

A Machine Learning Approach to Mitigating Radio Interference in Deep Wide-field Observations of the Lockman Hole

Jacob Edward Brooks



A THESIS SUBMITTED IN PARTIAL FULFILMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Jeremiah Horrocks Institute for Mathematics, Physics and Astronomy
University of Central Lancashire

Submitted: 31st March 2024

Declaration

Type of Award: Doctor of Philosophy

School: Engineering and Computing

I declare that while registered as a candidate for the research degree, I have not been a registered candidate or enrolled student for another award of the University or other academic or professional institution.

I declare that no material contained in the thesis has been used in any other submission for an academic award and is solely my own work.

No proof-reading service was used in the compilation of this thesis.

A handwritten signature in black ink, appearing to read 'J. Brooks', with a long horizontal flourish extending to the right.

Jacob Edward Brooks

31st March 2024

Abstract

A study of the Lockman Hole at 1.5GHz with the e-MERLIN radio telescope is performed, along with an experiment in using machine learning techniques to identify and remove interference in astronomical radio observations, and an investigation of a faint sub-kpc scale binary black hole candidate. It is found that manually flagged observations used as the basis for a training set does not produce a trained network that performs better than the current state of the art statistical tool, at least when using a GAN architecture. Though, it remains possible that refinements to the GAN model, the pre-processing method, and the training set could produce better results. Machine learning techniques remain as one of the most promising methods of dealing with interference in an increasingly difficult RFI environment populated by massive satellite networks that will impede the operation of next generation radio observatories, such as the SKA. In the investigation of the sub-kpc scale binary AGN candidate, two radio sources are detected that overlap with two previously identified stellar cores, both with a separation of $\sim 1''$. Specifically, the local epoch separation between the two radio sources is measured as 0.95 ± 0.29 kpc. However, follow up radio observations with greater resolving power did not produce a second detection. The source is very faint, so it is likely these observations were not sensitive enough to provide a second detection; further observation is recommended to reduce the noise level. Finally, a source catalogue of 78 sources in the Lockman Hole is produced from a wide field image with a mean sensitivity of $12 \mu\text{Jy}/\text{beam}$. Radio measurements from other studies at 140MHz and 3GHz are included for $\sim 50\%$ of sources, along with estimates of redshift for 35 sources. Five sources from this new catalogue are studied in detail and the differential source counts are calculated and compared with similar work. The source counts are consistent with an observed faint radio population with a suppressed SFG component at $\lesssim 1\text{mJy}$, which is expected at VLBI scales. The overall trend agrees with source counts calculated using a catalogue of VLA sources. The e-MERLIN image is an important step towards producing a deeper radio image of the Lockman Hole that can be used to classify sources within the field, and investigate the nature of the faint radio population.

Contents

| | |
|------------------------------------------------------|-----------|
| Declaration | 2 |
| Abstract | 3 |
| Acknowledgements | 6 |
| 1 Introduction | 9 |
| 1.1 Radio Interference | 11 |
| 1.2 Faint Radio Sources | 13 |
| 2 Techniques of Radio Astronomy | 17 |
| 2.1 Radio Interferometry | 18 |
| 2.2 The e-MERLIN Telescope | 20 |
| 2.3 The Two-Element Interferometer | 21 |
| 2.4 Aperture Synthesis | 26 |
| 2.5 Calibration | 31 |
| 2.5.1 Flagging | 31 |
| 2.5.2 Phase Calibration | 32 |
| 2.5.3 Amplitude Calibration | 33 |
| 2.5.4 Bandpass Calibration | 33 |
| 2.5.5 Self-calibration | 34 |
| 2.6 Constructing an Image | 35 |
| 2.7 The CLEANing Process | 39 |
| 2.8 Wide-field Imaging | 40 |
| 2.9 Primary Beam Correction | 41 |
| 3 Identifying RFI with Machine Learning | 43 |
| 3.1 Introduction to Machine Learning | 43 |
| 3.2 Fundamental Concepts | 45 |
| 3.2.1 The Convolutional Neural Network | 48 |
| 3.2.2 Generative Adversarial Networks | 51 |
| 3.3 Radio Frequency Interference | 56 |
| 3.3.1 Regulation | 58 |
| 3.3.2 Methods of Mitigation | 60 |
| 3.3.3 Types of RFI | 63 |
| 3.4 Machine Learning for RFI Mitigation | 65 |
| 3.4.1 IFlag | 66 |
| 3.5 Flagging RFI with a GAN | 68 |
| 3.5.1 Training Data | 72 |
| 3.5.2 Pre-processing | 73 |
| 3.5.3 Results | 75 |

| | | |
|----------|------------------------------------------------------------------------------------------------|------------|
| 3.5.4 | Discussion | 77 |
| 3.6 | Dark and Quiet Skies | 81 |
| 4 | The Faint Radio Sky | 83 |
| 4.1 | Active Galactic Nuclei | 84 |
| 4.1.1 | The Unified Model of AGN | 86 |
| 4.2 | AGN at Radio Frequencies | 90 |
| 4.2.1 | Synchrotron Radiation | 90 |
| 4.3 | Radio-loud and Radio-quiet AGN | 92 |
| 4.4 | Star Forming Galaxies | 94 |
| 4.5 | A Dual Population | 95 |
| 5 | Investigating a Binary AGN Candidate | 99 |
| 5.1 | Revealing dual radio sources in a sub-kpc-scale binary active galactic nucleus candidate | 100 |
| 5.1.1 | Abstract | 100 |
| 5.1.2 | Introduction | 100 |
| 5.1.3 | Observations | 102 |
| 5.1.4 | Methods & Results | 105 |
| 5.1.5 | Discussion | 108 |
| 5.1.6 | Conclusions | 111 |
| 5.2 | Follow up EVN Observations | 112 |
| 6 | Studying the Lockman Hole | 115 |
| 6.1 | The e-MERLIN Observations | 116 |
| 6.2 | Calibration & Imaging | 117 |
| 6.2.1 | Noise Structure | 118 |
| 6.3 | Source Catalogue | 124 |
| 6.4 | Source Studies | 132 |
| 6.4.1 | S51 + S52 | 132 |
| 6.4.2 | S56 | 134 |
| 6.4.3 | S29 | 135 |
| 6.4.4 | S39 + S40 | 138 |
| 6.4.5 | S6 + 7 | 139 |
| 6.5 | Differential Source Counts | 140 |
| 6.6 | Discussion | 142 |
| 7 | Conclusions | 145 |
| 7.1 | The RFI Experiment | 145 |
| 7.2 | The Binary Black Hole Study | 146 |
| 7.3 | The Lockman Hole Study | 147 |
| | Acronyms | 149 |
| | References | 151 |

Acknowledgements

It is tough to remember all the people I should be thanking for their part in driving me to do this PhD. It's something I've wanted to do since I was just a curious child in high school, so how could I possibly adequately acknowledge all of the teachers, friends, family members, colleagues, and even acquaintances who have supported me and my love of learning along the way.

I should start with all the many wonderful members of the Jeremiah Horrocks Institute, where I have spent the last (and best so far) seven years of my life. I'm not going to use the next several pages expounding on the many ways they have all touched my life, but suffice to say that I would not be the person I am today without their guidance and support, both personally and professionally. I will miss them all.

From Jodrell Bank, I sincerely thank Nick Wrigley and David Williams-Baldwin for supplying the data that made this experiment possible, and for their kindness, practical guidance, and fascinating discussions.

Of course, I would not have made it this far without my amazing supervisor Megan Argo, who was always willing to indulge my long rants on various topics, and my relentless questions to go along with them. She is an inspiration to me, as she is to so many others, for her intellect, love of learning, and her tireless efforts in spreading that love of learning. Megan, it has been a privilege to work with you for so long, and I look forward to working together again in the future.

Away from work, I have to thank my friends for always giving me a place to go for laughter, fun, joy, interesting discussion, and sometimes intense debate! I also give thanks to all the wonderful members of Preston Samaritans, for training and teaching me so well, showing me the value of true kindness, and for making me feel that there is always someone to offer help. They give me hope in a sometimes bleak world, and I'll miss them all dearly.

And finally, my parents, who have been so instrumental in getting me here. My father, who spent countless hours indulging the questions of a curious child, and gave me a love of learning that defines so much of my life to this day. And my mother, who was always there to pick me up, give me a chuckle, and then push me forward even when I didn't want to.

Look again at that dot. That's here. That's home. That's us. On it everyone you love, everyone you know, everyone you ever heard of, every human being who ever was, lived out their lives. The aggregate of our joy and suffering, thousands of confident religions, ideologies, and economic doctrines, every hunter and forager, every hero and coward, every creator and destroyer of civilization, every king and peasant, every young couple in love, every mother and father, hopeful child, inventor and explorer, every teacher of morals, every corrupt politician, every "superstar," every "supreme leader," every saint and sinner in the history of our species lived there – on a mote of dust suspended in a sunbeam.

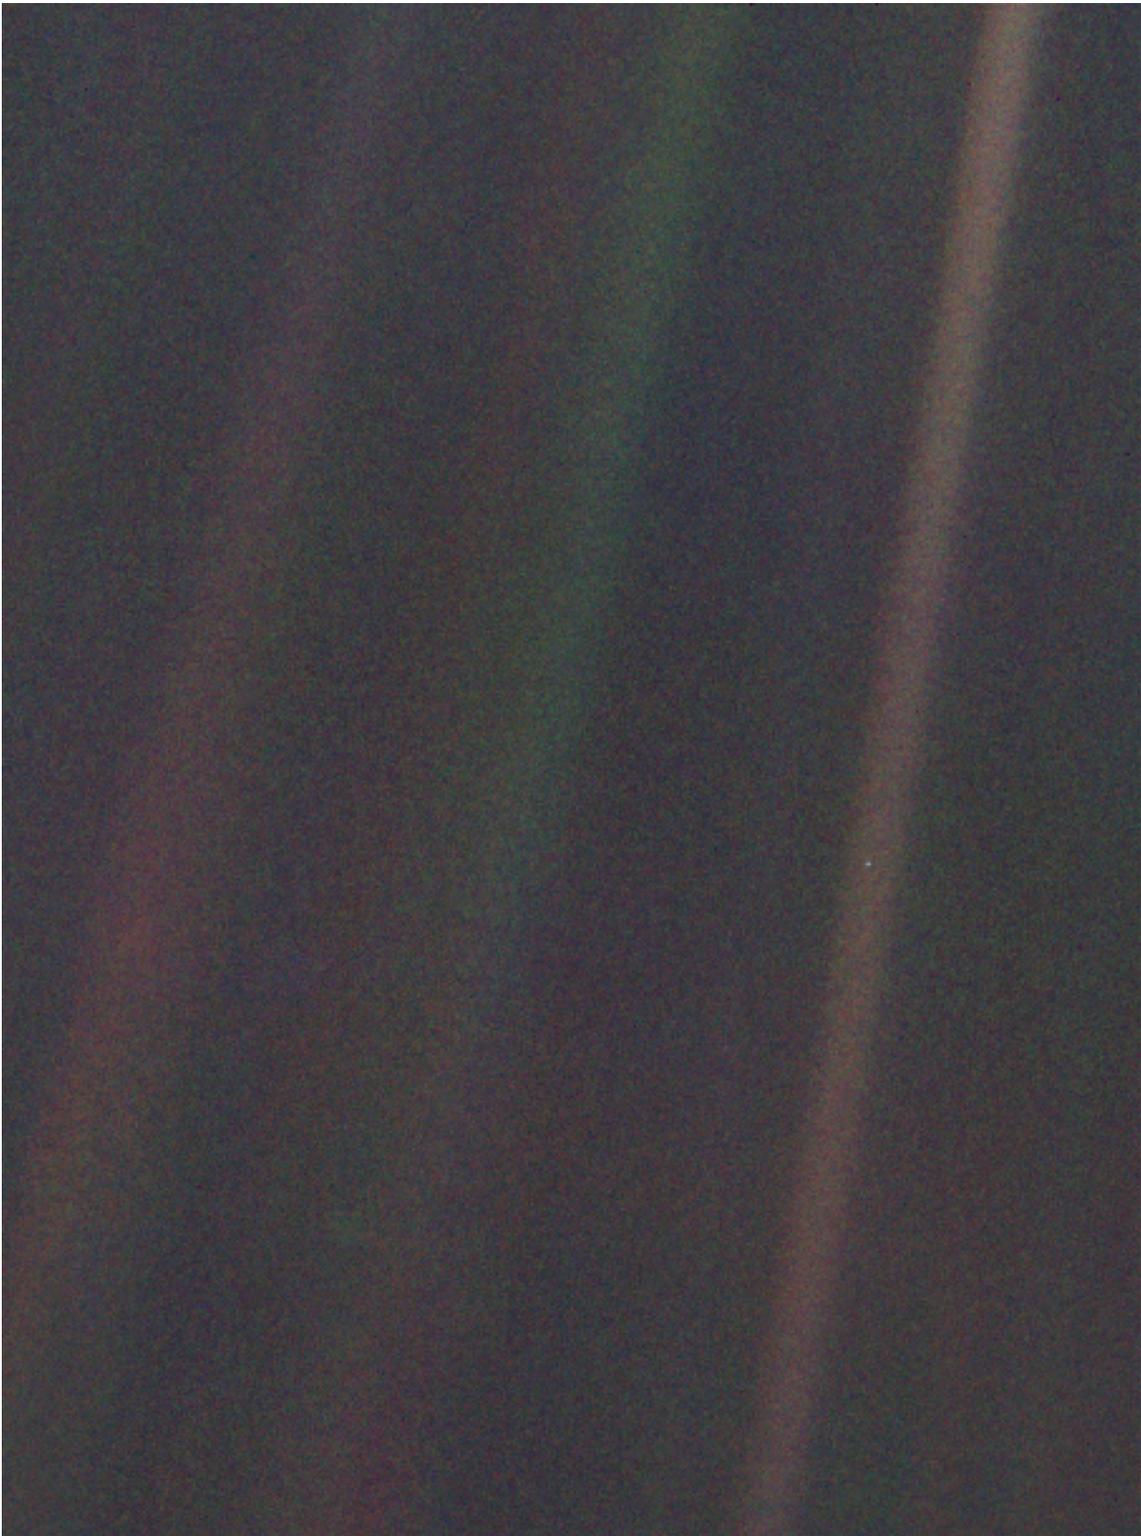
The Earth is a very small stage in a vast cosmic arena. Think of the rivers of blood spilled by all those generals and emperors so that, in glory and triumph, they could become the momentary masters of a fraction of a dot. Think of the endless cruelties visited by the inhabitants of one corner of this pixel on the scarcely distinguishable inhabitants of some other corner, how frequent their misunderstandings, how eager they are to kill one another, how fervent their hatreds.

Our posturings, our imagined self-importance, the delusion that we have some privileged position in the Universe, are challenged by this point of pale light. Our planet is a lonely speck in the great enveloping cosmic dark. In our obscurity, in all this vastness, there is no hint that help will come from elsewhere to save us from ourselves.

The Earth is the only world known so far to harbor life. There is nowhere else, at least in the near future, to which our species could migrate. Visit, yes. Settle, not yet. Like it or not, for the moment the Earth is where we make our stand.

It has been said that astronomy is a humbling and character-building experience. There is perhaps no better demonstration of the folly of human conceits than this distant image of our tiny world. To me, it underscores our responsibility to deal more kindly with one another, and to preserve and cherish the pale blue dot, the only home we've ever known.

– Carl Sagan,
describing the Voyager 1 image of Earth (overleaf)



Chapter 1

Introduction

In less than a century, the cutting edge of radio astronomy technology has gone from a metal and wire frame on wheels, to some of the largest, most complex equipment ever created for astronomical purposes (see [Figure 1.1](#)), giving astronomers a completely unique view of the universe. For the most part, this frenetic development has been driven by a desire for increased sensitivity and better imaging capability to keep up with other areas of astronomy, and to continue to explore the behaviour of astronomical sources at lower and lower radio flux densities with increasing accuracy. Furthermore, and arguably more importantly, radio astronomy reveals entirely unique phenomena including, but not limited to: the Cosmic Microwave Background (CMB), which is strongest at approximately 160 GHz; the famous 21 cm line, corresponding to a quantum state change in hydrogen atoms that produces energy of a precise frequency at approximately 1.4 GHz; supernova remnants, that are visible for long after their optical signature has faded away; and finally black holes, which are often surrounded by dust and gas that attenuates and obscures large portions of the radiation exiting from the central region, but is transparent to radio waves. The contribution of radio astronomers and observatories to the wider field of astronomy and astrophysics would take an entire text of its own to document, and is often not well understood by astronomers at large. Nevertheless, they provide a critical backbone in our understanding of our planet and the universe, from measuring tectonic plate shift, to studying the properties of spacetime.

The second driver of this massive progression is a desire for greater resolving power, to get an up-close, detailed view of the smallest and most compact objects, such as bound black hole systems. For traditional astronomy in the optical range the solution is *relatively* simple: build a bigger telescope. However, when observing radio waves this is rarely a



Figure 1.1: *Left:* Karl Jansky and his ‘Merry-go-Round’; one of the first radio instruments built for astronomical purposes. Image credit: NRAO¹. *Top right:* The VLA in New Mexico, United States. Image credit: NRAO². *Bottom right:* The Five-Hundred Aperture Spherical radio Telescope (FAST) in Guizhou Province, China. Image credit: NAOC³.

practical solution. To illustrate this, first recall that the resolving power of a telescope is generally described by the Rayleigh Criterion:

$$\theta = \frac{1.22\lambda}{D} \quad (1.1)$$

which represents the minimum angular spacing, θ , on the sky beyond which two point-like objects can be considered separate, when observing at wavelength λ , with a telescope aperture diameter of D . The Hubble Space Telescope (HST), for example, observes at wavelengths around 1 micrometre, and has a 2.4 m primary mirror. Therefore, the HST has a diffraction-limited resolving power of approximately $\theta = 0.1''$ in this configuration. Is it possible to achieve a similar resolving power at radio wavelengths? Again using [Equation 1.1](#), we can calculate the required telescope diameter for 1.4 GHz, which turns out to be $D = 504$ km, which would require about 82% of the entire land area of the United Kingdom! Suffice to say, building such a monolithic instrument is not practical, at least for now. While large single-dish radio telescopes, such as the Lovell telescope and the Effelsberg telescope, still exist and are still in use today, most high-resolution stud-

¹NRAO is the National Radio Astronomy Observatory. Image available at: <https://public.nrao.edu/gallery/karl-jansky-and-his-merrygoround>. Accessed 16th March 2024.

²Image available at: <https://public.nrao.edu/gallery/moonrise-over-the-very-large-array>. Accessed 16th March 2024.

³NAOC is the National Astronomical Observatories, Chinese Academy of Sciences. Image available at: <https://english.nao.cas.cn/research/researchdivisions/radioastronomy>. Accessed 16th March 2024.

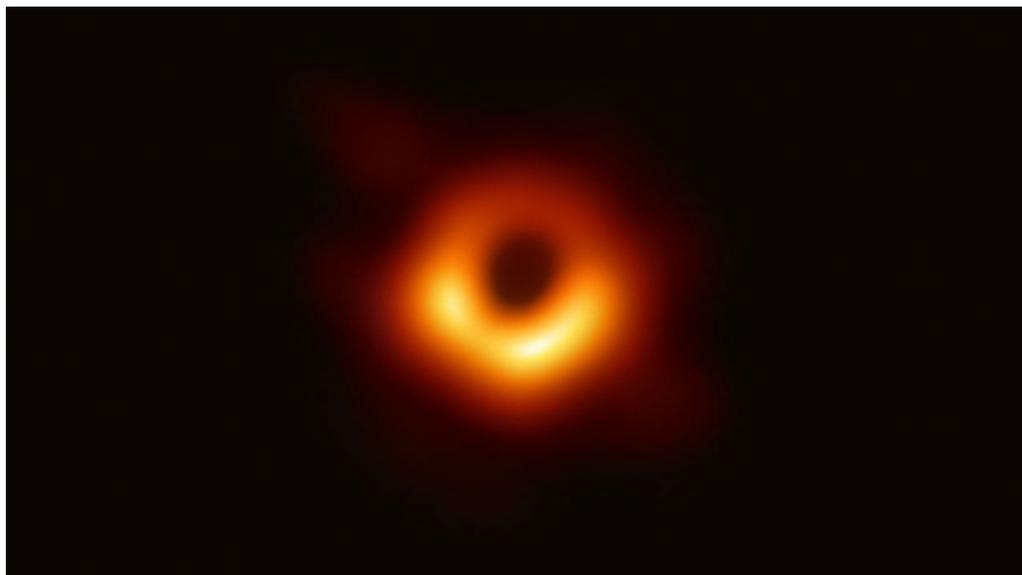


Figure 1.2: The Event Horizon Telescope (EHT) image of the black hole at the center of M87. The bright ring is the photon ring: incoming photons that have been curved towards us by the black hole's gravity. Image Credit: [Event Horizon Telescope Collaboration et al. \(2019\)](#)

ies require networks of smaller radio telescopes that link together to form interferometer arrays, like the Very Large Array (VLA, [Perley et al. 2011](#)). These telescope networks use the technique of aperture synthesis (see [chapter 2](#)) to create a representation of the radio sky with extremely high resolution that far surpasses instruments like the HST and its successor, the James Webb Space Telescope (JWST, [Gardner et al. 2006](#)). The technology has now developed so far that it has produced monumental scientific leaps, such as directly imaging the event horizon of the black hole at the centre of the galaxy Messier 87 (see [Figure 1.2](#), [Event Horizon Telescope Collaboration et al. 2019](#)). This landmark observation was performed by the EHT, which is a massive interferometer network that creates a radio telescope spanning the entire planet ([Doeleman et al., 2009](#)). The technique of radio interferometry continues to grow and develop, with pioneering institutions and observatories around the world preparing for the era of the Square-Kilometer Array (SKA), which will be the most powerful radio telescope in the world once completed ([Ekers, 2012](#)).

1.1 Radio Interference

During the course of a radio observation, which may take place over many hours, interfering signals are often detected from many different non-astronomical sources. These signals are usually quite strong compared to the desired sky signal, and create problems

during calibration of the observational data, and if left unaddressed the final image as well.

In the early days of radio astronomy, radio signals generated by humans that might interfere with an observation of the sky were comparatively rare, and given sufficient distance from civilisation one could conduct an astronomical radio observation without much interference. This began to change over the course of the late 20th century, and radio emissions became more widespread in and about populated regions. Still, it was common practice for radio astronomers to manually remove these interfering signals from their observations, scarce as they were. Now in the 21st century, radio signals have become so pervasive that manual intervention is (in general) no longer practical, and automated methods of interference removal are preferable and commonplace. The widespread adoption of radio emitting mobile phones (and the radio frequency networks that supply them), the massive increase in the Earth satellite population, the explosion of wireless internet, and the lasting existence of broadcast radio are all examples of the modern day radio environment that radio observatories must find a way to operate within.

The eternal challenge, however, of automatically removing interference in radio observations is to use or create a tool that operates with precision. The complexity of this task stems from the very blurry line that exists between interference and ‘not interference’ in the absence of full knowledge of the origin of the two signals. Therefore, how one decides to draw the line between these two is equivalent to finding the exact point at which the color white transitions to the color black on an infinitely continuous scale of gray, and is also the ultimate determining factor in the performance of the tool. Of course, it is possible to at least approximate this hypothetical boundary through the scientific method; using a combination of both *a priori* and *a posteriori* knowledge of the incoming radio signals, and the equipment used to detect them, to construct a method of filtering out unwanted signals. Achieving the ever sought after objective of close to 100% automatic removal of interfering signals requires intelligent methods and algorithms, even more so if one wishes the solution to work across all radio observatories. This particular desire for consistency across observatories comes from a need for scientific reproducibility: if two scientists use the same tools, their respective results become far more useful in comparison.

Currently, most observatories implement their own *ad hoc* measures, both in software and hardware, to reduce the impact of interference. After receiving data from the observatory, many radio astronomers will also want to fine-tune their observation by re-

examining the flagged data, and in this case there are a limited number of statistical tools available. The state of the art in automatic interference removal, and the most common generalised tool, is known as AOFLAGGER (Offringa et al., 2012b), which implements a form of adaptive thresholding. Among all tools though, including AOFLAGGER, there is still room for improvement in precision, as the designers of automated tools usually prefer to overestimate the presence of interference rather than risk allowing it through to the final image. This leads to an increased number of false positives that ultimately costs in effective observing time. Continuous refinement of these techniques to reduce the rate of false positives is necessary to prevent this needless loss of effective observing time.

