that the middle ground of showing Image Plane Only was enjoyed the least by volunteers. I therefore decided that the Image Plane Only workflow was unsatisfactory and did not pursue it further.

Between the LC+Best Detection and All Data workflows, the latter was slightly preferred by volunteers who tested it (more ‘very satisfied’ responses), although the number of responses is low. Deciding which of these to go for was slightly more difficult. I could have launched both, but then would have to make some decision about which data to feed into which workflow. I could have applied both workflows to all data, however this would be duplicating volunteer effort for the same fundamental science aims and so was to be avoided. Volunteers who responded having done the All Data workflow and said that the task was ‘Somewhat hard’ (4 testers) were asked why, with responses revolving around the subject being too complicated and taking too long to classify. Indeed, stepping through 26 images of a source to assess its variability was much more involved than directly using the light curve. This is reflected in the number of classifications made - for the LC+Best Detection workflow 752 responses were received during the beta test, whilst for the All Data scheme there were only 154 responses. Finally, as mentioned above, the data tested were from the *Swift* J1858 field, which has 26 epochs, whilst other datasets in Table 2.1 have many more observations than this. Requiring volunteers to look at 77 images in order to make a single classification is clearly a large barrier to making contributions. It might have been possible to split a single, long subject, into multiple sub-subjects, though how to do this for all the heterogeneous fields is not immediately apparent and again this somewhat duplicates volunteer efforts. Due to these considerations with respect to volunteer effort, classification time and difficulty, I decided to only go ahead with the LC+Best workflow, which is

detailed below. I will note, however, that it might be possible in the future to make a two-stage system wherein volunteers classify on the LC+Best Detection workflow and then candidate variables and transients (e.g. those where $> 50\%$ of respondents think something is a transient) are funnelled into the All Data workflow. This has not been done at this stage as the results from the simplest workflow already exceeded previous transient-finding efforts.

4.3.2 Final design

Taking on board feedback considerations from the beta-testers and the internal review team resulted in a single workflow in the final project design, with only one question. The final LC+Best Detection workflow can be seen in Figure 4.4. Volunteers are given a tutorial to familiarise them with the data and describe the figures shown in the project, as well as accounting for common pitfalls due to figure processing - namely there are a few visualisation issues that make classifications more difficult, discussed further in section 4.6. There are also description pages for the project detailing the team, the telescope and the kinds of objects we are searching for so that citizen scientists can learn more about astronomy and the work the ThunderKAT team do.

There is also a Field Guide (accessible via the main project interface) with examples to demonstrate the type of subjects intended for each class, as well as some help text describing the rough thought process behind each source. If volunteers feel a subject is particularly interesting, or they have questions on a particular source, they can create individual threads on the dedicated Talk forum for the project, where (citizen and project) scientists can discuss.

4.3.3 Data release

The first data release for the *Bursts from Space: MeerKAT Zooniverse* project ⁵ launched on 7th December 2021 with classifications concluding by early March 2022. During this time 1038 volunteers classified the sources using the workflow

⁵<https://www.zooniverse.org/projects/alex-andersson/bursts-from-space-meerkat>

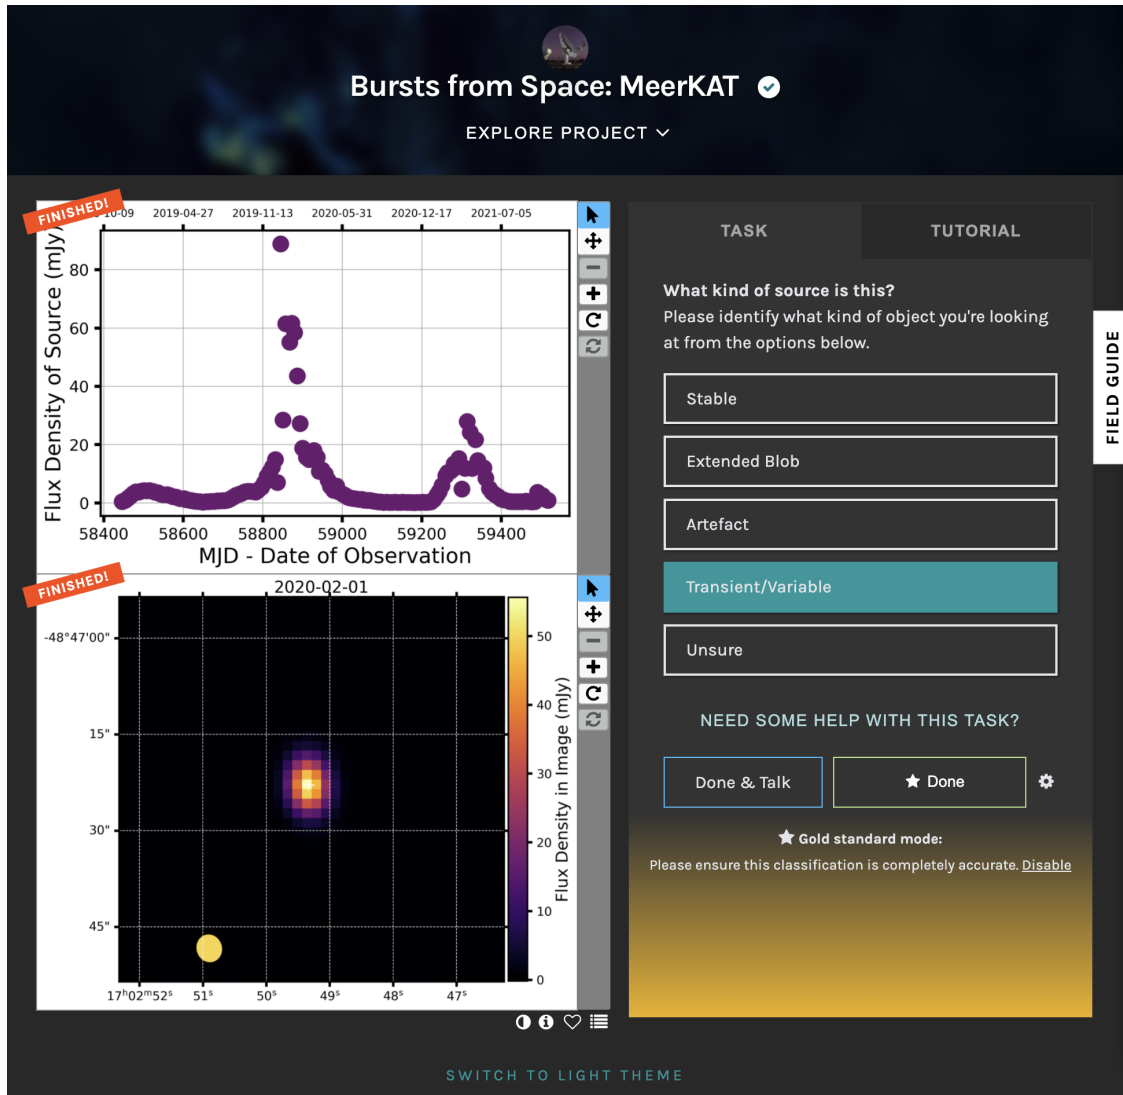


Figure 4.4: Classification workflow for *Bfs:MKT*, showing the light curve and local sky figures for GX339–4.

seen in Figure 4.4. Ten volunteers were required to classify a subject before it was considered classified, resulting in a total of 88740 classifications. These were classified over 90 days, or an average of 1 submission per 1-2 minutes. The histogram of all classifications for the project to date can be seen in Figure 4.5, showing the expected steep power law distribution of votes (Spiers et al., 2019). Many volunteers perform only a few classifications and a few volunteers perform many classifications. We can also see a number of ‘super users’ who have classified thousands of sources each. The median, mean and standard deviation of user classifications are 4.5, 86 and 490 respectively. The chosen retirement value of 10 is enough to mitigate outlier




Figure 4.5: Histogram of the number of classifications made per user on the *BfS:MKT* project. As might be expected, many users perform only a few classifications, whilst a few volunteers dedicate thousands of votes to the project.

votes, but not so high that it would take many months for a single subject to be fully classified. I note here that if further iterations of the project gain as much traction as this first batch of subjects I might be able to determine what the optimal trade off could be. Finally, the median and mean times for a classification to be made were 7 and 22 seconds respectively though I note that these are server times and could be affected by, for example, volunteers leaving their computer. Extrapolating from the mean, the total time spent is equivalent to 0.26 years of full-time employment.

Simple aggregation is performed for this one question workflow, using the standard aggregation scripts⁶, where I take the Boolean values for each classification and sum over every vote for each subject. This gives us 10 votes for a given source, from which I can calculate fractional classifications and determine how many subjects are deemed to be transient/variable by some number of citizen scientists (TF = transient fraction). In this process I chose not to weight volunteer classifications at all. This was partly motivated by the findings that most volunteer weighting schemes are heavily peaked at 1 (i.e. the same as no weighting applied; e.g. Lintott et al., 2008; Schwamb et al., 2012; Willett et al., 2013; Willett et al., 2017; Simmons et al.,

⁶<https://aggregation-caesar.zooniverse.org/index.html>

2017). These cited schemes were based on consistency, meaning that volunteers who disagree often with the majority of voters are down-weighted. However, this in effect reduces the number of classifications for a retired subject and so makes comparisons more difficult. Another way of weighting volunteers is by using some form of gold-standard data. Indeed, there do exist ‘correct’ classifications for the known variable and transients in the datasets (XRBs, for example) so it might have been possible to judge performance based on these. However, there are only 45 known variable sources in these data out of almost ~ 9000 light curves, and so relatively few volunteers would ever be voting on these. This is only compounded when considering the power law distribution of votes. A few of the highly contributing volunteers may end up voting for several of the known variables, but very few of the large majority of the volunteers would ever see a known variable. This would make it difficult to assess those who had not seen a known system and would make it perhaps unfair to judge only those who were spending more time engaging in the project. Finally, many of these known variables come from one single field, as found by Driessen et al. (2022). As mentioned earlier, the data are very heterogeneous, with different properties across all 11 observation fields. This adds yet another layer of complexity in that I would be extrapolating the potential votes from few volunteers on variables which may not be representative of the findings in different datasets. As a result, no weighting scheme was applied.

I set a threshold of 4/10 volunteers classifying a subject as transient/variable, reducing the sample from 8874 to 381 sources. This 0.4 threshold was chosen as a trade off between having many sources to vet and missing some low vote fraction variables. In essence, this provided a filtered list in which I could find transients at higher rates than without volunteer contributions. These 381 subjects were visually inspected by me to confirm or reject each source. This final vetting reduces the number of transient candidates to 168 sources (i.e. a true positive rate of $168/381 = 0.44$). When testing higher thresholds I agreed with a large majority of citizen classifications, meaning I was likely losing some below the cut, whilst a lower cut of 0.3 doubles the number of candidates to vet. Reasons for disagreement between

citizen and project scientists are numerous and include how much volunteers make use of error bars, PSFs and other parts of the figures, as well as the existence of systematic variability that is present in earlier MeerKAT observations (see e.g. the appendix of Driessen et al., 2022) that would only be noticeable to experts who have compared many light curves in a given field. I will discuss what information is obtainable from the false positives in section 4.6.

4.4 Results

Using citizen scientists to inspect a wealth of data from the MeerKAT telescope produces 168 variable and transient radio sources which have also been vetted by project scientists. Of these 168 variables and transients found by volunteers, 142 are not detailed in previous commensal or XRB work. This constitutes the largest sample of radio variable candidates to date from a single project, with their positions are listed in Table B. This table provides the TF, median 1.28 GHz flux density and η and V statistics described in section 2.2.1 for each of the sources, as well as the date on which they were detected at highest S/N.

In order to characterise these sources, I can examine the variability plane defined by η and V (equations 2.2 and 2.3) seen in Figure 4.6. As expected, the structure seen in this figure consists of many sources at low significance and amplitude, some spurious artefacts at high V and generally bright objects detected at high significance (η) but low variability amplitudes. Systematic differences between datasets due to their heterogeneous sampling and imaging are also present (e.g. the sources not voted as transients at $\log(V) \sim -0.6$). The most important thing to note is that there are many objects flagged by the project whose parameters are ‘normal’ i.e. non-anomalous. This means that, were one instead to take outliers above some $n\sigma$ threshold in η and V , sources could have missed that, upon individual inspection, appear variable or transient. So citizen scientists can find interesting variable radio sources that could have been missed by other techniques or without detailed analysis (see the other ThunderKAT commensal work for examples; Driessen et al., 2022; Rowlinson et al., 2022; Chastain et al., 2023). The previously known transient

sources (i.e. the circled sources of Figure 4.6) with high η and V are almost all recovered, as are many of those found in previous studies (see sections 4.4.1 and 4.4.2 for more detail), meaning that volunteers are able to recover or discover interesting sources across a range in statistical parameters. So our volunteers have been able to analyse large data volumes in just a few months and produce the largest sample to-date of variables from a radio telescope.

I can use these discovered (142) and recovered (26) variables, along with the other known transients (19) in our fields, to estimate that at least 2.1% of radio sources detected by the TRAP at 1.28 GHz with our approximately weekly observations are transient or variable for typical sensitivity limits of $\sim 30\mu\text{Jy}$. This rate is in line with previous work on radio transients - Driessen et al. (2022) find $\sim 2\%$ of sources in the GX339-4 field to be variable and Sarbadhicary et al. (2021) found $\sim 1\%$ of sources in the COSMOS field to be seen as variable, as seen by the VLA. It is worth noting that these different studies, including this work, probe slightly different timescales, though all of them ‘long’ (\geq a week). This means that our project can find comparable transient rates as previous work. This 2.1% is a lower limit on the number of variables and I can estimate what fraction of variables/transient are missed by assessing which known transients are not recovered. The fraction of previously known variables not recovered is 19/45, implying the true amount of variables in our fields could be as high as $\sim 5\%$. However, this estimate is only valid under the assumption that new variables are recovered at the same rate as the known transients in our field. In section 4.6 I discuss the selection effects that are evident in the different kinds of light curves selected by volunteers when compared to the known transients in our field.

4.4.1 Comparison to target sources

Of all the 8874 sources classified in this project, there are 45 known variable/transient objects published in the literature, including the 11 XRBs listed in Table 2.1. Of the 11 XRBs, 9 are classified as transient by citizen scientists. The only two that are missed are *Swift* J1858 and SAX J1808, whose light curves can be




Figure 4.6: Variability plane for the 168 sources identified by citizen scientists to be variable, along with those they find to not be so in grey. The outer histograms are the marginal distributions in η and V of volunteer identified variables (purple) and everything else (grey), normalised such that the area under each curve sums to one. The colour bar denotes the fraction of classifications as a transient/variable source. Known variable sources (e.g. XRBs) in our fields are circled. Imaging artefacts appear at low η and high V , whilst flux calibration uncertainties can produce high η , low V sources (due to lack of systematic uncertainty in equation 2.2). Most known transients are found by citizen scientists, whilst many new sources are identified and show a wide spread of values in this parameter space.

seen along with that of EXO 1846 for contrast, exactly as they were presented to volunteers, in Figure 4.7. It is likely that *Swift* J1858 was not classified as transient due to a combination of there being only one significantly bright data point, as well as the automated figure generation creating an overly large legend, which serves to distract the viewer. Similarly, SAX J1808’s misclassification can be understood as stemming from uncertainty surrounding so few data points (especially compared to other datasets). Indeed, SAX J1808 received 1 ‘Unsure’ vote which, had it instead been for ‘Transient/Variable’ would have pushed this subject above our classification threshold.

In addition to these central 11 sources, several of the XRBs also display discrete, resolved jet components (*MAXI* J1820+070, *MAXI* J1348–630, *MAXI* J1848–015 and 4U1543–47; see Table 2.1) that are counted as unique sources by the TRAP as they become resolved and move away from the core of the jet. In total there are 6 moving, transient jet components that are detected by the TRAP as unique, of which 3 are classified as variable by citizen scientists. Given that neither the software pipeline nor project classifications were designed to pick up moving point sources, it is interesting to note that there is room for unexpected discoveries even in such simple workflows.

4.4.2 Previous commensal studies

One of the largest works on variable sources in MeerKAT images to date is that of Driessen et al. (2022). In that work, 21 new long term variables (LTVs) are detailed, along with GX339–4, the first MeerKAT transient MKT J170456.2–482100 and the known mode-changing pulsar PSR J1703–4851 (both described in Driessen et al., 2020). Of these 24, excluding the XRB discussed earlier, MKT J170456.2–482100, PSR J1703–4851 and 10 of the LTVs are missed by classifications i.e. only 11 are marked as transient/variable by volunteers. One reason for this is that in the work of Driessen et al. (2022), each light curve is binned by a factor of 10 (i.e. every 10 data points are represented by their uncertainty-weighted mean) to make long term variability more apparent, something not done here. Furthermore, here I made no

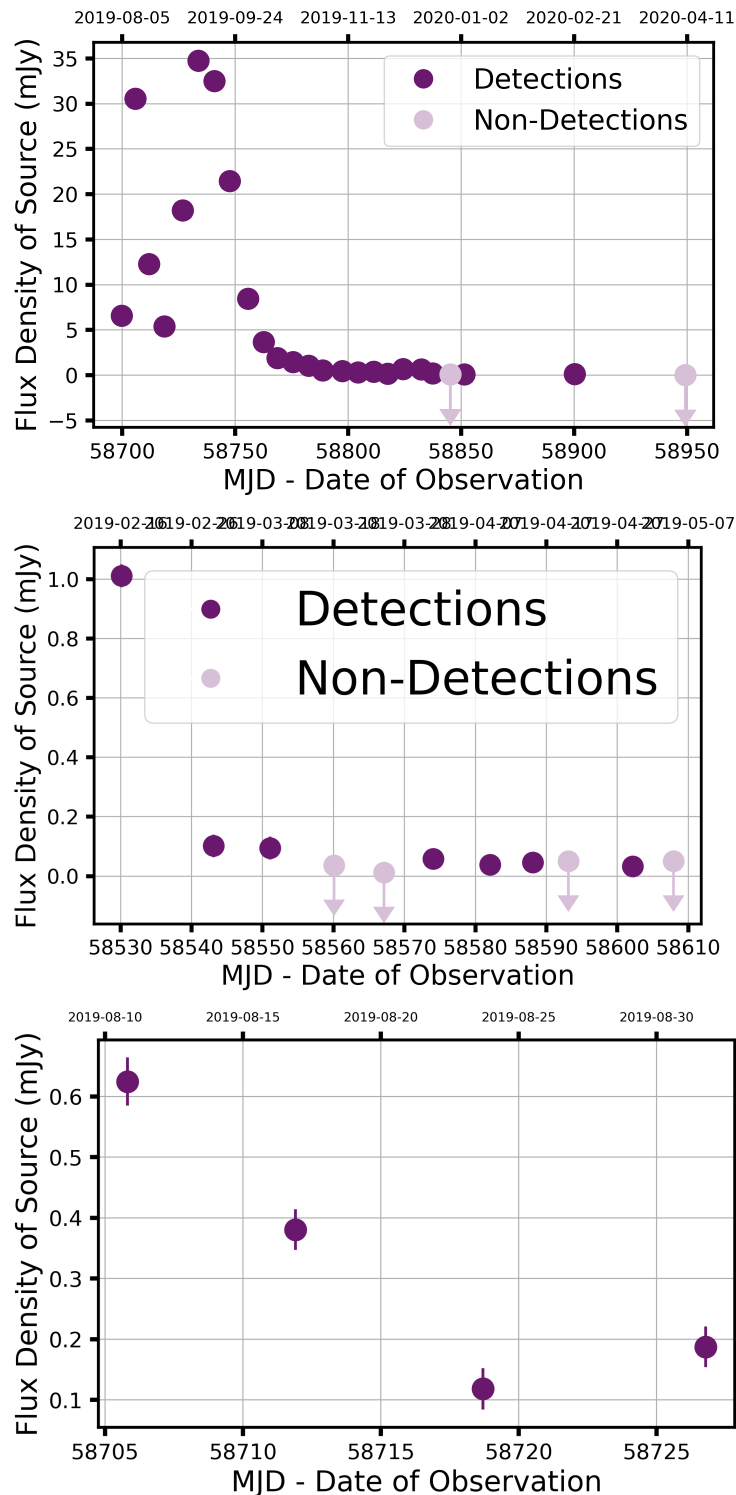


Figure 4.7: The TRAP light curves of EXO 1846-031, *Swift* J1858.6-0814 and SAX J1808.4-3658 (upper left, upper right and lower respectively), shown here exactly as they were to volunteers on *BfS:MKT*. The foremost was classified as transient by citizen scientists whilst the others were not. I believe that *Swift* J1858.6-0814 was misclassified due to the mis-scaled, automatically generated legend and the lack of apparent variability aside from the first data point. Similarly SAX J1808.4-3658 was likely misclassified due to the lack of data points on which to make a judgement.

use of a deep, stacked observation in order to follow sources through every epoch (see section 4.2.1's details on the `expiration` parameter), something that will result in different light curve shape and in one case resulted in a source not being detected at all in my set of images. In general, the LTVs not identified by volunteers were generally fainter and have less smooth light curve evolution.

Similarly, when comparing to the work of Rowlinson et al. (2022) we see that the three variables found therein - NVSS sources J181849+062843, J181752+064638 and J182029+063419 - were recovered by citizen scientists. However, the three sources found in their 'transient hunt' (where for Rowlinson et al., 2022, a transient is defined as sources not detected in a deep observation of the field, see their section 3.1) were not identified by volunteers. I therefore inspected the Zooniverse subjects for these three transients. Upon doing this I saw that the image of source MKT J182015.5+071455 (Rowlinson et al., 2022's source 2) provided to volunteers was automatically scaled to the brightest pixel in the image, not the central source of interest. As such, the source appears not to be present and was classified mostly as an artefact. The light curves of the other two transients were faint, with large uncertainties and figures that are more 'cluttered', with Rowlinson et al. (2022)'s pulsar receiving 3 votes as a transient/variable and 1 unsure classification. By contrast the three variables recovered have smooth light curve evolution over long timescales. These missed transients could be due to the heterogeneous nature of our classifiers, or it could reflect the uncertainty surrounding a less clear pattern in a light curve.

Finally, the M dwarf SCR 1746-3214 discussed in Chapter 3 was also shown to volunteers. In order to assess what light curves citizen scientists were most comfortable with classifying, I provided two light curves of this source - one with more data points and an additional detection and one with only two detections and two upper limits (see Figure 3.4). Interestingly, when provided with the shorter light curve of the source, the transient source was mostly classified as 'Unsure' or 'Stable'. However, when given the full light curve volunteers correctly identify the subject as transient.

In summary, aside from a few figure processing issues, several transients and variables were not recovered by volunteers. These tended to be from sources that were fainter, whose light curves had larger uncertainties and sources with fewer detections. Therefore this sample of transients and variables is likely biased towards brighter, slower evolving objects that occur in our most sampled fields.

4.4.3 Scintillation analysis

The majority of these 142 variables that our volunteers discover show long term variability, with light curves showing variations over weeks to months. As discussed in section 1.3.3.2, one cause of variability for radio sources is scintillation, caused by changes in electron density through intergalactic, interstellar, or interplanetary media (IGM, ISM and IPM respectively). As pointed out by Hancock et al. (2019) this is, in essence, the radio analogue of witnessing stars twinkle in the night sky. All of our observations are within latitude $|b| \lesssim 10^\circ$ of the Galactic plane (see Table 2.1 and Figure 2.2) so this may be a large contributor to the variability seen in our sample due to light from these variables traversing through (at least some part of) the Milky Way.

The theory behind scintillation is well described in the work of Rickett (1986) and Rickett (1990), which has been implemented into a user-friendly model by Hancock et al. (2019). Here I describe the most important aspects of Hancock et al. (2019)'s model and implementation, which builds on previous work. Firstly, it is assumed that the integrated effect over the line of sight between a telescope and a source can be modelled by a single thin scintillation screen at some distance D . It is then useful to define two particular length scales that together define weak and strong scattering regimes. The Fresnel scale r_F is defined as

$$r_F = \sqrt{\frac{\lambda D}{2\pi}} \quad (4.1)$$

and represents the transverse separation over which the phase delay is 1 rad. The next is the diffractive length scale r_{diff} , representing the transverse distance over

which the RMS phase variance is 1 rad. Under the assumption that the entirety of our scattering material is within the Milky Way then

$$r_{\text{diff}} \propto \left(\frac{\lambda}{1m} \right)^{-6/5} \left(\frac{SM}{10^{12}m^{-17/3}} \right) \quad (4.2)$$

where SM is the scattering measure, the integral along the line of sight of the amplitude of turbulence per unit length (analogous to the dispersion measure of pulsars and FRBs). The ratio of our two relevant length scales is then

$$\xi \equiv r_F/r_{\text{diff}} \propto \lambda^{17/10} SM^{3/5}. \quad (4.3)$$

Diffractive and refractive interstellar scintillation (DISS and RISS) are then the two relevant scattering effects caused by multipath propagation. The former occurs on rapid timescales (seconds to minutes), at high amplitude and over small fractional bandwidths, making it not relevant for our broad bandwidth observations taken over many weeks. RISS, however, occurs at lower amplitudes and over longer timescales and is caused by inhomogeneities in the ISM on large scales. The amplitude of RISS modulations is given by

$$m_R \equiv \frac{\sigma}{\mu} = \xi^{-1/3} \quad (4.4)$$

where μ and σ are the mean and standard deviation of our flux measurements - making this quantity identical to our previously defined V . Similarly the timescale for variations in RISS modulations is given by

$$t_R = \left(\frac{r_F^2}{r_{\text{diff}}} \right) \frac{1}{v} \quad (4.5)$$

where v is the relative velocity between the observer and a screen. In summary, if I have measurements or estimates of SM , v and D , along with wavelength λ , then I can make a prediction for both the amplitude and timescale of variability and compare this to the measured quantities on our observed light curves. The final missing piece of the puzzle, then, is an estimate of SM . These are folded into Hancock et al. (2019)'s model using an underlying $H\alpha$ map from Finkbeiner (2003). The underlying data come from three surveys (The Virginia Tech Spectral line

Survey, the Southern H-Alpha Sky Survey and the Wisconsin H-Alpha Mapper; Dennison, Simonetti, and Topasna, 1998; Gaustad et al., 2001; Reynolds, Haffner, and Madsen, 2002), forming a final sky mosaic with a resolution of $6'$. These data then allow for $H\alpha$ intensities to be mapped to values of SM for any line of sight. From this, and the equations above, I can then make our estimates of RISS-induced modulation and estimate the approximate timescale.

Importantly, I do not know what the appropriate distances or relative velocities are between the Earth and the modelled screen(s). Therefore I assume that the single phase screen is located halfway between the Sun and the ‘edge’ of the Milky Way, which is modelled as a flat disk, with a radius of 16 kpc, a height of 1 kpc, and with the Sun system sitting 8 kpc from the Galactic centre (Hancock et al., 2019). From equations 4.1, 4.3 and 4.4 we can see that $m_R \propto D^{-1/6}$ and so depends only weakly on this assumption.

There is no built in model for the relative velocity distribution of the Milky Way, so I use the default value of 10 km s^{-1} . Equation 4.5 shows how the timescales of variation are linearly dependent on this assumed velocity. As a result I must be careful in interpreting the calculated timescales. Nevertheless it is still an interesting exercise to see *if* there exist such velocities and distances to screens such that this RISS model can potentially explain the variations seen in some of my volunteers’ light curves.