While modern radio observatories are far more capable than their predecessors, that capability has come with a catch. The progress made in making observatories more sensitive to fainter radio emission has also made them more vulnerable to non-astronomical interference. Therefore, as our observing instruments develop, so too must our methods of handling interference in order to keep pace with an ever changing radio environment. This effort is the subject of [chapter 3](#) of this thesis. Appropriately dealing with radio interference is a critical step in conducting deeper, more sensitive observations of the radio sky, which in turn allows an exploration of faint and high-redshift source populations.

1.2 Faint Radio Sources

The progressive research and reinterpretation of the nature of radio sources is an excellent case study of the process of scientific discovery. Early radio telescopes being developed in the 1950s did not have the observing capability we expect from our instruments today; poor sensitivity, resolving power, and bandwidth (at least by modern standards) were normal. This means that experiments to discern the nature of radio sources were constrained by the number of sources they could detect above the sensitivity of the telescope being used. While this limitation always has been and always will be present to some degree, modern telescopes are, nevertheless, orders of magnitude better in key performance metrics. However, this change did not happen instantly; it is only because of the tireless efforts of astronomers and engineers over just seven decades to improve the technology they use to observe. Though necessary, this translated into a constant reinterpretation of the radio sky as new phenomena are observed, such as fast radio bursts (Lorimer et al., 2007; Petroff et al., 2019), and even entirely new objects are discovered, such as pulsars (Hewish et al., 1968). This effect is at least partly responsible for the semantic spaghetti

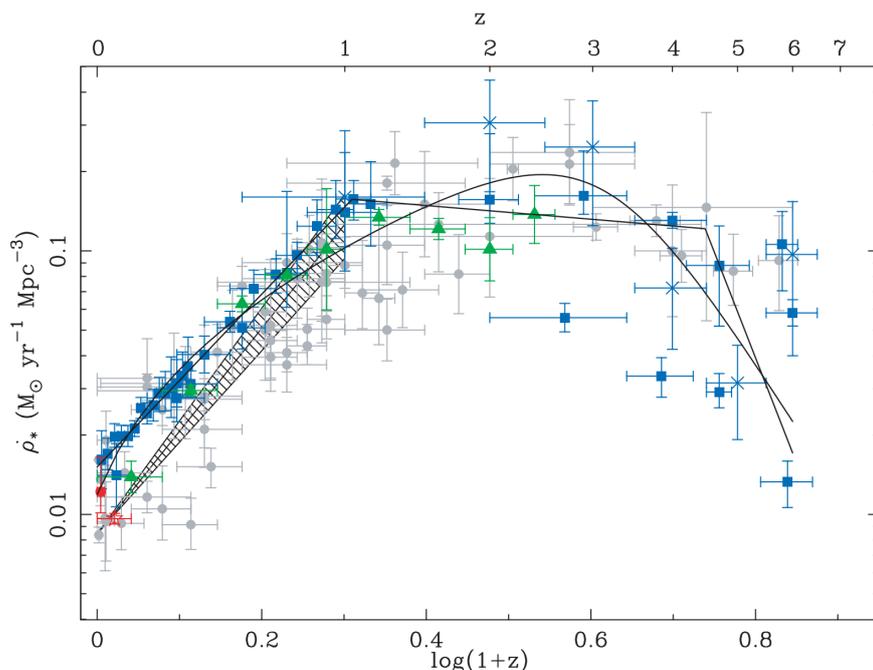


Figure 1.3: The evolution of the star formation rate density with redshift; the famous ‘Lilly-Madau plot’ (from Lilly et al. 1995 and Madau et al. 1996). The data points are from a variety of sources using data from across the electromagnetic spectrum. Note the peak in star formation rate density around $z \sim 2$ (~ 3 Gyr after the big bang), indicating that the majority of observable stellar mass was formed at redshifts $z \gtrsim 1$. Image Credit: Hopkins & Beacom (2006)

of Active Galactic Nuclei (AGN) that is only just being disentangled (see section 4.1).

One interesting aspect of the faint radio sky, defined in this work as the radio sky below $\lesssim 1$ mJy, is the relative distribution of the three main components presented in chapter 4: jetted and non-jetted AGN, and Star-forming Galaxies (SFGs). Jetted AGN dominate at flux densities $\gtrsim 400 \mu\text{Jy}$, while non-jetted AGN and SFGs share significant portions of the sky at $\lesssim 100 \mu\text{Jy}$ and decline in number at higher flux densities. The exact distribution of non-jetted active galaxies and star-forming galaxies, and how that changes with redshift, is an informative metric as it can be used to trace the cosmic star-formation rate which shows a distinct and curious peak at $z \sim 2$, and then declines in recent cosmic history (see Figure 1.3, see also Madau & Dickinson 2014).

Radio telescopes offer a unique view on this area of research. First and foremost, radio telescopes are not forced to use extinction corrections as they are insensitive to dust attenuation, which significantly reduces systematic uncertainties in any study. In addition, through combination with measurements at infrared wavelengths, AGN with jets that make up the majority of radio emission can be reliably classified using the FIR-Radio correlation (van der Kruit, 1973; Condon, 1992). This makes it relatively trivial, provided infrared measurements are available, to identify the group of non-jetted AGN and SFGs

in a sample. From this point, there are several methods of decomposing the AGN and SF components. If possible, it is always recommended to incorporate X-ray sources, since these are a clear indicator of the presence of an AGN. There are some (less-reliable) methods that can be carried out entirely at radio frequencies, which has the benefit of eliminating the limitations involved with multi-wavelength studies. One commonly used method is to measure the radio spectral slope Γ (as defined by $S(\nu) \propto \nu^\Gamma$)⁴, which can be a reliable indicator of AGN in particular in sources with $\Gamma > 0$ (Padovani, 2016). Another method is to track the radio morphology at varying angular scales; a method only available to radio interferometers through varied baseline lengths and their weightings.

The work presented in [chapter 6](#) seeks to provide evidence of the nature of the faint radio population in the Lockman Hole, as observed by the *enhanced* Multi Element Remotely Linked Interferometer Network (e-MERLIN) radio telescope. It is intended to be a complimentary study to the work of [Muxlow et al. \(2020\)](#), which investigates a different field, the Great Observatories Origins Deep Survey - North (GOODS-N) field, using e-MERLIN and VLA observations. By combining the conclusions of these two studies in two separate fields, any aspects of cosmic variance should be accounted for, and each can augment the others scientific value. Before a radio study like this can be conducted, however, there is work to be done on preparing and calibrating the observational data; introducing this process is the subject of [chapter 2](#). To maximise the sensitivity of the final radio survey, I investigate alternative methods of identifying and removing radio interference from the Lockman Hole observations. This work is presented in [chapter 3](#). In [chapter 4](#) I introduce the topic of the faint radio sky, its constituents, and the current research being conducted. A study of an extremely faint sub-kpc scale binary AGN candidate is presented in [chapter 5](#). I present the study of the Lockman Hole with e-MERLIN in [chapter 6](#), and then finish with some conclusions in [chapter 7](#) where I also discuss some recommendations for future work.

⁴Different authors will use different conventions for the sign of Γ . In this work I follow the recommendation of [Burke et al. \(2019\)](#) to use the positive convention.

Chapter 2

Techniques of Radio Astronomy

Studies involving radio telescopes often require some knowledge of how the observations are conducted, how the data is processed, and the assumptions that are made in each of these stages to properly assess the end conclusions. This chapter seeks to provide an introduction to the governing theory of radio interferometry, such that the results presented later in this work can be better understood. An overview of existing observatories, and the e-MERLIN radio telescope that is used for the majority of observations presented in this work, is given in [section 2.1](#) and [section 2.2](#) respectively. Then, the technique of interferometry that all of these observatories rely on is introduced in [section 2.3](#), by first starting at the basic example of an interferometer of two stations. This allows an explanation of the detection and correlation process that remains the same for even larger arrays. Once this foundation is laid, an explanation is given of the process of calibration, and the importance of ensuring a high-quality calibration that is not corrupted by unwanted radio signals. Then, the unique process of imaging at radio frequencies is described in [section 2.6](#), which is very different from more traditional imaging based on Charge-Coupled Devices (CCDs) used in optical imaging. The final sections will discuss some of the complexities introduced by imaging a ‘wide-field’, and the techniques used to overcome those challenges. Much of the following chapter was written using the combined works of [Wilson et al. \(2013\)](#), [Thompson et al. \(2017\)](#), and [Burke et al. \(2019\)](#) which are foundational texts in much of radio astronomy and interferometry.

2.1 Radio Interferometry

One of the many differences that distinguishes radio astronomy from other areas of astronomy is how observations of the sky are actually performed. Broadly speaking, there are two distinct methods: using either one large radio antenna, or using many smaller but connected radio antennas that detect the same incoming radio waves. The former offers a simpler and cleaner post-observation workflow but is limited in resolving power, the latter is more flexible and scalable, but has its own set of drawbacks. In this work I will exclusively focus on connected networks of smaller radio antennas, known as radio interferometers. These networks are further divided loosely into two groups based on the spacing between their constituent antennas (known as baselines), as this has major implications for not just the scientific output, but also the difficulty, complexity, and technology required for operation. Short baselines, for example, can rely on simple wire connections over short distances for transmitting the received signal. For higher resolution studies longer baselines are needed, and running a wire directly from the Australian outback to South Africa is not a practical solution. Longer baseline arrays therefore, use the Very-Long-Baseline Interferometry (VLBI) technique, which is an extension of the core interferometer technology but over continents instead of kilometres. Some of the most well known radio interferometer networks include:

- **Very Large Array - VLA**

In operation for nearly 50 years, the VLA is a leading instrument for \sim arcsecond resolution studies in the 1-50 GHz range. It is based in Socorro, New Mexico and is operated by the US National Radio Astronomy Observatory (NRAO)

See: [Perley et al. \(2011\)](#).

- **Atacama Large Millimeter Array - ALMA**

Situated on the Chajnantor Plateau in Chile, ALMA covers the higher frequency portion of the radio spectrum extending from 30-950 GHz, with a resolving power similar to that of the VLA. It is funded and operated by a collaboration of organisations from around the world.

See: [Wootten & Thompson \(2009\)](#)

- **Square Kilometer Array - SKA**

The SKA, when constructed, will be one of the most advanced radio interferometers on the planet. It will consist of 133 antennas located in South Africa (known as

SKA-mid), and just over 130000 smaller dipole antennas located in western Australia (known as SKA-low). These two parts will collectively cover a wide frequency range between 0.05-25 GHz. It is an ambitious project that drives much innovation and research into new and existing techniques, to ensure that astronomers can maximise its potential.

See: [Braun et al. \(2019\)](#)

- **Low-Frequency Array - LOFAR**

Covering the lower end of the frequency spectrum at 10-240 MHz, LOFAR instead uses dipole antennas to observe astronomical phenomena in a sparsely-explored window of the spectrum.

See: [van Haarlem et al. \(2013\)](#)

- **Multi Element Remotely Linked Interferometer Network - e-MERLIN**

Leading the UK radio astronomy effort is e-MERLIN, incorporating the nearly 70 year old 76 m Lovell telescope. Formerly known as MERLIN, the upgraded version, e-MERLIN, uses seven stations across the UK to produce ~milliarcsecond resolution images in the 1-25 GHz range.

See: [Garrington et al. \(2004\)](#). See also [section 2.2](#) for more detail.

- **European VLBI Network - EVN**

Bringing together stations primarily from around Europe (but also on other continents), the EVN produces some of the highest resolution images currently possible, primarily in the 1-25 GHz range.

See: [Venturi et al. \(2020\)](#)

The small sample of observatories listed above are roughly in order corresponding to the length of baselines within the array of stations they bring together, and each produces a different view of the universe. Astronomers seeking to use one of these facilities must carefully examine their science objective to determine which facility is most suited to them, as sources that appear bright and clear to one, might be invisible to another, even at the same frequency. This is because the length of the baselines in the array does not just determine the achievable resolving power, but also the angular scale the array is sensitive to. For example, arrays with an extremely high resolving power are not suitable for studying diffuse supernova remnants with a large angular size; the array simply will not detect the low surface brightness emission.



Figure 2.1: A picture of the Lovell telescope at Jodrell Bank Observatory. Image Credit: BBC¹.

2.2 The e-MERLIN Telescope

During World War II, the use of radar for both offensive and defensive purposes was pioneered by the British Royal Air Force. Most notably, the ‘Chain Home’ system of radar antennas spread across England was a critical factor in the Battle of Britain, giving time for people to escape to air-raid shelters, and for British fighters to scramble. At this time, Sir Bernard Lovell was developing the H2S radar system for nighttime air raids when he noticed that certain radar signals were ignored by operators, though they did not know their cause (Argo, 2012). After the war, Lovell returned to the University of Manchester where he continued his radio astronomy research, and later founded what came to be known as Jodrell Bank Observatory. Here, he constructed a massive 250 ft telescope; the largest in the world at the time and capable of observing any point in the sky. This telescope was later renamed in his honour giving us the Lovell Telescope, seen in Figure 2.1, which remains operational to this day.

During the 1970s and 80s, Lovell and his team also pioneered the technique of radio interferometry, and constructed the MERLIN array of antennas (Thomasson, 1986) based at Jodrell Bank Observatory. Since then, the MERLIN array, now upgraded to the e-MERLIN array (Garrington et al., 2004), has been a leading instrument in radio astronomy. The network connects together seven antennas spread across the United King-

¹Available at: <https://www.bbc.co.uk/news/uk-england-manchester-61630241>. Accessed: 15th March 2024

| | L-Band (1.5 GHz) | C-band (5 GHz) | K-band (22 GHz) |
|----------------------------|------------------|------------------|-------------------|
| Resolution | 0.2'' | 0.05'' | 0.02'' |
| Primary beam width | 30' | 7' | 2' |
| Sensitivity (from ~12 hrs) | 24 μ Jy/beam | 16 μ Jy/beam | 120 μ Jy/beam |

| Antenna | Mk2 Baseline Length (3D Geocentric) | Dish Size |
|----------------|-------------------------------------|-----------|
| Lovell (Lo) | 0.4 km | 76 m |
| Mark II (Mk2) | 0.0 km | 24 m |
| Pickmere (Pi) | 11.2 km | 25 m |
| Darnhall (Da) | 17.5 km | 25 m |
| Knockin (Kn) | 67.8 km | 25 m |
| Defford (De) | 126.6 km | 25 m |
| Cambridge (Cm) | 197.8 km | 32 m |

Table 2.1: Operating characteristics for the e-MERLIN interferometer array. The top table presents performance values for each of the available observing bands. The bottom table presents position and size values for each of the antennas in the array.

dom, including the Lovell telescope itself, via fiber-optic cable. The long baselines of e-MERLIN allow for high resolution observations at ~milliarcsecond scales. More details of e-MERLIN can be found in [Table 2.1](#).

2.3 The Two-Element Interferometer

At a fundamental level, a radio interferometer effectively replicates Young’s Double Slit experiment ([Young, 1804](#)), where two slits in a flat sheet of material are illuminated by an incident plane wave. Behind the sheet, an interference pattern (made up of a series of ‘fringes’) is produced on another surface from the interaction of light waves travelling through the two slits. The shape of the interference pattern contains information about the incident light and is also dependent on the setup of the experiment. Mapping the analogy to a radio interferometer, the slits can be thought of as radio antennas and the sheet with the interference pattern projected onto it is the ‘image plane’ (see [section 2.6](#)). The uv plane (see [section 2.4](#)) represents the illumination of the slits, which exhibits a Fourier transform relationship with the interference pattern. The shape of the resulting interference pattern is determined by the properties of the light entering the antennas (frequency, amplitude, phase), which all conspire to produce an interference pattern that is a direct representation of the sky. The comparison is most obvious in [Figure 2.6](#).

Radio interferometers with many antennas can, in one way or another, be effectively represented as a series of two-element interferometers, that is, a pair of antennas that produce a correlated radio signal. It is, therefore, a useful and commonly used starting

point for understanding how an interferometer works. To make effective use of the data produced by an interferometer it is critical to have a deep understanding of the mathematical significance of the received signal. What are its limitations? What effects can be corrected for? What effects can be safely ignored, and under what circumstances? All of these questions can be answered through a rigorous mathematical definition.

Currently, there are a few mathematical formalisms that solidify the interferometer response. In the latter half of the 20th century, much of the mathematics surrounding radio interferometers was relatively arcane, and contained many (at the time, necessary) assumptions that limited the scope of the observations. Works such as [Bracewell \(1956\)](#) and [Morris et al. \(1964\)](#) represented significant attempts to formalise the effects an incoming signal experiences between emission and correlation. However, the approach taken in these works involved extremely high level mathematics, which made it more difficult for the wider community to fully incorporate it. An alternative approach was put forward by [Hamaker et al. \(1996\)](#) which involved using 4x4 matrix algebra to condense the long-form mathematics into something more intuitive. At the same time, it allowed signal propagation effects to be neatly described by singular algebraic terms, representing a particular effect or transformation that occurs in the path to correlation. A similar approach, the one that is derived in this section, involves using 2x2 Jones matrices ([Jones, 1941](#)). This Jones formalism was alluded to by [Hamaker et al. \(1996\)](#), but fully encapsulated by [Smirnov \(2011\)](#) who demonstrated its advantages over previous formalisms, both in terms of tractability and extensibility.

Before considering the mathematics of the incoming radio signal, I will first provide a brief description of the journey of a radio wave from its source. Unlike many other electromagnetic waves, radio waves are largely undisrupted when travelling through the Interstellar Medium (ISM), so radio waves that arrive at the Earth's atmosphere are largely untouched since the time they were emitted. Similarly, the Earth's atmosphere is mostly transparent to radio waves. Since the distance over which the incoming wave is observable (i.e. less than the diameter of the Earth for Earth-based antennas) is much smaller than the distance it has travelled, a plane-wave approximation is appropriate². Therefore, a monochromatic electromagnetic signal propagating through space at an instantaneous

²While this is certainly true for observations of extragalactic objects, some intragalactic and intrasolar observations require more complex models of the incoming wave (e.g. [Sekido & Fukushima 2006](#)).

point in time can be described by a column vector of two complex cartesian components:

$$\boldsymbol{\alpha} = \begin{pmatrix} \alpha_x \\ \alpha_y \end{pmatrix} \quad (2.1)$$

The complex components α_x and α_y exist in a cartesian plane perpendicular to the direction of propagation. They are fully expressed by the classic complex representation of a wave:

$$\alpha = Ae^{i\theta} = A(\cos \theta + i \sin \theta) \quad (2.2)$$

In this representation, the real part contains amplitude information, whereas the complex part (or the real part of the exponent, θ) contains phase information.

Almost all radio antenna feeds possess two dipoles sensitive to each component of $\boldsymbol{\alpha}$, in order to preserve polarisation information. The signal $\boldsymbol{\alpha}$, therefore, induces a proportional voltage \boldsymbol{v} that can also be described by a column vector of two complex components representing the two dipoles:

$$\boldsymbol{v} = \begin{pmatrix} v_x \\ v_y \end{pmatrix} \quad (2.3)$$

An assumption of the Jones formalism is that all signal propagation effects can be described by one or more linear transformations of the original signal $\boldsymbol{\alpha}$. In linear algebra, this means that we can describe the conversion from a signal $\boldsymbol{\alpha}$ to a voltage \boldsymbol{v} by matrix multiplication with a 2x2 Jones matrix:

$$\boldsymbol{v} = \mathbf{J}\boldsymbol{\alpha} \quad (2.4)$$

In an example setup using two linear dipoles, there are two electronic gain factors associated with each dipole, so [Equation 2.4](#) expands to:

$$\boldsymbol{v} = \begin{pmatrix} v_x \\ v_y \end{pmatrix} = \begin{pmatrix} g_x & 0 \\ 0 & g_y \end{pmatrix} \begin{pmatrix} \alpha_x \\ \alpha_y \end{pmatrix} \quad (2.5)$$

In the case of a two-element interferometer, shown in [Figure 2.2](#), there are two antennas A and B that produce two voltage signals \boldsymbol{v}_A and \boldsymbol{v}_B , where:

$$\boldsymbol{v}_A = \begin{pmatrix} v_{Ax} \\ v_{Ay} \end{pmatrix} \text{ and } \boldsymbol{v}_B = \begin{pmatrix} v_{Bx} \\ v_{By} \end{pmatrix} \quad (2.6)$$

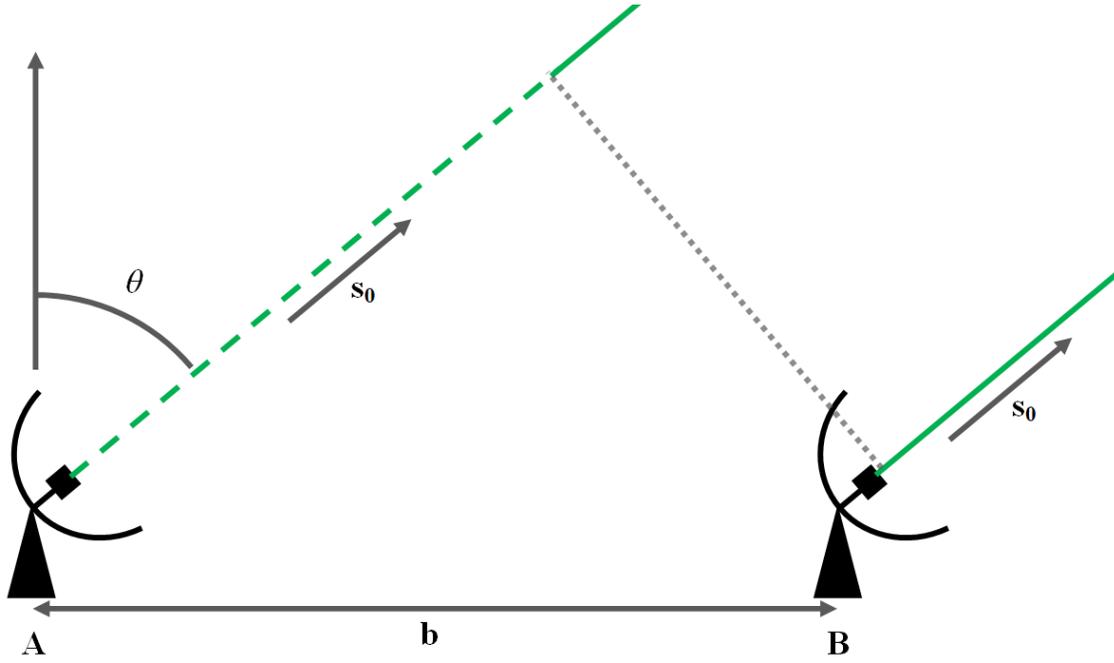


Figure 2.2: A diagram illustrating the simple case of an interferometer of only two antennas: the two-element interferometer. Two antennas A and B are separated by a distance in metres, b . The antennas both point in the direction of the unit vector s_0 , which makes an angle θ with the local zenith.

These voltage signals are sent to the correlator, a device (that exists in either software or hardware) which performs a cross-correlation operation between the four components over some time interval. The output of this operation is described by:

$$\mathbf{V} = 2 \begin{pmatrix} \langle v_{Ax} v_{Bx}^* \rangle & \langle v_{Ax} v_{By}^* \rangle \\ \langle v_{Ay} v_{Bx}^* \rangle & \langle v_{Ay} v_{By}^* \rangle \end{pmatrix} = 2 \left\langle \begin{pmatrix} v_{Ax} \\ v_{Ay} \end{pmatrix} \begin{pmatrix} v_{Bx}^* & v_{By}^* \end{pmatrix} \right\rangle = 2 \langle \mathbf{v}_A \mathbf{v}_B^H \rangle \quad (2.7)$$

where H represents a Hermitian transpose operation. Just as in [Equation 2.4](#), the quantity V , known as the *visibility*, can be represented as a linear transformation of the original signals from the two antennas α_A and α_B :

$$\mathbf{V} = 2 \langle \mathbf{J}_A (\alpha_A \alpha_B^H) \mathbf{J}_B^H \rangle \quad (2.8)$$

Note that each antenna has a different Jones transformation matrix because the transformation from electromagnetic signal to voltage could be different for each antenna. Following the notation of [Smirnov \(2011\)](#), the central quantity and the factor 2 in the above equation is redefined as the brightness matrix $\mathbf{B} = 2(\alpha_A \alpha_B^H)$, and the Jones matrices are assumed to be constant over the averaging interval, leading to:

$$\mathbf{V} = \mathbf{J}_A \mathbf{B} \mathbf{J}_B^H \quad (2.9)$$

which now simply describes the output of an ideal interferometer as an incoming electromagnetic signal \mathbf{B} detected by two separated antennas A and B with transformations to voltage described by \mathbf{J}_A and \mathbf{J}_B respectively.