I can now plug in the coordinates of all of our variable and transient sources to the above model and calculate a RISS-induced variability amplitude and timescale. I can then compare this predicted maximum to our measured V values, as seen in Figure 4.8. For the majority (131) of our new variables it can be seen that the predicted modulation is equal to or greater than that measured by TRAP. This means that the observed variability is consistent with (though not necessarily exclusively explained by) the scintillation of light passing along lines of sight through the Milky Way, having originated at extragalactic distances i.e. in AGN. By contrast, the known Galactic XRBs and their jets - whose variability is caused by shocks

and particle acceleration - show variability that is much greater than what would be expected due to refractive scintillation.

The predicted RISS variability is in some cases much higher than our observed V values. Part of this may be caused by the $H\alpha$ surveys that underlies the SM calculation. Namely, this map has relatively low resolution ($6'$) and an average uncertainty across our lines of sight of $\sim 12\%$, so it is natural to expect our values to not line up exactly. It is also true that the observed modulation will be a sum of all contributing causes (e.g. further propagation effects such as ionospheric scintillation but also any potential intrinsic variations). Therefore if a given light curve is sampled less frequently than a given variation timescale then the observed modulation will be reduced with respect to these predictions. Finally, the weighted average timescale of variation for our sources is found to be $\bar{t}_R = 8 \pm 4$ months, where the weights used are propagated through from the uncertainties in the underlying map and the quoted uncertainty is the standard deviation. This range of timescales of variation at 1.28 GHz matches well with the length of typical observations for our XRB fields, though as above I note caution in over-interpreting this calculation. Both these matching amplitudes and timescales of variability provide evidence for the majority of our transients being scintillating AGN or other point-like extragalactic radio sources.

4.5 Counterparts and associations

I have already shown in the previous chapter how multiwavelength data can be invaluable in characterising our variable and transient universe. Therefore here I again make use of the MeerLICHT optical telescope (Bloemen et al., 2016) to gain physical insight into possible source classes of our volunteers' transients. A typical observing schedule consists of 1 minute exposures of a given field, alternating between q -band (440–720nm) and each of 5 Sloan bands u , g , r , i and z . Whilst MeerLICHT operates in these 6 bands, here I use only the q -band due to its highest sampling rate and its broad wavelength coverage. I crossmatched with the MeerLICHT database at a $2''$ search radius - large enough to partially account for




Figure 4.8: The ratio between observed V and the variability predicted by Hancock et al. (2019)’s model for Refractive Interstellar Scintillation (RISS), for our sample of transients and variables. Most variability can be explained by this model of a scintillating, extragalactic source, apart from the known XRBs and jetted systems.

MeerKAT astrometric uncertainties (see Driessen et al., 2022) and some proper motion (e.g. of nearby stars) but not so large as to include many false matches – using the uncertainty-weighted mean position of each of our candidates as returned by the TraP. Running this crossmatch returns 25 counterparts in the MeerLICHT database and 143 3σ upper limits. This low rate of optical counterparts is not surprising as all of our XRB fields are within ~ 10 degrees of the Galactic plane and so many optical counterparts may be heavily extinguished.

I summarise a given entry to the MeerLICHT database by the mean of its light curve and compare to the average radio flux density, as for the purposes of classification I am only interested in the nature of counterparts and not necessarily their detailed optical behaviour.

The radio-optical plane for our variables can be seen in Figure 4.9, with comparison source types from Stewart et al. (2018)⁷. The first thing to note is that the majority of these cross-matches exist in the region where extragalactic sources have been detected - either quasars or GRBs. If most of our candidates are extragalactic in nature this agrees with previous studies, who find the vast majority of radio variables are extragalactic (Thyagarajan et al., 2011; Sarbadhicary et al., 2021). There are a few sources overlapping the ‘stellar’ region of the parameter space - these include the transients already reported by ThunderKAT (Driessen et al., 2020; Andersson et al., 2022). I note that moderate levels of extinction would displace some of these sources into the extragalactic part of parameter space and so perhaps an infrared version of this diagnostic plot is needed. Finally, it is important to note that by comparing to archival data in this way does not leave room for unknown astrophysical classes. By this I mean that just because a new variable or transient is consistent with some known source type, this does not mean that the transient exclusively must originate from the known source class and could instead be a novel, unknown class of object. Nevertheless, this figure still gives an indication of the overall distribution of multiwavelength counterparts.

4.5.1 Highlights

I also searched for counterparts at other wavelengths and with pre-existing classifications, with the aid of code made available by Driessen (2021). Crossmatching makes use of the `astroquery` package to search the SIMBAD⁸ database (Wenger et al., 2000) and several catalogues within Vizier⁹ (Ochsenbein, Bauer, and Marcout, 2000). I again searched at a 2'' radius to our radio sources. I searched for X-ray and gamma-ray counterparts to our radio variables. To do this I crossmatched with catalogs from the *Fermi* (Schinzel et al., 2014), *Chandra*, (Evans et al., 2010), *Swift* (Evans et al., 2020) and *XMM-Newton* (Traulsen et al., 2020) facilities. There were no counterparts for any of our sample, aside from the known transients of interest

⁷Code for reproducing this plot is available at <https://github.com/4pisky/radio-optical-transients-plot>

⁸<http://simbad.cds.unistra.fr/simbad/>

⁹<https://vizier.cds.unistra.fr/viz-bin/VizieR>




Figure 4.9: The mean optical and radio flux densities of our sample of radio variables, atop an underlying distribution of astrophysical classes (Stewart et al., 2018). Black crosses denote counterparts within the MeerLICHT database whilst grey triangles are upper limits. Diagonal lines denote a constant ratio between radio and optical flux density, whilst the A_R marker indicates the horizontal displacement caused by 5 magnitudes of optical extinction. The majority of our radio sources are likely extragalactic as they overlap in parameter space with quasars and GRBs.

(i.e. the XRBs). Below I note a few interesting objects returned in our search with IR or radio detections, or otherwise known counterparts.

4.5.1.1 OH Maser - BfS 80

This subject was classified by the following contributors to BfS:MKT - Newton Temoke, Najma Rahman, Sai Parthasarathi, Piilonen, Loïc Baert, Isabelle Whittle, Davide Iannone, Sauro Gaudenzi, Lizzeth Ruiz Arroyo and one anonymous volunteer.

We detected transient emission from a source in the EXO 1846 field which volunteers confidently classified as transient. Cross-matching reveals that the source in question is a known maser star OH 30.1–0.7 otherwise known as V1362 Aql, with radio observations stretching back almost 50 years (Evans et al., 1976), but no previous mentions of its variable nature. This asymptotic giant branch (AGB) star

is heavily dust-obscured at optical wavelengths, but very bright at IR wavelengths - $W1 = 6.9 \pm 0.1$ mag, or 279 Jy at 25 μm (Cutri et al., 2013; Gonidakis et al., 2014). A comparison of a MeerKAT radio detection (contours) and Spitzer Glimpse imaging can be seen in Figure 4.10, alongside its light curve. AGB stars are post-main sequence systems, whose low surface temperatures, radii of several hundred times the solar radius and stellar pulsations give rise to strong winds, expelling as much as $\sim 10^{-5} M_{\odot} \text{yr}^{-1}$ (Höfner and Olofsson, 2018). These winds create an oxygen-rich, dusty circumstellar environment that generate masers at 1612, 1655 and 1667 MHz as infrared photons pump OH molecules formed through photodissociation of water by interstellar radiation (see section 1.2.2.3).

The cause of the variability of OH 30.1–0.7 is not clear. Perhaps the OH maser emission is varying, due to stellar pulsations. The derived stellar period from the General Catalog of Variable Stars is $\sim 1730\text{d}$ (Samus' et al., 2017), with more recent estimates from the WISE W1 and W2 bands at 1950 ± 50 and 1520 ± 20 days respectively (Groenewegen, 2022). These timescales are much longer than the variability seen in the radio light curve here (of order a few months), so it is not clear if stellar pulsations are responsible for variable maser emission (averaged across the entire L-band from Jy to mJy levels). There could also be inhomogeneities at the site of the maser emission. The second cause for variability could be due to binary interactions - the system likely has a companion, as inferred from ALMA CO data by Decin et al. (2019). This, combined with the lack of any optical counterpart implies the system could be a dusty (D-type) symbiotic binary - all D-type symbiotics host Mira stars (Whitelock, 1987). These binaries consist of a windy red giant and a smaller companion onto which material is shed (Allen, 1984). Radio emission has been seen from symbiotic binaries (Seaquist, Taylor, and Button, 1984) and is known to vary in some sources, typically interpreted as optically thick emission from an inhomogeneous region in the red giant's wind that is ionised by its companion (Seaquist, 1988). The nature of OH 30.1–0.7's companion is as yet unknown due to the heavy extinction in this region and without evidence of high ionisation (e.g. HeII or [O]) I cannot claim that this is a symbiotic system. There is

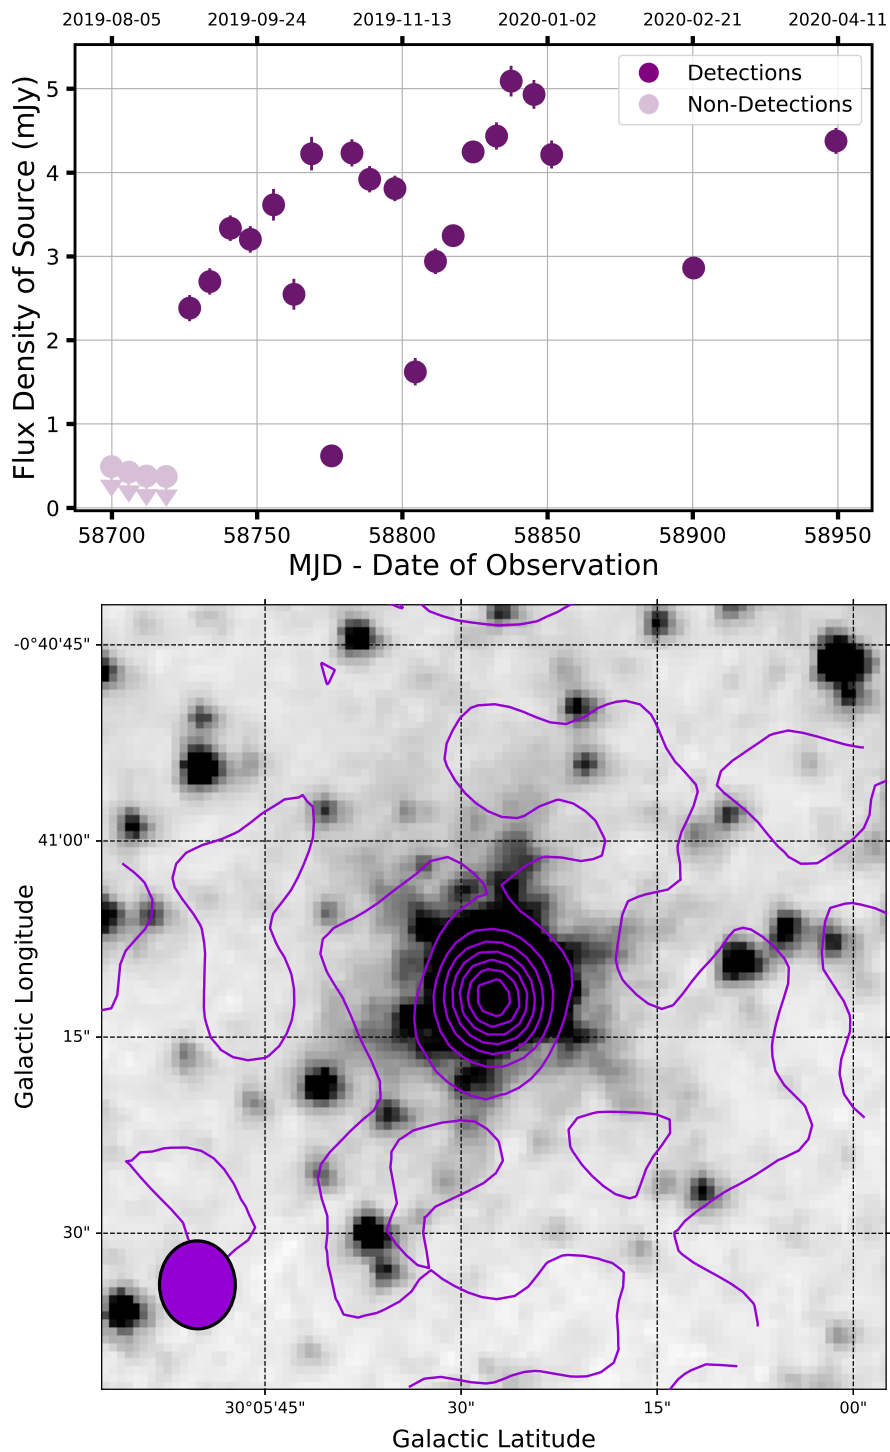


Figure 4.10: **Upper:** Light curve of OH Maser 30.1–0.7 picked up by volunteers. **Lower:** Overlay of MeerKAT radio contours over Spitzer imaging from the GLIMPSE survey (Benjamin et al., 2003). Contours are spaced linearly in 0.5 mJy increments from -0.5 to 3.5 mJy.

overlap between maser systems and Mira-type symbiotic binaries (Seaquist, Ivison, and Hall, 1995) and so the observed variability could be due to a combination of emission mechanisms discussed, or perhaps something else entirely. There are 4 non-detections before the system reached ~ 3 mJy and so the system must have changed significantly in roughly a week. Future radio observations could help determine what the nature of the variability is (including the initial non-detections), perhaps combined with spectroscopic searches for the nature of the companion (e.g. if it is a white dwarf, or for evidence for ionisation).

4.5.1.2 Pulsar - BfS 20

This subject was classified by Davide Iannone, Leticia Navarro, Kyle J. Melville, Chloe Tworek, Lizzeth Ruiz Arroyo, Tim Forsythe, Marianne De Sousa Nascimento, Karla Lahoz, Jeffrey Smith and Wilfried Domainko.

Pulsar PSR B1845-01 (J1848-0123), whose light curve can be found in Figure 4.11, was seen in the field surrounding XRB *MAXI* J1848. In the MeerKAT observations a mean flux density of ~ 15 mJy was measured, in close agreement with that measured by the Parkes radio telescope recently (15.2 ± 0.9 mJy Johnston and Kerr, 2018). However, reports of this pulsar's flux density have been as low as 8.9 ± 0.9 mJy at 1400 MHz (Hobbs et al., 2004), indicating long term changes in the received brightness from the source. The pulsar is brighter at lower frequencies (e.g. measured to be 79 ± 6 mJy at 408 MHz by Lorimer et al., 1995), with a spectral index α - defined at frequency ν as flux density $F \propto \nu^\alpha$ - of -1.3 ± 0.3 (McEwen et al., 2020).

The observed variability is consistent with RISS (see section 4.4), with a predicted modulation greater than the measured $V \sim 0.02$ and the observed timescale of variation matching the estimated 11 ± 2 months. Furthermore, this pulsar has a dispersion measure of 159.1 ± 0.2 pc cm $^{-3}$ (McEwen et al., 2020). This high dispersion (compared to typical DMs of order 10s of pc cm $^{-3}$ for pulsars with $|b| > 25^\circ$; Manchester et al., 2005) is known to be linked to long timescale, low amplitude refractive scintillation (Stinebring et al., 1990; Stinebring et al., 2000).




Figure 4.11: Light curve of PSR B1845–01 with MeerKAT. 7 of 10 volunteers voted for this as a transient/variable source. The low amplitude variability observed over a timescale of 100s of days is consistent with RISS.

Similarly, the pulsar’s location in the Galactic plane is in keeping with its radio emission traversing a large free electron content, hence the high DM and clear scintillation. When searching through the literature there were no other mentions of variability observed from this pulsar, with its most individually noteworthy feature being a triple-peaked pulse profile seen from 1.4 – 10.5 GHz, relevant for testing possible underlying emission mechanisms (Kramer, 1994; Shang et al., 2021). Finally, I note that B1845-01’s spin period of ~ 0.65943 s is much shorter than that of our observations (typically 15 minute epochs consisting of 8s correlator sampling) so this cannot be contributing to the observed variability.

Similar phenomena such as mode changing (Kramer et al., 2006) and pulse shape variations have been seen in the pulsar population (e.g. Lyne et al., 2010 or more recently Shaw et al., 2022) and there are known phenomena in rotating radio transients that can change on timescales from weeks up to years (Keane and McLaughlin, 2011). Whether the observed drift in flux density is intrinsic to the emission at the pulsar or a product of the interstellar medium is also not clear.

This pulsar adds to the diverse range of behaviours seen in pulsars spotted by MeerKAT in imaging data. Similar examples include the mode changing pulsar observed in the GX339–4 field (see Driessen et al., 2022) and, in the most extreme case, one of the slowest pulsars discovered (Caleb et al., 2022). Even though the variability seen here is only moderate, imaging of pulsars can be very useful for measuring their positions to high accuracy (e.g. Heywood, 2023).

4.5.1.3 VLASS1 J181955.28+074418.7 - BfS 146

This subject was classified by Matthew Baumann, B. Stewart, Jeffrey Smith, Rachel Ann Grenier, Chloe Tworek, Davide Iannone, Thorsten Eschweiler, Leticia Navarro, Lizzeth Ruiz Arroyo and Per Ola Jonas Andersson. The lattermost volunteer deserves a special, biased mention, as he is my dad.

Radio source VLASS1 J181955.28+074418.7 is an object spotted by our volunteers not only in our aggregated results (scoring 7/10 votes as a transient/variable) but also in our Talk board¹⁰. This source is in the field of MAXI J1820 but outside the 0.5 degree search radius set by Rowlinson et al. (2022). Its light curve can be seen to vary very smoothly in Figure 4.12, where it is worth noting that the non-detections are all due to drops in data quality in those images, which should be filtered out in future work. This variation is not similar to any nearby source in our data, nor is it correlated with PSF size or shape.

There is a counterpart to this source in the Very Large Array’s Sky Survey (VLASS QuickLook Epoch 1; Lacy et al., 2020; Gordon et al., 2021), with a 3 GHz flux density of 1.9 ± 0.3 mJy. Since the paper that forms the basis of this chapter was accepted for publication (late April 2023), this source has also been observed by the VAST project on ASKAP - with 6 observations between June and December 2023. VAST observed the source at a central frequency of 888 MHz, varying from 0.5 to 1.4 mJy over these 6 months. From the VAST measurements the source has (η, V) values of (2.9, 0.43) compared to the MeerKAT light curve values of (36.4, 0.37) i.e. roughly the same amplitude of variability is seen, though at lower significance

¹⁰<https://www.zooniverse.org/projects/alex-andersson/bursts-from-space-meerkat/talk/4567/2263526>

due to larger uncertainties on the ASKAP data. As flux density uncertainties enter into the denominator of equation 2.2 this is not a surprise, but it is interesting to note how the same source, seen at similar frequencies over comparable time periods, can produce such different statistics.

There are no counterparts to this source in any higher energy band, despite this field being our furthest from the Galactic plane ($b \sim 10^\circ$). This kind of source is very typical of our sample - scant extra data but an intriguing radio light curve. Further information at other wavelengths would help determine source type, as would simultaneous radio data - for example to determine a precise spectral index α and see if this points towards the source being an AGN or a pulsar. Our RISS analysis (section 4.4.3) produces a ratio of 1.17 between the observed V and predicted scintillation amplitude. This could indicate the source has some intrinsic variation however the scintillation parameters were not exhaustively tested, so this ratio could be explained by an incorrectly assumed distance to, or relative motion between, the model scattering screen and the observer.

4.6 Discussion

Using citizen science to discover transients is a fruitful endeavour, as shown by both the uptake of our Zooniverse project and the novel results produced. For our volunteers, the project provides new experiences with a branch of observations perhaps less familiar to their perception of astronomy (e.g. compared to galaxy morphologies or solar physics). The uptake of our project demonstrates a clear appetite for further development of citizen science for radio transients, as evidenced by >1000 citizen scientists applying their time to our science case. The time taken for an individual to check several years worth of data is much greater than the 3 months taken by our volunteers, so citizen scientists help project scientists analyse their observations more efficiently. In this work volunteers have been able to recover or discover a broad range of astrophysical transients that occur over several orders of magnitude in timescale, so the science results discussed here also give merit to developing *BfS:MKT* further. These results include flaring from




Figure 4.12: Light curve of source VLASS1 J181955.28+074418.7, showing smooth, near sinusoidal variations. The epochs denoted upper limits are in error and due to not filtering out low quality images in pre-processing. This light curve was described as “too perfect” (to be real) and “beautiful” by volunteers.¹⁰

stellar systems, the discrete and compact jets of XRBs, maser emission, pulsars, the long-term variation in likely AGN and potentially new source classes. Given these successes, I went on to plan and launch a second wave of classifications to the Zooniverse site, using different data from surveys on MeerKAT to explore the range of parameter space for both citizen and project scientists (see section 4.7). Of course this method is not restricted to MeerKAT data and any such set of (radio) images and light curves could be analysed in this way.

One of the most rewarding successes of the project for me was the discussion of the relevant science between users and beyond the immediate classifications. For example, there are instances where experienced volunteers would ask me a question on the project-specific talk forum, usually about how an edge-case or uncertain subject should be classified. This then sometimes had a cascading effect, where I would later see a new volunteer ask a similar question, only to have the experienced user answer the query. This proved to proliferate both the science aims of the

project, as well as the learning experiences and enthusiasm that volunteers had for this work. Furthermore, some of the more experienced users (many of whom share authorship on the relevant journal article) would take the metadata about a particularly interesting subject and go away to do their own research and report their findings back to me. This often constituted cross-matching the coordinates of a subject, as taken from the embedded metadata, with other surveys and reporting back on the lack of a counterpart or the mass of the black hole in a particular XRB system. This level of engagement was completely unexpected for me but one of the most rewarding parts of my entire DPhil. Finally I note that in the era of unprecedented data scales even the reduced, final list of candidate sources provided by volunteers can exceed the follow-up capabilities of a single science team. Perhaps this kind of community-generated interest and volunteer-driven study is a way to further investigate sources of interest, as well as encouraging the agency of the members who wish to take their citizen science efforts further.

However, citizen science is not without its challenges and some of these can be understood by comparing our findings to those of previous work from the ThunderKAT team - namely Driessen et al. (2022) and Rowlinson et al. (2022), as well as the published XRB work (see Table 2.1). As discussed in section 4.4.1, transients and variables were missed by volunteers due to issues with figure processing, including incorrect z-scaling of pixels and badly scaled figure labels. I will use these issues to improve our procedure for future batches of data, for example by manually setting the pixel scale in images, altering legend sizes etc.

Aside from these pre-processing issues, there were still some known transients not recovered, including the pulsars detected by Driessen et al. (2022) and Rowlinson et al. (2022) and the star from Chapter 3. When comparing the light curves of these transients to variables that are recovered we can see a clear trend - recovered sources show long term, smooth variations, and are typically brighter, with smaller uncertainties, resulting in clearer patterns on display to scientists. By contrast, the faint, transient pulsar light curves show very ‘noisy’ light curves, with less eye-catching patterns, despite being precisely the kinds of transients I want to

discover. We can perhaps use this to infer that citizen scientists were least unsure when classifying sources with more data points and less reliant on upper limits, as is the case for many of the new variables found in this study. I hope to alleviate some of this bias in future Zooniverse runs by emphasising the use of non-detections and by encouraging volunteers to label things as transient. I could also implement the transients and variables discussed in this work into the Field Guide for volunteers as more examples of the types of sources for which I am looking.

I can quantify how scalable our method is by comparing our observations to the MeerKAT MIGHTEE survey. MIGHTEE's observations produce ~ 6000 sources per square degree on the sky (Heywood et al., 2022a), for which our TRAP processing would produce a light curve. For comparison, in this work we produce ~ 1000 subjects from the ThunderKAT fields that are devoid of large diffuse structures or exceedingly bright sources. If we assume that only the same 1000 volunteers contribute to all future data releases as with this study, at the same classification rate, then it would take approximately 60 days to receive 10 classifications on every source at MIGHTEE's sensitivity (1σ RMS noise $\sim 1 \mu\text{Jy}$), per square degree. Multiplying this across the sky coverage of just the ThunderKAT fields used in this study (~ 1.5 square degrees over 11 fields) results in a required volunteer classification time of 990 days. This is far greater than the time taken for the observations (e.g. of order tens of 8-hour epochs) and which is needed to image the data and process them to form light curves. If the volunteer results were only analysed when all classifications were finished then this would also be too late for real-time follow-up of transients. To bring the classification time to that of this work (90 days) would require an order of magnitude more volunteers, which is achievable for Zooniverse projects, particularly when disseminated widely. For example, *Galaxy Zoo* variants receive many tens of thousands of volunteers, whilst *Gravity Spy* has had over 30,000 participants. However, for any considerable survey area the overall time required would again balloon to far larger timescales than reasonable, particularly if datasets and classifications are processed in batches once observations are complete. In addition, observations at the 8s integration time for MeerKAT produce far fewer

sources to classify compared to deeper images, however the imaging and processing time prior to volunteer classification increases hugely so the overall timescale remains long. So this kind of analysis will not scale easily to the most sensitive observations of MeerKAT fields (at 1.28 GHz), let alone those expected from the SKA, ngVLA (Hallinan et al., 2019) or DSA-2000 (Selina et al., 2018).

One way to alleviate this data deluge might be to develop machine learning methods to remove ‘bogus’ and ‘boring’ sources in favour of the rarer variables and transients for which I am searching, as has been done for supernovae and galaxy morphologies (Wright et al., 2017; Walmsley et al., 2022a). In transient astronomy, anomaly detection algorithms can uncover rare phenomena, finding unique stellar activity and flares (Giles and Walkowicz, 2019; Webb et al., 2020; Webb et al., 2021), extragalactic transients in current data streams (Muthukrishna et al., 2022; Malanchev et al., 2021; Aleo et al., 2021) and in preparing models for the Vera Rubin Observatory (Boone, 2019; Hložek et al., 2023; Villar et al., 2021). These kinds of interesting and powerful tools, along with the data deluge that citizen science efforts may come up against, motivated the work of Chapter 5.

4.7 Conclusions

In this chapter I have presented the first citizen science project for finding transients in image plane radio surveys. This was motivated by both the difficulties in finding radio transients and the previous success in transient astrophysics and radio astronomy on the Zooniverse platform. I used data taken originally as part of the ThunderKAT programme on MeerKAT to monitor XRBs and their jets. These data were processed comensally, producing ~ 9000 sources which were displayed to volunteers on the Bursts from Space: MeerKAT project page.

In this chapter I have detailed the process of review and beta-testing that resulted in the final design of the *BfS:MKT* project. Upon launch, the uptake of the project was very strong, with > 1000 volunteers taking part, demonstrating a healthy appetite for further Zooniverse data releases. I was also able to use the

known transients in our fields to understand some reasons why interesting sources may be missed and will fold this learning through to future iterations of the project.

Citizen scientists have uncovered a large sample of interesting transient and variable sources, some of which we may have not uncovered were it not for our volunteers' dedication. I provide the full catalogue of 168 radio transients and variables, which I believe to be the largest catalogue of candidate radio variables to date. It is my hope that these sources will be investigated by the community.

I made use of archival multiwavelength data, including the MeerLICHT telescope, to help classify the systems found. Many of our light curves have no multiwavelength counterpart, though this is not entirely unsurprising given the low Galactic latitude of our commensal observations. Analysis of the RISS properties of these variables indicates that many of our sources are likely extragalactic sources scintillating (i.e. AGN). Volunteers also found or recovered a broad range of physical phenomena including pulsars, radio-loud stars and XRBs. These results demonstrate the wealth of science possible with new radio facilities and the power of citizen science.

Moving forwards, I have already released a second batch of data to volunteers. These data came from the MIGHTEE survey on MeerKAT and constitute 19 deep ($\sim 2\mu\text{Jy}$) epochs of the Chandra Deep Field South. Classifications of the ~ 12000 sources in this field have already been made, again in only a few months, by our volunteers. Analysis of this data release has not been performed yet but initial observations on the Talk forum show many promising candidates. A third data release is planned, consisting of 8-s images of the same 19 observations as DR2. There are of order a few hundred sources detected consistently in the ~ 60000 images I have. These will be shown to volunteers in the near future. The combination of these deep, long timescale observations and shallow (though actually still quite deep) fast timescale images will provide a truly unique understanding of the multi-timescale radio variable sky.

Real stupidity beats artificial intelligence every time.

— Sir Terry Pratchett, from *Hogfather*

5

Anomaly detection

Contents

5.1	Introduction	118
5.1.1	Data-intensive astronomy	118
5.1.2	Machine learning in transient astronomy	119
5.2	Observations and classifications	125
5.3	Feature extraction	128
5.3.1	Baseline features	128
5.3.2	Feets features	129
5.3.3	Wavelet features	130
5.3.4	Feature space and clustering	134
5.4	Anomaly detection	135
5.4.1	Anomaly recall and purity	139
5.5	Active learning	144
5.5.1	Active learning improvement	146
5.6	Discussion	147
5.7	Conclusions	152

This chapter presents the first application of anomaly detection algorithms to the task of finding radio transients via their light curves. I demonstrate that fast and general-purpose algorithms are able to recover the majority of the transients, as found by citizen scientists in the previous chapter, in a minority of the dataset. Building these fast but effective algorithms into telescope backends could play a large part in dealing with the data rates expected from upcoming observatories.

5.1 Introduction

5.1.1 Data-intensive astronomy

Current and upcoming radio facilities promise to overwhelm astronomers with data, creating the problem of finding the needle of rare transients in a haystack of high-cadence, highly sensitive observations. For example, current deep image stacks (25 hour integration) with MeerKAT detect ~ 6000 sources/deg² (Heywood et al., 2022a). The SKA-Mid Array, when complete, will be able to reach the same depth in approximately 12hrs¹. If the array were to conduct non-overlapping 12hr pointings over an entire year, it would observe upwards of 4,000,000 sources (over less than 1/40th of the sky). As seen in the previous chapter, time-variable phenomena make up approximately 2% of all sources in a given radio image, meaning that the SKA-Mid will see up to $\sim 40,000$ transients and variables in its first year of operation. The PB-scale data coming from instruments like the SKA therefore necessitates the use of low-latency methods for finding transients, as ultimately not all data will be stored in perpetuity. An example of this can be already seen in the case of the MeerTRAP backend (More TRansients And Pulsars, Stappers, 2016) employed on MeerKAT. The MeerTRAP user supplied equipment taps into the MeerKAT datastream, commensally searching for fast transients such as FRBs and bright single pulses from pulsars. This effort requires real-time candidate processing and automatic filtering in order to reduce the data rate to reasonable volumes. Potential candidates from this system require additional confirmation or rejection by the team in order to not overwhelm storage capacity as, for example, repeating bursts from a known system can produce > 1 TB of data per day (Malenta et al., 2020; Rajwade et al., 2020). As a final note, data rates only grow when considering the many data products are required for a given science case. For example, the ThunderKAT observations used in this thesis are averaged from 32k frequency channels and two polarisation states to a single frequency-averaged, Stokes I image, reducing the data volume by over an order of magnitude. By contrast there are

¹Calculated from <https://sensitivity-calculator.skao.int/mid> on 13-04-2024

good examples from the past two years where the combination of spectral coverage, time baseline *and* polarimetry has been shown to provide a unique discovery space for pulsars and transients (Heywood, 2023; Smirnov et al., 2024).

This data-intensive regime is similar to that facing optical astronomers, who have spent several years preparing for the Vera C. Rubin Observatory to come online and deliver a Legacy Survey of Space and Time (LSST), consisting of millions of transient alerts per night in a cumulative volume of petabyte scale data (Ivezić et al., 2019). One avenue of research in this field has been into anomaly detection methods - these are often unsupervised (that is, untrained) machine learning models that seek to separate rare anomalies from everything else. Due to the large imbalance between transients and non-transients in my data I choose to apply such anomaly detection methods in order to understand how detectable these rare systems are *without* the need for citizen science verification (e.g. in a low-latency transient detection system). Ultimately, I hope to combine citizen science with machine learning methods in order to search for and characterise transients.

Below I describe broadly how machine learning algorithms work and present a case study in their application to the study of optical transients. This field holds many analogies to the searches for *radio* transients and several of the tools that I use in this chapter were first developed for the finding of (rare classes of) optical SNe.

5.1.2 Machine learning in transient astronomy

Machine learning refers to a large family of related algorithms and processes that have gained huge popularity in data-intensive sciences in the past 15 years. In general, machine learning consists of feeding data to algorithms in order to classify, predict, generate or infer some kind of label, value or structure. If this sounds vague and all-encompassing, that is because machine learning algorithms are a very general purpose set of tools. Fluke and Jacobs (2020) provide an overview of the maturity of a wide range machine learning applications in astrophysics. I will restrict discussion here to immediately relevant efforts in transient astronomy and in relation to citizen science. Nevertheless it is worth commenting on how machine learning

has rapidly been integrated into the toolkit of astronomers - what began as the occasional use of statistical techniques for morphological classifications of galaxies and quasars (Whitmore, 1984; Francis et al., 1992; Lahav et al., 1995) has grown to cover a wide range of applications, from the discovery of extrasolar planets (Pearson, Palafox, and Griffith, 2018; Shallue and Vanderburg, 2018), and the classification of transient objects (this chapter) to the categorisation of radio galaxies (Scaife and Porter, 2021; Slijepcevic et al., 2024), the forecasting of solar activity (Florios et al., 2018; Inceoglu et al., 2018), the assignment of photometric redshifts (Bilicki et al., 2018; Speagle and Eisenstein, 2017) and much more. In just the last 15 years, the number of annual, peer-reviewed entries to the NASA Astrophysical Data System² containing the phrase “machine learning” in the title has increased from < 200 to > 7,000 (between 2008 and 2023, queried in mid April 2024).

In the machine learning landscape there exist two well established families of algorithms, generally used for different categories of tasks, referred to as supervised and unsupervised learning (e.g. Chapter 5 of Goodfellow, Bengio, and Courville, 2016). Supervised learning refers to a class of processes that train on some known, labelled data - or, more precisely, a representation of that data - in order to infer something about new data. The output from a supervised machine learning model can be categorical (a class) or a numerical value. Supervised models can provide point estimates or be probabilistic (Murphy, 2022); Banerji et al. (2010), Gal, Islam, and Ghahramani (2017), and Walmsley et al. (2020) demonstrate how the uncertainty in *Galaxy Zoo* volunteer responses can be used to describe galaxy morphologies with sparse labels and with full probability distributions rather than point estimates.

Unsupervised learning, by contrast, involves *no* training labels. So in some sense, this is not really learning by being taught. Unsupervised learning is perhaps more accurately described as a request for a model to infer some structure, given the data, making these models an exercise in learning by exploring.

As a final note before discussing transient astronomy specifically, I have already made reference to learning from ‘the data’. The form of the data that is fed to

²<https://ui.adsabs.harvard.edu/>

a machine learning task is usually described as a set of features (parameters or variables) and the choice of feature set is critical to the success of a given model. These features might be extracted from the data based on some previously known importance (e.g. the peak magnitude of a supernova) or, for large and complex datasets, learnt from the data themselves. To learn the representation of e.g. an image directly from its input pixels requires much more training and generally more complex models, with multiple layers of processing, referred to as deep learning (LeCun, Bengio, and Hinton, 2015).

In time domain optical astronomy, scientists are concerned with the discovery of new transients such as supernovae. When an optical telescope takes a new picture of a given patch of sky, how can astronomers tell if there is a new transient that was not present in some previous, reference picture? The standard way for this to be done in optical surveys is by calculating ‘difference images’ - a subtraction between a reference picture and a new observation of the same part of the sky (Bond et al., 2001). Therefore any constant flux (e.g. a host galaxy) is removed, ideally leaving only a residual where a transient has appeared. However, difference imaging induces several false positives (Bloom et al., 2012). For example, observational differences between reference and new images can cause poor subtraction of host galaxies, whilst diffraction spikes around bright stars in a new night’s data might not be present in a reference image, producing a spurious detection. These issues have warranted the development of Real/Bogus machine learning classifiers that seek to retain real SNe and remove these artefacts and image processing issues (Bloom et al., 2012). The first examples of this were done by taking a set of features derived from images - e.g. the S/N of a source, or the FWHM of its pixels - and training a classifier to distinguish SNe from bogus detections. These efforts, such as Bailey et al. (2007), du Buisson et al. (2015), and Wright et al. (2015), classified their data using random forests, where a set of binary decision trees split subsets of data according to their training labels and then average over the entire forest, and support vector machines, which try to find a decision boundary that maximally separate each class in some parameter space (Bishop, 2006). Since then,

deep learning has become the paradigm of choice for real/bogus classifiers, where the features used to classify images are learnt from the difference images directly (LeCun, Bengio, and Hinton, 2015). This is generally done using convolutional neural networks (CNNs), where image pixels are convolved with some filter(s) before being fed to a set of non-linear operations to produce a final classification (Nielsen, 2015). Learning directly from the data can be a powerful way to extract meaningful features and these models often outperform their non-deep counterparts but this generally requires many more labelled examples on which to train and increases the computational cost of training. These CNNs have been built into the pipelines on programmes such as Pan-STARRS1 (Wright et al., 2017), the Gravitational-wave Optical Transient Observer (GOTO; Killestein et al., 2021), the Zwicky Transient Facility (ZTF; Duev et al., 2019) and MeerLICHT (Hosenie et al., 2021).

Once the real/bogus classifiers have eliminated most spurious detections, transient astronomers can then move on to the classification of a given object, typically with the aid of a spectrum to confirm its distance and its chemical composition. However, this presents another problem, as there are many more transients detected each night than can be observed spectroscopically. Depending on a given science case, it may be of interest to spectroscopically confirm every type Ia SNe, due to their standardisable nature and use as a distance measure in cosmology (Riess et al., 1998; Perlmutter et al., 1999). However, if one wishes to make discoveries of rarer transient classes such as kilonovae (Smartt et al., 2017) then getting spectra for many Ia's is not in one's interest (and vice versa). In either case, the classification of transients based on their light curves alone is one way of deciding which to send for further observations. Many photometric classifications were developed using the Photometric LSST Astronomical Time-Series Classification Challenge (PLAsTiCC; Boone, 2019; Hložek et al., 2023), as well as data streams from ZTF (Mahabal et al., 2019).

I draw specific attention to Lochner et al. (2016), where a novel decomposition of supernova light curves with wavelet functions provides accurate features for classification for a range of models, including random forests, support vector

machines and neural networks. These wavelet features have since been used at radio wavelengths along with contextual data (e.g. position with respect to the galactic plane or any multiwavelength coverage) for transient classification (Sooknunan et al., 2021). That these general wavelet based features have successfully been used in radio astronomy was part of the motivation for their use in section 5.3.

As with the previous step in the transient discovery pipeline, again deep learning provides data-driven feature extraction for the task of light curve classification. This has typically been done with a model known as the *recurrent* neural network (RNN), which features not only the primary, feed-forward element common to all NNs, but also a lateral sequence i.e. information about previous points of the input sequence, making them natural models for light curve classification. These RNNs allow for classifications on incomplete light curves, updating their output class probabilities as more data points are added, a significant improvement on previous models - knowing that a transient is rare whilst it is still observable is far more valuable than only finding out when it has faded. RNNs have been used for extracting features and classifying supernovae in both simulations (Charnock and Moss, 2017; Moss, 2018; Muthukrishna et al., 2019; Moller and de Boissiere, 2020) and observations (Villar et al., 2020; Burhanudin et al., 2021), as well as in the discovery of variable stars (Naul et al., 2018; Dékány and Grebel, 2020).

The above two examples of machine learning in transient astronomy have been supervised, where either simulations or previous observations inform a classifier. However, a classifier trained only on type Ia and type II SNe will fail to classify a type Ib/c. What then, should an astronomer do when searching for rare classes of transients? This is where unsupervised learning and specifically anomaly detection models have become useful tools for astronomers dealing with large datasets. For example, Pruzhinskaya et al. (2019) and Malanchev et al. (2021) demonstrate the applicability of isolation forests and local outlier factor algorithms (particular algorithms for anomaly detection, see section 5.4) to detecting rare supernovae classes and novel transients. Deep learning is again used for feature extraction concerning anomaly detection. For example, Villar et al. (2021) use a self-supervised

autoencoder framework, a kind of symmetric RNN architecture that forces the representation of the data through a small bottleneck, to extract 10 features that are then used for anomaly detection. However, deep learning architectures are not always the answer - for example, Muthukrishna et al. (2022) demonstrate that their overflexibility can make them less suitable for detection of anomalies when compared to parametric models.

One potential pitfall of these anomaly detection techniques is that a specific algorithmic definition of an anomaly may not match what the end-user (i.e. a science team) deem to be interesting. Practitioners can make use of active learning, in which an algorithm can query a human annotator (sometimes called an “oracle”) to label a small subset of data from which the learner can train. This querying can serve to update existing labels (for supervised tasks) or provide an additional recommendation layer on top of a standard unsupervised learning algorithm. If the machine learner can choose the data for the annotator to label, the hope is that performance will increase by labelling only a small but informative subset of the data (Settles, 2012; see also Walmsley et al., 2022b’s active learning implementation for *Galaxy Zoo*). By incorporating domain knowledge into machine learning models to refine their performance one can, for example, separate anomalous instrumental effects from bona fide astrophysical phenomena. This active learning for anomaly detection has been tested on simulated data in preparation for LSST and real data streams from ZTF and the Open Supernova Catalog (Ishida et al., 2021; Pruzhinskaya et al., 2023).

The `ASTRONOMALY` package (Lochner and Bassett, 2021) provides a general framework for discovering anomalies in astronomical datasets, which can be used to test different aspects of the discovery process. `ASTRONOMALY` has three main scientific steps consisting of feature extraction, anomaly detection and active learning, whilst also providing data management, pre-processing and visualisation tools. Exploration of how these different stages perform, with different features and algorithms applied to a novel dataset form sections 5.3–5.5 of this work, though I note that they are all built into `ASTRONOMALY`’s framework. `ASTRONOMALY`

has been used to discover variable and flaring stars in optical surveys (Webb et al., 2020; Webb et al., 2021), as well as rare and unique morphologies of optical galaxies, mergers and radio sources (Etsebeth et al., 2024; Mohale and Lochner, 2023; Lochner et al., 2023).

Despite the advances in appropriate feature use, anomaly detection and active learning in optical surveys, these tools have yet to be applied to the rapidly-growing field of wide-field radio transient surveys. In the previous chapter I have shown that many variables and transients are not outliers in the distributions of standard variability parameters used to detect these sources. As I am interested in finding such transients in large datasets (e.g. from the SKA), these findings were part of the motivation for studying and applying anomaly detection techniques. Anomaly detection models are well suited to large datasets as unsupervised learning does not require the costly and time-consuming labelling of data nor the training of large models. Therefore, by being fast and scalable, anomaly detection models could play an important role in both the finding of rare and interesting transients as well as in the handling of the PB-scale data to come from new surveys.

In this chapter I describe the first exploration of applying unsupervised anomaly detection techniques to large samples of light curves from radio telescopes. I make use of the same 8874 light curves as described in Chapter 4, along with their citizen science classifications. I have already mentioned that the lack of labels makes these models fast to implement. By the same token, however, it is hard to verify performance of unsupervised learning methods without some known ‘ground truth’ (see e.g. Giles and Walkowicz, 2019 for the use of Boyajian’s star as a ground truth). Therefore, I make use of the citizen science classifications as ground truth against which to compare anomaly detection methods from the `ASTRONOMALY` package.

5.2 Observations and classifications

This work builds on the data collected and presented in Chapter 4. Here I briefly summarise the observations and key data products used in this chapter.

As described in Chapter 2, the observational data consist of ~ 500 images taken by MeerKAT from 2018 - 2022 at 1.3 GHz. These observations were distributed across 11 fields of the sky and centred on XRBs of interest to the ThunderKAT programme and were typically taken at weekly cadence. The resulting set of images is heterogeneous in number of visits per field and this heterogeneity is entirely determined by the commensal nature of the observations - if an XRB faded into quiescence it was dropped from the weekly schedule. As a result some fields have only two months worth of observations, whilst fields around bright and active systems such as MAXI J1820+070, GRS 1915+105 and GX339-4 have over 50 visits.

Once imaged, light curves are extracted from the observations using the Transient Pipeline (TRAP, Swinbank et al., 2015). Examples of a bona fide transient, a false positive and a stable light curve are all visible in Figure 5.1, demonstrating the diversity of this data. I note that the middle, false positive, example looks qualitatively similar to an interesting transient. As described in Chapter 4, volunteers were also presented with images of each source, which aided in classifying something as an extended object or an imaging artefact. Combining the image data with light curve information would require the additional training of some CNN or else the extensive trialling of parametric image plane features. Given that the sources I am interested in are point sources and that the light curves are the descriptors of variability, I opted to work only with the time series data.

From running the TraP I produced 8874 light curves in total from the 11 XRB fields. As with the images, the number of light curves per field is not uniform, with some fields having many more than others. This is primarily governed by some fields having more diffuse emission, higher noise floors and worse u, v coverage due to their relative elevation from the MeerKAT site.

Each of the 8874 light curves used in this thesis was seen by 10 volunteers. From their votes I have already shown that there are 168 variables and transients found, consisting of 142 newly identified sources and 26 that had been previously studied. For the purposes of this work, I will treat these 168 variables and transients as




Figure 5.1: Example light curves of an XRB (MAXI J1820+070, top), a false positive caused by bad data (middle) and a stable light curve (bottom). My aim is for unsupervised machine learning algorithms to preferentially select light curves similar to the foremost over the latter two. The line and shaded region of each panel denote the mean and 1σ uncertainty of the computed Gaussian Process regression, detailed in section 5.3.3.

the ground truth sample. From hereon in, any reference to transients refers to this sample, representing $\sim 2\%$ of all light curves.

5.3 Feature extraction

In order to make use of anomaly detection tools, features must first be calculated that describe the input light curves. The choice of how to describe a dataset is crucial in machine learning tasks and can generally be split between features based on domain knowledge or experience and deep learning, as discussed in 5.1 In this work I use two domain-informed feature sets which have been used in machine learning searches for variability in previous studies. I chose to use these domain-informed features as opposed to those learnt from the data for several reasons. Firstly, automatic feature extraction methods can have many thousands of parameters (internal weights and biases) to update, requiring extensive training. With this moderate sample of $\sim 9,000$ light curves, taken from heterogeneously observed fields, I cannot expect this data to be representative of all possible underlying variability. Similarly, data-driven feature extraction can be computationally expensive, requiring many CPU- or GPU-hours to train. Finally, by using features tested in other regimes I can compare and contrast their performance between use cases. I will also compare the performance of the two feature sets to a baseline pair of parameters often used in transient detection and variable characterisation (see e.g. Chapters 3 and 4).

5.3.1 Baseline features

As a baseline feature set I make use of η and V as I have already done several times throughout this thesis. As a brief reminder, for a light curve of N flux density measurements $F_i \pm \sigma_i$, observed at frequency ν , these two variability statistics are defined as

$$\eta_\nu \equiv \chi_{N-1}^2 = \frac{1}{N-1} \sum_{i=1}^N \frac{(F_{i,\nu} - \overline{F_\nu}^*)^2}{\sigma_i^2} \quad (5.1)$$

and

$$V_\nu \equiv \frac{s_\nu}{\overline{F_\nu}} = \frac{1}{\overline{F_\nu}} \sqrt{\frac{N}{N-1} (\overline{F_\nu^2} - \overline{F_\nu}^2)} \quad (5.2)$$

where s , \overline{F} and $\overline{F_\nu}^*$ denote standard deviation, mean and weighted average of the flux density respectively.

Once calculated, I standardise each feature x by computing

$$z = \frac{x - \mu}{\sigma} \quad (5.3)$$

i.e. by subtracting the mean μ and normalising by standard deviation σ . This is a common approach in machine learning applications (e.g. Czech, Mishra, and Inggs, 2018) and is carried out so that features with large absolute ranges (or measured in different units) do not dominate. For example, if a model required computation of distances between datapoints in a 2-D dataset where one feature had a range $\mathcal{O}(10^4)$ whilst the other had a range $\mathcal{O}(1)$, the former would dominate the distance calculation. This standardisation is done for every feature used in this chapter.

5.3.2 Feets features

The next set of features used comes from the *feets* package (Cabral et al., 2018a; Cabral et al., 2018b), in turn inspired by the FATS code (Nun et al., 2015). These features have an implementation built into ASTRONOMALY's framework, making them ready to test for a given application out-of-the-box. The *feets* package combines over 40 features used for time series analysis, listed in Table 5.1, which can be split into approximately three subsets. First, some features are simple statistical quantities, such the mean, amplitude and standard deviation of a light curve. A second subset of features have been used for the classification of variable stars (Richards et al., 2011) and quasars (Kim et al., 2011) and include, for example, the skew, kurtosis and the power-law index of the light curve's structure function. Finally, several features are derived from a Lomb-Scargle periodogram (Lomb, 1976; Scargle, 1982) and its resulting best fit to each light curve. Specifically, each time series is modelled as a sum of components

$$y_i(t|f_i) = a_i \sin(2\pi f_i t) + b_i \cos(2\pi f_i t) + b_0 \quad (5.4)$$

with amplitudes a_i , b_i , offset b_0 and frequency f_i , which is searched for periodic variations by minimising

$$\chi^2 = \sum_k \frac{(F_k - y_i(t_k))^2}{\sigma_k^2}. \quad (5.5)$$

Defining

$$\chi_0^2 = \sum_k \frac{(F_k - \mu)^2}{\sigma_k^2}, \quad (5.6)$$

where μ is the weighted mean, the generalised Lomb-Scargle periodogram is then

$$P(f) = \frac{N-1}{2} \frac{\chi_0^2 - \chi_m^2(f)}{\chi_0^2} \quad (5.7)$$

where $\chi_m^2(f)$ is the χ^2 at frequency f , minimised with respect to a_i , b_i and b_0 . The amplitude and phase features listed in Table 5.1 refer specifically to the those of the first three dominant frequencies and their four harmonics i.e.

$$y(t) = ct + \sum_{i=1}^3 \sum_{j=1}^4 y_i(t|j f_i) \quad (5.8)$$

where, following Debosscher et al. (2007), each dominant frequency and its harmonics are found and then modelled out iteratively for subsequent values of i . The amplitudes and phases listed in Table 5.1 are calculated with the coefficients a_i and b_i of equation 5.8. Finally, I note that there are some features in this codebase that require colour information but as the data used here are monochromatic these features cannot be employed in this implementation and they are therefore dropped.

5.3.3 Wavelet features

Initial tests with just the *feets* features produced performance that showed room for improvement. Therefore a third and final set of features were used, called wavelet features, which have been used to describe and classify supernovae (Lochner et al., 2016; Alves et al., 2022). However, Sooknunan et al. (2021) successfully showed how wavelet features can be used to construct accurate classifiers for a wide range of light curves from different astrophysical classes (e.g. AGN or KNe), demonstrating

³[/www.mathworks.com/help/wavelet](http://www.mathworks.com/help/wavelet)

Feature	Description	Inputs	Reference
Amplitude	Half the difference between the median of the maximum 5% and the median of the minimum 5% fluxes	F_i	Richards et al., 2011
Anderson Darling test	A test statistic for non-Gaussianity	F_i	Kim et al., 2009
Auto correlation length	Length of linear dependence of a signal with itself at two points in time, taken when the ACF reaches $1/e$	F_i	Kim et al., 2011
Beyond1Std	Percentage of points beyond one standard deviation from the weighted mean	F_i, σ_i	Richards et al., 2011
Con	The number of three consecutive data points that are brighter or fainter than 2σ and normalized by $N - 2$	F_i	Kim et al., 2011
η^e	$\bar{w} \sum w_i (F_{i+1} - F_i)^2 / (\sigma^2 \sum w_i)$ where $w_i = 1/(t_{i+1} - t_i)^2$	F_i, t_i	Kim et al., 2014
Fourier Amplitudes	$A_{i,j} = \sqrt{a_{i,j}^2 + b_{i,j}^2}$ for $i = 1, 2, 3$ and $j = 1, 2, 3, 4$ (see text)	F_i, t_i	Debosscher et al., 2007
Fourier Phases	$\phi_{i,j} = \arctan\left(\frac{b_{i,j}}{a_{i,j}}\right) - \phi_{1,1}$ for $i = 1, 2, 3$ and $j = 1, 2, 3, 4$ (see text)	F_i, t_i	Debosscher et al., 2007
Gskew	Median-of-magnitudes based measure of the skew	F_i	Richards et al., 2011
Linear Trend	Slope of a linear fit to the light curve	F_i, t_i	Richards et al., 2011
Lomb-Scargle Period	The period of the largest periodogram peak	F_i, t_i	Kim et al., 2011
Lomb-Scargle η^e	η^e as applied to the phase folded light curve	F_i, t_i	Kim et al., 2011
Lomb-Scargle R_{CS}	R_{CS} as applied to the phase folded light curve	F_i, t_i	Kim et al., 2011
MaxSlope	Maximum absolute slope between two consecutive observations	F_i, t_i	Richards et al., 2011
Mean \bar{F}	The mean flux - $1/N \sum_i F_i$	F_i	Kim et al., 2014
Mean Variance	The ratio of the standard deviation to the mean flux - $\sigma/\bar{F} \equiv V$	F_i	Kim et al., 2011
Median Abs. Dev.	The median discrepancy of the data from the median - $\text{med}(F_i - \text{med}(F_i))$	F_i	Richards et al., 2011
Median Buffer Range %	Fraction of points with amplitude a tenth of the median flux	F_i	Richards et al., 2011
Pair Slope Trend	The fraction of increasing first differences minus the fraction of decreasing first differences	F_i	Richards et al., 2011
Percent Amplitude	Largest percentage difference between either the max or min flux and the median.	F_i	Richards et al., 2011
Q_{3-1}	The difference between the 3rd and 1st quarterlies	F_i	Kim et al., 2014
R_{CS}	Range of cumulative sum	F_i	Richards et al., 2011
Skew	$\frac{N}{(N-1)(N-2)} \sum (F_i - \bar{F}/\sigma)^3$	F_i	Richards et al., 2011
Slotted Auto Correlation Function Length	Slotted auto correlation length	F_i, t_i	Protopapas et al., 2015
Small Kurtosis	$E[(F - \bar{F}/\sigma)^4]$	F_i	Richards et al., 2011
Standard Deviation	Standard deviation of fluxes	F_i	Richards et al., 2011
Structure Function Index	The exponent of the $SF \propto \tau^\beta$ where $SF(\tau) = \langle [F(t) - F(t + \tau)]^2 \rangle$	F_i, t_i	Hughes et al., 1992

Table 5.1: Features used in this work as calculated by the *feets* package (Cabral et al., 2018b). Each light curve consists of $N = \sum_i$ measurements of flux density $F_i \pm \sigma_i$ taken at time t_i . A similar version, for the original FATS code is described by Webb et al. (2020).