One of the defining problems of any radio interferometer is that the incoming astronomical signal is detected by antenna B some finite time after it is detected by antenna A. This time delay between successive detections due to the placement of the antennas is known as the *geometric delay*, and causes a phase shift in the signal at B with respect to the signal at A that must be corrected for to prevent destructive addition of signals from the direction of observation. To calculate this phase shift and how it can be applied to the brightness matrix \mathbf{B} , consider the two-element interferometer in [Figure 2.2](#). The antennas A and B are separated by a distance b that produces a path length difference:

$$\Delta p = b \sin \theta \quad (2.10)$$

which corresponds to a time delay of:

$$\tau = \frac{\Delta p}{c} = \frac{b \sin \theta}{c} \quad (2.11)$$

which produces a phase delay relative to the signal at A of:

$$\phi = 2\pi\nu\tau = 2\pi\frac{c}{\lambda}\frac{b \sin \theta}{c} = 2\pi b_\lambda \sin \theta \quad (2.12)$$

where the delay, b , is changed to be expressed in wavelengths instead of metres. Note that the phase shift ϕ is dependent on both time τ and frequency ν , as θ will change as the source traverses the sky. In three dimensions, the phase difference ϕ can be expressed as the dot product of the pointing vector \mathbf{s}_0 and the antenna separation in wavelengths \mathbf{b}_λ :

$$\phi = 2\pi\mathbf{b}_\lambda \cdot \mathbf{s}_0 \quad (2.13)$$

To apply this quantity to \mathbf{B} , we arrange it into the scalar matrix \mathbf{K} :

$$\mathbf{K} = \begin{pmatrix} e^{-i\phi} & 0 \\ 0 & e^{-i\phi} \end{pmatrix} \quad (2.14)$$

and since \mathbf{K} is able to commute with other Jones matrices, it can then be applied to \mathbf{B}

through another matrix multiplication:

$$\mathbf{V} = \mathbf{J}_A \mathbf{K} \mathbf{B} \mathbf{J}_B^H \quad (2.15)$$

which describes the interferometer response with the geometric delay accounted for. However there are several limitations to the description of the visibility matrix in Equation 2.15. Firstly, the \mathbf{K} correction is defined in a coordinate frame centred on antenna A which, while valid, is inconvenient in an interferometer of more than two antennas. It would be more convenient to have the ability to define \mathbf{K} for all antennas from some arbitrary reference point (which may also be an antenna). Furthermore, there is actually a mismatch in coordinate frames between \mathbf{K} and \mathbf{B} , since \mathbf{K} is defined in a cartesian frame with z axis pointing towards the zenith, and \mathbf{B} is defined in a cartesian frame with z axis pointing towards the source direction. Secondly, we have assumed that the incoming signal α originates from a single perfectly point-like source in exactly the direction indicated by the pointing vector s_0 . In reality, the quantity α varies across the sky as it consists of many sources in all shapes and sizes. Lastly, it is entirely unclear how one could go from the quantity \mathbf{V} to a useful image of the radio sky. To achieve this, and to account for the limitations described above, there is further work that needs to be done in defining both the geometric delay and the incoming signal in a shared frame of reference.

2.4 Aperture Synthesis

The conceptual foundation on which radio interferometers are built is the uvw coordinate system, or uv coordinate system. The uv coordinate system, illustrated in figure Figure 2.3, is a cartesian coordinate system which defines a plane perpendicular to the direction of the source. The coordinates are most often defined in units of wavelengths, such that a point in the uv system at one frequency will have different coordinates at another frequency. The axis oriented towards the source, the w axis, is very often chosen to be parallel with the pointing vector s_0 , which also means that uv coordinates have a time dependence as the source moves across the sky. As per convention, the u and v axes are oriented east and north respectively. Antenna positions are defined in the uv plane, such that an antenna p operating at a certain time and frequency exists in the uv plane at a point $\mathbf{u}_p = (u_p, v_p, w_p)$.

Now, to represent an arbitrary point on the sky we can define an offset σ that is or-

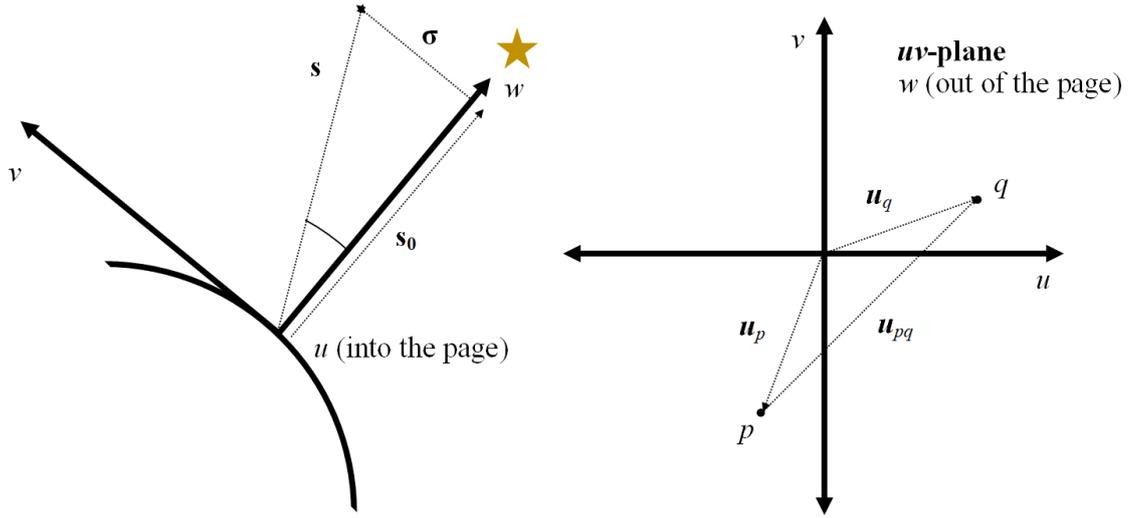


Figure 2.3: Representations of the uv coordinate system with w axis oriented towards the source indicated by the star. The left image views the system side-on, with the u axis projected into the page. The right image views the system from above, with the w axis pointing out of the page. The antennas p and q are defined in the uv plane, along with their separation.

thogonal to s_0 , which means we can define a new unit vector $s = s_0 \cdot \sigma$ that points to a location on the sky tangent plane. The vectors s and s_0 have three components, l , m , and n that are formally defined as direction cosines, and since s has a magnitude of unity, by definition, we can write n in terms of the other two components as $n = \sqrt{1 - l^2 - m^2}$. The components l and m exist on a plane that effectively projects onto the sky in the ‘far-field’, following the assumption about plane waves we made earlier. This also means that the l, m plane exists parallel to the uv plane.

Now we have defined the antenna position \mathbf{u}_p in the uv plane, we can write an equation for the phase difference due to the geometric delay at antenna p relative to $\mathbf{u} = 0$ in a signal coming from direction s_0 as:

$$\phi_p = 2\pi \mathbf{u}_p \cdot s_0 = 2\pi(u_p l_0 + v_p m_0 + w_p \sqrt{1 - l_0^2 - m_0^2}) = 2\pi w_p \quad (2.16)$$

as $s_0 = (l_0, m_0, n_0) = (0, 0, 1)$. Turning back to the two-element interferometer from [section 2.3](#), since we have now defined ϕ_p relative to $\mathbf{u} = 0$ instead of the exact location of antenna A, we require two delay corrections for both antennas A and B:

$$\mathbf{K}_A = \begin{pmatrix} e^{-i\phi_A} & 0 \\ 0 & e^{-i\phi_A} \end{pmatrix} \text{ and } \mathbf{K}_B = \begin{pmatrix} e^{-i\phi_B} & 0 \\ 0 & e^{-i\phi_B} \end{pmatrix} \quad (2.17)$$

which is applied like so:

$$\mathbf{V} = \mathbf{J}_A \mathbf{K}_A \mathbf{B} \mathbf{K}_B^H \mathbf{J}_B^H \quad (2.18)$$

This equation can be further generalised to the visibility from any two antennas p and q :

$$\mathbf{V}_{pq} = \mathbf{J}_p \mathbf{K}_p \mathbf{B} \mathbf{K}_q^H \mathbf{J}_q^H \quad (2.19)$$

Therefore, to summarise, Equation 2.19, represents a radio signal \mathbf{B} , coming from direction \mathbf{s}_0 , that is detected by two antennas separated by a distance $\mathbf{u}_p - \mathbf{u}_q$. The geometric delay relative to $\mathbf{u} = 0$ is corrected for by the phase delay matrices \mathbf{K}_p and \mathbf{K}_q , and the complex gains of the antenna instruments are corrected for by the gain matrices \mathbf{J}_p and \mathbf{J}_q .

Though, this is still a signal coming from a single, perfect point source at the so called ‘phase centre’ \mathbf{s}_0 . To relax this assumption, we must recognise that \mathbf{B} will vary over the ‘on sky’ coordinates l and m . This also means however that both \mathbf{K} and \mathbf{J} will vary over the sky as well.

It is simpler to understand the variation of \mathbf{K} by imagining the small deviations of $\boldsymbol{\sigma}$ from the pointing vector \mathbf{s}_0 in Figure 2.2. When the signal direction deviates from \mathbf{s}_0 the geometry of the setup changes, and therefore \mathbf{K} must also change to compensate, leading to an additional direction dependent delay correction $\mathbf{K}(\boldsymbol{\sigma})$. The total delay correction to the system therefore has two components: one component to correct for the geometry of the relative antenna positioning, and another to correct for deviations from the phase centre: $\boldsymbol{\sigma}$. If we assume that the primary geometric delay $\phi = 2\pi\mathbf{u} \cdot \mathbf{s}_0 = 2\pi w$ is handled directly in instrumentation (as it often is), then we are simply left with the residual delay due to varying signal direction $\mathbf{K}(\boldsymbol{\sigma})$. Redefining Equation 2.16 we have:

$$\phi_p = 2\pi\mathbf{u}_p \cdot \boldsymbol{\sigma} = 2\pi(\mathbf{u}_p \cdot \mathbf{s} - \mathbf{u}_p \cdot \mathbf{s}_0) = 2\pi(u_p l + v_p m + w_p(\sqrt{1 - l^2 - m^2} - 1)) \quad (2.20)$$

The direction dependence of \mathbf{J} is far more tricky, and represents one of the more recent challenges in radio interferometric imaging. Following the notation of Smirnov (2011), the \mathbf{J} term is decomposed into a direction dependent term, $\mathbf{E}_p(l, m)$, and a direction independent term, \mathbf{G}_p , for each antenna such that:

$$\mathbf{J}_p(l, m) = \mathbf{G}_p \mathbf{E}_p(l, m) \quad (2.21)$$

Historically, the direction dependent effects indicated by E have largely been ignored, or quietly (and sometimes unknowingly) approximated by other correction factors. This is largely caused by two factors: the computational complexity of dealing with it, and a lack of need to perform measurements at significant distances from the phase centre. It is only in the previous couple of decades or so that instruments capable of observing significant portions of the sky have emerged (instruments such as the LOw Frequency ARray (LOFAR)), and wide-field observations have become more common, making the E correction more important. Furthermore, the explosion in processing power within the 21st century has allowed it to become a practical option. For a detailed treatment of direction dependent effects and ways of dealing with them, see the works of Schwab (1984); Bhatnagar et al. (2008); Carozzi & Woan (2009); Smirnov (2011); Smirnov & Tasse (2015) and the references therein. In this work I will proceed in the traditional manner, and assume that the contribution of E is negligible over the field of view and ignore it (apart from the w term correction described in section 2.8, and the primary beam correction in section 2.9).

To finalise the transition from visibilities to an image, we make the assumption that visibilities from each point on the sky add up linearly, which means that we simply integrate B over all points l, m :

$$V_{pq} = G_p \left(\iint_{lm} K_p B K_q^H \frac{dl dm}{\sqrt{1-l^2-m^2}} \right) G_q^H \quad (2.22)$$

The scalar delay matrices may commute, so expanding the term $K_p K_q^H$ we have:

$$\begin{aligned} K_p K_q^H &= e^{-i\phi_p} e^{i\phi_q} \\ &= e^{-i(\phi_p - \phi_q)} \\ &= e^{-i2\pi(\mathbf{u}_q \cdot \boldsymbol{\sigma} - \mathbf{u}_p \cdot \boldsymbol{\sigma})} \\ &= e^{-i2\pi(u_{pq}l + v_{pq}m + w_{pq}(n-1))} \end{aligned} \quad (2.23)$$

and therefore:

$$V_{pq} = G_p \left(\iint_{lm} B e^{-i2\pi(u_{pq}l + v_{pq}m + w_{pq}(n-1))} \frac{dl dm}{\sqrt{1-l^2-m^2}} \right) G_q^H \quad (2.24)$$

The above equation is the first glimpse of how one might construct an image from the

visibilities \mathbf{V} , since it is actually a 3D Fourier Transform between the visibilities $V(u, v, w)$ and the sky map $B(l, m, n)$. However, a three-dimensional Fourier transform is incredibly expensive to compute, and the calculation is most often reduced to a two-dimensional transform by assuming *coplanarity* of the antennas. In other words, it is assumed that the term w_{pq} is essentially zero, yielding the much simpler calculation of a two-dimensional Fourier transform:

$$\mathbf{V}_{pq} = \mathbf{G}_p \left(\iint_{lm} \mathbf{B} e^{-i2\pi(u_{pq}l + v_{pq}m)} \frac{dl dm}{\sqrt{1 - l^2 - m^2}} \right) \mathbf{G}_q^H \quad (2.25)$$

with the explicit Fourier relationship:

$$V(u, v) = \mathcal{F}[B(l, m)], \quad B(l, m) = \mathcal{F}[V(u, v)] \quad (2.26)$$

This relationship is also known as the van Cittert-Zernike theorem ([van Cittert, 1934](#); [Zernike, 1938](#); [Thompson et al., 2017](#)), and is the cornerstone of producing a sky image with a radio interferometer. Naturally, directly calculating the Fourier transform would also be prohibitively time-consuming because of the $O(N^2)$ computational complexity, so [Equation 2.25](#) is most often performed using a Fast Fourier Transform (FFT) such as the Cooley-Tukey FFT ([Cooley & Tukey, 1965](#)). This requires a prior step known as ‘gridding’ to map the continuous visibilities onto a discrete two-dimensional image, which can then be transformed (see [Thompson et al. \(2017\)](#) for a full explanation of gridding). From the properties of the Fourier transform, and the knowledge that the resulting image of the sky must be real-valued, we can infer that each visibility actually contributes two points in the uv plane. This therefore means that the visibility \mathbf{V} is Hermitian symmetric:

$$\mathbf{V}_{pq}(u, v) = \mathbf{V}_{pq}^H(-u, -v) \quad (2.27)$$

This relation stems from the quantities u_{pq} , v_{pq} and w_{pq} in [Equation 2.23](#). There is an inherent ambiguity in the baseline separation vector \mathbf{u}_{pq} since it is exactly equivalent to $-\mathbf{u}_{qp}$, and the choice between the two is entirely arbitrary. As a result, the choice of sign in the exponent of the delay term in [Equation 2.25](#) is also entirely arbitrary, but convention assumes the negative. In any case, no further useful information is added to the uv plane because of [Equation 2.27](#).

The following sections will detail the practicalities of calculating the direction inde-

pendent correction factor G over time and frequency, as it often encompasses more than just one physical effect.

2.5 Calibration

Like most problems in radio interferometry, there is no general, universal solution to visibility calibration. This largely stems from a mixing of complexity and variable needs; there exists some perfect correction that exactly produces the ‘true’ visibilities for all spatial frequency coordinates u, v, w , frequency ν , and time t , but calculating a perfect correction is entirely unnecessary if you are only interested in the inner 10’ of the field-of-view, for example. In addition, the depth and complexity of the technique introduces significant room for ‘unknown unknowns’ which cannot, by definition, be compensated for. As a result, a ‘good’ radio interferometer calibration is more specifically a correction that is accurate enough for a specified set of research goals. That being said, there is a common set of calibrations and corrections that the vast majority of research goals require. This section will briefly cover these.

2.5.1 Flagging

The topic of radio interference and flagging is more deeply explored in [chapter 3](#), but it is introduced here to highlight its significance in the overall calibration process.

Calibrating the complex visibilities is mostly a mix of extrapolation and interpolation. There is a good understanding over which parameters certain quantities vary, and to what magnitude. Therefore, we can use this knowledge to make good, yet simple corrections to the data. The immediate problem is the presence of outliers; interpolating or extrapolating in the presence of outliers does not give good results.

As a radio interferometer is observing it will likely pick up radio emission that is entirely unrelated to the astronomical source being studied. Modern society uses radio waves for a wide range of applications, and to some extent our society could not exist as it is without using radio waves. Terrestrial radio signals are more likely to be very strong compared to the trickle of emission that reaches us from astronomical sources, so radio interferometers — which are designed to observe astronomical sources — can be blinded by these strong radio signals. Signals not relevant to the astronomical sources being studied are broadly labelled as Radio Frequency Interference (RFI).

The impact of RFI is very unforgiving, for without pre-correlation intervention, it overpowers the signal from the astronomical source for its duration, and the observing time is effectively lost. After correlation, the only thing left to do is to identify RFI signals in the visibilities, and mark them as such. This process is called flagging. Flagged visibilities are not included when extrapolating or interpolating calibration solutions as they would be a significant outlier and would skew the calculated solutions, so it is vital that as much RFI is flagged, and is done so with the best possible degree of accuracy. This ensures minimal data loss, and maximises the quality of the overall calibration to follow. Poor calibrations due to poorly flagged visibilities ultimately manifest as reduced image fidelity; the overall quality of the final image of the sky. Images with poor fidelity can contain artefacts that could be misinterpreted as real sources, especially by automated algorithms. Furthermore, the flux scaling of sources may be incorrect, leading to erroneous measurements and, therefore, erroneous conclusions.

2.5.2 Phase Calibration

As the incoming radio signal propagates through the atmosphere, the content and optical depth of the troposphere in between the receiver and the source is constantly changing over the duration of the observation (see [Figure 2.4](#)). At 1.5 GHz this primarily impacts the visibility phase, which must be known to a high level of accuracy to prevent any destructive interaction between Fourier components in the final image. Observations of the Lockman Hole and the candidate binary black hole documented in this thesis used the phase-referencing technique to correct for this effect.

Creating an optimal phase-referencing strategy is highly dependent on the characteristics of a particular interferometer. It requires the consideration of many factors such as atmospheric conditions, baseline length, antenna slewing speed, and so on ([Thompson et al., 2017](#)). However, the primary consideration is to select any unresolved source that is strong enough to produce clear phase measurements well above the system noise. Additionally this source should be located close enough to the target to ensure that the area of sky observed can be assumed to be the same within a single target-reference cycle. This allows phase solutions derived from the reference source to be extrapolated to the target signal. An unresolved source allows the calibration to ignore instrumental effects (such as variable baseline length) and purely track the atmospheric fluctuations, as a single unresolved source at the phase centre is easily modelled.

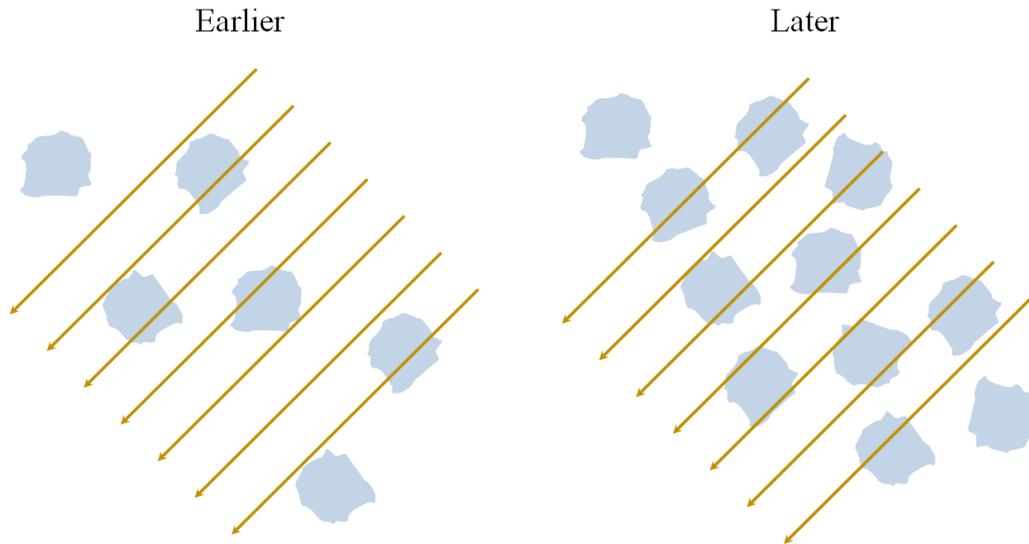


Figure 2.4: An illustration of how atmospheric pressure changes produce a time variable optical depth, which drives short timescale phase fluctuations in the received radio signal. This effect is corrected for by observing a phase calibrator source via ‘phase referencing’.

2.5.3 Amplitude Calibration

The signal that emerges from the correlator has lost physical meaning for a number of reasons. One reason is that, despite radio waves being disrupted much less by the atmosphere, there is still some attenuation that reduces the strength of the incoming signal. Another is that pointing errors in the antennas cause uncertainty in the output power. If an antenna is, in reality, pointing slightly off source, but the system reads as on source there is an error introduced to the output signal due to the directional gain of the antenna itself. Finally, and most importantly, the conversion of an incoming radio signal to an output voltage signal (as in Equation 2.4) is not an equal conversion; amplifiers and other electrical components introduce electronic gain. The net effect of these particular processes create a proportionality relationship between the incoming signal and the visibility amplitude, one that can be corrected by observing a well known and well modelled flux calibrator source, and deriving calibration solutions from the difference between the known flux from the calibrator and the measured visibility amplitude. The chosen source should ideally be a bright point source with minimal variability to minimise time dependent errors in the model.

2.5.4 Bandpass Calibration

One of the developments of modern interferometer arrays is a massively increased observing bandwidth. Fractional bandwidths (ν_{width}/ν_{obs}) between 50-100% are not uncommon,

enabling shorter observations with a higher sensitivity. In instrumentation, this is achieved by dividing up the bandwidth into spectral windows, which are then further subdivided into frequency channels. Because of instrumental effects, each antenna, and the interferometer as a whole will have a frequency dependent response that must be calibrated for. [Figure 2.5](#) shows an example bandwidth response for an e-MERLIN antenna. Since the bandpass response has minimal temporal variance, a short integration on a ‘bandpass calibrator’ is enough to generate calibration solutions for the whole observation. Thus, the bandpass calibrator is often also the flux calibrator referred to in [subsection 2.5.3](#), though in principle any bright source with a featureless spectrum can be used.