Figure 5.2: The first eight symlet wavelets, used in this work to build a feature set for describing light curves. Taken from the Mathworks wavelet introduction.³

their applicability to the data used here. These wavelet features arise from the understanding that time series $y(t)$ can be decomposed into a sum of basis functions

$$y(t) = \sum_i a_i \Phi_i(t) \quad (5.9)$$

with coefficients a_i , where $\Phi_i(t)$ might for example be a set of sines and cosines (i.e. a Fourier decomposition). I can then use the coefficients of these basis functions as a set of features for later use. Here, I use the symlet family of wavelets seen in Figure 5.2 as the basis functions due to their successful use in previous work (Lochner et al., 2016; Alves et al., 2022). Tests by Lochner et al. (2016) have shown that the choice of wavelet family makes minimal difference to final classification accuracies of supernovae classes.

These wavelet features are calculated without some pre-existing idea of canonical variability in mind and merely fit functional forms to the data. One might therefore hope that they describe the observed light curves well, whilst the *feets* features, whose calculations are all made with some existing class types in mind, may not describe the data in a way that is conducive to this anomaly detection task. Furthermore, wavelets functions localise power in both time and frequency, whilst for example a Fourier decomposition is infinite in time (Ivezic et al., 2014, Chapter

10). This makes them more suited for describing time-localised phenomena and not simply periodic signals.

I calculate the wavelet coefficients for the light curves using the SNMACHINE codebase developed by the LSST Dark Energy Science Collaboration (Lochner et al., 2016; Alves et al., 2022). In order to do this, I first interpolate my light curves with a Gaussian Process Regression (Rasmussen and Williams, 2006), allowing us to account for gaps in the uneven time series. To do this, all 8874 light curves are shifted to take place between 0 and t_{max} days, where for this dataset $t_{max} = 1300$ days, corresponding to the longest light curves in the GX339–4 field. The underlying procedure built into SNMACHINE fits Gaussian processes with a mean of zero and attempts two covariance kernels, keeping the kernel with the better fit. The two kernels tested are an exponential square kernel of the form

$$k_{\text{exp}}(t_i, t_j) = \exp\left(-\frac{|t_i - t_j|^2}{2}\right), \quad (5.10)$$

and a combination of $k_{\text{exp}}(t_i, t_j)$ with an exponential sine squared of scale parameter Γ and period P i.e.

$$k_{\text{sin}}(t_i, t_j) = k_{\text{exp}}(t_i, t_j) + \exp\left(-\Gamma \sin^2\left[\frac{\pi}{P}|t_i - t_j|\right]\right). \quad (5.11)$$

I then evaluate the resulting regression every ~ 5 days (i.e. at 276 instances), approximately matching the typical observational cadence. Example Gaussian process regression fits can be seen in Figure 5.1.

Once all light curves are interpolated by the regression, the wavelet transform is calculated, producing a set of 275 highly redundant coefficients. In order to reduce this high dimensionality, I run a Principal Component Analysis (PCA, Pearson, 1901), which transforms the data into an orthogonal basis. In this now-orthogonal data space, I preserve the 40 largest eigenvalues for each light curve, which capture 99.99995% of the summative variance of the data. This results in reducing the data space from 276 features per light curve to only 40, while retaining almost all of the variance. The choice of 40 principal components is a balance between reducing the dimensionality of the features whilst trying to

preserving enough information such that the anomalies, which constitute a small minority of all of the light curves, are retained.

It is worth noting that extracting these sets of features is the most computationally expensive part of the analysis. However, even the slowest step, the Gaussian process regression, runs in 10 minutes for the entire database on a standard desktop using a 3.5 GHz quad-core processor. All of these features, once calculated, can be saved and these routines then need not be run again. This is an important consideration when considering scalability for larger datasets or for a real-time transient detection system.

5.3.4 Feature space and clustering

An obvious question to ask is how well each feature set separates transients (anomalies) from the rest of the distribution. For a perfect feature set, one would want all transients to be separated cleanly from all non-anomalies. To inspect this I make use of a Uniform Manifold Approximation and Projection (UMAP; McInnes et al., 2018), a technique for dimensionality reduction. UMAP attempts to project data from a N-D space (where N is 45 and 40 for the *feets* and wavelet features respectively and 2 for η and V) down into 2-D space for visualisation, while preserving local structure. I stress here that I use UMAP *only* for the purposes of visualisation and some minor interpretation of qualitative behaviour and that the results of this work do not rely on its inner machinations, assumptions or calculations.

The *feets* feature space for the dataset can be seen in Figure 5.3, where stars are the citizen-scientist verified anomalies (transients) and the colour bar indicates to which observational field a given light curve is from (see section 5.2). It can be seen that sources in a given field are clustered together. This is likely due to all observational systematics varying on a field-by-field basis. For example, the SAX J1808 field only received 6 observations and therefore exists in a very different area of parameter space to other fields. Whilst this might be of concern, these 2-D representations are not used for downstream tasks, but rather as sanity checks and in order to judge qualitatively how well the data (anomalies/transients and everything

else) are described by a given feature set. Moreover, for a real-time or low-latency detection system, what is of more importance is how well transients separate from background data *within* a given field - as in such a system, it is only against sources in the same observation that one would compare. Indeed, we can see that anomalies are somewhat clustered and that, whilst they tend to lie in their respective field of light curves, they are often on the edges of each region. A demonstration of this can be seen for the GX339–4 field shown in the inset of Figure 5.3.

Similarly, the feature space of the wavelet coefficients in Figure 5.4 shows field-dependent clustering. However, in this case, there are clear strings of points in feature space that correspond to the observational fields with longer light curves (e.g. J1858, GX339–4). The majority of these sources have stable light curves and so perhaps they are constrained to existing on these strings due to requiring up to 40 features to describe their shape. By contrast, shorter light curves and variable sources are clustered in the same area of parameter space (inset), however as are some non-anomalous background sources. This distinct cluster of data points may give insight into the performance of the wavelet features in ASTRONOMALY’s models (see sections 5.4 and 5.5).

Finally, for the baseline feature set (η, V) , we know that many transients and variables are inliers, as demonstrated in Chapter 4. For comparison the feature space of these two parameters is shown in Figure 5.5. The presence of systematics in some fields (e.g. MAXI J1820) can be seen, along with the difficulty in separating the transients (anomalies) from all background sources. This was one of the main motivations for testing other feature sets for anomaly detection.

5.4 Anomaly detection

Having calculated sets of features for the light curves, I now move on to the core anomaly detection step of ASTRONOMALY. To do this I pass the features into two common anomaly detection models, detailed below. These models provide a ranked list of the sources from most-to-least anomalous. I can then use the positions of our citizen-science verified transients in our anomaly-ordered list as a metric of success

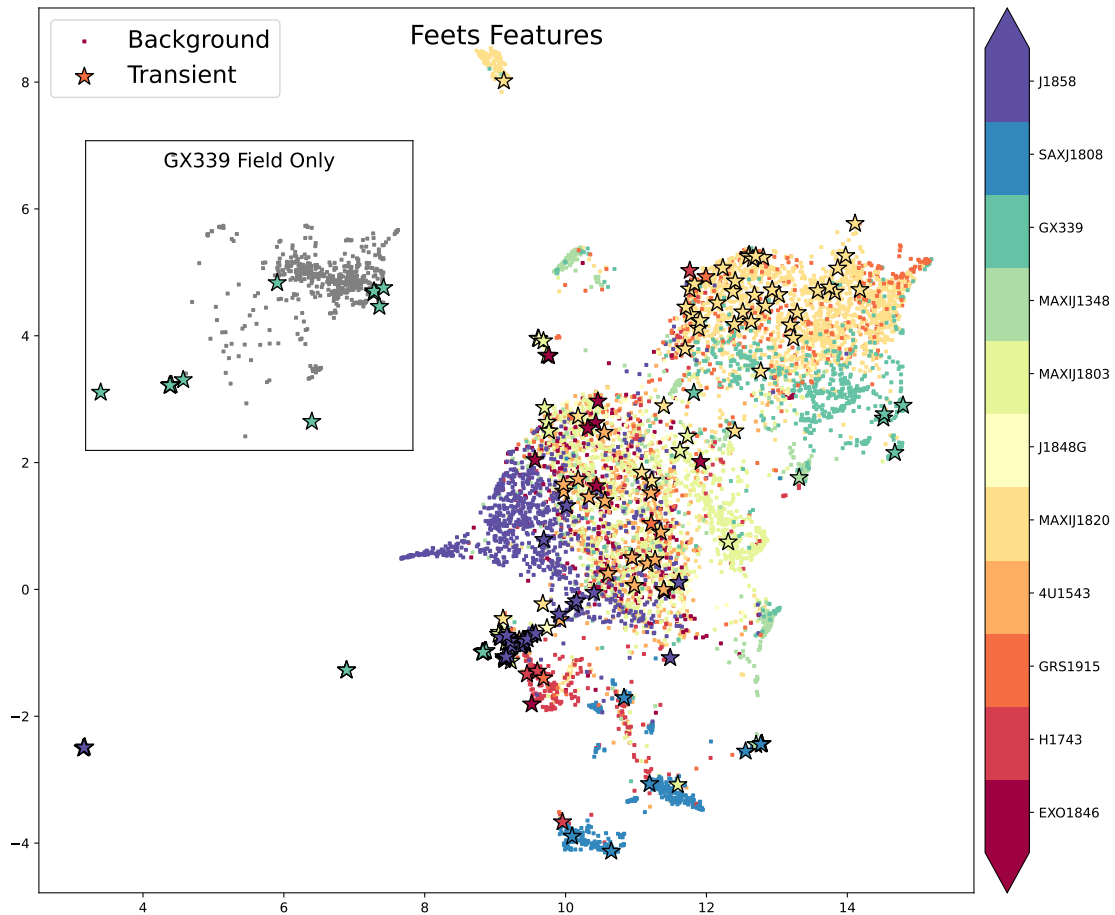


Figure 5.3: The UMAP of the *feets* feature space, colour-coded by observational field. Stars indicate volunteer-verified anomalies (transients). We see that clustering is mainly a function of observational field, aligning with the heterogeneous nature of the dataset. I note that the axes are unlabelled as these numbers are arbitrary combinations of many features and have no physical meaning. The inset is a demonstration of how the data of one particular field (GX339) is distributed, as mentioned in the text.

- the more transients that are ranked as highly anomalous, the better. Throughout this section I am considering all of our volunteer-labelled transients as anomalies.

Anomaly detection is typically done by each algorithm having a definition of what is normal and assigning subjects that are far in parameter space from normal as highly anomalous, without any user labels or training. In this work I explore two anomaly detection algorithms from the `scikit-learn` package (Pedregosa et al., 2011). Once anomalousness has been ranked by a given algorithm, `ASTRONOMALY` scales this to a score between 0 to 5, where 5 is most anomalous. It is important to point out that `ASTRONOMALY` never decides whether or not something is anomalous

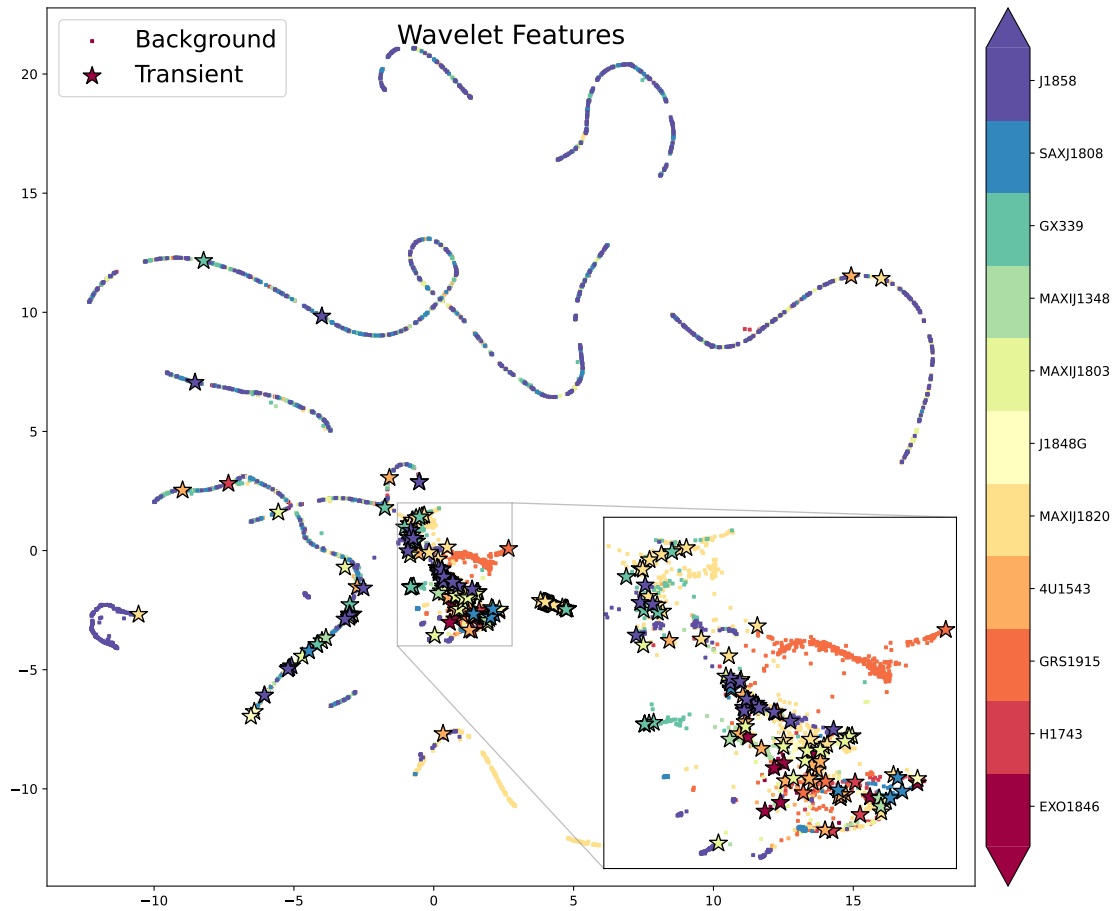


Figure 5.4: The UMAP of the wavelet coefficient feature space, colour-coded and labelled as in Figure 5.3. In this instance, the observational fields with longer light curves (e.g. J1858, GX339–4) show distinct patterns, whilst most other sources are in a very dense region of parameter space, as shown by the inset zoom.

and simply orders the entire list by anomaly score. It is only thanks to the volunteer classifications that I can make statements about, for example, how many true anomalies are ranked highly in this list.

The first of the two anomaly detection methods used herein is the Local Outlier Factor (LOF, Breunig et al., 2000). The LOF is a density-based method that measures the deviation of the density of a given sample with respect to some user-defined number of nearest neighbours, k . I set $k = 100$ throughout this work, approximately 1% of the sample. I tested values of k between 20 (the default) and 200, which produced changes in performance at the sub-percent level. I note that without labelled data it would be difficult to optimise this (and other) hyperparameters. The LOF is local in that the resulting anomaly score depends on




Figure 5.5: The (η, V) feature space, colour-coded by observational field as in previous figures. This is the same as Figure 4.6 but with a different colour-coding.

how isolated the object is with respect to the surrounding neighbourhood. Formally, the local density of A , for a set of k nearest neighbours $N_k(A)$, is defined as

$$\rho_k(A) = 1 / \left(\frac{\sum_{B \in N_k(A)} d(A, B)}{|N_k(A)|} \right) \quad (5.12)$$

where $d(A, B)$ is the distance between A and B . The LOF is then

$$LOF_k(A) = \frac{\sum_{B \in N_k(A)} \rho_k(B)}{|N_k(A)|} \frac{1}{\rho_k(A)} \quad (5.13)$$

i.e. the LOF of datum A is a measure of the average local density of all k neighbours divided by A 's own local density. One advantage of this local approach is that it might pick up anomalies missed by global density measurements. That is, a point only a small distance from a very dense cluster might be an anomaly, whereas a point in a sparse cluster might still have the same density as its neighbours and so

not be an outlier despite its low absolute density value. For more on the distinction between global and local anomalies, see the discussion of Rogers et al. (2024).

The second method used is the Isolation Forest (IF; Liu, Ting, and Zhou, 2008). IF takes inspiration from the supervised random forest algorithm (RF; Breiman, 2001) and operates on the core principles that anomalies are both few in number and different compared to normal data. In an IF an ensemble of decision trees randomly split random features, growing each ‘branch’ in the tree until all data considered are isolated. Given that anomalies should be few and different, they are easier to isolate than normal data that are close to each other in feature space. Therefore data with shorter decision branches are more likely to be anomalous, whilst those that take many splits to isolate are more similar to the rest of the sample. The final anomaly score is an average score over all trees in the forest. More specifically, $h(x)$ is the number of edges that datum x traverses to reach a terminal node i.e. the path length in a single decision tree. For a forest of average path length across all trees of L , the anomaly score is then calculated as

$$s(x) = 2^{-\frac{E(h(x))}{L}} \quad (5.14)$$

where $E(h(x))$ is the average path length of x from a collection of decision trees. For a large average path length $s(x) \rightarrow 0$, whilst for a small path length $s(x) \rightarrow 1$. This model is particularly applicable to high dimensional data as there is no costly calculation of distances and therefore has a linear time complexity (i.e. twice the input data takes twice the runtime). By comparison, LOF can take up to $\mathcal{O}(N^2)$ (Breunig et al., 2000).

5.4.1 Anomaly recall and purity

In order to assess the performance of our models one might ask how many anomalies are returned within some n sources in the ranked list? This quantity is known as the recall and is defined as

$$R = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}, \quad (5.15)$$

where for us the denominator is the total number of transients (anomalies) in the data, i.e. 168. This recall value, whilst simple, corresponds well with preferable performance consisting of many more anomalies high in the list for the user to see quickly. I can calculate the recall R as a function of n until the end of the dataset is reached ($n = 8874$), by which time all sources in the list will have been recovered. I define R_{10} and R_{30} as the recall of transients in the top 10 and 30% of the list respectively, which I will use as quantitative measures of performance throughout this section.

The recall curves for each feature set applied to both anomaly detection models can be seen in Figure 5.6. These are normalised on both axes i.e. $x=0.1$ corresponds to the top 10% most anomalous sources and $y=0.5$ corresponds to half of the 168 variables found by volunteers. Better performance corresponds to curves with higher gradient at low n , appearing pushed up into the upper left corner. It can be clearly seen that (η, V) show worse performance on this dataset, essentially matching random chance, in line with Chapter 4's findings. Contrastingly, it can be seen that the new feature sets are able to recover a majority of volunteer-verified transients (anomalies) in a small subset of the data. This is the first demonstration of anomaly detection as applied to radio time series.

When comparing performance between anomaly detection models (i.e. for the same feature set), minimal variation and qualitatively similar behaviour between recall curves can be seen. For the *feets* and wavelet features a 'broken power-law' behaviour can be seen, with an approximately constant gradient initially, followed by a knee and then worse performance. Table 5.2 lists all R_{10} and R_{30} values for model-feature set pairs. The largest discrepancies between models occur at the 'top' of the list, with R_{10} values varying by 0.1 for *feets* features and 0.14 for wavelets. However, R_{30} values differ from 0.03 to 0.07, indicating that model choice most affects the highly anomalous sources in the dataset. It is worth noting that several iterations of each feature-model pair were tested, with performance never varying by more than $\sim 1\%$, so I have rounded all values to this precision.




Figure 5.6: Upper panel: The recall curves for both of the anomaly detection algorithms as applied to three sets of features derived from the radio light curves. The dotted line in each corresponds to $y = x$ i.e. random chance. **Lower panel:** A combined plot of all six recall lines for better comparison between feature sets. Dashed lines make comparisons (see text) for recall rates at 10 and 30% of the dataset. The values in the legend correspond to the area under each recall curve (AUC).

	<i>feets</i>		Wavelets		(η, V)	
	LOF	IF	LOF	IF	LOF	IF
R_{10}	0.52, 0.52	0.42, 0.45	0.24, 0.32	0.38, 0.42	0.14, 0.16	0.12, 0.11
R_{30}	0.74, 0.73	0.78, 0.78	0.91, 0.91	0.84, 0.84	0.34, 0.32	0.31, 0.30
P_{10}	0.11, 0.11	0.09, 0.10	0.05, 0.06	0.07, 0.08	0.03, 0.03	0.02, 0.02
P_{30}	0.05, 0.05	0.05, 0.05	0.06, 0.06	0.05, 0.05	0.02, 0.02	0.02, 0.02

Table 5.2: Specific recall (R) and purity (P) values, rounded to two decimal places, for each model-feature set pair, measured after ten and thirty percent of the most anomalous sources have been ‘seen’. Each pair of values listed are the default, **active** iteration of that feature-model instance.

Differences in feature set have a much greater effect than changing the anomaly detection method - for the same model R_{10} changes by factors of ~ 3.5 between best- and worst- performing feature sets. Similarly, R_{30} for the baseline parameters reaches only ~ 0.32 , whilst the wavelet features can recover up to 91% of volunteer-verified transients. This underscores the importance of using features that best describe a given dataset.

I can also calculate an aggregate recall performance by simply integrating the area under each curve (AUC) of each model-feature set’s recall curve. The corresponding AUC values for each model-feature set instance are listed in the legend of Figure 5.6. As already noted, differences between feature sets are greater than differences between anomaly detection models. These AUC values show that the wavelet-IF combination has the best recall performance at an AUC of 0.83. However, this is a summary over all the data and does not consider the actual finding of transients in a live setting. By this I mean that one must consider how many anomalies a human will be able to find in a reasonable amount of time. As discussed above and can be seen in the lower panel of Figure 5.6, if a user can only inspect the top 10% of the dataset, then the feets-LOF analysis performs the best, with $R_{10} = 0.52$. However, if there are resources available to inspect 30% of the dataset (computation, citizen scientists, a team of experts) then the LOF performance on wavelet features recalls 91% of transients, 7% better than the next-best feature-model pair ($R_{30} = 0.84$).

We can also ask what fraction of the top n sources are transients. This is

known as the precision or purity, defined as

$$P = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}. \quad (5.16)$$

This can be thought of as a measure of the hit-rate of our anomalies; if an individual were to look through the top 10 or 30% of the data, how many transients (anomalies) would they find? I note that, by definition, as n approaches the total length of the dataset (8874 sources), the precision will tend towards the aggregate rate of transients in the data i.e. $\sim 2\%$. The precision values for all feature-model pairs can be found in Table 5.2, again looking at 10 and 30% of the anomaly ordered lists. For the *feets* and wavelet feature sets, the precision is always at least a factor of 2 better than the corresponding baseline performance.

In general, the same trends hold as with the recall values - differences between feature sets are greater than differences between anomaly detection models. Likewise, we can see that the *feets* features produce the best results when looking at the top 10% of the data, with a hit-rate of over 1 in 10, compared to the baseline of 1 in 50. The results 30% of the way through the anomaly ordered lists, as seen with the recall values, show that the wavelet-LOF instance performs the best, but only marginally. Nevertheless, we can see that for either feature set, as applied to both of the anomaly detection algorithms, the hit-rate of anomalies in the data is always a factor of 2 or greater when compared to the baseline. This directly translates into a reduction in the number of sources that would need to be sifted out before finding a new transient.

In summary, the choice of model used makes a small difference, though LOF is marginally preferred. When considering feature sets, the baseline features perform the worst, recovering transients almost uniformly through the anomaly-ordered lists. The *feets* features perform the best at the beginning of each anomaly-ordered list, recovering 52% of transients in 10% of the data and increasing the purity above random by a factor of 5. However, the wavelet features perform better when considering a larger fraction of each anomaly-ordered list. Using the wavelet features over 90% of anomalies are recovered in 30% of the data, at a purity level that is three times higher than the baseline features.

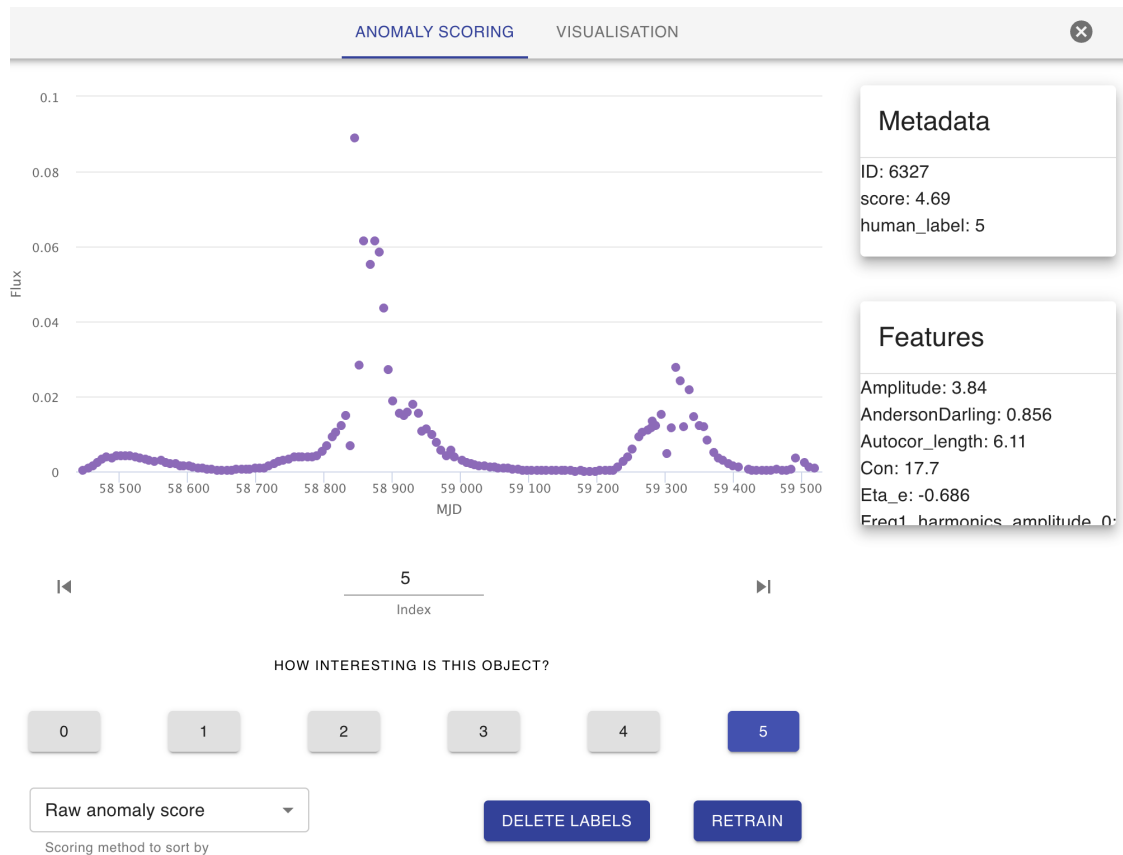


Figure 5.7: The Astronomy GUI, showing subjects in anomaly-order. Using this interface, I can easily inspect the data and provide labels on some subset of objects in order to perform active learning.