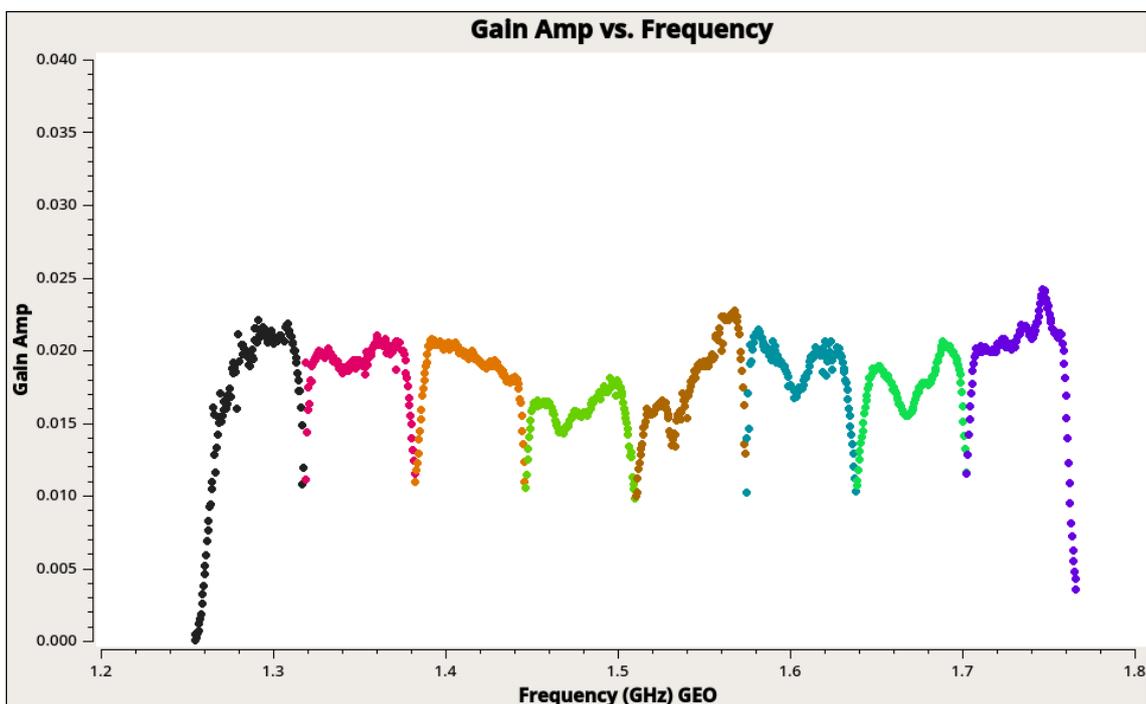


Figure 2.5: An example bandpass response calibration taken from a single observation of the Lockman Hole used in this thesis. The data indicates the time-constant corrections applied to the 1024 channels spread across eight spectral windows at L-band (21 cm) for e-MERLIN.

2.5.5 Self-calibration

Inevitably, residual errors will exist in the visibility phases and amplitudes. When using phase-referencing, the periodic phase measurement will only give an approximation of the phase during integration on the target, and is biased in its correction towards large-scale phase fluctuations. Small-scale phase errors will persist on scales of a minute or less. This residual error can be corrected through self-calibration.

Recall that a baseline produces the incoming signal \mathbf{B} scaled by some pair of gain

factors that is unique to each antenna:

$$\mathbf{V}_{pq} = \mathbf{G}_p \mathbf{B} \mathbf{G}_q^H \quad (2.28)$$

At the point where self-calibration is an available option, the visibilities will normally have already gone through some other calibration such that an approximate model of the sky $\mathbf{B}_M(\boldsymbol{\sigma})$ can be made. Alternatively, the patch of sky being observed may have already been observed previously. This sky model has corresponding visibilities $\mathbf{V}_{pq,M}$ on the baseline pq , calculated from the Fourier transform relationship. By leveraging the sky model, one can generate further complex gain corrections that minimise the gap between the model visibilities and the observed visibilities, $\Delta \mathbf{V}_{pq}$. This is done through a least squares approach:

$$\Delta \mathbf{V}_{pq} = |\mathbf{V}_{pq} - \mathbf{V}_{pq,M}|^2 = |\mathbf{V}_{pq} - \mathbf{G}_p \mathbf{B}_M \mathbf{G}_q^H|^2 \quad (2.29)$$

With perfect knowledge of every complex gain factor \mathbf{G} , and a perfect sky model such that $\mathbf{V}_{pq,M} = \mathbf{V}_{pq}$, the difference $\Delta \mathbf{V}_{pq}$ will be zero in all cases. By iteratively updating $\mathbf{V}_{pq,M}$ with iteratively improved complex gains and an improved sky model, residual errors in the visibilities can be significantly reduced. The process of generating \mathbf{B}_M is described in [section 2.6](#) and [section 2.7](#).

2.6 Constructing an Image

An observation by a radio interferometer is simply a process of sampling many visibilities over an extended period. Each collected visibility samples a certain uvw coordinate at a specific time t , and frequency ν . Furthermore, because of the rotation of the Earth, the interferometer naturally samples different points in the uv plane as time passes; this is the fundamental mechanism of aperture synthesis. In the l, m plane, known as the *image* plane, when the visibility is Fourier transformed, a single visibility sample corresponds to a single Fourier component. Adding further visibility samples at different points in the uv plane adds further Fourier components to the image, and progressively improves its quality. This is best understood by visualising the changes in the image that occur as more visibilities are added. [Figure 2.6](#) shows the result of transforming a single visibility (with a further mirrored visibility from the Hermitian relation in [Equation 2.27](#)) on a shorter baseline, and a longer baseline some time later. The illustration shows how sampling different points in the uv plane produces different outcomes in the image plane. Longer

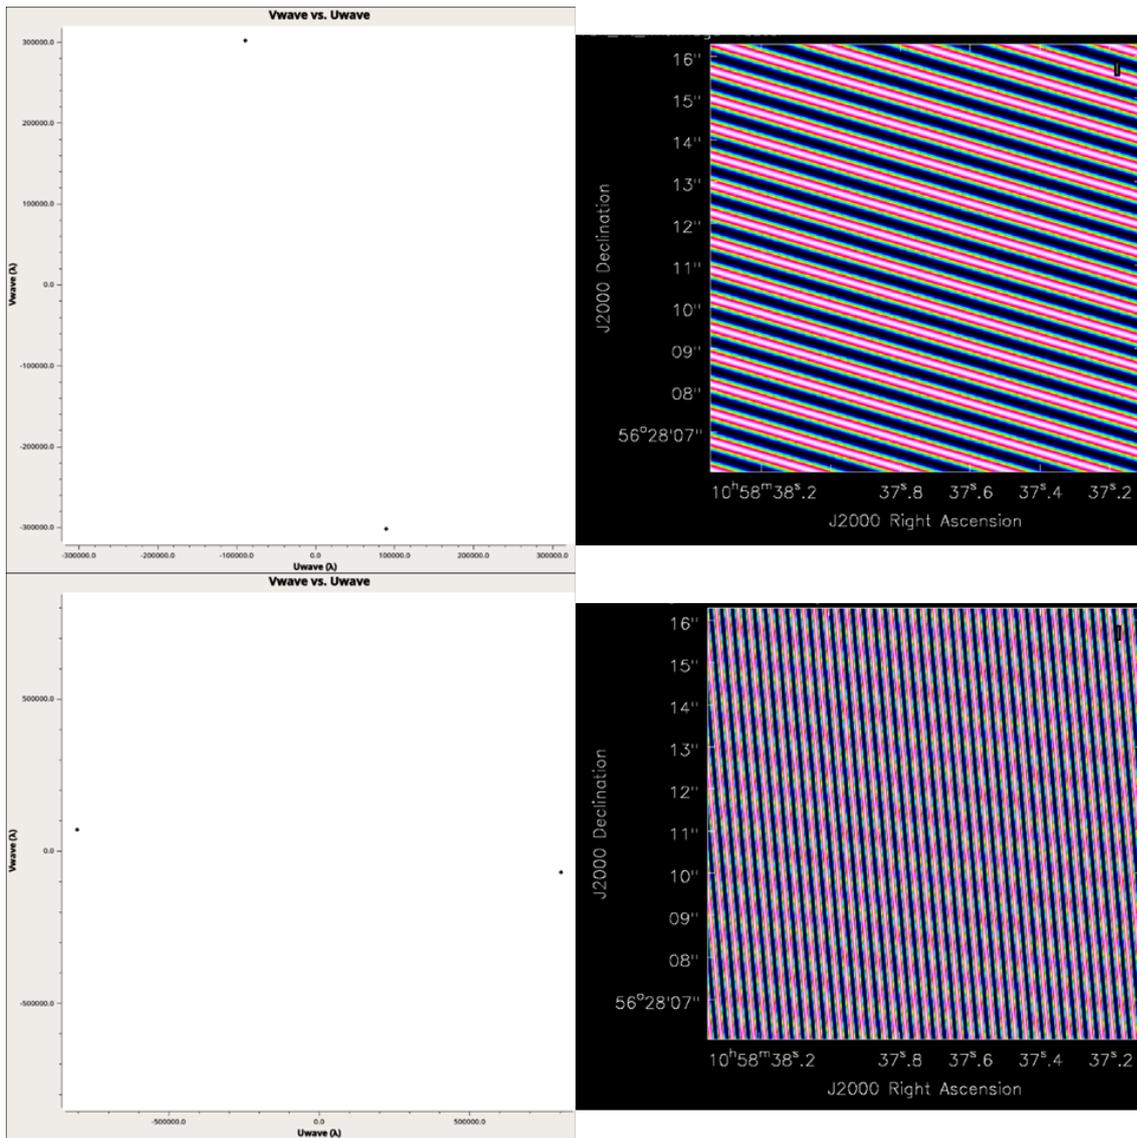


Figure 2.6: Images showing the different scales sampled by a shorter baseline (top) and a longer baseline (bottom). Each image shows a single visibility taken at a different time, and includes the mirrored visibility indicated by Equation 2.27. Since only a single visibility is used to construct the output image (right), it contains only a single Fourier component.

baselines produce components with shorter spacings in the image plane, and therefore sample smaller scale structures. As an interferometer continues to observe a source, the antennas appear to move their position in the uv plane because of the rotation of the Earth, so simply by the passage of time an interferometer can sample many different parts of the uv plane. This is further illustrated in Figure 2.7, which shows the effect of continued integration on a source. As more visibilities are added to the uv plane the image of the target becomes progressively better. Figure 2.7 also shows the significant improvement in image quality from observing at multiple frequencies simultaneously (recall that the spatial frequency coordinates u, v, w are frequency dependent). Each additional frequency observed acts as a multiplier on the total Fourier components added to the final image

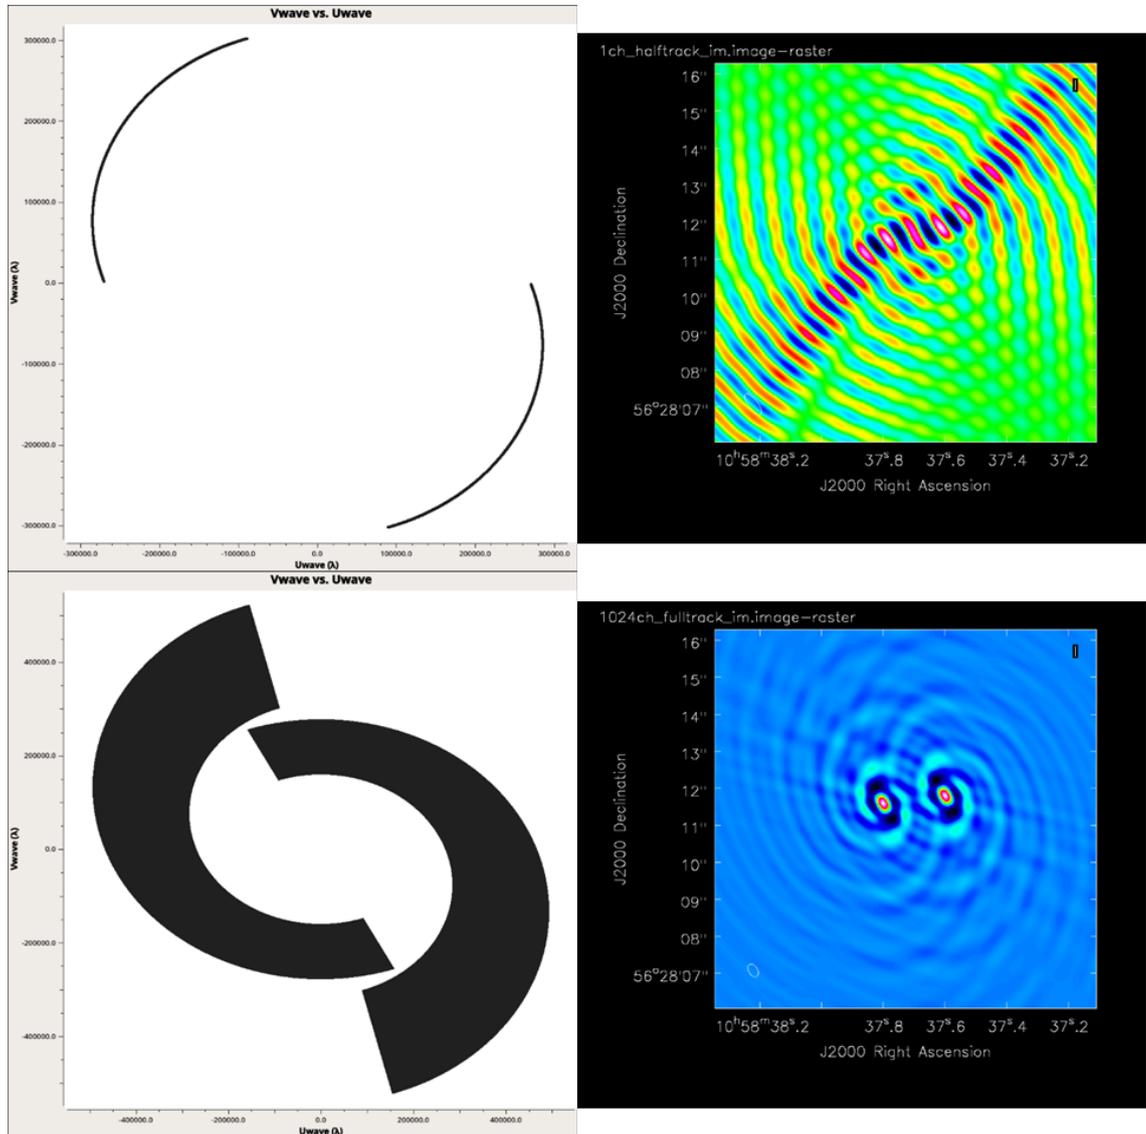


Figure 2.7: Two images showing the dramatic increase in image quality from an increased bandwidth and increased exposure time. The top image shows a six-hour observation observing a single frequency. The bottom image shows a twelve-hour observation of 1024 separate frequency channels. Since each time and frequency step fills the uv plane at a new location, further structural information is added to the final image (right) as unique Fourier components are progressively added to the image.

at each timestep. Further consider that the image shown in [Figure 2.7](#) uses only two antennas forming a single baseline, but modern interferometer arrays contain many more antennas. The VLA for example contains 27 antennas, giving well over 300 baselines. The number of baselines is a further multiplier for the number of Fourier components added to the final image at each timestep! Because of the wide bandwidths observed by potentially hundreds of baselines over many hours, measuring millions of visibilities in a single observation is readily achievable. The combination of wide bandwidths, and an ever increasing number of radio antennas in a given array, is the heart of why modern radio observations can produce such high quality images of the radio sky.

The Fourier transform of the filled uv plane produces an image of the sky with one remaining error: the point-source response. This raw sky map, known as the dirty image, is effectively the true sky map convolved with the Point Spread Function (PSF) of the associated interferometer, which causes image artefacts due to the sidelobes of the PSF. In a filled-aperture telescope the point-source response generally consists of a single relatively simple component. In contrast, the PSF of an interferometer, which does not use a filled aperture, can have many different components of variable complexity. In certain cases, it can even produce the illusion of sources that do not actually exist. Mitigating the effect of the PSF requires a uv plane with as few gaps as possible, i.e. an interferometer with as close to complete aperture coverage as possible. When imaging, the effect of the PSF is removed by the deconvolution algorithm CLEAN and its successors.

2.7 The CLEANing Process

The CLEAN algorithm (Högbom, 1974) provided a significant jump in radio image deconvolution. The central idea is to consider all emission in the dirty image as the result of a superposition of many point sources, arranged according to the sky brightness distribution $\mathbf{B}(\sigma)$. Areas of the image with little or no emission consequently have little or no sidelobe artefacts (from *local* emission specifically), and naturally areas with strong emission have strong sidelobe artefacts. The uncorrupted synthesised beam with the PSF sidelobes removed, known as the clean beam, is easily obtained by fitting the central peak of the PSF. It is possible, therefore, to remove the effect of the dirty beam (the PSF) on the image by finding and marking as many of these point sources as possible, since they collectively form the entire image. This is done through an iterative, top-down ap-

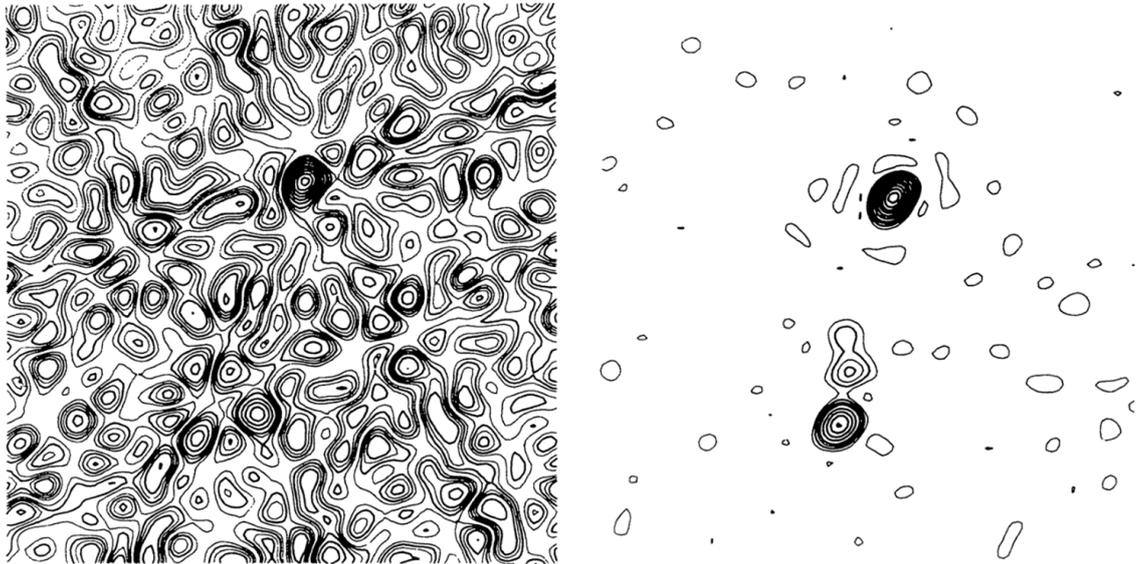


Figure 2.8: A radio image before (left) and after (right) processing by the CLEAN deconvolution algorithm. The structure that is removed from the left image is the result of the interferometer PSF. Image credit: Högbom (1974)

proach: the strongest point in the image is multiplied by a Kronecker delta function and then convolved with the dirty beam. The result is subtracted from the dirty image and the process is repeated until some criteria, usually defined by the user, is met. After the criteria are met, all of the delta functions found during the iterations are multiplied by the clean beam, and added back to the residuals of the dirty image. This produces an image with the ‘real’ emission preserved, and the sidelobe artefacts removed. Great care must be taken to ensure that delta functions are placed at positions of real emission, otherwise CLEAN can introduce sources of emission that do not actually exist. A common problem with images generated from sparsely filled uv planes is making sure that CLEAN does not

see the sidelobes originating from another strong, nearby, source as genuine independent sources. Using the example of [Figure 2.8](#), it is possible that CLEAN may identify one of the sidelobe artefacts as a point of real emission, and mistakenly introduce a fake source to the cleaned image. It is also possible that sidelobes from multiple discrete sources may overlap and add together, creating the illusion of a brighter source. Advanced cleaning algorithms contain as many safeguards as possible to prevent this from happening, but it remains an inherent problem of the technique.

There are some caveats that apply to an image produced by CLEAN, most notably that it is a best-case *representation* of the sky. Fundamentally, it is not a direct image because we have replaced the point source response of the interferometer with a model of the clean beam. Secondly, though the Högbom CLEAN works for extended sources, its effectiveness begins to diminish relative to its performance as the complexity of the source increases ([Cornwell, 2008](#)). Numerous other variations of the CLEAN algorithm have been designed and implemented though most, if not all, share the same fundamental iterative approach of identifying real emission, and using that information to separate out the contribution of the PSF sidelobes.

2.8 Wide-field Imaging

In [section 2.4](#) we assumed that the w_{pq} term in [Equation 2.24](#) simplifies to $w_{pq} = 0$. However there are several scenarios where this is no longer a valid simplification. In general, the error introduced to the final image by assuming $w_{pq} = 0$ scales with distance from the phase centre, or the direction s_0 . The magnitude of the error is clearly proportional to how much the true value of w_{pq} exceeds $w_{pq} = 0$. Therefore, there are two ways of avoiding the error introduced by the simplification. Most often radio astronomers simply accept that there is a limit on the field-of-view of their observation; as the magnitude of the w_{pq} error scales with distance from the phase centre $|\sigma(l, m)|$ they limit their measurements to within some calculated uncertainty range. Many observations are targeted directly at the source of interest so in these cases the astronomer does not need large fields-of-view. Another more limiting approach is to constrain the length of the baselines in the antenna array, since this will limit the fundamental value of w_{pq} regardless of pointing. Other designs constrain the placement of the antennas along a straight line, normally East-West, and align the w axis with the rotation axis of the Earth (see [section 3.1.2](#) of [Thompson et al. \(2017\)](#) for a detailed explanation of this technique). This is the approach taken by

the Westerbork Synthesis Radio Telescope (WSRT, Baars & Hooghoudt 1974; Hogbom & Brouw 1974).

For many applications including the work conducted in this thesis, wide fields-of-view are desired, so the error introduced by simplifying w_{pq} must be accounted for. The most common approach is to utilise the so-called ‘w-stacking’ technique, which calculates multiple images using a subset of values for w_{pq} . By taking the inverse Fourier relationship of Equation 2.24, this time for the full antenna array, and assuming perfect gain we can write³:

$$\mathbf{B}(l, m, n) = n \iiint \mathbf{V}(u, v, w) e^{i2\pi(ul+vm+w(n-1))} dudvdw \quad (2.30)$$

To simplify the three-dimensional integral we instead sum multiple two-dimensional integrals together (in this case images), and apply a further w -correction to each image over a discrete range of values for w :

$$\mathbf{B}(l, m) = n \sum_{w_{\min}}^{w_{\max}} e^{i2\pi w(n-1)} \iint \mathbf{V}(u, v, w) e^{i2\pi(ul+vm)} dudv \quad (2.31)$$

This is the method implemented by the WSCLEAN imager (Offringa et al., 2014) for wide-field imaging, which is used throughout this thesis. Less formally, this process can be written as the construction of an image of the sky from a series of fast Fourier transforms of the discretized visibilities:

$$\mathbf{B}(l, m) = \sum_{w_{\min}}^{w_{\max}} e^{i2\pi w(n-1)} \times \text{FFT}[V_{\Pi}(u, v, w)] \quad (2.32)$$

where the values of w over which the summation (and the FFTs) is calculated is set by the user, though WSCLEAN itself calculates an optimal set of w values by default.

2.9 Primary Beam Correction

Another direction dependent effect that is universal amongst interferometers (or all beam-forming radio antennas for that matter) is the direction dependent power response of an antenna known as the *primary beam*. The primary beam describes the reduction in power sensitivity of the antenna to signals coming from directions other than the pointing direction. Analytically deriving the frequency dependent primary beam response, $\mathbf{P}_p(l, m)$, of a

³Although the 3D Fourier transform technically produces a sky *sphere*, the sky signal of interest is recovered from an infinitely thin spherical shell with a radius equal to 1 in the l, m, n domain.

given antenna is extremely difficult, and requires the consideration of many factors including, but not limited to, antenna design, shadowing from the receiver support structure, and dish size. Thus, \mathbf{P} is often calculated numerically from observations of known sources, or by mapping the antenna surface via holography (e.g. [Hunter et al. 2011](#)). In practice, the primary beam is often generated and applied to the final image in the l, m plane, with the Full Width at Half Maximum (FWHM) acting as the limit of uncertainty for the image. The primary beam response calculated for the Lockman Hole field presented in [chapter 6](#) is shown in [Figure 2.9](#).