5.5 Active learning

Once anomaly scores have been calculated for all data, ASTRONOMALY allows for interactive visualisation by use of a GUI, seen in Figure 5.7. This GUI displays sources listed in anomaly order, displaying the raw data, its features and any associated metadata. From this interface a user can label sources on a range of 0 to 5. This allows for active learning, whereby an additional model balances the raw anomaly score with user input and calculating a trained anomaly score. This then acts as a recommendation engine, displaying sources with higher trained anomaly scores preferentially to the user, speeding up anomaly detection.

The core of this active learning process is described by ranking objects with a new score

$$\hat{S} = S \tanh(\delta - 1 + \operatorname{arctanh}(\tilde{U})) \quad (5.17)$$

where S is the default anomaly score, δ represents a distance penalty term and \tilde{U} describes a relevance score given by

$$\tilde{U} = \epsilon_1 + \epsilon_2 \left(\frac{U}{U_{\max}} \right). \quad (5.18)$$

Here ϵ are normalisation constants and the user input score is U , normalised by $U_{\max} = 5$. Since U will only ever be input by the user for a small fraction of a dataset, its value for all remaining sources is estimated by a random forest regression calculated with 100 estimators. This regression will have uncertainty, quantified by δ , which is large when an object is far from a human-labelled neighbour and small when near to labelled data. In summary this means that, in regions far from human-labelled data ($\delta \gg 1$), \hat{S} reverts to its original score as $\tanh \rightarrow 1$. Similarly, if $U \sim U_{\max}$ then S will not change due to high human ranking. Conversely, regions of feature space near labelled data ($\delta \ll 1$), with low user scores will be downweighted, corresponding to relegating uninteresting data to lower in the anomaly-ranked list.

In this work I actively label the top $\sim 2\%$ of each feature-model pair, which takes of order 10 minutes. That is, examples of clear variability or known transients were ranked highly ($U = 4$ or 5), whilst low-significance light curves, or those showing clear false positive behaviour - e.g. where one epoch shows large changes in flux due to a different calibrator or worse image quality - were ranked at 1 or 0. However, it is worth noting that this labelling is inherently subjective and dependent on a user's end goal. An instrument scientist interested in flux calibration would vote differently to an astronomer interested in stellar radio flares. Active learning attempts to quantify this and balance an algorithmic definition of anomalous with a user-dependent idea about what is an *interesting* anomaly. Finally I note that the regressor used to model user interest is by default a RF but that using a Gaussian process provides an improvement on this for some datasets, which will be incorporated into future versions of ASTRONOMALY (Walmsley et al., 2022b).

5.5.1 Active learning improvement

To quantify how much active learning changes the results, I subtract each updated recall curve from those in section 5.4.1. The residual curves for all feature-model pairs can be seen in Figure 5.8. We can immediately see that the wavelet feature instances improve dramatically, increasing recall by as much as 10% within the most anomalous region of feature space. Active learning applied to the *feets* features improves anomaly recall for the most anomalous sources, but once $\sim 15\%$ of the data have been analysed, minimal improvement is seen. This is reflected in the respective R_{10} and R_{30} values seen in Table 5.2 which remain almost unchanged to two decimal places. This difference in improvement between the *feets* and wavelets features may be related to the clustering seen in Figures 5.3 and 5.4. The baseline features show no improvement and in some cases perform worse than before. This is in part because the active learning procedure involves labelling the sources at the top of the anomaly-ordered list. Given that the anomaly detection performance using the baseline features essentially matches random chance and that transients make up $\sim 2\%$ of the sample, labelling only 2% of the data will not provide the regressor with many positive anomalies from which to learn.

The precision values of Table 5.2 show minimal improvement following active learning. Only the P_{10} values improve by, at most, 1%, whilst all P_{30} values remain the same to two decimal places. This is in part due to the asymptotic behaviour of these precision calculations as mentioned earlier, namely that as n reaches the total sample size, the precision must approach the overall rate of transients in the data of $\sim 2\%$. As a result, it is hard to improve the P_{30} values much when there are so few anomalies to find in the dataset.

In terms of AUC values, minimal changes are seen compared to the default anomaly detection procedure, all *feets* and (η, V) values varying at the sub-per cent level and wavelet instances improving by only 0.01. This matches what we see in Figure 5.8 where after $x=0.3$ all curves vary only minimally around no change ($y=0$). As before I stress that this metric does not match the user-driven experience of finding anomalies preferentially in as small a data volume as possible. As before,




Figure 5.8: Changes in performance due to active learning for each model-feature pair, where $y = 0$ indicates no improvement. This plot has been truncated at $x = 0.4$ as, after this point, all curves vary minimally about 0. I show two figures for clarity when comparing between curves.

the feets-LOF iteration provide the most transients in 10% of the data, though the wavelet instances perform almost as well following the active learning.

5.6 Discussion

In this work I have demonstrated the first application of anomaly detection to the task of finding transients with radio interferometers. In general we see that appropriate features, models and domain knowledge via active learning allow us to recover over half of the volunteer-verified transients (anomalies) in one tenth of the total data volume. Similarly, I have shown that the number of sources a human must inspect in order to find transients can as reduced by a factor of 5 compared to baseline performance. The choice of model used makes a small but noticeable difference, with the LOF algorithm generally performing better. The

choice of feature or representation is a much stronger determinant of which and how many anomalies can be recovered. For example, when dealing with the large and heterogeneous sample of data used in this work, it is clear that a simple 2-parameter setup fails to represent the data in a way that allows the subsequent anomaly detection models to succeed. This is not a surprise, given that η and V are intended for intra-field comparisons i.e. between light curves in a single set of observations. Indeed these are still useful statistics that have aided in the discovery in numerous transients from the ThunderKAT observations.

The *feets* features show good promise at recovering transients and can recall over half of the anomalies in less than 10% of the data. The wavelet features, after active learning, perform almost as well as the *feets* features with respect to R_{10} and are by a significant margin the best feature set when comparing R_{30} values. The large improvement seen in the wavelet features may be related to the clear clustering of anomalies seen in Figure 5.4, where the majority of the transients are located in a relatively small part of this representation. Therefore we might conclude that if active learning sufficiently promotes sources from this region of feature space, performance will increase.

I note that one could spend a great deal longer curating the perfect feature set for a given task and there is no guarantee that the features tested here form the optimal choices. However, that pre-existing methods work well in this regime is a significant step towards finding radio transients in real-time surveys.

With regards to scalability, all processing has been performed on a standard CPU with no attempt at optimisation. Despite this, all feature extraction and anomaly detection took less than 10 minutes, with the Gaussian process regression being the rate-limiting step. Even when considering the additional step of producing light curves from images in the first place, the overall run time is low enough to allow for a low-latency online transient detection system. Commensal real-time searches already exist on some interferometers, such as the *realfast* system on the Karl G Jansky’s Very Large Array (Law et al., 2018) and on the Murchison Widefield Array (see Methods in Hurley-Walker et al., 2023). This work shows that

the incorporation of anomaly detection techniques would be a flexible addition to such a backend system, whereby the anomaly ranking and user-defined preference would allow for rapid discovery without ballooning data rates.

However, additional steps will be required in order to maximise potential discoveries from current and upcoming data streams. For example, several recent discoveries have been made using fast images (that is, sub-integrations of a single long observation) that warrant further investigation (Hurley-Walker et al., 2022; Caleb et al., 2022; Wang et al., 2023), whilst herein the images are typically separated by a week. I plan to apply these anomaly detection methods to sets of 8s-images taken by MeerKAT in the near future. Similarly, I have not made use of any spectral or polarimetric information inherent in the data, both of which can provide crucial additional information and reduce contaminants (Heywood, 2023). However, it is impossible to know on which timescales and at which frequencies to search in a manner that is completely agnostic to all the possible underlying astrophysics. Therefore it is worth thinking about the specific science case in mind and how to best leverage these data products - for example, if one is searching for radio flaring stars, the use of Stokes V imaging is a powerful tool that eliminates many contaminants from images (e.g. Pritchard et al., 2021; Pritchard et al., 2024). These additional pipeline steps would, however, add to the computational time of finding these anomalies. It is therefore down to the given science case whether data rates are such that only the bare minimum information can be extracted, or whether the astrophysical phenomena necessitate additional processing.

Finally, I note again that all of the discoveries used as the ground truth were found by volunteers. With upcoming (big) data projects such as the SKA and LSST, it is worth discussing the role of individuals and groups with regards to discovery science. A crude but immediate comparison between human discoveries and anomaly detection can be seen in Figure 5.9, where the transient vote fraction is how many out of 10 volunteers that saw each source voted for it as a transient/variable. The anomaly ranking is a source's position in the anomaly-ordered list for the wavelet-LOF instance, though other feature-model pairs were tested with qualitatively

similar results. It can be seen that the majority of volunteer transients are ranked highly by the anomaly detection model, matching the recall curves seen earlier. However, we can also see that there are known transients and variables missed by volunteers (i.e. below Chapter 4’s threshold of 0.4) - these are typically low significance sources or those that show variability that was difficult for volunteers to identify but were found in previous studies i.e. Driessen et al. (2020), Rowlinson et al. (2022), and Driessen et al. (2022). It is encouraging then that the anomaly detection methods rank these anomalies as highly as many of the new transients/variables confirmed by volunteers, as this provides a method of recovering sources that may be missed by other techniques. Similarly, there are sources ranked low in the anomaly detection algorithms (even with active learning) that were clearly identified by a majority of volunteers - these are the data points about the $y = x$ line in Figure 5.9. So these are anomalies that volunteers find “more easily” than the algorithms. It is therefore pertinent to leverage each search method in a complementary manner (see also Wright et al., 2017). For example, one could perform a first pass with anomaly detection procedures, which would find the highly anomalous sources easily and then pass the sources that are not ranked highly (e.g. beyond the knee of a given recall curve) to citizen scientists. However, it is also important to value the time taken by individual volunteers and consider whether asking them to find relatively few sources in large data volumes is an optimal or even fair division of labour. There exist Zooniverse projects where the hit-rate is lower than the one presented here (e.g. *Planet Hunters*’ $<0.7\%$ Schwamb et al., 2012, compared to our $\sim 2\%$) which are engaged with by many thousands of volunteers. Therefore, as long as the odds are clearly communicated to contributors, volunteers have the agency to decide if they want to contribute. In any case, the combination of rapid computation and the advantages of volunteer-led searches in concert with expert analysis and, as demonstrated here, anomaly detection techniques, hold great promise for future transient hunting efforts.




Figure 5.9: A comparison between how many citizen scientists voted for a source as a transient and its ranking in the anomaly detection work, where a rank of 1 indicates the most anomalous source. That is, the most anomalous source in my anomaly score ordered list has a ranking of 1, whilst the least anomalous source has a ranking of 0. Sources ‘not recovered’ refers to known variables/transients in the light curves that fall below the 4/10 threshold used in Chapter 4 and were therefore regarded as missed by volunteers. The background points are all remaining sources, coloured by the density of points. The anomaly list used here corresponds to the wavelet features applied to the LOF algorithm, but results are consistent for all feature-model pairs.

5.7 Conclusions

In this chapter I have presented the first analysis of anomaly detection models and active learning for finding radio transients via their light curves. This was motivated by the rare but interesting transients and variables found in radio surveys. I have shown how the volunteer contributions on the *Bursts from Space: MeerKAT* citizen science project can act as a ground truth sample against which to test these unsupervised machine learning techniques incorporated into the `ASTRONOMALY` package.

I have shown how choices in feature sets result in different clustering of sources and how the appropriate representation of data has a strong impact on the ability to recover anomalies. I have also shown that choice of anomaly detection model is a less significant determinant in performance. Also, for any combination of feature set and model, the hit-rate of transients is always at least a factor of two greater than previous methods, for highly anomalous sources. This is highly encouraging for the necessary acceleration of transient searches with modern (radio) telescopes.

Active learning on a small subset of the data produces better recall of transients, particularly for those ranked as highly anomalous. This enables users to focus on anomalies that are interesting for their given science case. The increases from active learning were not so large as to be an absolute necessity, meaning that instances without the possibility for such active learning will still gain the majority of the benefit from these methods.

These promising methods, in combination with the volunteer-led efforts and commensal search engines being built into telescope backends will allow for unprecedented discovery of rare and novel transients that will drive science cases for upcoming facilities such as the SKA. In the future I will be applying these techniques to surveys coming from MeerKAT, including in combination with upcoming citizen science data releases.

*I don't know where I'm going from here, but I promise
I won't bore you*

David Bowie

6

Conclusions and Future Work

Contents

6.1	Stellar activity	153
6.2	Citizen science and transient searches	155
6.3	Anomaly detection with large surveys	157
6.4	Final remarks	158

This thesis aimed to show how commensal data and novel methods can be used to find interesting radio transients with MeerKAT. Below I summarise the research discussed in this thesis and how each science chapter relates to these goals. I also describe how I will continue working in each area moving forwards and make some concluding remarks.

6.1 Stellar activity

In Chapter 3 I discussed how the detection of stars is relatively rare in widefield radio searches, but these nearby sources of transient emission are vital probes of the magnetic fields and plasma densities of stellar systems. They also provide clues to their local space weather, relevant to, for example, habitability studies of hosted exoplanets. This is particularly important for M dwarfs, due to their relative abundance, proximity and high activity rates.

During my studies I found a new radio flaring star as part of a project to find serendipitous transients in images from the MeerKAT radio telescope. I did this by making use of commensal observations from the ThunderKAT survey, searching for transients in a set of test observations. This led me to investigate the multiwavelength counterpart of this transient, making use of dedicated facilities such as MeerLICHT as well as widefield surveys such as TESS. Spectroscopic follow-up with SALT confirmed the source to be a magnetically active, rotating M dwarf, at a distance of 12pc. The lack of circular polarisation measurements limited the conclusions I could make about the emission processes occurring on this star. At the time of discovery, this radio-loud star was the third radio-loud stellar system found by ThunderKATs. Since then, at least four more active stellar systems have been uncovered thanks to the unique capabilities of MeerKAT, as demonstrated in Appendix A.

The wealth of data available in the MeerKAT archives are prime hunting grounds for further detections of stellar activity. I plan to do this by making use of the data stored by the Breakthrough Listen User Supplied Equipment (BLUSE) on MeerKAT, in which $\sim 1,000,000$ nearby stars are being observed commensally (Czech et al., 2021). The Breakthrough Listen project's main focus is the search for extraterrestrial intelligence via the detection of communication technology and technosignatures. This BLUSE backend will observe almost all stars within ~ 2500 pc that are visible to MeerKAT. I plan to use the skills and knowledge gained from my work on stellar activity to make use of these data for finding radio-loud stars. It should be possible to take the data products of the BLUSE and form images at the positions of all nearby stars observed by MeerKAT, which could provide radio detections or upper limits on the entire local volume of stars visible from the southern hemisphere. As mentioned above, the use of circular polarisation can be invaluable in determining the emission processes at play and so I will fold this learning through to detected these SETI star searches. Surrounding these stellar targets there are approximately 2600 known exoplanets¹, each of which will be affected by any space weather caused

¹Queried at <https://exoplanet.eu/> on 18-04-2024

by a host star. Understanding how the activity of a host star affects its host planet's atmosphere and habitability therefore makes such a search relevant to exoplanetary studies, in addition to the search for technosignatures. I also note that imaging products from this backend would also be a useful form of RFI rejection as part of technosignature searches. By this I mean that RFI appears significantly different in images to real far-field point sources and so imaging with this backend would be of use to several SETI-related goals. Finally, the BLUSE swatches and imaging products could be integrated into citizen science efforts, though there already exist SETI-focused citizen science projects². This characterisation of the local radio sky will be unprecedented and of relevance to our understanding of stellar activity, coherent emission processes and in the search for technosignatures. I have already shown in this thesis how modest sampling of the radio sky with MeerKAT already opens up this search area and so future searches look to be very promising.

6.2 Citizen science and transient searches

In order to scale-up commensal searches with MeerKAT I planned, built, beta-tested and launched my own citizen science project on the Zooniverse platform, as presented in Chapter 4. *Bursts from Space: MeerKAT* has engaged over 1000 members of the public to date in finding novel transients. The first data release (DR1) was analysed at a rate of 1 click/minute for 90 days, i.e. much faster than an individual scientist could manage. This has produced a catalogue of ~ 150 new transient and variable sources that had never been classified before, including transient OH maser systems, image-plane detections of pulsars and many variable AGN. Furthermore, I have shown that simple statistical cuts on the data fail to find these systems, showing that the transients that volunteers find are an under-studied population. To date, this is the first and only example of citizen science as applied to radio images for transient astrophysics. Working alongside the thousands of volunteers who have taken part in *BfS:MKT* has been a uniquely joyful experience.

²www.zooniverse.org/projects/ucla-seti-group/are-we-alone-in-the-universe

It has allowed for both novel science as well as engagement with volunteers all around the world, some of which is described in Appendix A.

I am involved in several current and upcoming transient searches moving forwards. Firstly, as part of my citizen science project, we have had over 120,000 classification of deep, long-timescale observations of the LADUMA field, forming our second data release (DR2). The third citizen science data release will consist of $\sim 60,000$ images of the same LADUMA field taken at 8-s cadence (DR3), allowing for unprecedented searches for both faint, long-timescale variables and fast, bright transients, all in one of the most commonly visited extragalactic fields (e.g. by Chandra; Giacconi et al., 2002). It has been demonstrated that fast imaging on timescales of 2- and 8-seconds with MeerKAT can detect and immediately localise all FRBs with fluences similar to those published by the CHIME telescope (Rhodes et al., 2023a; Driessen et al., 2024) making this data promising for FRB detections with excellent host galaxy information. An 8s snapshot with MeerKAT reaches the typical sensitivity of the VLA's FIRST survey (0.15 mJy; Becker, White, and Helfand, 1995) whilst even a 2s exposure reaches ~ 0.45 mJy (equivalent to the NVSS; Condon et al., 1998).

I am also part of the MeerKAT-Max Planck Galactic Centre Survey, which is utilising the new S-band receiver to image the Galactic Centre. I will be directly involved in using anomaly detection methods to search for transients from this dataset, including looking for image-plane detections of pulsars and other coherent emitters (e.g. Hyman et al., 2005). This project will search for transients on timescales from 8s to months, in contrast to the weekly ThunderKAT images used here, as well as making use of circular polarisation calibration, two features that have not been present in this work. Finally, I am also leading a search for transients in MeerKAT data of five bow shock and star-forming complexes (Proposal code MKT-22079). The recent image plane discoveries of ultra-long period systems, including in the field of a bow shock (Caleb et al., 2022) has generated interest in these complex and young parts of the Milky Way. Therefore searches in both the Galactic Centre and in young complexes could be fruitful grounds for building the population of these new and rare systems. The transient search methods that I have carried out

in this thesis provide a set of methods that could be applied to these and other datasets to search for new sources. In short, there is no end in sight for the kinds of searches that have, in recent years, uncovered a range of interesting radio transients.

6.3 Anomaly detection with large surveys

Modern astronomical observatories are becoming data-intensive machines, as exemplified by both the upcoming Vera C. Rubin Observatory’s Legacy Survey of Space and Time (Ivezić et al., 2019) and the existing surveys on MeerKAT. These big data projects create an enormous needle-in-a-haystack problem which is only exacerbated when looking ahead to the SKA. Many of the relevant machine learning tools have been developed for the LSST (e.g. SNMACHINE; Alves et al., 2022), which promises to deliver millions of alerts every night to astronomers. In Chapter 5 I described how I have applied such anomaly detection algorithms and feature extraction routines to a set of ~ 9000 radio light curves, all of which have volunteer data. The ASTRONOMALY framework (Lochner and Bassett, 2021) has allowed me to apply the isolation forest and local outlier factor algorithms to these feature sets in order to test anomaly detection methods. I have shown that, depending on the choice of feature and anomaly detection model, over half of all transients identified by volunteers can be recovered in less than 10% of the data, within 10 minutes of light curve ingest. Furthermore, the active learning within ASTRONOMALY produces a 10% improvement on our recall of anomalies when labelling only 2% of the data. This work provides a crucial step for real-time discovery of transients with radio telescopes and is the first exploration of anomaly detection in this regime.

Moving forwards, I think anomaly detection techniques will be incorporated into the transient searches at all wavelengths. As I have shown, algorithms such as IFs and LOFs are computationally very inexpensive but provide excellent agreement with citizen science votes, assuaging fears of missing interesting transients. One challenge may be ensuring that the data are described appropriately and, as with any unsupervised learning task, it is impossible to know a priori what the optimal features might be. I also do not see anomaly detection entirely replacing humans

from the discovery process. As shown with ASTRONOMALY, active learning is able to provide improvements on basic anomaly detection, honed for a specific science case. I therefore hope that active learning strategies are developed further in astronomy, allowing for smarter (more efficient) labelling of data, as well as more interesting data with which to engage with volunteers. Future work may investigate whether, for example, autoencoder architectures (e.g. Villar et al., 2021) provide a way to extract meaningful features directly from the radio time series.

In the future, some part of the Rubin Observatory’s datastream will be plugged into the Zooniverse architecture. This will provide a wealth of observational data in which to search for rare anomalies in an overwhelming stream of millions of alerts per night. The multi-epoch, multi-band photometry provided by Rubin will revolutionise the search for weird and wonderful time variable phenomena. These searches could be aimed at rare interstellar objects such as ‘Oumuamua (Meech et al., 2017), of which LSST will be poised to find hundreds, with anomaly detection predictions already being investigated (Rogers et al., 2024). In addition to solar system anomalies, Rubin data is naturally well suited for other space oddities such as rare SNe classes. Using the Lasair broker (Smith et al., 2019), led by collaborators in Oxford, would be an obvious area to dovetail these unique methods with the unprecedented survey to come.

6.4 Final remarks

In this thesis I have presented work on the finding and characterisation of transients with the MeerKAT telescope, the development of a citizen science project in order to find further transients and the first application of anomaly detection to this field. I will continue to work in these three areas - finding (stellar) transients, citizen science and anomaly detection - and their intersection moving forwards. One of the broadest and hopefully clearest conclusions I arrive at after my studies is that using new facilities and applying novel methodologies to a problem can produce valuable insight into a range of interesting phenomena. It is therefore an exciting time to be working in these areas of astrophysics when facilities such MeerKAT

and ASKAP are already finding interesting transients and as the Vera C. Rubin Observatory and the SKA are due to come online. Of course, only post hoc will it be possible to say if we are truly transitioning from an antediluvian world into a sea of interesting data. Nevertheless, the future looks bright. The aphorism “the harder [we] work, the luckier [we] get” springs to mind, where in the astronomer’s case the “work” constitutes the accruing of experiences and tools that allow for systematically serendipitous discoveries, whilst the lucky payoff is the opportunity to contribute, alongside fellow stargazers, to literally the largest puzzles of them all.

Appendices



Transient searches and citizen science in South Africa

Contents

A.1 Transient Searches	161
A.2 Engagement in Carnarvon, Cape Town and online . .	163

During the course of my DPhil studies I was lucky enough to spend a significant amount of time in South Africa, working alongside colleagues at the University of Cape Town (UCT), the University of the Western Cape (UWC) and at the South African Radio Astronomy Observatory (SARAO). Below I touch on some of the work I did alongside students at UCT to help develop commensal ThunderKAT transient searches, as well as engagement work centred on my citizen science project alongside SARAO, UWC and UCT.

A.1 Transient Searches

During my DPhil I was lucky enough to spend three months living in Cape Town, working at UCT and with collaborators at UWC (this collaboration forms the core work of Chapter 5). At UCT I worked with three Masters and three Honours students, each of whom had a set of MeerKAT observations in which to find

transients. These included MeerKAT fields from ThunderKAT though others were from other large survey projects, including MIGHTEE and MHONGOOSE (Jarvis et al., 2016; de Blok et al., 2016).

As described in Chapter 3, the use of the Transients Pipeline (TraP; Swinbank et al., 2015) underpins the majority of the search for commensal image transients with MeerKAT. Even though, as I have shown, the variability statistics computed by the TraP struggle to exhaustively find all transients in our data, they still provide a good jumping off point and can be suitable when working with a single dataset. Furthermore, the TraP provides a convenient source finder and light curve builder for then developing one's own analysis (e.g. a citizen science platform or an anomaly detection pipeline).

Using to my experiences working with the TraP I ran three, 2 hour workshops with these MSc and Hons students - affectionately known as the UCTTraPers - detailing the kinds of astrophysical transients we can find with these observations, the actual hands-on running of the TraP and the analysis of its output. Over the rest of my time at UCT and since then, I have worked with these students when they have had thoughts or roadblocks on the way through their investigations.

As of the time of writing this thesis, these students have all made key contributions to discoveries of radio transients with MeerKAT. I leave all detailed analysis to the theses and papers to come from these students but wanted to highlight in this appendix the joys of working collaboratively with other early career astronomers to make discoveries. One of these students led the radio observations of GRB 190114C, seen by the MAGIC telescope (MAGIC Collaboration et al., 2019) and made use of the TraP for characterising the field of observations, building the light curve for the GRB and searching for any additional transients. The other students have found or helped to find radio transient emission from flaring stars. From these students, including Peters (2023), Mlangeni et al. *in prep*, Nyamai et al. *in prep*, along with the work of Driessen et al. (2020), Driessen et al. (2022), Chastain et al. (2023), and Andersson et al. (2022) we now have a small sample of stellar transients from the ThunderKAT team and so I here display briefly the results of this collection.

In Figure A.1 the optical-radio plot of Stewart et al. (2018) is shown, overlaid with the stellar transients found by the ThunderKAT team. The majority of these are M dwarfs, with subtypes M0 - M5 covered, along with the probable (but not confirmed) RS CVn binary of Driessen et al. (2020) and the giant star found by Chastain et al. (2023). These line up nicely with the stellar region of this parameter space, pushing towards fainter in both the radio and optical directions. Furthermore, Figure A.2 shows a comparison between these stars found by ThunderKAT with those detected by ASKAP, as described by Pritchard et al. (2024). These ASKAP stars span a wide range of types, including dMes, dK stars, Algol and RSCVn binaries, young stellar objects and white dwarfs though I have not labelled all these types for the sake of clarity. Furthermore, this work utilised circular polarisation to identify many of these active systems over very large areas of the sky (thousands of square degrees). By comparison, the ThunderKAT observations, though only scanning (relatively) small areas of the sky, have been able to see fainter sources i.e. further away for the same luminosity. To me, this highlights the complementary power of observing both broad areas of the sky and smaller, deeper fields - akin to the tiered, wedding cake strategies of extragalactic surveys (see e.g. Grogin et al., 2011). It is my hope that over the next few years the community will be able to populate this parameter space, making use of facilities such as ASKAP and MeerKAT.

In summary, the experience of working with others and making novel discoveries in this space has been a richly rewarding part of my DPhil and it would therefore have been remiss of me not to include it in my thesis. The other ThunderKATs with whom I have worked have made interesting discoveries and have gone / are going on to do further great science and I look forward to collaborating further in the future.

A.2 Engagement in Carnarvon, Cape Town and online

During my studies I had the opportunity to use my citizen science platform, as described in section 4, to engage with communities in South Africa. Specifically, along with a team from SARA0 and translators from UWC I had the opportunity to




Figure A.1: The mean optical and radio flux densities of the radio stars found by the ThunderKAT team, atop an underlying distribution of astrophysical classes from (Stewart et al., 2018).