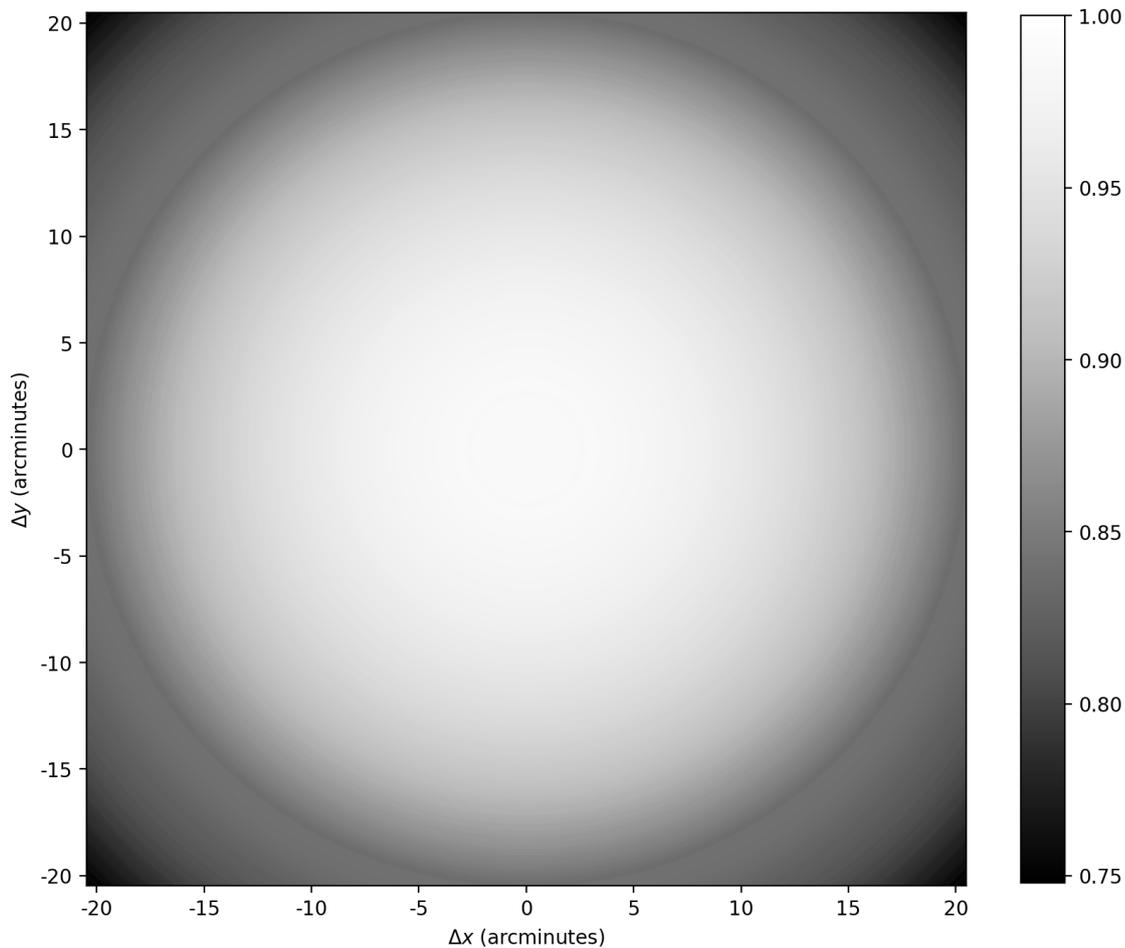


Figure 2.9: An approximation of the e-MERLIN primary beam correction applied to the Lockman Hole field.

Chapter 3

Identifying RFI with Machine Learning

This chapter will introduce the concepts behind the field of machine learning, and explain the specific model architecture used in this experiment. This will lead directly into the discussion of [section 3.3](#) and onwards, which outlines how machine learning was used for identifying radio interference in astronomical observations.

3.1 Introduction to Machine Learning

To the majority of non-experts the field of machine learning can seem unreasonably opaque, with words that seem to have different meanings depending on the author, the time of day, and perhaps even the current lunar phase. This is made worse by the plethora of machine learning models that all seem to have their own set of acronyms¹, nouns, and verbs to describe its function. This chapter will attempt build an understanding of the wider field of machine learning from the ground up, with a focus on the methods used in this piece of research and without using esoteric language. It will then outline the research conducted using the methods described, and analyse their success.

First, consider the cliché question: what exactly *is* machine learning? Broadly, the term describes a methodology where some kind of pattern exists - normally represented in a computer dataset - and one seeks to create a model of this pattern. The model they seek to create of this pattern is contained within a machine learning model, that is then used to identify and/or predict a similar pattern in some unknown dataset. Before the model can be used for prediction, it must ‘learn’ the pattern from the original dataset. The vast majority of machine learning models ‘learn’ through an iterative approach, that can

¹Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Gated Recurrent Units (GRUs), Long Short-Term Memory (LSTM) Networks; just a handful of offences

be most simply described as a series of educated guesses. At each iteration a ‘loss’ is calculated to represent how good or bad the current iteration was at guessing the pattern; a good guess has a low loss. Each iteration has access to the performance of every iteration that came before it, and one can use this knowledge to steadily converge towards an optimal model of the desired pattern. The strength of this methodology is that it is independent of any analytical constraint and entirely based on empirical evidence; any pattern that exists in reality can be modelled given a sufficient level of starting knowledge. This strength must be tempered by caution however, as no new information is actually created and the model is only as good as the data it ‘learns’ from. There are many scenarios that must be considered, for example, what if the data provided to the model was incorrect or misleading in some way? What if multiple patterns exist in the same dataset? What if the model learns a pattern in the dataset, but it is not the pattern you want to extract? What if, between the time of training and prediction, the pattern has changed? Answering questions like this and accounting for the numerous possibilities is one of the fundamental problems of machine learning. The model itself stores the pattern in a collection of operations that work together to map the input data to the desired output. A simplified version of this process is described in [section 3.2](#). The most advanced machine learning models that exist today (such as GPT-4, see [OpenAI et al. 2023](#)) still use this fundamental concept, albeit to a significantly larger and more complex degree.

Given that machine learning models are fundamentally empirical models, it’s no surprise that they have been extensively used for observational astronomy, and any task that seeks to extract information from large blocks of data such as what will be produced by the SKA and the Rubin Observatory ([Ivezić et al., 2019](#)). Applications of machine learning in the wider field of astronomy include, but are not limited to:

- Radio transient detection (e.g. [Bethapudi & Desai 2018](#))
- Solar Flare prediction (e.g. [Qahwaji & Colak 2007](#))
- Identification of near-Earth objects (e.g. [Lieu et al. 2019](#))
- Galaxy classification (e.g. [Aniyan & Thorat 2017](#))
- Exoplanet identification (e.g. [Shallue & Vanderburg 2018](#))
- Search for Extraterrestrial Intelligence (SETI) (e.g. [Ma et al. 2023](#))

3.2 Fundamental Concepts

The question of “How can a machine *learn*?” is best understood by considering a simple classification problem. Classification is something humans do instinctively, we see the picture of a dog in [Figure 3.1](#) and immediately know that we are looking at a dog. But we only know this because we have seen dogs for our entire lives; if an alien were to visit us from another planet it’s likely they would have no clue what this mysterious creature is. We are also able to store past experiences in our brains and *remember* that something of this shape, size, colour, smell (and so on) is a dog. In addition, humans have evolved to develop the biological mechanisms necessary to detect and measure all of those properties in the first place, and then cognitively condense all that information into a single concept: a dog.

From this point of view, learning can be understood more specifically as the ability to absorb, store, and later recall information. Computers already have the ability to store and recall information, and modern technology allows them store quite a significant amount of information. Therefore, the problem of learning or, to invoke the relevant name, *machine learning* is narrowed to a problem of information absorption. To go back to [Figure 3.1](#), the question becomes “How can the information in this image be condensed such that it represents the concept of a dog?”. This is a problem of classification.

An intuitive approach to classifying the image would be to use the information in the image to calculate a single number representing the ‘dogness’ of the image. A high ‘dogness’ score means the image likely contains a dog, a low score indicates the opposite. The simplest way to calculate such a score can be done by one of the simplest machine learning methods: a linear classifier.

A linear classifier implements this scoring system via a single, simple multiplication between two matrices \mathbf{x} , which represents the image [Figure 3.1](#), and \mathbf{W} , which is a weighting factor representing the translation from image data to a ‘dogness’ score. This operation can be written like so:

$$f(\mathbf{x}, \mathbf{W}) = \mathbf{W} \cdot \mathbf{x} \quad (3.1)$$

where the output f is a single number representing the ‘dogness’ of \mathbf{x} . The weight matrix \mathbf{W} applies a weighting to each pixel in the image in an attempt to yield the most accurate ‘dogness’ score for the input image. For example, pictures of dogs tend to have



Figure 3.1: A picture of a golden retriever from the Stanford Dogs benchmark dataset (Khosla et al., 2011).

them located right in the middle of the image (they are wonderful after all) so a better \mathbf{W} would likely downweight the contribution of edge pixels to the ‘dogness’ score, since dogs will rarely be in the edge of the image. Conversely, if we were instead trying to detect something like the Loch Ness Monster in an image, that would presumably be cryptically located towards the edge of most images, and the best weight matrix would upweight those pixels for the classification instead.

However, manually crafting an optimal weight matrix is laborious and not a good use of time. It is instead selected through an automatic, iterative, trial-and-error-like process that sits at the heart of machine learning. To implement this method, we first construct a way of quantifying exactly how bad the weight matrix \mathbf{W} has performed. Note, that during training of the linear classifier we have access to the ground truth of our data, so we know that Figure 3.1 contains a dog and we can use that information to improve \mathbf{W} . Assessing the performance of \mathbf{W} is done through a ‘loss function’. There are an infinite number of loss functions, but generally speaking the machine learning community uses a select few that have been proven to perform well and are simple and reliable. An example of one of these loss functions is known as the L1 loss.

To flesh out the example a bit more, imagine that the linear classifier in Equation 3.1 produces a ‘dogness’ score where zero indicates not a dog, and positive one indicates a dog. We have passed in to the classifier the image in Figure 3.1 which we’re going to label as x_1 . The output $f(x_1, \mathbf{W})$ is labelled as y_1 . We know that there is a dog in the image, so we also know that y_1 should at least be as close to the number 1 as possible. We can assess the performance of our selection of \mathbf{W} by scoring it based on how far away the output y_1

is from the ground truth t_1 . The L1 loss function implements this using Manhattan or Taxicab distance:

$$L_{L1}(y) = |t - y| \quad (3.2)$$

where t is the ground truth of the image x . Imagine that our classifier produced a poor score of -3.5 for the image in [Figure 3.1](#). Then the L1 loss would be:

$$L_{L1} = |t_1 - y_1| = |t_1 - f(x_1, \mathbf{W})| = |1 - (-3.5)| = 4.5$$

which is a direct quantisation of how badly the particular \mathbf{W} performed. Now imagine that the selection of \mathbf{W} was good, and the score was instead $y_1 = 1.04$. Now the loss would be:

$$L_{L1} = |1 - 1.04| = 0.04$$

The output of the loss function can therefore be used to inform how much the weight matrix needs to change to improve the next attempt; a larger loss means a worse prediction and greater change needs to be made. How that change is made to \mathbf{W} is a process known as optimisation, which today is mostly implemented through variants of stochastic gradient descent ([Robbins & Monro, 1951](#)). A detailed explanation of that method is well beyond the scope of this document, but see [Kingma & Ba \(2014\)](#) for a detailed explanation of the popular *Adam* optimiser. Broadly speaking, the process of optimisation involves the calculation of, either numerically or analytically, the value of $\frac{dL}{d\mathbf{W}}$ multiplied by some scale factor (also known as the learning rate), and using that to update the values within \mathbf{W} .

Finally, to make the linear classifier a tool of prediction, we iterate over many images $x = x_1 \dots x_N$, and for each image go through a continuous loop of scoring, loss calculation, and weight matrix optimisation. Through these iterations, the weight matrix \mathbf{W} will begin to ‘learn’ and store the characteristics of what a dog looks like, and be able to use that information to give a ‘dogness’ score to an image it hasn’t seen before through a simple multiplication of the image data. The method described above now fulfils the concept of learning described at the beginning of this section: the linear classifier has the ability to absorb, store, and later recall information.

However, the method of training and the linear classifier described above have several limitations, a few of which I will describe. Firstly, the linear classifier would be restricted to classifications of dogs; it would be unable to classify anything other than dogs (it would likely struggle to even classify different breeds of dogs, since they vary in appearance and

size). This limits its scope of application and fundamentally makes the model less useful.

Secondly, the method of using a simple multiplication by a single weight matrix \mathbf{W} means the information stored in \mathbf{W} is whatever set of values produces the smallest total loss over all inputs x_n . Aggregating over such a large amount of data in such a simplistic way dampens the accuracy of the classifier because there is a great deal of unnecessary or unwanted information in each image. For example, many of the images in the Stanford Dogs dataset contain backgrounds of grass or sand, but this background is entirely irrelevant to the desired classification: whether a dog is on grass or sand, it is still a dog. Overfitting to this redundant information only reduces the ability of the model to generalise with no added benefit.

Lastly, we have dedicated all our input images to training the network, so we don't have a way of fairly testing how it is performing at each stage of training. Normally, training datasets are split into three groups: training, validation, and test data. Training data is the largest partition, and is used to actually train the model as described above. Validation data is used to further inform the direction the training should go, specifically to tune the model's hyperparameters. It effectively allows an *ad hoc* measure of performance as the model is being trained that is indicative of model quality instead of how well it is learning from the data (which is indicated instead by the loss). Test data is typically used as the final performance test, since the model will never actually encounter it at all. The overall objective of the training process is to produce a model with the best performance when operating on the test data.

For tasks such as image classification, significant performance gains can be realised by increasing the depth and complexity of the function f . Indeed, much of machine learning research is based loosely around this goal. By introducing many more parameters to f and combining them in a variety of configurations, accuracies upwards of 90% can be achieved for the task of classifying all dog breeds in the Stanford Dogs dataset ([Do et al., 2022](#)).

3.2.1 The Convolutional Neural Network

One of the key advancements in machine learning was the discovery and quick proliferation of the Convolutional Neural Network (CNN). Inspired by processes occurring in the visual cortex, the 'Neocognitron' proposed by [Fukushima \(1980\)](#) was the first to use a series of convolutions arranged as a neural network to detect patterns in images inde-

Figure 3.2. The focus of this iteration is on the pixel X_{11} , so the kernel is placed over this pixel and the multiplication and summation is performed over all the overlapping pixels including and surrounding it. In full, the output pixel in the convolution image, C_{11} , would be:

$$C_{11} = W_5X_{11} + W_6X_{12} + W_8X_{21} + W_9X_{22}$$

Normally input images have pixels inserted around them so the convolution can occur even at the edges (known as padding), but these surrounding pixels are often set to zero, hence the omission from the above summation. The second iteration on the right of Figure 3.2 now focuses on the next pixel over, X_{12} . Using the same procedure, the output pixel C_{12} is:

$$C_{12} = W_4X_{11} + W_5X_{12} + W_6X_{13} + W_7X_{21} + W_8X_{22} + W_9X_{23}$$

This procedure is continued over every pixel in the input image until a full output image has been constructed; a so called ‘activation map’.

A convolution layer is commonly followed by a downsampling operation known as max pooling. This operation reduces the resolution of the image by taking a contiguous group of pixels in the input image, which would be the activation map from the associated convolution layer, and producing one output pixel using the maximum value of the group of input pixels. Since it is operating on an activation map, and it is selectively preserving pixels with a larger value, features detected by the previous convolution layer are further emphasised in every convolution deeper in the network. An example of the entire process is given in Figure 3.3.

Through training, the weights of the convolution kernels can be automatically tuned to minimise the network loss function, instead of being manually created. Therefore, while it is more common to see simpler features at shallow layers, ultimately the network is left to decide which features are important at what point and which are not, based on loss evaluations during training. This ability to learn and store complex features is the essence of what makes CNNs so powerful. Architectures using CNNs can be very sophisticated, which allows them to learn increasingly complex features. A prime example of this is the UNet (Ronneberger et al., 2015), which was developed to address the problem of ‘image segmentation’: classifying pixels in an image on an individual basis rather than classifying the image as a whole. An example application of the UNet might be to separate which

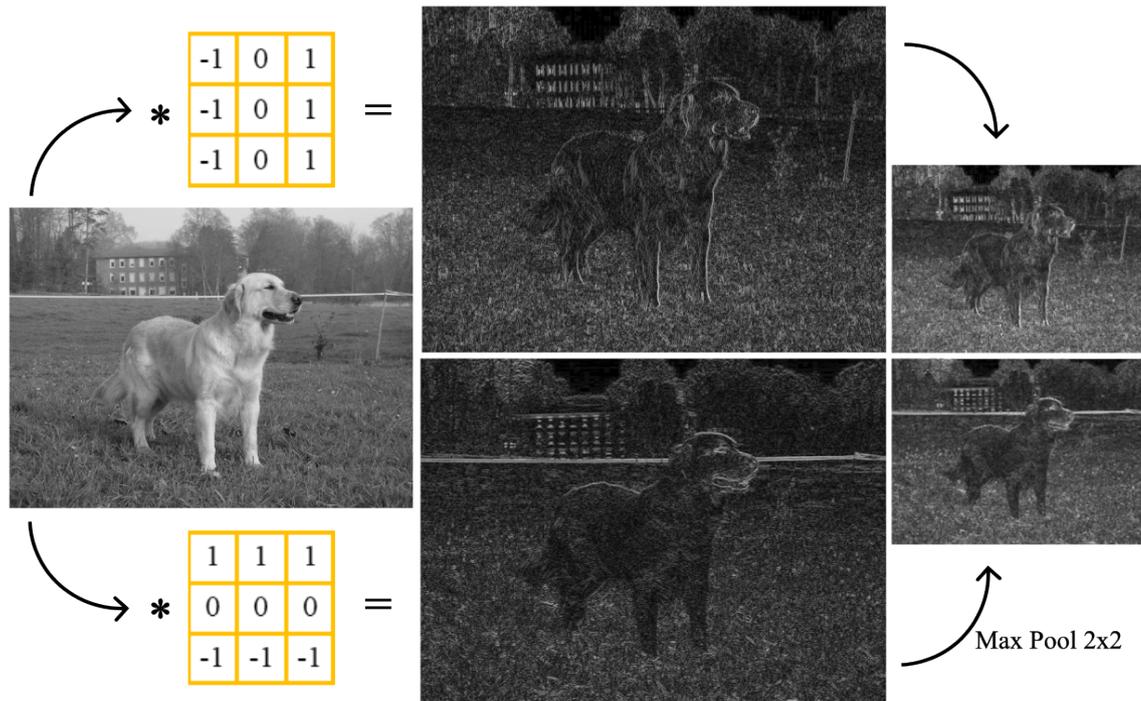


Figure 3.3: Two examples of the convolution operation followed by a max pool operation. The convolution kernels are manually crafted to detect vertical features (top), and horizontal features (bottom).

pixels in [Figure 3.1](#) belong to the dog, and which do not. Another example of a modern and powerful application of CNNs is the Generative Adversarial Network (GAN).

3.2.2 Generative Adversarial Networks

The GAN architecture is different from most other, more traditional, architectures because it borrows and utilises one key concept from the broad field of game theory: the ‘Nash equilibrium’. This term describes a specific state within a zero-sum game where each player can choose a strategy to win the game. Among all the strategies the player can choose, there exists one strategy that provides the best chance of winning against the strategies of all other players. The same can be said for all players, and so there exists a state where each player has selected their associated optimal strategy, and the game has reached a type of equilibrium. In practice, the equilibrium will likely be reached after many iterations of the game, where each player improves and perfects their strategy after each iteration until a change to their strategy would provide no additional chance of success, or even lessen their chance of success. The concept can be applied to a huge range of problems not just limited to mathematics, and can be useful to model any type of system where parties engaging in an interaction are pursuing mutually exclusive outcomes.

Goodfellow et al. (2014) introduced this concept to machine learning by modifying the conventional training process and replacing the loss function with another neural network that fulfils the same criteria. Their objective was to use this technique to generate realistic images of a certain object. For example, one could train a GAN to learn the characteristics of golden retrievers, and then generate entirely synthetic images of them. It does this by creating a type of game between the two neural networks, and tailoring two competing objectives such that one of the players becomes extremely proficient at generating realistic images.

Discriminator

The network that replaces the conventional loss function is known as the discriminator. The discriminative side of the model is responsible for classifying input images as either real or fake, with real meaning that it is a genuine image from the training set, and fake meaning that it is a synthetic image created by the opposing side of the GAN. Each of the two players in the GAN still have their own analytic loss functions that use the output of the other network, so that they can learn from each other. The objective of a discriminator, D , is to output probabilities as close as possible to the following:

$$D(x) \rightarrow 1 \text{ and } D(G(z)) \rightarrow 0 \quad (3.3)$$

where x is an image from the training set, and $G(z)$ is a synthetic image created by the opposing network (described in [section 3.2.2](#)). This details a discriminator loss function, L_D of the form:

$$L_D = (1 - D(x)) + D(G(z)) \quad (3.4)$$

since the objective of the discriminator is to minimise both components of the summation. In practice, the loss function is commonly evaluated using logarithms, but the premise is the same. CNNs are a good choice to fulfil this role since they are capable of learning the features of what makes an image ‘real’ and what makes an image ‘fake’.

Generator

The other side of the GAN is known as the generator, G , which is responsible for generating a synthetic image $G(z)$ from an input z . In direct opposition to [Equation 3.3](#), the

generator seeks a scenario where:

$$D(x) \rightarrow 1 \text{ and } D(G(z)) \rightarrow 1 \quad (3.5)$$

and it is able to trick the discriminator into classifying the synthetic images as real. However, because of the tension between that goal and the competing goal of the discriminator, the actual ideal situation for the overall system is closer to:

$$D(x) \rightarrow 0.5 \text{ and } D(G(z)) \rightarrow 0.5 \quad (3.6)$$

which effectively means that the discriminator has the best chance of success by simply guessing, which also indicates that the output $G(z)$ is almost indistinguishable from x . [Equation 3.5](#) describes a loss function for the generator of the form:

$$L_G = 1 - D(G(z)) \quad (3.7)$$

which again retains the underlying premise of the practical implementation.

Training

The exact training process for a GAN depends on the specific type of GAN being used (see [section 3.2.2](#)). However, the training process tends to follow a *similar* path regardless of type, so I will provide a brief explanation of the training process for the conventional Deep Convolutional Generative Adversarial Network (DCGAN) first presented by [Goodfellow et al. \(2014\)](#). Training loops often take place over many *epochs*, usually defined as one full loop over the entire training set. Each epoch consists of many *steps* where each step takes a finite number of samples (called a *batch*) from the training set and passes them to the network.

It is common practice to start the training step with the discriminator. The discriminator processes the batch of samples from the training set and produces an output score for each one. Then, the generator is tasked with creating an equal sized batch of fake samples derived from a type of random seed vector referred to as a ‘latent vector’. This batch is also evaluated by the discriminator. These two outputs from the discriminator are then used to evaluate the discriminator loss function (e.g. [Equation 3.4](#)) and perform the optimisation process. Now the training step moves on to the generator, but since it has

already created a batch of samples in the previous step its job is mostly completed. All that is left is for the discriminator to produce a further batch of scores, since it has just been updated with an optimisation step, from the fake batch of samples which, along with the real batch of samples from the training set, can be used to evaluate the generator loss (e.g. [Equation 3.7](#)) and perform optimisation.

The above process is repeated over many steps until the full training set has been sampled and processed. Then the whole loop is repeated over many epochs. How many samples are used in a batch and how many epochs are performed are both examples of *hyperparameters*, which can have a significant impact on the overall performance of the trained network and the time taken for training. These parameters are selected through a further method referred to as *tuning*, which essentially implements some method (e.g. grid search) of dynamically selecting a hyperparameters value by analysing the performance of the network on the validation data.

Variations of GANs

The distinctive feature of the original DCGAN described above is the form of the input z . It is a so-called ‘latent vector’, which is a vector of random numbers sampled from a ‘latent space’ which is a random distribution of numbers. The latent vector provides seed information to the GAN which, once trained, allows it to generate new images, not contained in the training set, but ones that are still incredibly realistic. One downside though is that without a good understanding of the GANs interpretation of the latent vector, it isn’t possible to specify the type of output you want. As an example, imagine that a GAN has been trained to generate realistic images of all kinds of dogs using a random vector as a seed. But if you wanted a synthetic image of specifically a golden retriever there would be no way, without doing a comprehensive study of the latent space and its relation with the GAN, to predetermine that a golden retriever would be in the synthetic output image.