Figure A.2: Detections of radio stars found in ThunderKAT observations, compared with recent work from the ASKAP telescope. The dotted, dashed and solid lines correspond to 5σ detection limits for the data from Pritchard et al. (2024), 15 minutes with MeerKAT and 1 hour with the SKA-Mid array, i.e. 1.25 mJy, 60 and 22 μ Jy respectively.

run a three day workshop with students aged ~ 16 from Carnarvon, the nearest town to the MeerKAT site in the Northern Cape. SRAO, the organisation that runs MeerKAT, already have existing links with many operations in Carnarvon, including in the High School (CHS). As stated in the SRAO Anniversary Report (Schröder et al., 2024), “schools in this region... serve financially constrained communities that cannot afford to pay school fees. The schools therefore rely completely on the funds they receive from the Department of Education, which only covers the most basic costs. In addition, due to the remote location of the schools, it is an ongoing challenge for schools to attract and retain qualified educators, especially for Mathematics and Physical/Natural Sciences.”

The opportunity to provide direct access to data coming from the telescope on the proverbial doorstep of the school was immediately of interest to me and, with SRAO and their human capacity development team, I devised, developed and delivered a workshop centred around the second data release on *BfS:MKT*. DR2 consisted of approximately 12,000 sources detected in the Chandra Deep Field South images taken by the MIGHTEE survey, one image of which can be seen in Figure A.3.

12 students took part in the workshop, essentially those who had opted to take mathematics and physics to the end of their secondary education at CHS. This workshop took place in Cape Town and learners got the opportunity to look at how MeerKAT is operated, what science can be done with radio telescopes and visited the historic optical observatory across the road at SAAO. The workshop sessions themselves began with the fundamentals of radio astronomy and discussing the kinds of transient-focused science that can be done with MeerKAT (see Chapter 1), before describing my citizen science project and the Zooniverse more generally. Over the three days, the students did over 3000 classifications as part of DR2, which I hope indicates that they were engaged with and enjoyed the project, as well as each taking part in a range of other Zooniverse projects. I also invited three students from UCT to discuss their backgrounds and paths into science with the learners.

In addition to this focused project, DR2 was launched and marketed more broadly as part of South Africa’s National Science Week, with significant help from




Figure A.3: One epoch of the Chandra Deep Field South, taken as part of the MIGHTEE survey with MeerKAT. It is light curves of these sources that volunteers classified as part of DR2.

the Department of Science and Innovation. As part of this push, I spoke to local radio stations about the project and how to get involved with citizen science. I also got the opportunity to be heard on episodes from two podcasts that were released around this time - *Looking Up*¹ and *The Cosmic Savannah*². This was a great chance to share *BfS:MKT* A histogram of the combined number of classifications taken over the first two months of DR2 are shown in Figure A.4, where it can be seen that

¹<https://www.fmr.co.za/podcasts/looking-up/>

²<https://thecosmicsavannah.com/>

every opportunity to share *BfS:MKT* resulted in an uptick in classifications. The project was also assisted by a team from UWC and the IAU's Office of Astronomy for Development, who provided translations from English into 5 of South Africa's 11 official spoken languages - IsiXhosa, Sepedi, Sesotho, Tshivenda and Xitshonga. This was a significant undertaking, for which I am very grateful and I hope that the volunteers who have taken part in the project in their languages have found the experience enjoyable. Finally I note that the website traffic data shows that for DR1 there were fewer than 32 page visits from South African IP addresses to the classification workflow for *BfS:MKT*. By contrast, DR2 received 572 visits from 185 different IP addresses in South Africa, showing that hundreds of visits were made and that they were made repeatedly.

In summary I was lucky enough to be able to share my Zooniverse project with learners from Carnarvon High School and with the public more broadly via radio broadcasts and podcasts online. This was an immensely rewarding experience and has been essential for the development of my communication and outreach skills. More importantly, the learners with whom I engaged and statistics on page visits and classifications to the project show that there is an appetite for citizen science in these communities and I hope that it has been an engaging experience for those who have taken part and with whom I have had the privilege of working. I will continue this kind of engagement work with all of my future citizen science efforts.

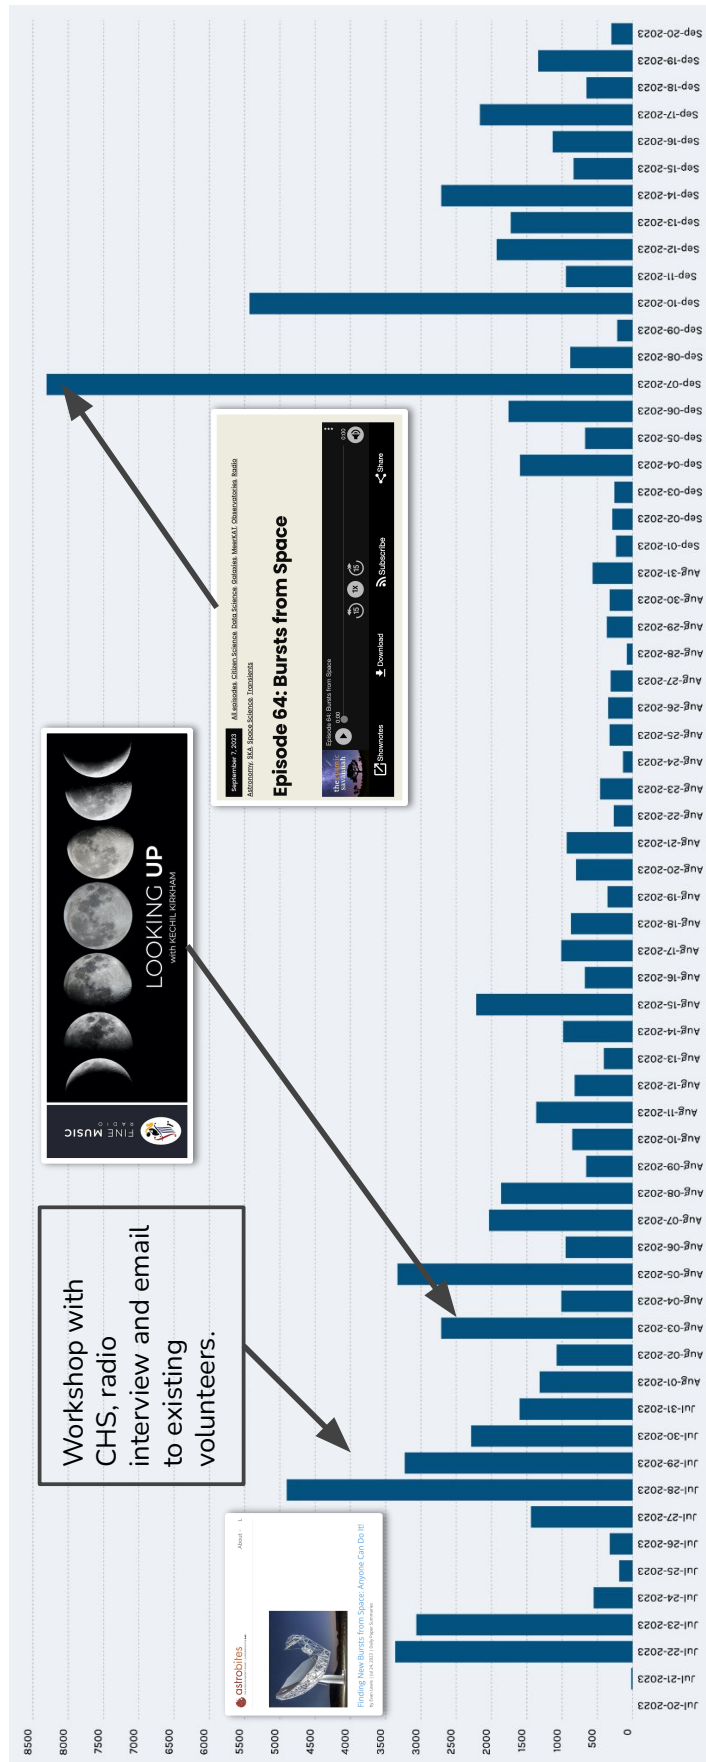


Figure A.4: Number of classifications per day over the first two months of DR2 of *Bfs:MKT*. Each uptick corresponds to the release of an article, podcast or engagement session.

B

Bursts from Space:MeerKAT results table

Table B.1: Table of transients and variables found by volunteers during BfS:MKT, with their associated positions, median flux densities (F_{med}) and variability parameters η and V . The observed date given here is that in which the source was detected at highest S/N by the TRAP, whilst the distance recorded is how far an entry is from the pointing centre of that observation. The Transient Fraction (TF) is the fraction of classifications as a transient/variable, from 10 volunteers. Note that their Zooniverse subject ID is a unique identifier from which one can find the image and light curve online and freely accessible at <https://www.zooniverse.org/projects/alex-andersson/bursts-from-space-meerkat/talk/subjects/ID>, replacing ID with the numeric value below. This table will be available in machine-readable form online.

Name	Subject ID	RA ($^{\circ}$)	Dec. ($^{\circ}$)	Date Obs.	TF	F_{med} (mJy)	η	V	Distance ($'$)	Known?
BfS 0	70785323	255.424	-48.776	14/04/2018	0.4	13.2	83.7	0.05	11.2	0
BfS 1	70785357	255.364	-48.970	14/04/2018	0.9	11.2	395.4	0.11	17.3	1
BfS 2	70785384	254.710	-48.877	14/04/2018	0.5	4.7	26.8	0.07	39.7	0
BfS 3	70785398	255.441	-48.675	14/04/2018	0.6	0.7	1.8	0.11	12.5	1
BfS 4	70785401	255.369	-48.499	14/04/2018	0.4	0.4	0.6	0.14	22	1
BfS 5	70785440	255.156	-48.946	14/04/2018	0.5	1.7	6.6	0.09	23.6	1
BfS 6	70785461	255.634	-48.850	14/04/2018	0.9	0.6	4.3	0.19	4.6	0
BfS 7	70785641	255.291	-48.597	14/04/2018	0.7	3.9	100.6	0.15	20.1	1
BfS 8	70785723	255.117	-48.429	14/04/2018	0.6	0.6	1.1	0.11	31.9	1
BfS 9	70785821	256.988	-48.972	14/04/2018	0.4	8.8	130.5	0.09	51.8	0
BfS 10	70785881	255.917	-48.670	14/04/2018	0.4	0.5	0.9	0.11	11.1	1
BfS 11	70785898	255.706	-48.790	01/02/2020	0.7	2.5	22427.6	1.86	0	1
BfS 12	70785951	255.983	-48.932	14/04/2018	0.6	3.6	62.6	0.13	13.9	1
BfS 13	70785952	255.557	-48.560	14/04/2018	0.8	1.4	3.6	0.08	15	1
BfS 14	70785973	256.443	-48.806	21/03/2020	0.5	0.2	0.6	0.23	29.2	1
BfS 15	70786097	289.084	10.836	28/02/2021	0.9	0.8	7.6	0.32	18	0
BfS 16	70786101	288.964	10.940	05/04/2021	0.4	0.2	8.7	0.49	9.8	0
BfS 17	70786128	288.798	10.946	24/07/2020	0.6	37	66097.3	1.51	0.3	1
BfS 18	70786366	288.400	10.354	08/12/2018	0.4	8.1	8002.1	0.24	42.3	0
BfS 19	70786703	282.208	-1.497	08/11/2021	0.4	4.8	37	0.46	0.3	1
BfS 20	70786771	282.098	-1.400	28/02/2021	0.7	15.2	5	0.05	8.6	0
BfS 21	70786879	282.207	-1.499	08/11/2021	0.9	15.7	984.2	0.67	0.4	1
BfS 22	70786931	282.080	-1.679	04/10/2021	0.5	80.4	12.5	0.02	13.6	0
BfS 23	70789583	271.374	-30.120	22/05/2021	0.4	9.1	23	0.02	36.4	0
BfS 24	70789643	270.424	-29.801	31/07/2021	0.6	0.5	1	0.07	17.6	0
BfS 25	70789787	271.225	-29.740	07/05/2021	0.6	1.9	2.9	0.03	24.7	0
BfS 26	70789902	270.687	-30.534	12/06/2021	0.6	2.1	2.1	0.02	42.6	0
BfS 27	70789937	271.029	-30.293	15/08/2021	0.5	127.3	99.8	0.02	31.2	0

Continued on next page

Table B.1 – continued from previous page

Name	Subject ID	RA ($^{\circ}$)	Dec. ($^{\circ}$)	Date Obs.	TF	F_{med} (mJy)	η	V	Distance ($'$)	Known?
BfS 28	70790055	270.934	-29.951	12/06/2021	0.5	1.6	2.4	0.04	11.7	0
BfS 29	70790141	270.681	-29.466	12/06/2021	0.4	2	3.8	0.03	22	0
BfS 30	70790174	270.699	-29.561	07/05/2021	0.8	2.4	3.5	0.03	16.2	0
BfS 31	70790384	271.383	-30.411	12/06/2021	0.4	6.2	20.9	0.03	47.7	0
BfS 32	70790438	271.424	-30.144	07/05/2021	0.5	8	11.9	0.02	39.3	0
BfS 33	70790547	270.539	-29.637	19/06/2021	0.7	1.9	8.5	0.05	16.2	0
BfS 34	70790666	270.559	-29.467	05/09/2021	0.4	3	1.8	0.02	24	0
BfS 35	70790672	271.167	-29.189	19/06/2021	0.5	6.2	36.7	0.04	43.7	0
BfS 36	70790754	270.762	-29.830	15/05/2021	0.4	0.1	13026.9	2.38	0.2	1
BfS 37	70790813	270.399	-29.923	12/07/2021	0.4	1	2.2	0.05	19.7	0
BfS 38	70790848	270.833	-29.289	23/10/2021	0.5	0.8	2.4	0.07	32.5	0
BfS 39	70790883	271.545	-30.141	27/05/2021	0.7	2.1	4.5	0.03	44.9	0
BfS 40	70790918	270.650	-29.852	05/06/2021	0.5	2.8	14.3	0.05	6	0
BfS 41	70790962	270.488	-30.022	05/09/2021	0.5	0.3	6.6	0.28	18.4	0
BfS 42	70790975	270.723	-29.525	14/11/2021	0.8	1.7	3.5	0.04	18.2	0
BfS 43	70791011	270.369	-29.511	15/08/2021	0.4	2.5	4.1	0.03	27.9	0
BfS 44	70791023	270.899	-29.293	15/05/2021	0.4	1.5	3.4	0.04	32.8	0
BfS 45	70791355	272.603	-36.434	31/08/2019	0.5	2.3	6.4	0.04	40.3	0
BfS 46	70791434	272.517	-37.488	16/08/2019	0.5	14.9	64.9	0.03	36.1	0
BfS 47	70791686	273.023	-36.568	23/08/2019	0.4	17.7	200.6	0.06	50.1	0
BfS 48	70791761	271.248	-37.247	31/08/2019	0.4	3	9.9	0.05	44.5	0
BfS 49	70791852	272.860	-37.199	31/08/2019	0.4	2.1	6.2	0.04	38	0
BfS 50	70792117	272.207	-37.403	31/08/2019	0.4	2.7	1.5	0.02	25.8	0
BfS 51	70792231	237.067	-48.184	31/07/2021	0.5	1	10.4	0.1	32.6	0
BfS 52	70792375	235.628	-47.398	17/10/2021	0.4	0.9	3.3	0.07	49.7	0
BfS 53	70792445	236.614	-47.322	31/10/2021	0.4	20.4	27.5	0.02	22.2	0
BfS 54	70792463	236.839	-47.273	17/10/2021	0.4	0.3	1.8	0.18	24.1	0
BfS 55	70792507	236.701	-47.932	19/06/2021	0.5	0.7	2.9	0.1	15.9	0
BfS 56	70792542	235.884	-47.378	17/10/2021	0.4	4.7	29.7	0.05	40.6	0
BfS 57	70792578	236.787	-47.672	27/09/2021	0.5	0.2	12.3	0.76	0.2	1
BfS 58	70792640	236.045	-47.552	17/10/2021	0.4	4.2	8.3	0.04	30.8	0
BfS 59	70792656	236.271	-47.244	07/08/2021	0.4	4.3	5.7	0.03	33.1	0
BfS 60	70792673	237.774	-47.166	31/07/2021	0.6	1.3	3.1	0.05	50.4	0
BfS 61	70792680	237.816	-47.732	31/07/2021	0.6	2.6	8.3	0.04	41.8	0
BfS 62	70792689	236.974	-48.085	23/10/2021	0.4	5.8	21.4	0.03	25.8	0
BfS 63	70792709	236.908	-47.147	23/10/2021	0.4	4.1	11.6	0.04	32	0

Continued on next page

Table B.1 – continued from previous page

Name	Subject ID	RA (°)	Dec. (°)	Date Obs.	TF	F_{med} (mJy)	η	V	Distance (')	Known?
BfS 64	70792753	237.413	-47.324	31/10/2021	0.6	8	83	0.06	33	0
BfS 65	70792803	237.946	-47.945	31/07/2021	0.6	0.6	1.4	0.07	49.6	0
BfS 66	70792836	237.318	-48.068	19/06/2021	0.4	0.1	0.7	0.18	32	0
BfS 67	70792856	235.627	-47.468	17/10/2021	0.6	0.6	4.1	0.13	48.5	0
BfS 68	70792874	236.501	-47.879	31/07/2021	0.7	4	31.8	0.06	16.8	0
BfS 69	70792923	235.831	-47.516	05/09/2021	0.5	0.7	0.9	0.05	39.7	0
BfS 70	70792950	237.244	-48.086	23/10/2021	0.6	5.7	23.6	0.03	30.9	0
BfS 71	70792982	237.261	-47.217	31/10/2021	0.6	0.5	6.4	0.2	33.5	0
BfS 72	70792990	237.207	-47.712	31/07/2021	0.4	4	2.9	0.02	17.2	0
BfS 73	70793015	235.650	-47.575	04/10/2021	0.4	0.9	2.7	0.06	46.3	0
BfS 74	70793035	235.482	-47.815	31/10/2021	0.6	1.9	18.9	0.08	53.2	0
BfS 75	70793087	282.321	-3.065	07/09/2019	0.6	1	9528.9	1.53	0	1
BfS 76	70793134	282.315	-2.990	30/11/2019	0.4	74.9	11.7	0.02	4.5	0
BfS 77	70793169	282.957	-3.192	10/04/2020	0.4	5.9	8	0.04	38.9	0
BfS 78	70793209	282.732	-2.533	30/11/2019	0.5	6.7	7.6	0.03	40.3	0
BfS 79	70793222	282.107	-3.352	19/10/2019	0.4	37.6	8.1	0.01	21.5	0
BfS 80	70793281	282.175	-2.841	07/12/2019	0.4	3.6	54.5	0.32	16	0
BfS 81	70793303	282.491	-3.292	10/04/2020	0.4	0.8	3.2	0.21	17	0
BfS 82	70793382	282.355	-3.728	07/09/2019	0.4	13.8	9.9	0.03	39.8	0
BfS 83	70793482	266.565	-32.234	21/09/2018	0.6	0.4	146.3	1.09	0	1
BfS 84	70793565	266.151	-32.564	27/10/2018	0.4	0.6	1.5	0.22	28.9	0
BfS 85	70793663	266.671	-32.234	19/10/2018	0.4	0.2	26.8	1.18	5.4	1
BfS 86	70793688	266.761	-32.828	27/10/2018	0.5	0.8	1.6	0.12	37	0
BfS 87	70793773	266.127	-32.060	28/09/2018	0.4	1	6.7	0.23	24.6	0
BfS 88	70793847	284.523	-9.106	05/08/2019	0.5	2.1	8.9	0.05	52.6	0
BfS 89	70793894	284.524	-7.907	05/08/2019	0.6	3.3	6.4	0.03	21.1	0
BfS 90	70793933	284.260	-8.211	05/08/2019	0.7	2.4	8.1	0.04	22.9	0
BfS 91	70794074	285.144	-8.774	05/08/2019	0.6	0.8	9.2	0.11	43.7	0
BfS 92	70794096	283.910	-8.295	05/08/2019	0.4	0.9	4.2	0.08	43.8	0
BfS 93	70794108	284.531	-8.706	05/08/2019	0.7	3.1	14.9	0.04	28.9	0
BfS 94	70794119	284.483	-8.148	05/08/2019	0.4	5.6	16.9	0.02	11	0
BfS 95	70794176	284.696	-7.772	05/08/2019	0.4	0.6	1	0.07	28.1	0
BfS 96	70794327	284.412	-7.658	05/08/2019	0.6	0.4	1.1	0.08	37.4	0
BfS 97	70794440	284.789	-8.391	05/08/2019	0.4	0.3	0.7	0.11	12.5	0
BfS 98	70794613	285.194	-8.346	05/08/2019	0.5	1.5	5.5	0.05	33.2	0
BfS 99	70794684	284.574	-8.078	05/08/2019	0.7	3.3	30.7	0.05	10.5	0

Continued on next page

Table B.1 – continued from previous page

Name	Subject ID	RA ($^{\circ}$)	Dec. ($^{\circ}$)	Date Obs.	TF	F_{med} (mJy)	η	V	Distance ($'$)	Known?
BfS 100	70794685	284.858	-7.609	05/08/2019	0.7	1.1	55.5	0.23	39.8	0
BfS 101	70794791	284.623	-8.379	05/08/2019	0.9	2.3	21.8	0.05	8.6	0
BfS 102	70794807	285.090	-8.988	05/08/2019	0.8	0.5	5.1	0.14	52.1	0
BfS 103	70794826	284.444	-8.611	05/08/2019	0.9	1.8	12.5	0.06	25.4	0
BfS 104	70794918	285.358	-8.161	05/08/2019	0.6	1.2	6.6	0.07	42.5	0
BfS 105	70794964	284.783	-8.853	05/08/2019	0.4	1.1	2.3	0.05	37.8	0
BfS 106	70795034	284.820	-8.057	05/08/2019	0.8	1.1	6.7	0.07	15	0
BfS 107	70795054	285.000	-8.120	05/08/2019	1	2.3	21.6	0.07	22.2	0
BfS 108	70795071	285.019	-8.570	05/08/2019	0.4	0.5	2.9	0.12	29.8	0
BfS 109	70795100	284.877	-8.091	05/08/2019	0.7	0.7	5	0.13	16.3	0
BfS 110	70795197	284.812	-8.673	05/08/2019	0.8	1.9	8.4	0.06	27.9	0
BfS 111	70795265	284.676	-8.695	05/08/2019	0.4	1.8	2.1	0.03	27.5	0
BfS 112	70795290	284.448	-7.598	05/08/2019	0.4	2.9	14.3	0.05	40.1	0
BfS 113	70795327	284.871	-8.073	05/08/2019	0.8	9	103	0.05	16.6	0
BfS 114	70795334	284.800	-7.904	05/08/2019	0.4	1.8	1.9	0.03	22	0
BfS 115	70795365	284.708	-8.128	05/08/2019	0.8	3.3	15.3	0.03	7.5	0
BfS 116	70795426	284.341	-7.904	05/08/2019	0.6	0.5	2.4	0.11	27	0
BfS 117	70795476	284.619	-7.956	05/08/2019	0.7	1.7	3.4	0.04	17	0
BfS 118	70795913	207.308	-62.922	17/11/2019	0.5	2	2.9	0.07	22.2	0
BfS 119	70795974	207.569	-63.676	09/04/2019	0.4	20.9	50.9	0.04	27.8	0
BfS 120	70796049	207.055	-63.273	09/04/2019	0.7	0.2	434.9	1.72	0.1	1
BfS 121	70796442	207.054	-63.274	01/02/2019	0.4	0.5	2936.6	1.88	0	1
BfS 122	70796444	207.993	-63.692	01/03/2019	0.4	2.3	3.7	0.07	35.6	0
BfS 123	70796614	274.548	7.188	14/10/2018	0.4	4.2	434.7	0.24	32.3	0
BfS 124	70796677	275.168	7.804	14/10/2018	0.6	8	487.9	0.24	37.4	0
BfS 125	70796699	274.869	6.912	14/10/2018	0.4	0.2	1.4	0.29	21.1	0
BfS 126	70796701	274.509	7.811	14/10/2018	0.6	1.3	81.5	0.26	51.1	0
BfS 127	70796797	275.120	6.698	14/10/2018	0.5	0.7	22.5	0.26	29.3	0
BfS 128	70796831	274.919	6.673	14/10/2018	0.4	0.2	3.4	0.4	32.4	0
BfS 129	70796842	274.773	7.007	14/10/2018	0.8	0.6	21.6	0.31	21.7	0
BfS 130	70796845	274.883	7.940	14/10/2018	0.4	0.2	3.3	0.31	47	0
BfS 131	70796863	274.588	7.085	14/10/2018	0.4	0.3	8.9	0.3	30.6	0
BfS 132	70796900	275.511	7.878	14/10/2018	0.4	0.8	28.2	0.31	48.5	0
BfS 133	70796970	274.727	7.523	14/10/2018	0.4	0.9	39.5	0.26	29.7	0
BfS 134	70796978	275.166	7.056	14/10/2018	0.6	0.6	14.5	0.31	8.9	0
BfS 135	70797098	275.110	6.915	14/10/2018	0.4	1.4	63.7	0.27	16.3	0

Continued on next page

Table B.1 – continued from previous page

Name	Subject ID	RA (°)	Dec. (°)	Date Obs.	TF	F_{med} (mJy)	η	V	Distance (')	Known?
BfS 136	70797121	275.105	7.526	14/10/2018	0.4	0.5	6.1	0.28	20.5	0
BfS 137	70797142	275.235	6.442	14/10/2018	0.4	0.2	5.6	0.36	45.4	0
BfS 138	70797182	274.707	6.479	14/10/2018	0.6	2.4	322.8	0.36	48.2	1
BfS 139	70797198	275.717	7.629	14/10/2018	0.5	0.3	4.8	0.31	45.7	0
BfS 140	70797292	274.609	6.771	14/10/2018	0.5	0.4	9.2	0.32	38	0
BfS 141	70797295	274.819	6.766	14/10/2018	0.5	0.4	3.3	0.24	29.9	0
BfS 142	70797361	275.038	7.447	13/11/2018	0.7	0.5	14.2	0.32	16	0
BfS 143	70797376	274.875	7.784	14/10/2018	0.6	8.7	1052.6	0.26	38.2	0
BfS 144	70797391	275.126	7.338	14/10/2018	0.4	0.4	6.3	0.3	9.4	0
BfS 145	70797458	275.649	6.797	14/10/2018	0.7	0.8	8.2	0.32	40.5	0
BfS 146	70797482	274.980	7.739	14/10/2018	0.7	0.7	36.4	0.37	33.8	0
BfS 147	70797595	275.125	6.572	14/10/2018	0.7	3.8	520.9	0.28	36.9	1
BfS 148	70797626	274.716	7.064	14/10/2018	0.6	0.1	2	0.32	23.5	0
BfS 149	70797723	274.544	7.789	14/10/2018	0.4	0.3	8.7	0.28	48.7	0
BfS 150	70797725	275.532	7.768	14/10/2018	0.5	1.1	35.8	0.26	43.7	0
BfS 151	70797752	275.091	7.186	05/10/2018	0.7	0.2	1546.9	1.78	0	1
BfS 152	70797802	275.526	6.549	14/10/2018	0.6	0.6	10.2	0.27	46.1	0
BfS 153	70797809	275.818	7.101	14/10/2018	0.5	0.5	3.2	0.28	43.6	0
BfS 154	70797837	274.955	7.512	14/10/2018	0.4	0.9	18.1	0.25	21.2	0
BfS 155	70797847	275.930	7.263	14/10/2018	0.6	0.7	28.6	0.3	50.2	0
BfS 156	70797867	274.860	7.798	14/10/2018	0.7	0.6	6.8	0.24	39.2	0
BfS 157	70797932	275.014	7.750	14/10/2018	0.5	1.9	79.4	0.26	34.2	0
BfS 158	70798003	275.227	7.283	14/10/2018	0.4	12	853.4	0.26	10	0
BfS 159	70798081	275.063	7.475	14/10/2018	0.4	0.8	13.2	0.24	17.5	0
BfS 160	70798095	274.520	6.958	14/10/2018	0.6	0.9	25.4	0.24	36.7	0
BfS 161	70798104	274.469	6.777	10/04/2020	0.7	31.5	1686.3	0.3	44.4	1
BfS 162	70798126	274.372	7.327	14/10/2018	0.4	0.4	13	0.3	43.6	0
BfS 163	70798130	274.748	7.214	14/10/2018	0.6	1.3	81.3	0.27	20.5	0
BfS 164	70798239	275.506	7.480	14/10/2018	0.5	0.3	4.4	0.31	30.4	0
BfS 165	70798392	274.511	7.213	14/10/2018	0.5	1.1	96.5	0.3	34.6	0
BfS 166	70798400	275.445	7.451	14/10/2018	0.5	0.1	1.3	0.38	26.4	0
BfS 167	70798559	275.166	7.145	14/10/2018	0.6	8.9	409.4	0.24	5	0

Bibliography

- Abbott, B. et al. (2009). *Phys. Rev. D* 79 (2), p. 022001.
- Abbott, B. P. et al. (2016). *Physical Review Letters* 116.6, 061102, p. 061102.
- Agazie, G. et al. (2023). *The Astrophysical Journal Letters* 951.1, L8, p. L8.
- Airapetian, V. S. et al. (2017). *The Astrophysical Journal* 836.1, p. L3.
- Aleo, P. D. et al. (2021). *New Astronomy* 96.
- Alexander, K. D. et al. (2020). *Space Science Reviews* 216.5, 81, p. 81.
- Allen, D. A. (1984). *Astrophysics and Space Science* 99.80, p. 101.
- Alves, C. S. et al. (2022). *The Astrophysical Journal, Supplement* 258.2, p. 23.
- Andersson, A. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 513.3, pp. 3482–3492.
- Andersson, A. et al. (2023). *Monthly Notices of the Royal Astronomical Society* 523.2, pp. 2219–2235.
- Anglada-Escudé, G. et al. (2016). *Nature* 536.7617, pp. 437–440.
- Bahramian, A. et al. (2023). *The Astrophysical Journal Letters* 948.1, L7, p. L7.
- Bailer-Jones, C. A. L. et al. (2018). *The Astronomical Journal* 156.2, p. 58.
- Bailey, S. et al. (2007). *The Astrophysical Journal* 665.2, pp. 1246–1253.
- Banerji, M. et al. (2010). *Monthly Notices of the Royal Astronomical Society* 406.1, pp. 342–353.
- Bannister, K. W. et al. (2011). *Monthly Notices of the Royal Astronomical Society* 412.1, pp. 634–664.
- Bastian, T. S. (1990). *Solar Physics* 130.1-2, pp. 265–294.
- Beall, J. H. (2016). In: *Frontier Research in Astrophysics II (FRAPWS2016)*, 53, p. 53.