This problem is addressed by a version of GANs known as Conditional Generative Adversarial Networks (CGANs, [Mirza & Osindero 2014](#)), which take in an input of a class label in addition to a random vector. By including a class label in the training process, which in practice just provides a different input vector distribution per class label, the output of the GAN can have conditions placed on it, hence the name.

The desire to place conditions on the output of the GAN can be extended even further

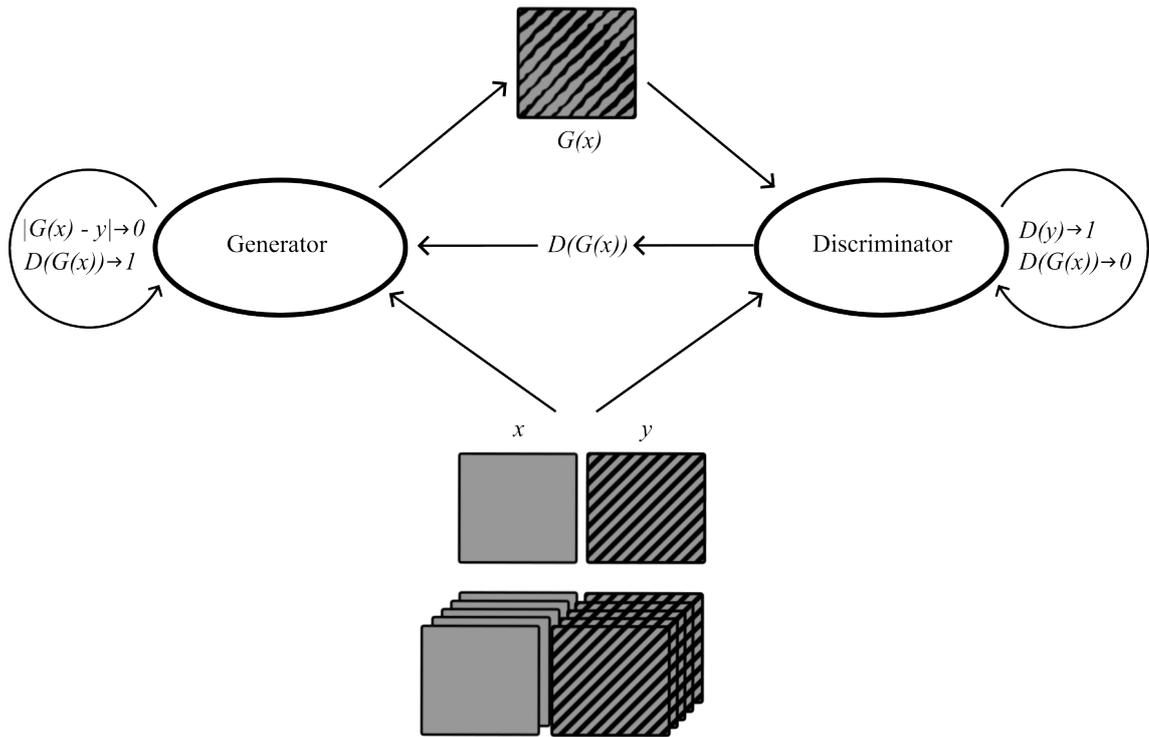


Figure 3.4: A diagram of the pix2pix model and the procedure for training. The training set is shown at the bottom of the image, from which an image and mask is taken. The generator produces a fake mask $G(x)$ from the image, and the discriminator calculates a probability $D(G(x))$. Each side of the model optimises itself towards the quantities given in the circular arrows.

through the pix2pix model (see [Figure 3.4](#), [Isola et al. 2016](#)), which modifies the CGAN model to produce a pixel-by-pixel classification of an input image. To achieve this, the training set must contain pairs of images where one image, x , may contain any number of arbitrary objects belonging to any set of classes, such as a picture of a dog and a cat. This image is accompanied by its ground truth y ; the pixel level classification of objects in the image. During training the network learns to reproduce this classification on new, unclassified images of the same objects. This modification effectively transforms the GAN use case from an image generation tool to an image translation tool, where the desired translation can be specified in the training set. How the pix2pix model is used in this work for creating interference masks is detailed in [section 3.5](#).

Regardless of type, all GANs share similar characteristics because of the unconventional training process. With optimal training on a high quality training set, GANs are known to quickly produce accurate and realistic images that certainly surpassed many models that came before it. Additionally, the GAN architecture and training process are quite generalisable; it only requires a generator that will create some output, and a discriminator that will criticise that output. This means that GAN based models can be applied to a huge variety of applications depending on the needs of a particular problem. Fur-

thermore, it means that the generator and discriminator themselves can be made more or less complex, or even changed into something completely different, again depending on the needs of a given problem. The GAN framework is certainly an important tool in the machine learning toolbox.

However, GANs are not without their downsides. The most notable issue with GANs is how unstable they can be during training. The entire process is reliant on the tension between the generator and the discriminator, but this tension can sometimes lead to divergent behaviour, instead of the overall network converging on an optimal solution. There can be a few causes of this, but most commonly it is because the generator often has a fundamentally more difficult task. It can take longer for the generator to start producing even somewhat accurate results, and all the while the discriminator is making larger strides in its ability to perform. Eventually, the gap between the discriminator and the generator may become too large, the generator cannot ever get close to an accurate result, and the process diverges from equilibrium. Using the analogy of a chess game, this is effectively where the skill gap between two players is so large, one player is unable to learn from the other because the game is over too quickly, so they simply get up and walk away!

This issue has a further secondary effect. In the case of the pix2pix model the complexity of the whole network, combined with the tension in the training process, means that small changes to the underlying architectures can have a very large, sometimes detrimental impact on whether the GAN can converge on an optimal state. This means that users of the pix2pix model cannot easily diverge from the original configuration which has been rigorously tested.

3.3 Radio Frequency Interference

As discussed in [chapter 1](#), radio observatories hold incredibly sensitive equipment for detecting astronomical radio emission. This equipment must be more sensitive than most other astronomical observatories, since many astronomical objects tend to emit less overall energy at radio wavelengths. To further illustrate this, consider the radio source Cassiopeia A. It is one of the strongest sources in the radio sky, and has a flux density of approximately $2000 \text{ Jy} = 2000 \times 10^{-26} \text{ Wm}^{-2}\text{Hz}^{-1}$ at 1.4 GHz ([Findlay et al., 1965](#)). In a typical 12-hour radio observation, a colossal 500 m radio telescope (the size of FAST, [Li & Pan 2016](#)) sensitive to radio waves at 1.4 GHz would collect a total of approximately $4.3 \times 10^{-16} \text{ J}$ of energy. This is just enough energy to power an average handheld calcula-

lator for about $4.3 \text{ ps} = 4 \times 10^{-12} \text{ s}$. Though, many radio observatories today have a large observing bandwidth around 1 – 2 GHz if observing at 1.4 GHz. Assuming a flat Spectral Energy Distribution (SED) over a 2 GHz receiving bandwidth, enough power is actually collected for about 8.6 ms of calculator time.

At the same time, at least for now, radio observatories operate almost exclusively on the Earth. This is problematic because of how pervasive radio emissions are on a planet swarming with technologically advanced humans. Radio transmissions are fundamental to a wide range of practices and technologies, including but not limited to:

- Broadcast radio
- Television
- Long range communication
- Radar
- Remote operation
- Satellite communication

One of the fundamental challenges of radio astronomy is how to reduce or completely remove the impact of these radio signals on the incredibly sensitive equipment, while simultaneously preserving the comparatively weak astronomical radio emission. Radio signals that are not related to astronomical sources are broadly labelled as Radio Frequency Interference (RFI) by radio astronomers.

Single-dish radio observations are impacted more by RFI than interferometric observations, since they have no other reference point for a detected signal in the way that interferometer baselines do. Since interferometers are correlating two signals from two separated points in space, an RFI signal originating from a single point near to one antenna is less likely to be detected by both antennas in a baseline, and therefore is less likely to be correlated. This effect only increases with baseline length, so long-baseline VLBI observations suffer from RFI much less than single-dish observations. Despite this, RFI signals are so prolific that even the largest VLBI arrays will encounter interference to some extent. This is a growing concern given the massive increase in air and space-borne radio emitters (see [section 3.6](#)).

Even if RFI is not correlated, the extra power entering the receiver can contribute, sometimes quite significantly, to warming up the receiver. This raises the noise level of the baseline output, which degrades image quality even though the RFI itself was not

actually correlated! Furthermore, RFI can have a deleterious impact on spectroscopic radio observations, since RFI coincident in frequency with the desired spectral lines can irrecoverably overwrite the measured data for the duration the RFI is present. This prevents the astronomer from carrying out their research, and costs the observatory in lost observing time.

The following sections will first outline the current state of radio astronomy in terms of dealing with RFI. It will explore the techniques currently in use, and then introduce the experiment conducted on flagging RFI using a GAN.

3.3.1 Regulation

Modern society is very much dependent on the efficient use of radio waves for wireless communication. As a result, a great deal of regulation surrounds the process of emitting radio waves, to ensure that all parties involved can fulfil their objectives without interfering with one another. Internationally, these regulations are set forth by the United Nations specialist agency known as the International Telecommunications Union (ITU), which is broadly accepted to be the arbiter for radio related issues. Every four years, the ITU produces the Radio Regulations (ITU-R, 2019) informed by the World Radio Communication Conference and the radiocommunications sector of the ITU (known as ITU-R). This document is incorporated into the constitution of the ITU which, in turn, is ratified by UN member states. While the UN places the importance of national sovereignty over UN regulation, it is generally beneficial for member states to agree to these rules, to facilitate economic and social development. As a result, these regulations are broadly respected and implemented.

As part of their regulatory framework, the ITU allocates frequencies for specific purposes to prevent any competition or conflict over frequency usage. Radio astronomy is formally recognised by the ITU as a passive service, and so is afforded the same rights as other recognised services. The relevant ITU material is condensed into the ITU-R Handbook on Radio Astronomy², and a review of the regulations from the perspective of radio astronomy is given by Baan (2019).

Perhaps the most important allocation for radio astronomy allows undisturbed observation of the local Hydrogen 21 cm line, which occurs in L-band at around 1400 MHz (see Figure 3.5). Naturally, the proportion of frequencies allocated to radio astronomy is quite

²Available at: <http://handle.itu.int/11.1002/pub/809847c8-en>

| Allocation to services | | |
|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|----------|
| Region 1 | Region 2 | Region 3 |
| 1 300-1 350 | RADIOLOCATION AERONAUTICAL RADIONAVIGATION 5.337 RADIONAVIGATION-SATELLITE (Earth-to-space) 5.149 5.337A | |
| 1 350-1 400 FIXED MOBILE RADIOLOCATION 5.149 5.338 5.338A 5.339 | 1 350-1 400 RADIOLOCATION 5.338A 5.149 5.334 5.339 | |
| 1 400-1 427 | EARTH EXPLORATION-SATELLITE (passive) RADIO ASTRONOMY SPACE RESEARCH (passive) 5.340 5.341 | |

Figure 3.5: An excerpt from ITU Radio Regulations (ITU-R, 2019), outlining the international allocation of 1400-1427 MHz specifically to the Radio Astronomy service, to allow unobstructed access to the Hydrogen 21 cm line.

small and not always exclusive, and the range of frequencies in a particular allocation tend to be narrow in comparison to modern observational bandwidths. Furthermore, the frequency allocations to radio astronomy are intended to protect local ($z \sim 0$) receiving, and do not account for Doppler shifted emission.

This creates a problem for radio astronomers. State of the art radio observatories regularly conduct experiments covering fractional bandwidths ≥ 0.5 (the ratio of observing bandwidth to central frequency $\Delta\nu/\nu$). This means that observations at L-band with a wideband receiver can also pick up transmissions outside the 21 cm ITU allocation. Of course, the spectrum allocation cannot reasonably be expected to conform to the desires of every radio astronomer, so robust solutions that allow a coexistence with neighbouring allocations must be developed. This is doubly true, since there are numerous situations where the ITU regulations do not explicitly protect the radio astronomy service. Firstly, some allocations are shared with other active services, so transmissions can still be picked up from those even indirectly (such as being reflected off an aircraft). Secondly, so called ‘out of band’ emissions: emissions received in one allocation from another separate allocation, are allowed up to a limit. Finally, accidental emissions are protected under ITU regulations, so interference in protected bands may originate from an accidental source. These are just a few examples of the difficulties of operating a passive service, and illustrate why effective methods of interference management will always be critical to radio astronomy even in a favourable regulatory environment.

3.3.2 Methods of Mitigation

Methods of mitigating the presence or impact of RFI can be roughly categorised based on when they are implemented relative to the time of observation. There are, generally speaking, three such categories: preventative measures, flagging during observation, and post-correlation flagging. The methods within each of these categories form a layered defence against corruption from RFI; the earlier the RFI is caught and removed, the better it is for the astronomer.

Preventative Measures

Protecting an observatory against RFI is a process that starts right at the time of construction. Given a choice, astronomers will position their antennas far away from emitting sources, which normally involves groups of humans. As optical astronomers position their telescopes in part based on light pollution, radio astronomers position their antennas in part based on RFI.

Often this isn't enough to completely eliminate the problem. A radio antenna can, and often is, surrounded by a so-called 'radome' but this does nothing more than protect the sensitive equipment underneath from extreme weather. In rare cases, an enforced Radio Quiet Zone (RQZ) can be put in place in areas surrounding antennas to reduce the likelihood of RFI detection. Examples of RQZs include the United States National Radio Quiet Zone surrounding the Green Bank Observatory ([Prestage et al., 2009](#)) in West Virginia, the FAST RQZ ([Li & Pan, 2016](#); [Zhang et al., 2019](#)) in China's Guizhou Province, and the Australian Radio Quiet Zone WA (ARQZWA) surrounding several radio telescopes in Western Australia ([DeBoer et al., 2009](#)), including the future site of SKA-Low. As RQZs require political will to instantiate and sustain, they are not an option open to all astronomers and even observatories with a functioning RQZ still require other ways of addressing RFI.

A more direct approach is to integrate band-stop filters (or 'notch' filters) into the antenna setup, which attenuate signals over a range of frequencies. If a nearby interference source is known to produce strong emissions within a specific frequency range, then these filters can be permanently installed at the receiver to attenuate those frequencies, and eliminate the interference. Obviously, this still prevents observation of the filtered frequencies, but it is preferable if the interfering signal is strong enough to saturate the receiver, which can prevent useful data being gathered across the entire range of frequencies

under observation.

Finally, most preventative measures require long-term monitoring of the local RFI environment to be most effective. Detailed knowledge of emissions entering the receiver allows the observatory to adapt quickly to changes in the RFI environment, and provides an evidence base from which the observatory can make decisions about how best to combat interference. Therefore most observatories, regardless of the levels of protection already in place, will include regular monitoring of local emissions as part of regularly scheduled operations³.

Pre-correlation Flagging

There are many different methods of flagging in real-time, but they all broadly implement one of two solutions: separating out the RFI signal, or preventing the detection of RFI altogether. The former is more complex, as it requires the decomposition of the astronomical signal from the interfering signal with minimal error, the latter is easier (and often cheaper) to implement, but comes with an associated data loss. Which methods are implemented comes down to a cost-benefit analysis made by the observatory, informed primarily by the characteristics that define the telescope itself (antenna count and design, antenna separation, frequency range, operational windows, location, etc). For example, if the observatory manages one or two closely spaced antennas, then it becomes far more practical to implement a complex system of real-time detection and subtraction of interference. If instead the observatory consists of many antennas with large separations (perhaps hundreds of kilometres) then this system is less viable, and simply accepting a certain level of data loss may be more practical.

An example of a real-time method of separating out the RFI signal is a technique known as ‘spatial filtering’ (Boonstra & van der Tol, 2005), which leverages the ability of interferometers to create direction-dependent antenna gains (known as ‘beam forming’). If a strong signal is known to originate from a certain direction, or a strong signal is detected during an observation, the response pattern of a particular antenna can be mutated such that the nulls in the response pattern coincide with the direction of the interference. This allows the antenna to continue observing the astronomical source, while being blind to the interfering source. This particular technique can be practically difficult, and is only useful in arrays with short baselines (Fridman & Baan, 2001) and wide primary beams,

³The VLA, for example, conducts regular surveys for RFI: see <https://science.nrao.edu/facilities/vla/docs/manuals/obsguide/rfi>

such as LOFAR or SKA-Low.

An example of a real-time method of removing corrupted signals is to simply ‘blank’ the incoming signal for the duration of the interfering signal (Mitchell et al., 2010). This is effectively a form of real-time flagging, with the advantage that it can be done before any averaging of the data takes place, and is highly flexible in how the blanking threshold(s) can be selected.

Systems based on software correlators lend themselves well to pre-correlation flagging methods as, with sufficient computing power, the methods can be implemented entirely in software as part of the correlation process; hardware components are not necessary. Though, given the extremely large data volumes of pre-correlation signals, even the simplest method of RFI mitigation may add significant overhead to the correlation pipeline.

Post-correlation Flagging

Post-correlation flagging is what is generally referred to when astronomers discuss flagging, since it is a standard step in all post-observation data pipelines. Since the observation has already been conducted and the data correlated (for interferometer arrays), it is very difficult to disentangle the RFI signal from the astronomical. Therefore, almost all post-correlation flagging methods involve some level of data loss, and the challenge of good flagging algorithms is to maximise RFI removal while minimising the loss of astronomical signal.

If the volume of RFI encountered is small, then post-correlation flagging may be done manually either by direct visual inspection, simple thresholding (such as clipping), or implementing flags based on information gathered during observation. These approaches were standard in the early days of radio astronomy and are still used for extremely remote observatories, or VLBI observatories. Though this quickly becomes impractical as the volume of RFI increases along with increased bandwidths and additional antennas; naturally, automated solutions become preferable.

The distinguishing feature of RFI is that it is, in most cases, orders of magnitude more powerful than astronomical signals. There are several reasons for this, one is that artificial emissions originate from sources closer to the receiver. As a result, many post-correlation RFI detection methods rely on this, and implement algorithms based on thresholding. This threshold is often calculated per time-frequency bin; how it is calculated is the fundamental challenge. Current techniques perform thresholding on a time-frequency image

of the output of the antenna (or antennas), an example of which is shown in [Figure 3.6](#).

One of the most popular software solutions is called AOFlagger ([Offringa et al., 2012b](#)), which utilises a form of adaptive thresholding, combined with extension of the resulting mask. As a general purpose post-correlation flagging tool, the current iteration of AOFlagger represents the state of the art in interference removal. It was first introduced by [Offringa et al. \(2010b\)](#) as a prospective flagging tool for the LOFAR pipeline. The core of the technique implements the SumThreshold algorithm ([Offringa et al., 2010a](#)), which iterates over each pixel in an input time-frequency image constructed using data from a single antenna or baseline. By examining the surrounding pixels in time or frequency, the SumThreshold method deduces the likelihood of the current pixel being contaminated using a variable threshold. Following this, AOFlagger implements a scale-invariant rank operator, which extends (or ‘dilates’) the flag mask constructed by the SumThreshold method to include remaining false negatives in the flags.

One of the problems with statistical methods of flagging is that they often require many parameters to be specified by the user to achieve an optimal set of flags that minimise RFI corruption and data loss for a specific observation, by a specific observatory. This makes deploying these solutions tricky, and often large amounts of ‘good’ data can be wasted because of an improper setup. An alternative approach that is becoming more and more popular in research (though explicitly *not* in practice), is to train a neural network to perform the flagging. This has the advantage of reducing the amount of user-specific customisation to almost nothing, while (ideally) maintaining a high level of precision. The downside is the requirement for large amounts of accurate information on RFI to form an effective training set, which can be a laborious task to produce.

3.3.3 Types of RFI

Though the sources of RFI are mostly unique to a particular telescope (with a few exceptions such as mobile phone signals and satellite transmissions), the reason that many post-correlation RFI removal tools all choose to operate in the time-frequency space is that, for the most part, RFI presents itself in somewhat regular morphologies in this plane. This reduces the complexity of the problem, and allows some assumptions to be made about the likely morphology of the RFI being detected. This section will describe and present these regularly encountered types of RFI at the commonly used frequencies for radio astronomy between ~ 0.1 -100 GHz.

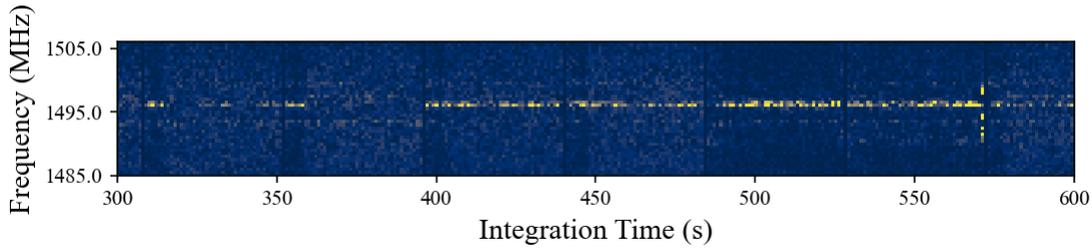


Figure 3.6: An example of narrow band RFI found in the Lockman Hole data used in this experiment. It lasts ≥ 4 minutes and occupies a narrow frequency range of ~ 1 MHz.

Narrow band RFI refers to a signal that is narrow in frequency. Most often, narrow band RFI presents as a strong signal occupying a small number of channels and will persist for somewhere between several minutes up to several hours. It may turn on and off at irregular intervals throughout the observation and may also fluctuate in strength, causing it to sink below the sensitivity limit of the receiver. The principal method of removing narrow band interference is to constrain the data removal to the corrupted channel, but in certain cases the signal is strong enough to cause a leakage effect that corrupts adjacent channels, which may also have to be removed. [Figure 3.6](#) shows an example of a narrow band interference signal within the Lockman Hole data that is presented in [chapter 6](#).

If left unflagged, narrow band RFI can raise the noise level in the final image by introducing artefacts exceeding the true visibility noise. In the worst case, these artefacts can be picked up during CLEANing (see [section 2.7](#)) and introduced to the sky model, which degrades the effectiveness of self-calibration, and causes radio sources to appear where there are none.

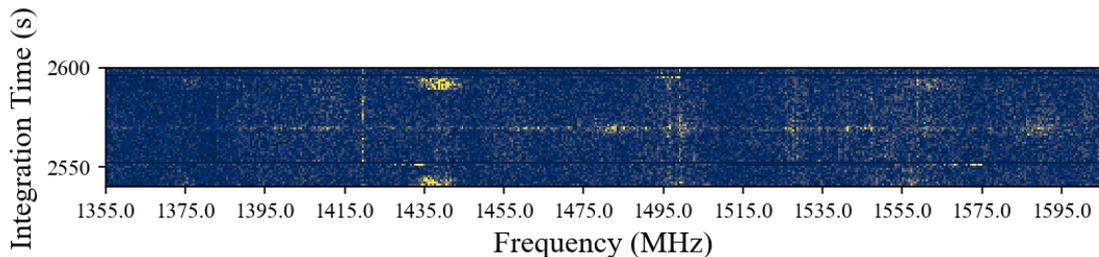


Figure 3.7: An example of a weak, wide band interference signal found in the Lockman Hole data used in this experiment. It only lasts for ~ 8 s, but occupies a wide frequency range ≥ 300 MHz.

Wide band RFI, like that shown in [Figure 3.7](#), occupies a wide range of frequencies that may span the entire bandwidth of the receiver. It does not normally last as long as narrow band RFI. Given that it occurs over a wider range of frequencies, wide band RFI can be more detrimental to the quality of an observation since it effectively overwrites

more visibilities, but it is also far more conspicuous. The primary impact of unflagged wide band interference is on the calibration of the bandpass. Normally, it is assumed the frequency response of the receiver does not vary significantly over the course of a single observation, so the calibration is calculated per-source using all available visibilities. If a strong wide-band interference signal is present at this stage, the calculated bandpass solutions will incorporate the RFI, and then use that ‘incorrect correction’ to correct the rest of the visibilities, which skews the frequency information in the corrected visibilities and in the final image.

3.4 Machine Learning for RFI Mitigation

Currently, there is no universal method or technique to remove RFI from radio astronomical data. As described in [subsection 3.3.2](#), observatories will implement their own measures to mitigate or remove RFI, but the sheer volume of data involved means the process must be automatic and there is always some amount of interference that persists in the data that is delivered to the principal investigator. This situation does not seem like it is going anywhere, so it is vital that effective solutions are developed both at the observatory, and for the end user. Since its explosion in both popular culture and research in the 2010s, machine learning has been explored as an alternative solution to the problem of RFI.