- Becker, R. H., White, R. L., and Helfand, D. J. (1995). *The Astrophysical Journal* 450, p. 559.
- Bell, A. R. (1978). *Monthly Notices of the Royal Astronomical Society* 182, pp. 147–156.
- Bell, M. E. et al. (2014). *Monthly Notices of the Royal Astronomical Society* 438.1, pp. 352–367.
- Benjamin, R. A. et al. (2003). *Publications of the Astronomical Society of the Pacific* 115.810, pp. 953–964.
- Benz, A., Aref, W., and Guedel, M. (1995). *Astronomy and Astrophysics* 298, pp. 187–192.
- Benz, A. and Guedel, M. (1994). *Astronomy and Astrophysics* 285, pp. 621–630.
- Benz, A. O. (2017). *Living Reviews in Solar Physics* 14.1, 2, p. 2.
- Berger, E. et al. (2001). *Nature* 410.6826, pp. 338–340.
- Berger, E. (2002). *The Astrophysical Journal* 572.1, pp. 503–513.
- Berger, E. et al. (2010). *Astrophysical Journal* 709.1, pp. 332–341.
- Bertin, E. and Arnouts, S. (1996). *Astronomy and Astrophysics, Supplement* 117, pp. 393–404.
- Bilicki, M. et al. (2018). *Astronomy and Astrophysics* 616, A69, A69.
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
- Bloemen, S. et al. (2016). In: *Ground-based and Airborne Telescopes VI*. Vol. 9906. International Society for Optics and Photonics, p. 990664.
- Bloom, J. S. et al. (2012). *Publications of the Astronomical Society of the Pacific* 124.921, p. 1175.
- Blyth, S. et al. (2016). In: *MeerKAT Science: On the Pathway to the SKA*. Ed. by A. R. Taylor et al., 4, p. 4.
- Bochenek, C. D. et al. (2020). *Nature* 587.7832, pp. 59–62.
- Bond, I. A. et al. (2001). *Monthly Notices of the Royal Astronomical Society* 327.3, pp. 868–880.
- Boone, K. (2019). *The Astronomical Journal* 158.6.
- Bopp, T. (1997). *Earth, Moon, and Planets* 79.1, pp. 307–308.

- Bower, G. C. et al. (2007). *The Astrophysical Journal* 666.1, pp. 346–360.
- Boyd, M. R. et al. (2011). *Astronomical Journal* 142.1, p. 10.
- Bramall, D. G. et al. (2012). In: *Ground-based and Airborne Instrumentation for Astronomy IV*. Vol. 8446. International Society for Optics and Photonics, 84460A.
- Braun, R. et al. (2015). In: *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 174, p. 174.
- Breiman, L. (2001). *Machine Learning* 45.1, pp. 5–32.
- Breunig, M. M. et al. (2000). *ACM SIGMOD Record* 29.2, pp. 93–104.
- Bright, J. S. et al. (2018). *Monthly Notices of the Royal Astronomical Society* 475.3, pp. 4011–4019.
- Bright, J. S. et al. (2020). *Nature Astronomy* 4.7, pp. 697–703.
- Buckley, D. A. H., Swart, G. P., and Meiring, J. G. (2006). In: *Ground-based and Airborne Telescopes*. Vol. 6267. 23. International Society for Optics and Photonics, 62670Z.
- Budding, E. et al. (2004). *Astronomy and Astrophysics* 417, pp. 263–268.
- Burgasser, A. J. and Putman, M. E. (2005). *The Astrophysical Journal* 626.1, pp. 486–497.
- Burhanudin, U. F. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 505.3, pp. 4345–4361.
- Burke, B. F. and Franklin, K. L. (1955). *Journal of Geophysics Research* 60.2, pp. 213–217.
- Burke, B. F. and Graham-Smith, F. (2014). *An Introduction to Radio Astronomy*. Cambridge University Press.
- Cabral, J. B. et al. (2018a). *Astronomy and Computing* 25, pp. 213–220.
- Cabral, J. B. et al. (2018b). *Astrophysics Source Code Library*, ascl:1806.001.
- Caleb, M. et al. (2022). *Nature Astronomy* 6.7, pp. 828–836.
- Callingham, J. R. et al. (2021). *Nature Astronomy*.
- Camilo, F. et al. (2018). *The Astrophysical Journal* 856.2, p. 180.

- Cardamone, C. et al. (2009). *Monthly Notices of the Royal Astronomical Society* 399.3, pp. 1191–1205.
- Carignan, C. et al. (2013). *The Astronomical Journal* 146.3, 48, p. 48.
- Carotenuto, F. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 504.1, pp. 444–468.
- Carrington, R. C. (1859). *Monthly Notices of the Royal Astronomical Society* 20, pp. 13–15.
- Castelli, J. P., Aarons, J., and Michael, G. A. (1968). *The Astrophysical Journal* 153, p. 267.
- Charnock, T. and Moss, A. (2017). *The Astrophysical Journal* 837.2, p. L28.
- Chastain, S. I. et al. (2023). *Monthly Notices of the Royal Astronomical Society* 526.2, pp. 1888–1903.
- Chatterjee, S. et al. (2017). *Nature* 541.7635, pp. 58–61.
- Chen, K. and Ruderman, M. (1993). *The Astrophysical Journal* 402, p. 264.
- Chevalier, R. A. (1982). *The Astrophysical Journal* 259, pp. 302–310.
- Chevalier, R. A. (1998). *The Astrophysical Journal* 499.2, pp. 810–819.
- Chevalier, R. A. and Fransson, C. (2006). *The Astrophysical Journal* 651.1, pp. 381–391.
- CHIME/FRB Collaboration et al. (2020). *Nature* 587.7832, pp. 54–58.
- Chiti, A. et al. (2016). *The Astrophysical Journal* 833.1, p. 11.
- Christy, C. T. et al. (2022). *Publications of the Astronomical Society of the Pacific* 134.1032, pp. 1–19.
- Clubbe, C. et al. (2018). *South Georgia & the South Sandwich Islands Marine Protected Area 5-year Review*. Tech. rep.
- Colom i Bernadich, M. et al. (2023). *Astronomy and Astrophysics* 678, A187, A187.
- Condon, J. J. et al. (1979). *The Astronomical Journal* 84, pp. 1–11.
- Condon, J. J. et al. (1998). *The Astronomical Journal* 115.5, pp. 1693–1716.
- Coppejans, D. L. and Knigge, C. (2020). *New Astronomy Review* 89, 101540, p. 101540.

- Cordes, J. M. and Lazio, T. J. W. (2002). *arXiv e-prints*, astro-ph/0207156, astro-ph/0207156.
- Coriat, M. et al. (2011). *Monthly Notices of the Royal Astronomical Society* 414.1, pp. 677–690.
- Cram, L. E. and Giampapa, M. S. (1987). *The Astrophysical Journal* 323, p. 316.
- Cram, L. E. and Mullan, D. J. (1979). *The Astrophysical Journal* 234, p. 579.
- Crause, L. A. et al. (2014). In: *Ground-based and Airborne Instrumentation for Astronomy V*. Vol. 9147. International Society for Optics and Photonics, 91476T.
- Crawford, S. M. et al. (2016). In: *Ground-based and Airborne Instrumentation for Astronomy VI*. Vol. 9908. International Society for Optics and Photonics, p. 99082L.
- Croft, S. et al. (2010). *The Astrophysical Journal* 719.1, pp. 45–58.
- Croft, S. et al. (2011). *The Astrophysical Journal* 731.1, 34, p. 34.
- Cutri, R. M. et al. (2003). *VizieR On-line Data Catalog: II/246*. Originally published in: *2003yCat.2246....0C*, pp. II/246.
- Cutri, R. M. et al. (2013). *VizieR On-line Data Catalog: II/328*. Originally published in: *IPAC/Caltech (2013)*, pp. II/328.
- Czech, D., Mishra, A., and Inggis, M. (2018). *Astronomy and Computing* 25, pp. 52–57.
- Czech, D. et al. (2021). *Publications of the Astronomical Society of the Pacific* 133.1024, 064502, p. 064502.
- D’Elia, V. et al. (2013). *Astronomy and Astrophysics* 551, p. 142.
- Damasso, M. et al. (2020). *Science Advances* 6.3, eaax7467.
- Davies, R. D. et al. (1976). *Nature* 261.5560, pp. 476–478.
- Dawson, K. S. et al. (2013). *The Astronomical Journal* 145.1, p. 10.
- de Blok, W. J. G. et al. (2016). In: *MeerKAT Science: On the Pathway to the SKA*. Ed. by A. R. Taylor et al., 7, p. 7.
- de Bruyn, A. G. et al. (2009). In: *The Low-Frequency Radio Universe*. Ed. by D. J. Saikia et al. Vol. 407. Astronomical Society of the Pacific Conference Series, p. 3.
- de Ruiter, I. et al. (2023). *Monthly Notices of the Royal Astronomical Society* 523.1, pp. 132–148.

- de Villiers, M. S. and Cotton, W. D. (2022). *The Astronomical Journal* 163.3, 135, p. 135.
- Debosscher, J. et al. (2007). *Astronomy and Astrophysics* 475.3, pp. 1159–1183.
- Decin, L. et al. (2019). *Nature Astronomy* 3.5, pp. 408–415.
- Dékány, I. and Grebel, E. K. (2020). *The Astrophysical Journal* 898.1, p. 46.
- Dennison, B., Simonetti, J. H., and Topasna, G. A. (1998). *Publications of the Astronomical Society of Australia* 15.1, pp. 147–48.
- Dewdney, P. E. et al. (2009). *IEEE Proceedings* 97.8, pp. 1482–1496.
- Dicke, R. H. et al. (1965). *The Astrophysical Journal* 142, pp. 414–419.
- Dong, D. Z. et al. (2021). *Science* 373.6559, pp. 1125–1129.
- Douglas, S. T. et al. (2014). *Astrophysical Journal* 795.2, p. 161.
- Driessen, L. N. et al. (2020). *Monthly Notices of the Royal Astronomical Society* 491.1, pp. 560–575.
- Driessen, L. N. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 000, pp. 1–10.
- Driessen, L. N. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 512.4, pp. 5037–5066.
- Driessen, L. N. et al. (2024). *Monthly Notices of the Royal Astronomical Society* 527.2, pp. 3659–3673.
- Driessen, L. (2021). *Zenodo*, Record 4460968.
- Du Buisson, L. et al. (2015). *Monthly Notices of the Royal Astronomical Society* 454.2, pp. 2026–2038.
- Duev, D. A. et al. (2019). *Monthly Notices of the Royal Astronomical Society* 489.3, pp. 3582–3590.
- Dulk, G. A. and Marsh, K. A. (1982). *The Astrophysical Journal* 259, pp. 350–358.
- Dulk, G. A. (1985). *Annual Review of Astronomy and Astrophysics* 23.1, pp. 169–224.
- Einstein, A. (1916). *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp. 688–696.

- Einstein, A. (1918). *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp. 154–167.
- EPTA Collaboration et al. (2023). *Astronomy and Astrophysics* 678, A50, A50.
- Esipov, V. F. et al. (1999). *Astronomy Letters* 25.10, pp. 672–677.
- Espinasse, M. et al. (2021). *The Astronomer’s Telegram* 14607, p. 1.
- Etsebeth, V. et al. (2024). *Monthly Notices of the Royal Astronomical Society* 529.1, pp. 732–747.
- Evans, I. N. et al. (2010). *Astrophysical Journal, Supplement Series* 189.1, pp. 37–82.
- Evans N. J., I. et al. (1976). *The Astrophysical Journal* 206, pp. 440–442.
- Evans, P. A. et al. (2014). *Astrophysical Journal, Supplement Series* 210.1, p. 8.
- Evans, P. A. et al. (2020). *The Astrophysical Journal Supplement Series* 247.2, p. 54.
- Fender, R. (2006). “Jets from X-ray binaries”. In: *Compact stellar X-ray sources*. Vol. 39, pp. 381–419.
- Fender, R., Belloni, T. M., and Gallo, E. (2004). *Monthly Notices of the Royal Astronomical Society* 355.4, pp. 1105–1118.
- Fender, R. et al. (2015). In: *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 51, p. 51.
- Fender, R. et al. (2016). In: *MeerKAT Science: On the Pathway to the SKA*. Ed. by A. R. Taylor et al.
- Fermi, E. (1949). *Physical Review* 75.8, pp. 1169–1174.
- Fijma, S. et al. (2024). *Monthly Notices of the Royal Astronomical Society*.
- Filippazzo, J. C. et al. (2015). *Astrophysical Journal* 810.2, p. 158.
- Finkbeiner, D. P. (2003). *The Astrophysical Journal Supplement Series, Volume 146, Issue 2, pp. 407-415*. 146.2, p. 407.
- Florios, K. et al. (2018). *Solar Physics* 293.2, 28, p. 28.
- Fluke, C. J. and Jacobs, C. (2020). *WIREs Data Mining and Knowledge Discovery* 10.2, e1349, e1349.
- Francis, P. J. et al. (1992). *The Astrophysical Journal* 398, p. 476.

- Gaia Collaboration et al. (2016). *Astronomy and Astrophysics* 595, A1.
- Gaia Collaboration et al. (2018). *Astronomy and Astrophysics* 616, A1.
- Gal, Y., Islam, R., and Ghahramani, Z. (2017). *arXiv e-prints*, arXiv:1703.02910, arXiv:1703.02910.
- Gal-Yam, A. et al. (2006). *The Astrophysical Journal* 639.1, pp. 331–339.
- García-Sánchez, J., Paredes, J. M., and Ribó, M. (2003). *Astronomy and Astrophysics* 403, pp. 613–624.
- Gasealahwe, K. et al. (2022). *The Astronomer’s Telegram* 15584, p. 1.
- Gaustad, J. E. et al. (2001). *Publications of the Astronomical Society of the Pacific* 113.789, pp. 1326–1348.
- Giacconi, R. et al. (2002). *The Astrophysical Journal Supplement* 139.2, pp. 369–410.
- Giles, D. and Walkowicz, L. (2019). *Monthly Notices of the Royal Astronomical Society* 484.1, pp. 834–849.
- Gillon, M. et al. (2017). *Nature* 542.7642, pp. 456–460.
- Ginzburg, V. L. and Zhelezniakov, V. V. (1958). *Soviet Astronomy* 2, p. 653.
- Gizis, J. E. et al. (2017). *The Astrophysical Journal* 845.1, p. 33.
- Goldreich, P. and Julian, W. H. (1969). *The Astrophysical Journal* 157, p. 869.
- Gonidakis, I. et al. (2014). *Monthly Notices of the Royal Astronomical Society* 443, pp. 3819–3838.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- Gordon, Y. A. et al. (2021). *Astrophysical Journal, Supplement* 255.2, 30, p. 30.
- Granot, J. and van der Horst, A. J. (2014). *Publications of the Astronomical Society of Australia* 31, e008, e008.
- Groenewegen, M. A. (2022). *Astronomy and Astrophysics* 659.
- Grogin, N. A. et al. (2011). *The Astrophysical Journal Supplement* 197.2, 35, p. 35.
- Güdel, M. (2002). *Annual Review of Astronomy and Astrophysics* 40, pp. 217–261.
- Guedel, M. and Benz, A. O. (1993). *The Astrophysical Journal* 405, p. L63.

- Gulati, A. et al. (2023). *Publications of the Astronomical Society of Australia* 40, e025, e025.
- Günther, M. N. et al. (2020). *The Astronomical Journal* 159.2, p. 60.
- Hall, D. S. (1976). In: *IAU Colloq. 29: Multiple Periodic Variable Stars*. Ed. by W. S. Fitch. Vol. 60. Astrophysics and Space Science Library, p. 287.
- Halley, E. (1716). *Philosophical Transactions (1683-1775)* 29, pp. 245–262.
- Halley, E. and Senex, J. (1715). *Institute of Astronomy Library*.
- Hallinan, G. et al. (2019). In: *Astro2020: Decadal Survey on Astronomy and Astrophysics*. Vol. 51. 7, p. 255.
- Hancock, P. J. et al. (2019). *arXiv*, arXiv:1907.08395.
- Hankins, T. H. et al. (2003). *Nature* 422.6928, pp. 141–143.
- Hardegree-Ullman, K. K. et al. (2019). *The Astronomical Journal* 158.2, p. 75.
- Harvey, P. M. et al. (1974). *The Astrophysical Journal Supplement* 27, pp. 331–357.
- Harwit, M. (2003). *Physics Today* 56.11, pp. 38–43.
- Haslam, C. G. T. et al. (1982). *Astronomy and Astrophysics Supplement* 47, p. 1.
- Henry, T. J. (2007). In: *Binary Stars as Critical Tools & Tests in Contemporary Astrophysics*. Ed. by W. I. Hartkopf, P. Harmanec, and E. F. Guinan. Vol. 240, pp. 299–299.
- Herman, J. and Habing, H. J. (1985). *Physics Reports* 124.4, pp. 255–314.
- Hewish, A., Scott, P. F., and Wills, D. (1964). *Nature* 203.4951, pp. 1214–1217.
- Hewish, A. et al. (1968). *Nature* 217.5130, pp. 709–713.
- Heywood, I. et al. (2022a). *Monthly Notices of the Royal Astronomical Society* 509.2, pp. 2150–2168.
- Heywood, I. et al. (2022b). *The Astrophysical Journal* 925.2, 165, p. 165.
- Heywood, I. (2020). *Astrophysics Source Code Library*, ascl:2009.003.
- Heywood, I. (2023). *Monthly Notices of the Royal Astronomical Society* 525.1, pp. L76–L81.

- Hjellming, R. M. et al. (1979). *The Astronomical Journal* 84, pp. 1619–1631.
- Hložek, R. et al. (2023). *The Astrophysical Journal Supplement* 267.2, 25, p. 25.
- Hobbs, G. et al. (2004). *Monthly Notices of the Royal Astronomical Society* 352.4, pp. 1439–1472.
- Höfner, S. and Olofsson, H. (2018). *The Astronomy and Astrophysics Review 2017 26:1* 26.1, pp. 1–92.
- Högbom, J. A. (1974). *Astronomy and Astrophysics, Supplement* 15, p. 417.
- Hosenie, Z. et al. (2021). *Experimental Astronomy*.
- Hotan, A. W. et al. (2021). *Publications of the Astronomy Society of Australia* 38, e009, e009.
- Hovatta, T., Lehto, H. J., and Tornikoski, M. (2008). *Astronomy and Astrophysics* 488.3, pp. 897–903.
- Hovatta, T. et al. (2008). *Astronomy and Astrophysics* 485.1, pp. 51–61.
- Hubble, E. P. (1936). *Realm of the Nebulae*. Yale University Press.
- Hughes, P. A., Aller, H. D., and Aller, M. F. (1989). *The Astrophysical Journal* 341, p. 54.
- Hughes, P. A. et al. (1992). *The Astrophysical Journal* 396, p. 469.
- Hugo, B. V. et al. (2022). In: *Astronomical Society of the Pacific Conference Series*. Ed. by J. E. Ruiz, F. Pierfederici, and P. Teuben. Vol. 532. Astronomical Society of the Pacific Conference Series, p. 541.
- Hurley-Walker, N. et al. (2022). *Nature* 601.7894, pp. 526–530.
- Hurley-Walker, N. et al. (2023). *Nature* 619.7970, pp. 487–490.
- Hyman, S. D. et al. (2005). *Nature* 434.7029, pp. 50–52.
- Hyman, S. D. et al. (2009). *The Astrophysical Journal* 696.1, pp. 280–286.
- Inceoglu, F. et al. (2018). *The Astrophysical Journal* 861.2, 128, p. 128.
- Ishida, E. E. et al. (2021). *Astronomy & Astrophysics* 650, A195.
- Itagaki, K. (2023). *Transient Name Server Discovery Report* 2023-1158, p. 1.

- Ivezic, Ž. et al. (2014). *Statistics, Data Mining, and Machine Learning in Astronomy: A Practical Python Guide for the Analysis of Survey Data*. Princeton University Press.
- Ivezić, Ž. et al. (2019). *The Astrophysical Journal* 873.2, p. 111.
- Jacobson-Galán, W. V. et al. (2023). *The Astrophysical Journal Letters* 954.2, L42, p. L42.
- Jaeger, T. R. et al. (2012). *Astronomical Journal* 143.4, p. 96.
- Jansky, K. G. (1933). *Nature* 132.3323, p. 66.
- Jarvis, M. et al. (2016). In: *MeerKAT Science: On the Pathway to the SKA*. Ed. by A. R. Taylor et al., 6, p. 6.
- Jenkins, J. M. et al. (2016). In: *Software and Cyberinfrastructure for Astronomy IV*. Vol. 9913. Proceedings of the SPIE, 99133E, 99133E.
- Johnston, S. et al. (2007). *Publications of the Astronomical Society of Australia* 24.4, pp. 174–188.
- Johnston, S. and Kerr, M. (2018). *Monthly Notices of the Royal Astronomical Society* 474.4, pp. 4629–4636.
- Jonas, J. and MeerKAT Team (2016). In: *MeerKAT Science: On the Pathway to the SKA*. Ed. by A. R. Taylor et al., 1, p. 1.
- Keane, E. F. (2018). *Nature Astronomy* 2, pp. 865–872.
- Keane, E. F. and McLaughlin, M. A. (2011). *Bulletin of the Astronomical Society of India, Vol. 39, No. 3, p. 333-352* 39.3, p. 333.
- Keel, W. C. et al. (2012). *Monthly Notices of the Royal Astronomical Society* 420.1, pp. 878–900.
- Kellermann, K. I. and Pauliny-Toth, I. I. K. (1969). *The Astrophysical Journal Letters* 155, p. L71.
- Kenyon, J. S. et al. (2018). *Monthly Notices of the Royal Astronomical Society* 478.2, pp. 2399–2415.
- Kesseli, A. Y. et al. (2017). *The Astrophysical Journal Supplement Series* 230.2, p. 16.
- Khodachenko, M. L. et al. (2007). *Astrobiology* 7.1, pp. 167–184.
- Killestein, T. L. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 503.4, pp. 4838–4854.

- Kilpatrick, C. D. et al. (2023). *The Astrophysical Journal Letters* 952.1, L23, p. L23.
- Kim, D.-W. et al. (2009). *Monthly Notices of the Royal Astronomical Society* 397.2, pp. 558–568.
- Kim, D.-W. et al. (2011). *The Astrophysical Journal* 735.2, p. 68.
- Kim, D.-W. et al. (2014). *Astronomy and Astrophysics* 566, A43, A43.
- Kirkpatrick, J. D., Henry, T. J., and McCarthy, D. W. (1991). *The Astrophysical Journal Supplement Series* 77, p. 417.
- Kochanek, C. S. et al. (2017). *Publications of the Astronomical Society of the Pacific* 129.980, p. 104502.
- Kramer, M. (1994). *Astronomy and Astrophysics, Supplement* 107, pp. 527–539.
- Kramer, M. et al. (2006). *Science* 312.5773, pp. 549–551.
- Kuzmin, A. D. (2007). *Astrophysics and Space Science* 308.1-4, pp. 563–567.
- Lacy, C. H., Moffett, T. J., and Evans, D. S. (1976). *The Astrophysical Journal Supplement Series* 30, p. 85.
- Lacy, M. et al. (2020). *Publications of the ASP* 132.1009, 035001, p. 035001.
- Lacy, M. et al. (2019). In: *Astronomical Data Analysis Software and Systems XXVII*. Ed. by P. J. Teuben et al. Vol. 523. Astronomical Society of the Pacific Conference Series, p. 217.
- Lahav, O. et al. (1995). *Science* 267.5199, pp. 859–862.
- Lainela, M. (1994). *Astronomy and Astrophysics* 286, pp. 408–414.
- Lammer, H. et al. (2007). *Astrobiology* 7.1, pp. 185–207.
- Lane, C. et al. (2007). *The Astrophysical Journal* 668.2, pp. L163–L166.
- Lang, K. R. (2009). *The Sun From Space*. Springer Berlin, Heidelberg.
- Lanza, A. F. (2005). *Monthly Notices of the Royal Astronomical Society* 364.1, pp. 238–246.
- Law, C. J. et al. (2018). *The Astrophysical Journal, Supplement* 236.1, p. 8.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). *Nature* 521.7553, pp. 436–444.