Taking a step back, the first fundamental goal of any RFI mitigation solution is to accurately identify an occurrence of a strong signal that ‘sticks out’ of the underlying signal in some way. RFI mitigation methods then go one step further and decide if and how to actually flag or remove the signal, but the initial goal of identification is common to a much wider range of applications. Technologies such as speech-to-text and speech recognition must identify a strong signal (the user’s voice) amongst what may be a significant noise signal (background conversations, passing cars, etc). Speech recognition methods, therefore, may also be of interest to radio astronomers (e.g. [Hannun et al. 2014](#)). Another application is the search for extraterrestrial intelligence (SETI), which also attempts to find what is assumed to be a stronger signal amongst noise. In fact, a key issue in SETI applications is the differentiation of candidate extraterrestrial signals with RFI, which also disrupts their operation ([Enriquez et al., 2017](#)).

Many related fields have found success in applying machine learning methods, so naturally radio astronomers have been examining the effectiveness of using these methods

to flag RFI. [Agarwal et al. \(2020\)](#) find success in training CNNs for the identification of Fast Radio Bursts (FRBs), which is functionally almost identical to identifying RFI signals. [Czech et al. \(2018\)](#) also use CNNs but with an additional Long Short-Term Memory (LSTM) network for identifying RFI signals in one-dimensional signal data. [Long et al. \(2019\)](#) utilise a modified U-Net which is also used as part of the GAN model used in this work.

Generative Adversarial Networks (see [subsection 3.2.2](#)) are an interesting option for automatically generating accurate flag masks. As of the time of writing, only [Li et al. \(2021\)](#) have implemented the GAN method for flagging interference in radio observations. They train a GAN using exclusively simulated data from the HI Data Emulator package (HIDE, [Akeret et al. 2017](#)) with the objective of identifying RFI in data from FAST. The work presented in this thesis uses a similar GAN model to the one proposed by [Li et al. \(2021\)](#), but instead trains and tests it against manually flagged real VLBI data (from e-MERLIN) instead of simulated single-dish data, or data with automatically generated flags. The reasoning behind this is that by training on simulated data, the network is not exposed to the ‘real world factor’ of data from an actual telescope. Real radio astronomy data contains many imperfections that are not easily simulated unless specific attention is paid to it, such as bandpass attenuation, antenna dropouts, and complex RFI morphologies. If a machine learning solution is to be viable, it must have some way of dealing with these. Furthermore, training the network on data that is automatically flagged simply encourages the network to reproduce the output of the automated method, negating the need to use machine learning in the first place, while simultaneously replicating any errors or inefficiencies in the automated method. By training on manually flagged real data, these issues are sidestepped.

3.4.1 IFlag

As mentioned in earlier sections, most flagging is not done manually due to the sheer amount of data produced by a modern radio interferometer. To make this concrete, consider the case of the European VLBI Network (EVN) interferometer array ([Venturi et al., 2020](#)). It is an international VLBI facility that can operate with as many as 30 stations. Imagine a hypothetical EVN observation that produces a total of 45 correlated baselines from only 10 separate stations; there would be a minimum of three sources observed (the target, the phase calibrator, and the flux calibrator) over the course of a single 12 hr

observation. Most observatories measure polarisation at both antennas in a baseline, so each baseline produces a minimum of 4 time-frequency images per source. So far, this means that our hypothetical observation would produce $45 \times 4 \times 3 = 540$ separate time-frequency images. However, this assumes one continuous scan for each source, which isn't possible due to the need to periodically observe the phase calibrator in between target scans. Assuming that the full twelve hour observation, including all sources, is instead broken up into scans of a maximum of 10 minutes each, we would instead have $45 \times 4 \times (\frac{12 \times 60}{10}) = 12960$ separate time-frequency images! A typical VLBI observation may average the data over 4 s, and may have somewhere in the region of 500 frequency channels, which means that each of these images would be 500x150 pixels in size. So, manually flagging this hypothetical observation would mean that nearly 13000 images of size 500x150 would have to be examined and flagged individually.

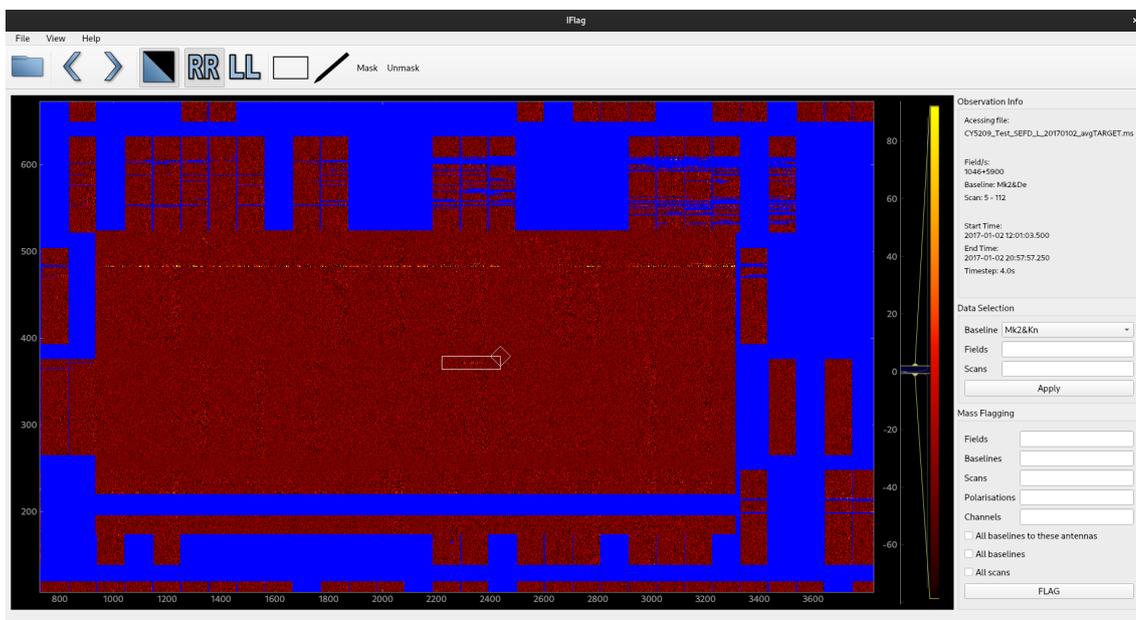


Figure 3.8: An example of the IFlag user interface. Buttons along the top allow, from left to right, opening a radio measurement set, moving between images, toggling the mask overlay (shown in blue), switching between polarisations, toggling the box drawing tool, toggling the pen drawing tool, and masking/unmasking with a drawn box. Input fields along the right side of the interface allow the user to show a specific image, or flag multiple parts of the measurement set matching the input criteria.

This is, obviously, a laborious task, and is similarly challenging for many modern observatories. Unfortunately, the current software environment in radio astronomy does not offer any reasonable solutions. Existing tools are command line interfaces that are not efficient for bulk manual operations or interfaces that allow viewing the time-frequency data but do not facilitate quick manual flagging through the interface. Flagging functionality exists in the commonly used Astronomical Image Processing System (AIPS) library

(Greisen, 2003) in the form of routines such as UVFLG and RFLAG, but again these are not designed for precise, manual flagging of large amounts of data.

Therefore, to make this challenge more accessible and less time-consuming, a graphical user interface was developed specifically for the purpose of enabling quick manual flagging of large amounts of radio astronomy data⁴. This enabled the efficient production of a training set containing manually flagged radio data. This interface is referred to as IFLAG throughout this thesis. IFLAG is developed in pure PYTHON, and builds on top of the CASA software stack (CASA Team et al., 2022) to interface with radio astronomy data contained in the Measurement Set (MS) format. It uses PYTHON bindings to the Qt GUI framework to construct the interface. Each observed field is presented as one contiguous image to the user since many RFI signals occur on timescales longer than a single scan, but there are many different options to view different slices of the data. Figure 3.8 presents the IFLAG interface.

The GUI allows quick manual flagging by giving the user the ability to draw over parts of the image deemed to be RFI, by using either the ‘pen’ mode or the ‘boxing’ mode, that allows drawing or boxing respectively. The use cases for such an application include: fine-tuning an already flagged observation, annotating radio astronomy data for use in machine learning, or even to manually flag small observations. At present, it is the only application I know of to provide this functionality. It was utilised in this work to create the training set used to train the GAN.

3.5 Flagging RFI with a GAN

This section presents the work conducted on using the GAN model described in subsection 3.2.2 to automatically create interference masks. Full details of the generator and discriminator networks used in this experiment can be found in Table 3.1 and Table 3.2. The networks were constructed using the PYTORCH toolkit (Paszke et al., 2019)⁵. The networks were designed by attempting to mimic the ‘RFI-NET’ presented by Li et al. (2021), but also to fit into the constraint of available processing power. Unfortunately, this limitation means that the GAN setup used in this experiment is essentially a scaled down version of the RFI-NET. Such is the nature of machine learning architecture design, the final configuration presented here was ultimately settled on through repeated trial and er-

⁴The source code is available on request

⁵The source code is available at <https://github.com/JakeEBrooks/BrooksRFIGAN>

ror experiments, and the hyperparameters were selected through a sparse grid search. The key differences between this design and the design of [Li et al. \(2021\)](#) are:

- **Input Shape:** The U-Net presented here takes an input array of size $batchsize \times 128 \times 1024 \times 1$. In [Table 3.1](#) and [Table 3.2](#) the batch size dimension is omitted as this is generally selected for reasons related to computation time and not model performance. The U-Net model presented by [Li et al. \(2021\)](#) takes an input shape of $8 \times 256 \times 128 \times 1$ ⁶. In theory, the input shape should not have a significant impact on the performance of the network, and is generally something selected for useability or computing reasons.
- **Number of Convolution Filters:** Traditional U-Net models double the number of filters at each layer of the U-Net, to match the reduction in image size through max pooling. This is the approach adopted by [Li et al. \(2021\)](#). The number of convolution filters is one of the biggest drivers of computation time, so to reduce this the convolution filters in the model presented here were increased additively rather than multiplicatively. This undoubtedly has an impact on the performance of the model.
- **Discriminator Layers:** [Li et al. \(2021\)](#) present a discriminator with four convolution layers. The model presented here uses only three to mirror the decreased number of filters in the generator and for computational reasons. Furthermore, GAN designs tend to intentionally provide the discriminator with less capability than the generator, since the generator has a fundamentally more difficult task.

⁶Note that the [Li et al. \(2021\)](#) model has an extra channel in the array passed to the discriminator than what is passed to the generator. It is unclear what purpose this extra channel serves. Furthermore, there is a contradiction in the number of convolution operations performed in the generator between [Figure 2](#) and [Table 1](#). [Table 1](#) suggests three convolution operations per layer, whereas [Figure 2](#) suggests two convolutions on the encoding side, and three convolutions on the decoding side; an uncommon approach. It seems most likely that [Table 1](#) is a more accurate description, in which case the model presented here performs one less convolution in each layer in keeping with the traditional U-Net. It is unknown what effect this has on model performance.

| Layer Type | Kernel Size | Output Shape | Filter No. |
|---------------------------|-----------------|------------------|------------|
| Input | - | 128 x 1024 x 1 | - |
| Convolution + BN + ReLU | 3 x 3 | 128 x 1024 x 64 | 64 |
| Convolution + BN + ReLU | 3 x 3 | 128 x 1024 x 64 | 64 |
| Max Pool | 2 x 2, Stride 2 | 64 x 512 x 64 | - |
| Dropout | - | 64 x 512 x 64 | - |
| Convolution + BN + ReLU | 3 x 3 | 64 x 512 x 128 | 128 |
| Convolution + BN + ReLU | 3 x 3 | 64 x 512 x 128 | 128 |
| Max Pool | 2 x 2, Stride 2 | 32 x 256 x 128 | - |
| Dropout | - | 32 x 256 x 128 | - |
| Convolution + BN + ReLU | 3 x 3 | 32 x 256 x 192 | 192 |
| Convolution + BN + ReLU | 3 x 3 | 32 x 256 x 192 | 192 |
| Max Pool | 2 x 2, Stride 2 | 16 x 128 x 192 | - |
| Dropout | - | 16 x 128 x 192 | - |
| Convolution + BN + ReLU | 3 x 3 | 16 x 128 x 256 | 256 |
| Convolution + BN + ReLU | 3 x 3 | 16 x 128 x 256 | 256 |
| Max Pool | 2 x 2, Stride 2 | 8 x 64 x 256 | - |
| Dropout | - | 8 x 64 x 256 | - |
| Convolution + BN + ReLU | 3 x 3 | 8 x 64 x 320 | 320 |
| Convolution + BN + ReLU | 3 x 3 | 8 x 64 x 320 | 320 |
| Deconvolution + BN + ReLU | 3 x 3, Stride 2 | 16 x 128 x 256 | 256 |
| Concatenation | - | 16 x 128 x 512 | - |
| Dropout | - | 16 x 128 x 512 | - |
| Convolution + BN + ReLU | 3 x 3 | 16 x 128 x 256 | 256 |
| Convolution + BN + ReLU | 3 x 3 | 16 x 128 x 256 | 256 |
| Deconvolution + BN + ReLU | 3 x 3, Stride 2 | 32 x 256 x 192 | 192 |
| Concatenation | - | 32 x 256 x 384 | - |
| Dropout | - | 32 x 256 x 384 | - |
| Convolution + BN + ReLU | 3 x 3 | 32 x 256 x 192 | 192 |
| Convolution + BN + ReLU | 3 x 3 | 32 x 256 x 192 | 192 |
| Deconvolution + BN + ReLU | 3 x 3, Stride 2 | 64 x 512 x 128 | 128 |
| Concatenation | - | 64 x 512 x 256 | - |
| Dropout | - | 64 x 512 x 256 | - |
| Convolution + BN + ReLU | 3 x 3 | 64 x 512 x 128 | 128 |
| Convolution + BN + ReLU | 3 x 3 | 64 x 512 x 128 | 128 |
| Deconvolution + BN + ReLU | 3 x 3, Stride 2 | 128 x 1024 x 64 | 64 |
| Concatenation | - | 128 x 1024 x 128 | - |
| Dropout | - | 128 x 1024 x 128 | - |
| Convolution + BN + ReLU | 3 x 3 | 128 x 1024 x 64 | 64 |
| Convolution + BN + ReLU | 3 x 3 | 128 x 1024 x 64 | 64 |
| Output Convolution | 1 x 1 | 128 x 1024 x 1 | 1 |

Table 3.1: Full details of the generator network constructed in PYTORCH.

| Layer Type | Kernel Size | Output Shape | Filter No. |
|-------------------------|-----------------|----------------|------------|
| Input | - | 128 x 1024 x 1 | - |
| Convolution + BN + ReLU | 3 x 3, Stride 2 | 64 x 512 x 64 | 64 |
| Convolution + BN + ReLU | 3 x 3, Stride 2 | 32 x 256 x 128 | 128 |
| Convolution + BN + ReLU | 3 x 3, Stride 2 | 16 x 128 x 192 | 192 |
| Flatten | - | 393216 | - |
| Linear | - | 192 | - |
| Linear + Sigmoid | - | 1 | - |

Table 3.2: Full details of the discriminator network constructed in PYTORCH.

3.5.1 Training Data

The signal received during a radio observation can be broken down into a series of components that represent different sources. The correlated signal between two antennas for a single polarisation, $V(\nu, t)$, can broadly be described as:

$$V(\nu, t) = A(\nu, t) + R(\nu, t) + \sigma(\nu, t) + \epsilon(\nu, t) \quad (3.8)$$

where $A(\nu, t)$ is the astronomical signal detected at the antenna. In most observations, this is a slowly varying weak signal compared to the strength of the majority of RFI. $R(\nu, t)$ is the RFI component, which is typically the strongest overall signal received. The objective of any RFI flagging technique is to remove this component, or eliminate its effect. The noise component, $\sigma(\nu, t)$, is a weak signal ideally described by a gaussian probability distribution. The final component, $\epsilon(\nu, t)$, is the ‘real world’ factor, which could include a dependence on any of the other three components. It represents the characteristics of a particular receiver, or collection of receivers, and their properties during an observation (their temperature, for example). It may also describe events such as antenna dropouts and component failure. Therefore, the ϵ component is, at best, specific to a particular correlator output, and not easily characterised by simulations without paying specific attention towards reconstructing it. If a machine learning model is trained on simulated visibilities that do not include the information contained within $\epsilon(\nu, t)$, then once the network begins to operate on ‘real’ visibilities it will encounter these ‘real world’ problems and may behave unpredictably. By using manually flagged ‘real’ visibilities as a training set a model can be trained that is, in theory, capable of predictably identifying $R(\nu, t)$ amongst the net effect of $\sigma(\nu, t)$ and, most importantly, $\epsilon(\nu, t)$.

Three separate L-band observations of the Lockman Hole ([Lockman et al., 1986](#)) were conducted by e-MERLIN. The observations used the phase-referencing technique, so another nearby point-like source was observed intermittently during the course of the observations. The same source, the phase calibrator, was used in all three observations and is the sole source making up the training and test sets. The specific phase calibrator used is a BL Lac type AGN located at $\alpha = 10^{\text{h}}58^{\text{m}}37.73^{\text{s}}$ and $\delta = +56^{\circ}28'11.18''$. The two observations making up the training set occurred on 16th June 2017, and 1st January 2018. Total observing time for the phase calibrators in each of these observations was approximately three hours, making a training set that uses just six hours of observing time. The observation used as a test set occurred on 2nd January 2018. All three observations were manually

flagged using the IFLAG interface (see [subsection 3.4.1](#)) to approximate knowledge of RFI within each observation.

3.5.2 Pre-processing

Operating on raw time-frequency images does not achieve desirable results. This is a common problem across all flagging techniques, not just those based on machine learning. In the case of a GAN, the generator can be overloaded by redundant yet prominent features such as dead pixels at the start of a scan and bandpass attenuation at the edges of each spectral window. Moreover, the absence of these flagged features are a clear indicator to the discriminator of a fake image, which causes the generator to put wasted effort into reproducing them. Therefore, these features should be removed or suppressed before the image is passed to the generator, so that the network can focus on identifying the correct features associated with RFI. Many different pre-processing methods could be used to enhance the ability of the network to identify RFI, and there is a large scope for experimentation, the only limitation is that the training images must be representative of the distribution of images to be assessed by the trained network. Consequently, ideal pre-processing operations will give consistent results even when the input is unknown. In this work, three modifications are made to the training data and appear to sacrifice little in terms of similarity to the input but dramatically boost the efficacy of the network. These steps are outlined below in the order they are applied. Note, all of the following steps are performed sequentially on an image representing a single baseline, a single polarisation, and the full spectral range, yielding a single image of size 128x1024.

Most observations will have a trivial set of flags at the start (and sometimes the end) of a scan, usually as the antennas in a baseline are not yet simultaneously observing the same pointing. This is dealt with by so-called ‘quack’ flagging, which simply flags the first and/or last n seconds in a scan where n is configurable. Observatories also normally have a standard set of flags that they apply to specific frequency channels at the edges of each sub-band (or spectral window) which are normally highly attenuated (see [Figure 2.5](#)). For the purposes of RFI identification it is easier to force the network to ignore the attenuation by fully removing it from the training set. When making predictions using the trained model, both of these features are removed and then added back to the flag mask after inference.

Once these features have been removed, it is useful to limit any extreme values. To

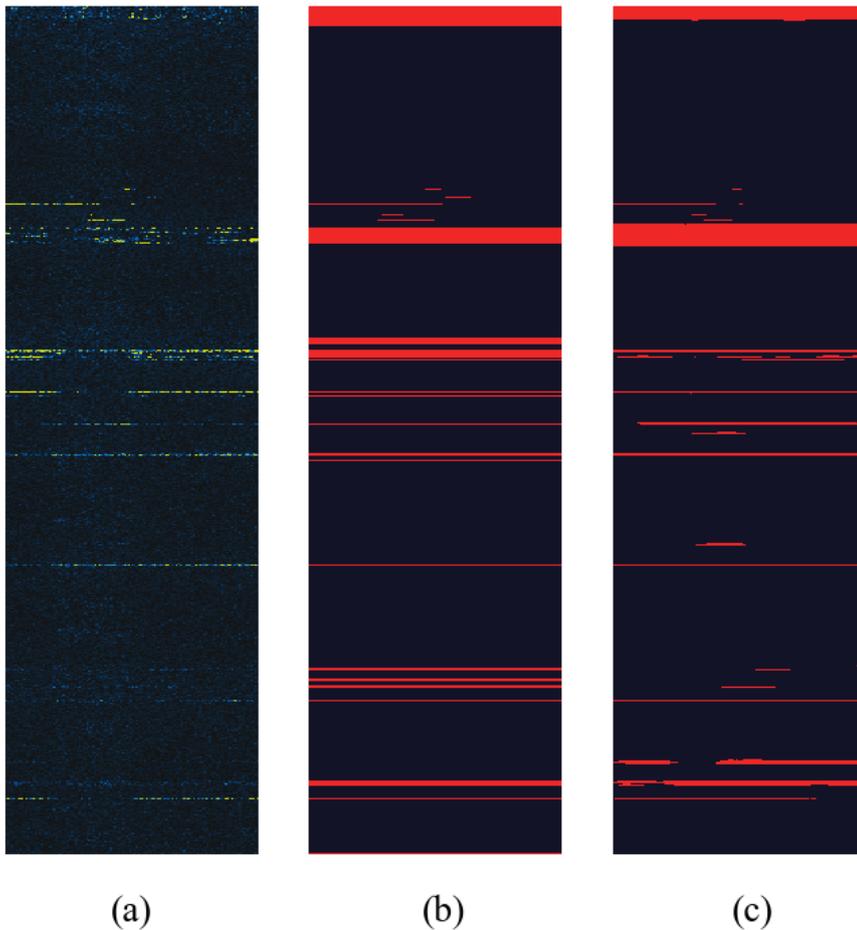


Figure 3.9: (a): An example of an image just before it is passed to the generator and just after it has been pre-processed. (b): The mask produced by manually flagging image (a), and the ground truth that would be used during training. (c): The mask produced by the trained generator operating on image (a).

counter the effect of these extreme values the absolute visibilities $|V(v, t)|$ are clipped (or, more specifically, winsorized) to the top 1% of values. Truncating extreme values in this way significantly reduces the mean of the distribution, and causes RFI to appear more prominent without disturbing the low-level, astronomical signal. However this method is maximally effective when the strength of RFI in an image is significantly greater than the strength of the astronomical signal, which is true in the vast majority of cases. There are rare cases where there is very little, or only very faint RFI in an image. In these cases, winsorizing the image can cause the dynamic range to be further reduced which may impact the efficacy of the network. Though, since these are rare cases, and there is only a small amount of RFI in these images anyway, this issue is not addressed any further.

After winsorizing, a two-dimensional time-frequency surface is fitted to the visibilities with the objective of removing the slowly varying astronomical signal, and leaving the RFI untouched. The process of computing this surface is very similar to that described

in section 2.2 of [Offringa et al. \(2010a\)](#), and is essentially a gaussian blur operation. The result of these pre-processing stages is shown in [Figure 3.9](#).

3.5.3 Results

Accurately assessing the performance of the GAN is extremely difficult due to the inherent problem in classifying real cases of interference: there is no definitive distinction between RFI and ‘not RFI’. Semantically this isn’t the case, but representing the semantic distinction in data is not realistic. Despite this inherent limitation, it is still useful to use the standard set of binary classification metrics, such as:

- **True Positive Rate:** Also known as the *recall*, the True Positive Rate (TPR) is defined as $TPR = TP/P$; the number of correctly identified (true) positives as a fraction of the number of actual positive cases.
- **True Negative Rate:** Also known as the *specificity*, the True Negative Rate (TNR) is defined as $TNR = TN/N$; the number of true negatives as a fraction of the number of actual negative cases.
- **Positive Predictive Value:** Also known as the *precision*, the Positive Predictive Value (PPV) is defined as $PPV = TP/PP$; the number of true positives (correct positive predictions) as a fraction of the number of predicted positives (all positive predictions).
- **F₁-Score:** The F₁-score is a commonly used performance metric in the field of machine learning, and is defined as the harmonic mean of the PPV and the TNR:

$$F_1 = \frac{2}{PPV^{-1} + TNR^{-1}} = \frac{2TP}{2TP + FP + FN} \quad (3.9)$$

where FP is the number of false positives; the number of incorrectly identified positive cases, and FN is the number of false negatives; the number of incorrectly identified negative cases.