- Levinson, A. et al. (2002). *The Astrophysical Journal* 576.2, pp. 923–931.
- Lin, C.-L. et al. (2021). *The Astronomical Journal* 162.1, p. 11.
- Link, B., Epstein, R. I., and Lattimer, J. M. (1999). *Physical Review Letters* 83.17, pp. 3362–3365.
- Lintott, C. J. et al. (2008). *Monthly Notices of the Royal Astronomical Society* 389.3, pp. 1179–1189.
- Lintott, C. J. et al. (2009). *Monthly Notices of the Royal Astronomical Society* 399.1, pp. 129–140.
- Liu, F. T., Ting, K. M., and Zhou, Z. H. (2008). *Proceedings - IEEE International Conference on Data Mining, ICDM*, pp. 413–422.
- Lochner, M. et al. (2023). *Monthly Notices of the Royal Astronomical Society, Advance Access* 000, pp. 1–7.
- Lochner, M. and Bassett, B. A. (2021). *Astronomy and Computing* 36.
- Lochner, M. et al. (2016). *The Astrophysical Journal, Supplement* 225.2, p. 31.
- Lomb, N. R. (1976). *Astrophysics and Space Science* 39.2, pp. 447–462.
- Longair, M. S. (2011). *High Energy Astrophysics*. Cambridge University Press.
- Lonsdale, C. J. et al. (2009). *IEEE Proceedings* 97.8, pp. 1497–1506.
- Lorimer, D. R. and Kramer, M. (2004). *Handbook of Pulsar Astronomy*. Vol. 4. Cambridge University Press.
- Lorimer, D. R. et al. (1995). *Monthly Notices of the Royal Astronomical Society* 273.2, pp. 411–421.
- Lorimer, D. R. et al. (2007). *Science* 318.5851, p. 777.
- Lovell, B. (1969). *Nature* 222.5199, pp. 1126–1129.
- Lovell, B., Whipple, F. L., and Solomon, L. H. (1963). *Nature* 198.4877, pp. 228–230.
- Lower, M. E. et al. (2024). *Nature Astronomy*.
- Luger, R. et al. (2017). *Nature Astronomy* 1, p. 0129.
- Lynch, C. et al. (2016). *Monthly Notices of the Royal Astronomical Society* 457.2, pp. 1224–1232.

- Lynch, C. R. et al. (2017). *The Astrophysical Journal Letters* 836.2, L30, p. L30.
- Lyne, A. and Graham-Smith, F. (2012). *Pulsar Astronomy*. 4th ed. Cambridge Astrophysics. Cambridge University Press.
- Lyne, A. et al. (2010). *Science* 329.5990, pp. 408–412.
- Macquart, J. P. et al. (2020). *Nature* 581.7809, pp. 391–395.
- MAGIC Collaboration et al. (2019). *Nature* 575.7783, pp. 459–463.
- Mahabal, A. et al. (2019). *Publications of the Astronomical Society of the Pacific* 131.997, p. 27.
- Malanchev, K. L. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 502.4, pp. 5147–5175.
- Malenta, M. et al. (2020). In: *Astronomical Data Analysis Software and Systems XXIX*. Ed. by R. Pizzo et al. Vol. 527. Astronomical Society of the Pacific Conference Series, p. 457.
- Manchester, R. N. et al. (2005). *The Astronomical Journal* 129.4, pp. 1993–200.
- Marcote, B. et al. (2017). *The Astrophysical Journal Letters* 834.2, L8, p. L8.
- Markwardt, C. B. and Swank, J. H. (2003). *The Astronomer's Telegram* 133, p. 1.
- Marshall, P. J., Lintott, C. J., and Fletcher, L. N. (2015). *Annual Review of Astronomy and Astrophysics* 53, pp. 247–278.
- Masters, K. et al. (2015). In: *IAU General Assembly*. Vol. 29, 2225804, p. 2225804.
- McConnell, D. et al. (2020). *Publications of the Astronomical Society of Australia* 37, e048, e048.
- McEwen, A. E. et al. (2020). *The Astrophysical Journal* 892.2, 76, p. 76.
- McInnes, L. et al. (2018). *Journal of Open Source Software* 3.29, p. 861.
- McMullin, J. P. et al. (2007). *ASPC* 376, p. 127.
- Meech, K. J. et al. (2017). *Nature* 552.7685, pp. 378–381.
- Melrose, D. B. (2009). In: *Universal Heliophysical Processes*. Ed. by N. Gopalswamy and D. F. Webb. Vol. 257, pp. 305–315.
- Melrose, D. B. (2017). *Reviews of Modern Plasma Physics* 1.1, 5, p. 5.

- Méndez, B. J. H. (2008). In: *EPO and a Changing World: Creating Linkages and Expanding Partnerships*. Ed. by C. Garmany, M. G. Gibbs, and J. W. Moody. Vol. 389. Astronomical Society of the Pacific Conference Series, p. 219.
- Mirabel, I. F. and Rodríguez, L. F. (1994). *Nature* 371.6492, pp. 46–48.
- Mohale, K. and Lochner, M. (2023). *arXiv e-prints*, arXiv:2311.14157.
- Mohan, N. and Rafferty, D. (2015). *Astrophysics Source Code Library*, ascl:1502.007.
- Moller, A. and de Boissiere, T. (2020). *Monthly Notices of the Royal Astronomical Society* 491.3, pp. 4277–4293.
- Monet, D. G. et al. (2003). *The Astronomical Journal* 125, pp. 984–993.
- Mooley, K. P. et al. (2013). *The Astrophysical Journal* 768.2, 165, p. 165.
- Mooley, K. P. et al. (2016). *The Astrophysical Journal* 818.2, p. 105.
- Moss, A. (2018). *arXiv*.
- Motta, S. E. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 503.1, pp. 152–161.
- Murphy, K. P. (2022). *Probabilistic Machine Learning: An introduction*. MIT Press.
- Murphy, T. et al. (2013). *Publications of the Astronomical Society of Australia* 30.e006.
- Murphy, T. et al. (2017). *Monthly Notices of the Royal Astronomical Society* 466.2, pp. 1944–1953.
- Murphy, T. et al. (2021). *Publications of the Astronomical Society of Australia* 38, e054, e054.
- Mutel, R. L. et al. (1987). *The Astronomical Journal* 93, p. 1220.
- Muthukrishna, D. et al. (2019). *Publications of the Astronomical Society of the Pacific* 131.
- Muthukrishna, D. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 517.1, pp. 393–419.
- Naul, B. et al. (2018). *Nature Astronomy* 2, pp. 151–155.
- Newton, E. R. et al. (2017). *The Astrophysical Journal* 834.1, p. 85.
- Nielsen, M. (2015). *Neural Networks and Deep Learning*. Determination Press.

- Norris, R. P. (2017). *Publications of the Astronomical Society of Australia* 34, e007, e007.
- Nun, I. et al. (2015). *arXiv*, arXiv:1506.00010.
- Nyamai, M. M. et al. (2023). *Monthly Notices of the Royal Astronomical Society* 523.2, pp. 1661–1675.
- Ochsenbein, F., Bauer, P., and Marcout, J. (2000). *Astronomy and Astrophysics, Supplement* 143, pp. 23–32.
- Ofek, E. O. et al. (2010). *The Astrophysical Journal* 711.1, pp. 517–531.
- Ofek, E. O. et al. (2011). *Astrophysical Journal* 740.2, p. 65.
- Offringa, A. R. and Smirnov, O. (2017). *Monthly Notices of the Royal Astronomical Society* 471.1, pp. 301–316.
- Offringa, A. R. (2010). *Astrophysics Source Code Library*, ascl–1010.
- Offringa, A. R., Van De Gronde, J. J., and Roerdink, J. B. (2012). *Astronomy and Astrophysics* 539, A95.
- Offringa, A. R. et al. (2014). *Monthly Notices of the Royal Astronomical Society* 444.1, pp. 606–619.
- Olausen, S. A. and Kaspi, V. M. (2014). *The Astrophysical Journal Supplement* 212.1, 6, p. 6.
- Osten, R. A. (2007). In: *Proceedings of Science*. Vol. 56. Sissa Medialab Srl.
- Osten, R. A. and Bastian, T. S. (2006). *The Astrophysical Journal* 637.2, pp. 1016–1024.
- Osten, R. A. et al. (2006). *The Astrophysical Journal* 637.1, p. 518.
- Oswald, L. S. et al. (2023). *Monthly Notices of the Royal Astronomical Society* 520.4, pp. 4961–4980.
- Pacholczyk, A. G. (1970). *Radio astrophysics. Nonthermal processes in galactic and extragalactic sources*. WH Freeman & Company.
- Padovani, P. et al. (2017). *Astronomy and Astrophysics Reviews* 25.1, 2, p. 2.
- Page, M. J. et al. (2012). *Monthly Notices of the Royal Astronomical Society* 426.2, pp. 903–926.
- Pass, E. K. et al. (2023). *Astronomical Journal* 166.1, 16, p. 16.

- Paudel, R. R. et al. (2018). *The Astrophysical Journal* 858.1, p. 55.
- Pearson, K. (1901). *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11, pp. 559–572.
- Pearson, K. A., Palafox, L., and Griffith, C. A. (2018). *Monthly Notices of the Royal Astronomical Society* 474.1, pp. 478–491.
- Pedregosa, F. et al. (2011). *Journal of Machine Learning Research* 12, pp. 2825–2830.
- Penzias, A. A. and Wilson, R. W. (1965). *The Astrophysical Journal* 142, pp. 419–421.
- Perlmutter, S. et al. (1999). *The Astrophysical Journal* 517.2, pp. 565–586.
- Peters, N. (2023). *The first transient search with automatically created MeerKAT images*. MSc Thesis, University of Amsterdam.
- Philippov, A. and Kramer, M. (2022). *Annual Review of Astronomy and Astrophysics* 60, pp. 495–558.
- Pickering, E. C. and Fleming, W. P. (1896). *The Astrophysical Journal* 4, pp. 369–370.
- Pietka, M., Fender, R. P., and Keane, E. F. (2015). *Monthly Notices of the Royal Astronomical Society* 446.4, pp. 3687–3696.
- Pineda, J. S. and Villadsen, J. (2023). *Nature Astronomy* 7, pp. 569–578.
- Pintaldi, S. et al. (2022). In: *Astronomical Society of the Pacific Conference Series*. Ed. by J. E. Ruiz, F. Pierfederici, and P. Teuben. Vol. 532. Astronomical Society of the Pacific Conference Series, p. 333.
- Planck Collaboration et al. (2020). *Astronomy and Astrophysics* 641, A6, A6.
- Platts, E. et al. (2019). *Physics Reports* 821, pp. 1–27.
- Pleunis, Z. et al. (2021). *The Astrophysical Journal* 923.1, 1, p. 1.
- Polisensky, E. et al. (2016). *The Astrophysical Journal* 832.1, 60, p. 60.
- Pritchard, J. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 502.4, pp. 5438–5454.
- Pritchard, J. et al. (2024). *Monthly Notices of the Royal Astronomical Society* 529.2, pp. 1258–1270.
- Protopapas, P. et al. (2015). *The Astrophysical Journal Supplement* 216.2, 25, p. 25.

- Pruzhinskaya, M. V. et al. (2019). *Monthly Notices of the Royal Astronomical Society* 489.3, pp. 3591–3608.
- Pruzhinskaya, M. V. et al. (2023). *Astronomy and Astrophysics* 672, A111.
- Rajpurohit, A. S. et al. (2013). *Astronomy and Astrophysics* 556.
- Rajwade, K. et al. (2020). In: *Ground-based and Airborne Instrumentation for Astronomy VIII*. Ed. by C. J. Evans, J. J. Bryant, and K. Motohara. Vol. 11447. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 114470J, 114470J.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge.
- Ravi, V. et al. (2022). *The Astrophysical Journal* 925.2, 220, p. 220.
- Rea, N. et al. (2024). *The Astrophysical Journal* 961.2, 214, p. 214.
- Readhead, A. C. S. (1994). *The Astrophysical Journal* 426, p. 51.
- Reardon, D. J. et al. (2023). *The Astrophysical Journal Letters* 951.1, L6, p. L6.
- Rees, M. J. (1988). *Nature* 333.6173, pp. 523–528.
- Reid, M. J. (2007). In: *Astrophysical Masers and their Environments*. Ed. by J. M. Chapman and W. A. Baan. Vol. 242, pp. 522–529.
- Remazeilles, M. et al. (2015). *Monthly Notices of the Royal Astronomical Society* 451.4, pp. 4311–4327.
- Reynolds, R. J., Haffner, L. M., and Madsen, G. J. (2002). In: *Galaxies: the Third Dimension*. Ed. by M. Rosada, L. Binette, and L. Arias. Vol. 282. Astronomical Society of the Pacific Conference Series, p. 31.
- Rhodes, L. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 513.2, pp. 2708–2718.
- Rhodes, L. et al. (2023a). *Monthly Notices of the Royal Astronomical Society* 525.3, pp. 3626–3632.
- Rhodes, L. et al. (2023b). *Monthly Notices of the Royal Astronomical Society* 521.1, pp. 389–395.
- Richards, J. W. et al. (2011). *The Astrophysical Journal* 733.1, p. 10.
- Richards, M. T. et al. (2003). *The Astrophysical Journal Supplement* 147.2, pp. 337–361.

- Ricker, G. R. et al. (2015). *Journal of Astronomical Telescopes, Instruments, and Systems* 1, 014003, p. 014003.
- Rickett, B. J. (1986). *The Astrophysical Journal* 307, p. 564.
- Rickett, B. J. (1990). *Annual Review of Astronomy and Astrophysics* 28.1, pp. 561–605.
- Rickett, B. J. (2007). *Astronomical and Astrophysical Transactions* 26.6, pp. 429–439.
- Rickett, B. (2001). *Astrophysics and Space Science 2001 278:1* 278.1, pp. 5–10.
- Ridder, M. E. et al. (2023). *Monthly Notices of the Royal Astronomical Society* 519.4, pp. 5922–5930.
- Riess, A. G. et al. (1998). *The Astronomical Journal* 116.3, pp. 1009–1038.
- Robinson, P. A. and Melrose, D. B. (1984). *Australian Journal of Physics* 37.6, pp. 675–704.
- Robinson, R. D. (1986). *Solar Physics* 104.1, pp. 33–39.
- Rogers, B. et al. (2024). *arXiv e-prints*, arXiv:2401.08763, arXiv:2401.08763.
- Roulston, B. R., Green, P. J., and Kesseli, A. Y. (2020). *The Astrophysical Journal Supplement Series* 249.2, p. 34.
- Rowlinson, A. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 517.2, pp. 2894–2911.
- Ruderman, M. A. and Sutherland, P. G. (1975). *The Astrophysical Journal* 196, pp. 51–72.
- Rybicki, G. B. and Lightman, A. P. (1986). *Radiative Processes in Astrophysics*. Wiley Publishing.
- Salas-Matamoros, C. and Klein, K.-L. (2020). *Astronomy and Astrophysics* 639, A102, A102.
- Samus', N. N. et al. (2017). *Astronomy Reports* 61.1, pp. 80–88.
- Sarbadhicary, S. K. et al. (2021). *The Astrophysical Journal* 923.1, 31, p. 31.
- Scaife, A. M. M. (2020). *Philosophical Transactions of the Royal Society of London Series A* 378.2166, 20190060, p. 20190060.
- Scaife, A. M. M. and Porter, F. (2021). *Monthly Notices of the Royal Astronomical Society* 503.2, pp. 2369–2379.

- Scargle, J. D. (1982). *The Astrophysical Journal* 263, p. 835.
- Schinzell, F. K. et al. (2014). *The Astrophysical Journal, Supplement* 217.1, p. 4.
- Schröder, A. et al. (2024). *SARAO 20th Anniversary Report*. Tech. rep. SARAO.
- Schroeder, G. et al. (2023). *arXiv e-prints*, arXiv:2308.10936, arXiv:2308.10936.
- Schwamb, M. E. et al. (2012). *The Astrophysical Journal Letters* 754.2, 129, p. 129.
- Scoville, N. et al. (2007). *The Astrophysical Journals* 172.1, pp. 1–8.
- Seaquist, E. R., Ivison, R. J., and Hall, P. J. (1995). *Monthly Notices of the Royal Astronomical Society* 276.3, p. 867.
- Seaquist, E. R., Taylor, A. R., and Button, S. (1984). *The Astrophysical Journal* 284, p. 202.
- Seaquist, E. (1988). *International Astronomical Union Colloquium* 103, pp. 69–75.
- Sebastian, D. et al. (2021). *Astronomy and Astrophysics* 645, p. 100.
- Selina, R. J. et al. (2018). In: *Astronomical Society of the Pacific Conference Series*. Vol. 517, p. 15.
- Serra, I. et al. (2020). In: *EGU General Assembly Conference Abstracts*. EGU General Assembly Conference Abstracts, 8763, p. 8763.
- Settles, B. (2012). *Active Learning*. Synthesis lectures on artificial intelligence and machine learning. Morgan & Claypool.
- Shallue, C. J. and Vanderburg, A. (2018). *The Astronomical Journal* 155.2, 94, p. 94.
- Shang, L. H. et al. (2021). *The Astrophysical Journal* 916.2, 62, p. 62.
- Shappee, B. J. et al. (2014). *Astrophysical Journal* 788.1 (1), p. 48.
- Shaw, B. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 513.4, pp. 5861–5880.
- Shui, Q. C. et al. (2023). *The Astrophysical Journal* 943.2, 165, p. 165.
- Sieber, W. (1973). *Astronomy and Astrophysics* 28, p. 237.
- Simmons, B. D. et al. (2017). *Monthly Notices of the Royal Astronomical Society* 464.4, pp. 4420–4447.

- Skrutskie, M. F. et al. (2006). *The Astronomical Journal* 131.2, pp. 1163–1183.
- Slee, O. B., Willes, A. J., and Robinson, R. D. (2003). *Publications of the Astronomical Society of Australia* 20.3, pp. 257–262.
- Slijepcevic, I. V. et al. (2024). *RAS Techniques and Instruments* 3.1, pp. 19–32.
- Smartt, S. J. et al. (2017). *Nature* 551.7678, pp. 75–79.
- Smirnov, O. M. et al. (2024). *Monthly Notices of the Royal Astronomical Society*.
- Smith, K. W. et al. (2019). *Research Notes of the AAS* 3.1, p. 26.
- Smith, R. K. et al. (2001). *The Astrophysical Journal* 556.2, pp. L91–L95.
- Smolčić, V. et al. (2017). *Astronomy and Astrophysics* 602, A2, A2.
- Sooknunan, K. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 502.1, pp. 206–224.
- Speagle, J. S. and Eisenstein, D. J. (2017). *Monthly Notices of the Royal Astronomical Society* 469.1, pp. 1186–1204.
- Spiers, H. et al. (2019). *Journal of Science Communication* 18.1, A04.
- Spitler, L. G. et al. (2016). *Nature* 531.7593, pp. 202–205.
- Spitzer Science Center (2009). *Vizier*, pp. II/293.
- Stappers, B. (2016). In: *MeerKAT Science: On the Pathway to the SKA*. Ed. by A. R. Taylor et al., 10, p. 10.
- Starrfield, S. (1989). In: *Classical Novae*, pp. 39–60.
- Stassun, K. G. et al. (2019). *The Astronomical Journal* 158.4, p. 138.
- Stelzer, B. et al. (2013). *Monthly Notices of the Royal Astronomical Society* 431.3, pp. 2063–2079.
- Stewart, A. J. et al. (2016). *Monthly Notices of the Royal Astronomical Society* 456.3, pp. 2321–2342.
- Stewart, A. J. et al. (2018). *Monthly Notices of the Royal Astronomical Society* 479.2, pp. 2481–2504.
- Stinebring, D. R. et al. (1990). *The Astrophysical Journal* 352, p. 207.

- Stinebring, D. R. et al. (2000). *The Astrophysical Journal* 539.1, pp. 300–316.
- Stroh, M. C. et al. (2021). *The Astrophysical Journal Letters* 923.2, L24, p. L24.
- Swarup, G. et al. (1991). *Current Science* 60, p. 95.
- Swinbank, J. D. et al. (2015). *Astronomy and Computing* 11.PA, pp. 25–48.
- Tan, C. M. et al. (2018). *The Astrophysical Journal* 866.1, 54, p. 54.
- Tasse, C. et al. (2018). *Astronomy and Astrophysics* 611, A87.
- Taylor, A. R. and Gregory, P. C. (1983). *The Astronomical Journal* 88, pp. 1784–1809.
- Taylor, J. H. and Weisberg, J. M. (1982). *The Astrophysical Journal* 253, pp. 908–920.
- Thompson, A. R., Moran, J. M., and Swenson George W., J. (2017). *Interferometry and Synthesis in Radio Astronomy, 3rd Edition*.
- Thyagarajan, N. et al. (2011). *The Astrophysical Journal* 742.15pp, p. 49.
- Tingay, S. J. et al. (2012). In: *Proceedings of Science*. Vol. 2012-April. Sissa Medialab Srl.
- Traulsen, I. et al. (2020). *Astronomy & Astrophysics* 641, A137.
- Tremou, E. et al. (2020). *Monthly Notices of the Royal Astronomical Society* 493.1, pp. L132–L137.
- Tremou, E. et al. (2021). *The Astronomer’s Telegram* 14432, p. 1.
- Turatto, M. (2003). “Classification of Supernovae”. In: *Supernovae and Gamma-Ray Bursters*. Ed. by K. W. Weiler. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 21–36.
- Van Haarlem, M. P. et al. (2013). *Astronomy and Astrophysics* 556, A2.
- Varghese, S. S. et al. (2019). *The Astrophysical Journal* 874.2, p. 151.
- Vedantham, H. K. et al. (2020). *Nature Astronomy* 4, pp. 577–583.
- Villadsen, J. and Hallinan, G. (2019). *The Astrophysical Journal* 871.2, p. 214.
- Villar, V. A. et al. (2020). *The Astrophysical Journal* 905.2, 94, p. 94.
- Villar, V. A. et al. (2021). *The Astrophysical Journal Supplement Series* 255.2, p. 24.

- Walker, M. A. (1998). *Monthly Notices of the Royal Astronomical Society* 294, pp. 307–311.
- Walkowicz, L. M., Hawley, S. L., and West, A. A. (2004). *Publications of the Astronomical Society of the Pacific* 116.826, pp. 1105–1110.
- Walmsley, M. (2021). *Advancing machine learning in astrophysics*. DPhil Thesis, University of Oxford.
- Walmsley, M. et al. (2020). *Monthly Notices of the Royal Astronomical Society* 491.2, pp. 1554–1574.
- Walmsley, M. et al. (2022a). *Monthly Notices of the Royal Astronomical Society* 509.3, pp. 3966–3988.
- Walmsley, M. et al. (2022b). *Monthly Notices of the Royal Astronomical Society* 513.2, pp. 1581–1599.
- Wang, Y. et al. (2021a). *Monthly Notices of the Royal Astronomical Society* 502.3, pp. 3294–3311.
- Wang, Y. et al. (2023). *Monthly Notices of the Royal Astronomical Society* 523.4, pp. 5661–5680.
- Wang, Z. et al. (2021b). *The Astrophysical Journal* 920.1, p. 45.
- Wang, Z. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 516.4, pp. 5972–5988.
- Webb, D. F. and Howard, T. A. (2012). *Living Reviews in Solar Physics* 9.1, 3, p. 3.
- Webb, S. et al. (2020). *Monthly Notices of the Royal Astronomical Society* 498.3, pp. 3077–3094.
- Webb, S. et al. (2021). *Monthly Notices of the Royal Astronomical Society* 506.2, pp. 2089–2103.
- Weiler, K. W. et al. (2002). *Annual Review of Astronomy and Astrophysics* 40, pp. 387–438.
- Welch, J. et al. (2009). *IEEE Proceedings* 97.8, pp. 1438–1447.
- Wells, D. C., Greisen, E. W., and Harten, R. H. (1981). *Astronomy and Astrophysics, Supplement* 44, p. 363.
- Wenger, M. et al. (2000). *Astronomy and Astrophysics Supplement Series* 143.1, pp. 9–22.

- Werthimer, D. et al. (2001). In: *The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*. Ed. by S. A. Kingsley and R. Bhathal. Vol. 4273. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 104–109.
- West, A. A., Walkowicz, L. M., and Hawley, S. L. (2005). *Publications of the ASP* 117.833, pp. 706–710.
- Westphal, A. J. et al. (2006). In: *AGU Fall Meeting Abstracts*. Vol. 2006, P52B-08, P52B-08.
- Whitelock, P. A. (1987). *Publications of the Astronomical Society of the Pacific* 99, p. 573.
- Whitmore, B. C. (1984). *The Astrophysical Journal* 278, pp. 61–80.
- Willett, K. W. et al. (2013). *Monthly Notices of the Royal Astronomical Society* 435.4, pp. 2835–2860.
- Willett, K. W. et al. (2017). *Monthly Notices of the Royal Astronomical Society* 464.4, pp. 4176–4203.
- Williams, D. R. A. et al. (2022). *Monthly Notices of the Royal Astronomical Society* 517.2, pp. 2801–2817.
- Williams, D. R. A. et al. (2020). *Monthly Notices of the Royal Astronomical Society: Letters* 491.1, pp. L29–L33.
- Winnberg, A., Brand, J., and Engels, D. (2024). *arXiv e-prints*, arXiv:2403.00535, arXiv:2403.00535.
- Winnberg, A. et al. (2008). *Astronomy and Astrophysics* 482.3, pp. 831–848.
- Wright, D. E. et al. (2015). *Monthly Notices of the Royal Astronomical Society* 449.1, pp. 451–466.
- Wright, D. E. et al. (2017). *Monthly Notices of the Royal Astronomical Society* 472.2, pp. 1315–1323.
- Xu, H. et al. (2023). *Research in Astronomy and Astrophysics* 23.7, 075024, p. 075024.
- Yamada, M., Kulsrud, R., and Ji, H. (2010). *Rev. Mod. Phys.* 82 (1), pp. 603–664.
- Yao, J. M., Manchester, R. N., and Wang, N. (2017). *The Astrophysical Journal* 835.1, 29, p. 29.
- Zarka, P. (1998). *Journal of Geophysics Research* 103.E9, pp. 20159–20194.

Zarka, P. (2007). *Planetary Space Science* 55.5, pp. 598–617.

Zhang, X. et al. (2021). *The Astronomer's Telegram* 14878, p. 1.

Zhao, J.-H. et al. (1992). *Science* 255.5051, pp. 1538–1543.

Zic, A. et al. (2019). *Monthly Notices of the Royal Astronomical Society* 488.1, pp. 559–571.

Zic, A. et al. (2020). *The Astrophysical Journal* 905.1, p. 23.

Zic, A. et al. (2024). *Monthly Notices of the Royal Astronomical Society* 528.4, pp. 5730–5741.