These values are used to assess the performance of the trained GAN using manually flagged observations as test data, as described in [subsection 3.5.1](#). These test data are observations taken at a similar time, of the same target, and flagged using the exact same method as the two training data sets, to ensure a fair comparison. Using these flag masks

as the ground truth, the trained GAN achieves a TPR of 36.4%, and a TNR of 94.5% when operating on the test dataset. The GAN achieves an F_1 -score of 40.0% when operating on the test observation; significantly lower than what one would expect from an effective flagging method.

Test Simulations

It is difficult to draw solid conclusions from the results of operating on the test dataset as they are calculated using manually produced flag masks. As can be seen in [Figure 3.9](#), even these are not perfect. This is especially true when producing such a large amount of flag masks since, for reasons of practicality, the process of manually flagging radio data is a tradeoff between speed and precision. A more robust test can be conducted by using a simulated dataset for testing, where full and complete knowledge of RFI is possible. To this end, the GAN is also tested using simulated RFI images against AOFLAGGER ([Offringa et al., 2012a](#)); the current state of the art in fast, automatic flagging.

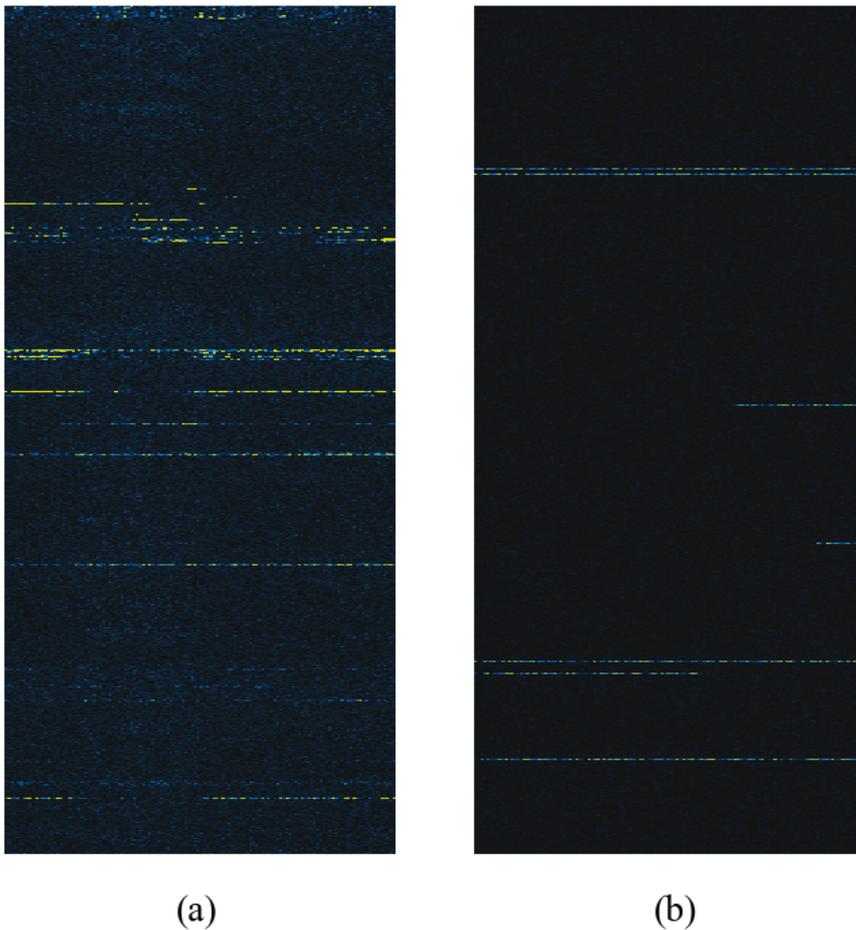


Figure 3.10: A comparison of a real time-frequency image (a) from the test dataset and a simulated image (b) generated via the method described in [section 3.5.3](#).

To create the simulated time-frequency images, the following method is used. An image containing only RFI is generated by first iterating over twice the chosen observation length. At each timestep, there is a $p_{\text{RFI}} = 0.01$ chance of an RFI signal beginning at that timestep, and a further $p_{\text{WB}} = 0.05$ chance of the RFI being wide-band interference, as opposed to narrow-band interference. A duration is randomly generated based on the RFI type: if it is wide-band it is likely to last only a few timesteps, if it is narrow-band it is likely to last for much longer. The strength of each RFI event is randomly drawn from a normal distribution with mean μ_{RFI} and $\sigma_{\text{RFI}} = 3$. Finally, a correction factor is applied across the RFI image such that all RFI signals appear to randomly fluctuate between full and $\frac{1}{10}$ power, to approximate the fluctuations seen in a real image. Finally, gaussian noise with $\mu_{\text{noise}} = 0$, $\sigma_{\text{noise}} = 0.1$ is added to this image to approximate the signal background after pre-processing surface subtraction. The simulated RFI image is rudimentary, but it does provide the basis for a fair experiment using each method. A comparison of the real and simulated images is shown in [Figure 3.10](#).

To fully explore the capability of the GAN over a wide distribution of dynamic ranges (controlled by μ_{RFI}), a $N = 500$ Monte Carlo experiment is performed to map the expected TPRs, TNRs, and F_1 -scores in the range $1 \leq \mu_{\text{RFI}} \leq 100$. At each iteration, 50 cutouts of size 128×1024 , like those in [Figure 3.10](#), are generated and flagged by both the GAN and AOFLAGGER. The results of this test are presented in [Figure 3.11](#). Before the sharp decline at $\mu_{\text{RFI}} \sim 10$, the GAN achieves a mean F_1 -score of 64.3%. It is difficult to determine exactly what causes the drop in performance around $\mu_{\text{RFI}} = 10$, though some possible explanations are discussed in [subsection 3.5.4](#).

3.5.4 Discussion

The process of flagging RFI has become extremely complex as astronomers are now almost entirely reliant on automated tools. These tools attempt to find a dividing line between RFI and ‘not RFI’, which can never be known with one hundred percent certainty, and attempts to move this dividing line must be done with great care as minor changes can produce dramatically different end results for key measurements such as source flux, and background noise. To avoid any RFI leaking through the flagging process, automated methods usually side with caution, so there will always be some number of false-positives, that, in general, will scale with the volume of RFI encountered in a single observation. This becomes problematic as RFI is more frequently encountered over time due to an in-

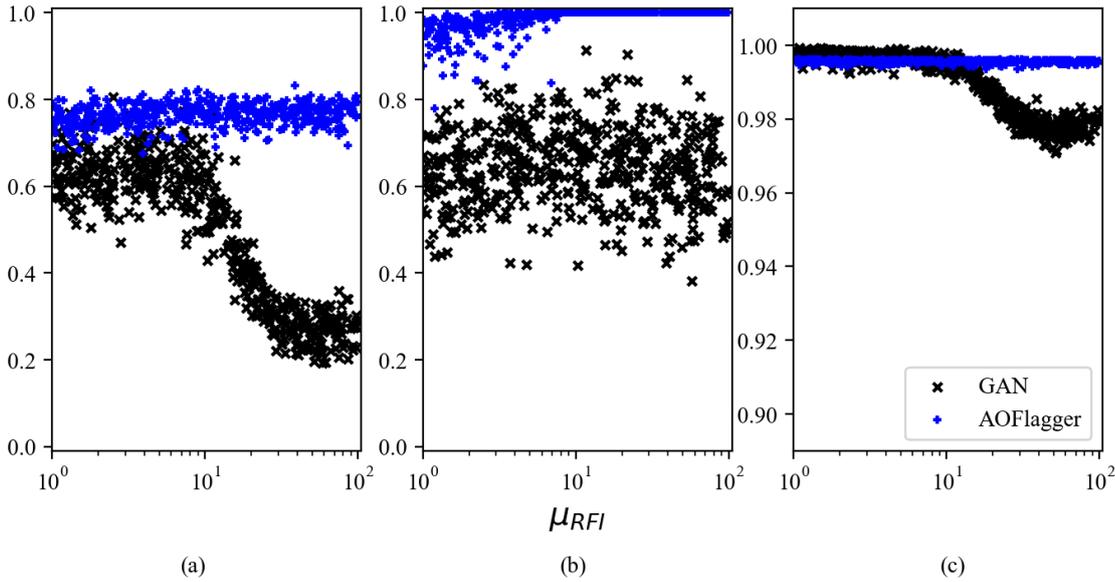


Figure 3.11: The results of the Monte Carlo experiment. A total of $N = 500$ tests were conducted, each operating on 50 simulated images. Each test involved flagging the batch of images with both the trained GAN and AOFLAGGER. Image (a) presents the F₁-scores, image (b) presents the TPRs and image (c) presents the TNRs.

creasing use of electronic communication, and we therefore see an increase in the amount of ‘not RFI’ being classified as RFI and being essentially wasted. For an individual experiment, this is not a primary concern. But for the entire output of an observatory this loss of data can be significant. For example, consider an observatory that observes on 300 of 365 days in a year. Each day a 12 hour observation is made, and 10% of each observation is corrupted by RFI. The flagging method has a modest false-positive rate of 2%, meaning that 11.8% of each observation is removed in practice, with 1.8% being ‘not RFI’. This means that, over the course of a single year, the observatory will have produced 64.8 hours of visibilities that are never used in an experiment. From this perspective, reducing the false-positive rate during flagging is an important area of research.

Another issue is that current techniques often require parameter optimisation, sometimes for a single observation. Having a wide selection of parameters is what allows these techniques to generalise across observatories and observations. However, the cost of not performing this optimisation step shows as a further increase in the false-positive rate, or an increase in the number of false negatives which can severely degrade the fidelity of the final sky map. The GAN mitigates this issue by effectively baking in the parameter optimisation to the process of constructing the training set, so rather than ensuring many different parameters, perhaps across many different methods, are optimal for a particular observation or set of observations one only needs to ensure two conditions: that

the training data is representative of future inputs, and that the training data is accurately describing RFI. If these two conditions are maintained, the GAN technique (in theory) performs optimally. However, there are currently few software tools that are designed to allow the production of a large amount of annotated visibilities quickly and efficiently, so accumulating a large and accurate training set from many different sources to allow for a generalised tool, without using simulations, is still an outstanding problem.

The measurements presented at the start of [subsection 3.5.3](#) immediately suggest two things. Firstly, the GAN is capable of identifying the *majority* of instances of RFI. This is further shown by the results of operating on simulated data (see [Figure 3.11](#)) that show a TPR of $>50\%$ at $\mu_{\text{RFI}} < 10$. Additionally, maintaining the low false positive rate seen below $\mu_{\text{RFI}} \sim 10$ is important for preserving as much valuable observing time as possible. Secondly, the results suggest a discrepancy exists between what the GAN classifies as RFI, and what the training data classifies as RFI. It must be remembered that the training data itself only offers an approximation of the RFI detected. Given the ability of this particular machine learning method to infer details from the training set rather than rigorously copying it, and that the flags contained in the training set are an approximation of perfect knowledge of RFI, it should not be expected that a well-trained GAN achieves a high TPR when operating on the test set. On the other hand, by simple inspection of the flag masks it creates it is obvious the GAN does not flag all instances of RFI. One example of this is that it repeats a similar error found in a number of thresholding methods where pixels nearby or directly adjacent to RFI that have a significantly lower power are not flagged while still technically being corrupted. Most methods will mitigate this problem by ‘extending’ the initial flag mask ([Offringa et al., 2012a](#)), and it is possible the GAN could also benefit from this.

The exact cause of the change in behaviour around $\mu_{\text{RFI}} \approx 10$ is unknown, but it is possible that the GAN has learned to flag pixels surrounding very strong instances of RFI (which is a common practice in manual flagging), leading to an increased false-positive rate when operating on the simulated instances of RFI which do not simulate this effect. Alternatively, the change in behaviour could be a combination of two other factors. Firstly, the clipping method described in [subsection 3.5.2](#) is based on percentiles rather than a hard cutoff. The number and strength of the pixels at the higher end of the visibility distribution will vary from image to image, therefore the threshold where the clip is applied will also change. Secondly, the e-MERLIN correlator produces visibilities with strengths capped at a certain limit, so strong instances of RFI will cluster around

an upper limit, which is then further winsorized as previously described. I hypothesize that this manifests in the training data as a maximum SNR ~ 10 in each image, which means that the GAN is effectively untrained on RFI instances exceeding this limit. In any case, it is unlikely that the network would regularly encounter RFI with $\mu_{\text{RFI}} > 10$ after normal pre-processing in a realistic scenario, and robust safeguards could be constructed to handle this eventuality.

More broadly, this experiment describes a technique with significant potential, but one that requires further refinement before it can be classified as reliable. A more diverse and expansive training set would likely have yielded better results, though realising this is limited by the practical difficulties of manually flagging large collections of visibilities. It is possible, by leveraging the ability of this particular architecture to infer properties of RFI, to train the network on visibilities flagged by some other automated method. However, care must be taken that the GAN does not learn to simply replicate the output of that other method. The decision to use manually flagged data in this experiment was driven by a desire to avoid this, and at the same time replicate the precision of human intervention, which is notoriously difficult with statistical tools. In any case, it seems clear from this and other work that human intervention in the training data should and will play a critical role in producing the most effective machine learning models of RFI (Agarwal et al., 2020; Vafaei Sadr et al., 2020; Pinchuk & Margot, 2022).

Potential systematic improvements include a more complex pre-processing stage, and modifications of the architecture itself. The overall objective of pre-processing is to emphasise RFI signals and de-emphasise all other signals, to make corrupted visibilities easier to identify. Therefore, any deterministic technique that satisfies this objective is viable. The architectures themselves may be improved by utilising *a posteriori* knowledge of RFI morphology, which is often quite consistent in a particular RFI environment. The U-Net architecture was designed with the specific intention of tracking cells in biomedical imaging, where the cells can be somewhat arbitrary in shape and size. RFI morphology is very different from this; it often extends narrowly along either the time or frequency dimension, and appears as two-dimensional box or circular shapes less often. Furthermore, the U-Net in this experiment was tasked with a binary classification problem, that is to identify RFI and ‘not RFI’. It’s possible that broadening to a multi-class problem might give better performance, where the U-Net is instead given the task of identifying and flagging different types of RFI such as narrow-band and wide-band RFI. One could even include slightly more abstract classes such as ‘unsure’ or ‘bordering RFI’ which could be

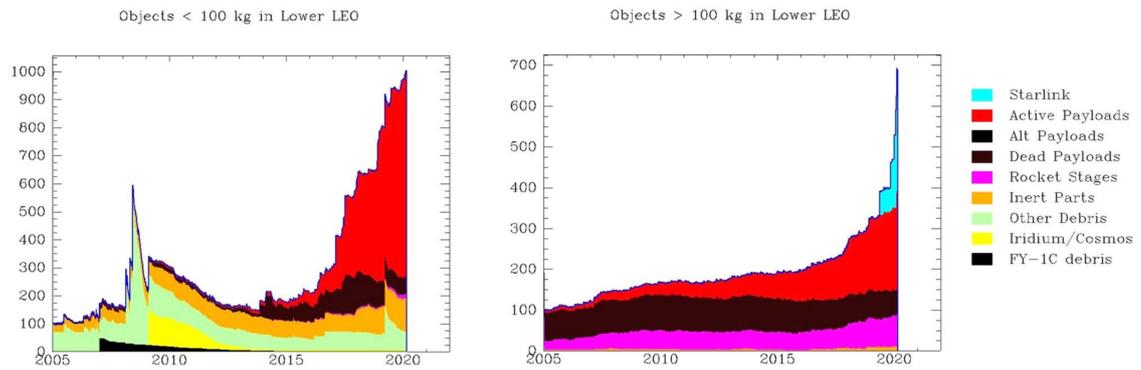


Figure 3.12: The number of objects in low earth orbit over time. These objects are divided between objects < 100 kg (left), and > 100 kg (right). Notice the spike in objects >100 kg due to the Starlink satellite program. Image credit: McDowell (2020).

useful information in a post-processing pipeline.

3.6 Dark and Quiet Skies

Looking forward, the greatest emerging threat to radio astronomy is the development of massive satellite networks. In approximately the last decade, there has been a significant increase in the volume of satellites in orbit around the Earth, and these satellites are capable of emitting very strong radio waves for communication. For maximum global coverage at all times, it has become a common objective for these satellites to be organised into ‘constellations’: hundreds of satellites distributed across the sky. Examples of these constellations in use today are the Starlink network, the OneWeb network, and the Iridium network.

The Starlink network in particular has been the greatest cause for concern so far and has prompted much of the action being taken by radio astronomers to manage satellite interference. While satellites have always been a concern, the Starlink launch program produced a massive increase in the volume of radio emitting satellites, such that ~50% of all objects >100 kg in low Earth orbit are reported to be Starlink satellites (see [Figure 3.12](#), [McDowell 2020](#)). As of March 2024 at least 5000 Starlink satellites have been deployed, but as many as 12000 are currently being proposed. These Starlink satellites have a typical visual magnitude of ≤ 6 , so wide-field optical surveys such as those that will be produced by the upcoming Vera C. Rubin observatory must find ways of handling the image artefacts introduced by Starlink and other satellites ([Tyson et al., 2020](#)).

Since many satellite networks operate outside the jurisdiction of any one state, it is difficult to regulate the transmission of signals from them as the ITU-R Radio Regulations

intend. Furthermore, because the cost of running these networks is high, the businesses that operate them have an incentive to stretch the limit of their obligations to ensure maximum monetary value. Negotiations between astronomers and the handful of private enterprises that operate these networks are ongoing, to try and come to a mutually acceptable arrangement.

In any case, the impact on radio astronomy will be significant, as a shared sky filled with radio transmitting satellites breaks the assumption that many RFI signals will not be correlated at both antennas in a baseline. It also breaks the assumption among many radio astronomers that RFI tends to originate from or near the horizon exclusively. Already there are some early warning signs of problems to come in radio astronomy, such as the data presented by [Di Vruno et al. \(2023\)](#) and [Grigg et al. \(2023\)](#). At the more extreme end of concern, there is the ever-present worry of an exponential increase in space debris caused by the increased volume of objects in low earth orbit; the so-called ‘Kessler Syndrome’ ([Kessler, 1991](#)). Furthermore, astronomers have proposed the construction of a lunar radio observatory as the radio environment of the surface of the moon allows radio observations at frequencies <30 MHz, something not possible on the Earth due to the RFI environment and the reflection of these low frequency waves off the ionosphere ([Jester & Falcke, 2009](#)). However, recent experiences with these satellite constellations has already prompted concern over the protection of potential lunar sites from interference ([Le Conte et al., 2023](#)).

These oncoming challenges only increases the need for more sophisticated methods of dealing with RFI, since even the most remote observatories on Earth will experience significant increases in the volume of RFI encountered. This will translate into greater amounts of data lost to flagging, and a reduction in effective observation time, which ultimately raises the noise level in any final data product. Finding and experimenting with new methods of precise RFI removal is, therefore, a critical area of research.

Chapter 4

The Faint Radio Sky

As discussed in [section 1.2](#), the understanding of the nature of the faintest radio sources has, naturally, evolved alongside the capabilities of our radio observatories. But only in the last few decades has this capability allowed observations at $S_{1.4\text{GHz}} \lesssim 100\mu\text{Jy}$, and only in the last decade or so has this been extended to $\lesssim 10\mu\text{Jy}$. Until these observations were possible, understanding the nature of the faint radio population was a challenge of theoretical knowledge, tested by effective simulations. Some attempts to model and simulate the radio sky include [Dunlop & Peacock \(1990\)](#); [Hopkins et al. \(2000\)](#); [Jarvis & Rawlings \(2004\)](#); [Wilman et al. \(2008\)](#) and [Wilman et al. \(2010\)](#); [Mancuso et al. \(2015\)](#), and [Bonaldi et al. \(2019\)](#). For the most part, these models have been successful at predicting the radio source counts above $\sim 100\mu\text{Jy}$ (see [Figure 4.1](#)), but there is not much consensus below this point as current instruments can be limited by confusion in the μJy region. Furthermore, adding more context to the counts by classifying the objects normally requires multi-wavelength information, which isn't always available for radio sources.

Nevertheless, understanding these faint radio sources is crucial to many areas of astronomy such as mapping the cosmic star formation rate ([Madau & Dickinson, 2014](#)). Radio surveys naturally sample higher redshift objects, so deep high resolution radio surveys also allow probing active and star forming galaxies at extremely high redshifts, which in turn provides evidence for current cosmological theory. This chapter will introduce the current theory of the faint radio sky and its two main constituents. It will also introduce the questions this experiment aims to explore, and their importance in the wider field of radio astronomy. This will be followed by a brief review of research into the radio sky at brightnesses below $\sim 1\text{ mJy}$.

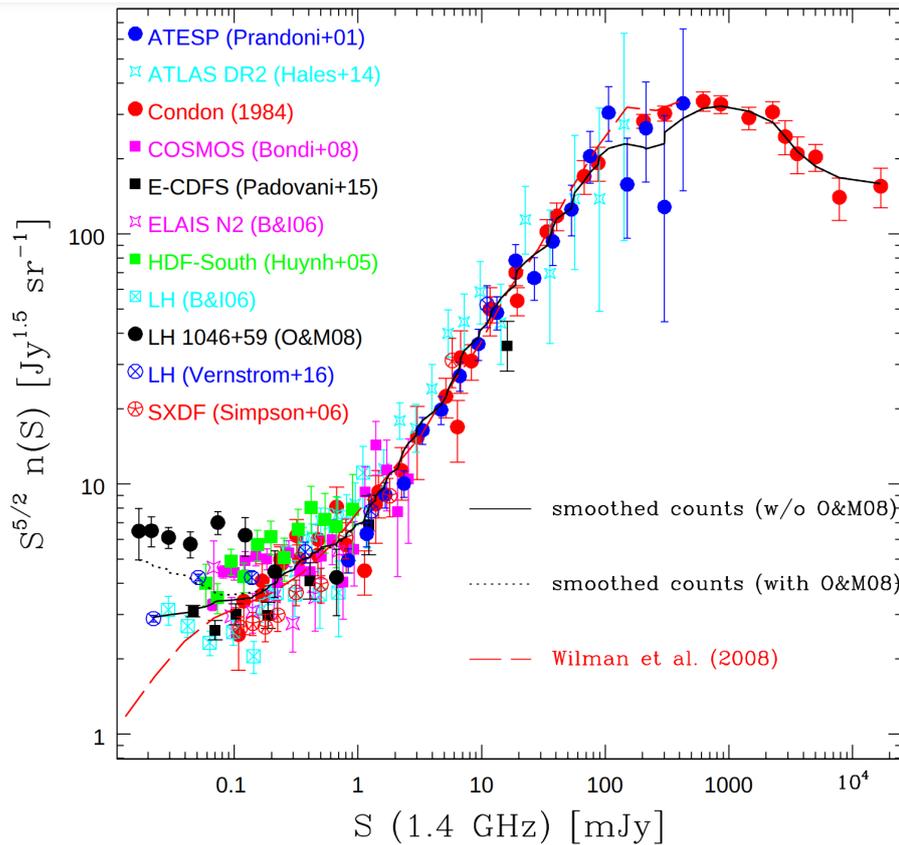


Figure 4.1: The 1.4 GHz source counts down to $\sim \mu\text{Jy}$ from both observational data and predictive models. Image Credit: [Padovani \(2016\)](#).

4.1 Active Galactic Nuclei

Of all the wondrous objects in the universe, black holes are certainly among the most impressive. To this day, they remain a constant source of enquiry and subsequent confusion, and are the ultimate testing ground of any theory of the universe. Until recently, black holes were still technically only theoretical (though universally accepted) since there existed no direct observational evidence for them. The EHT observation of the event horizon of the black hole in the center of M87, shown in [Figure 1.2](#) (see also [Event Horizon Telescope Collaboration et al. 2019](#)), was the first direct indication of a singularity, and all but cemented general relativity as a pillar of theoretical astrophysics. That observation further provided evidence for Supermassive Black Holes (SMBHs) as the gravitational core of the vast majority of, if not all, galaxies.

SMBHs are also theorised to be the mechanism that drives the observational phenomena known as AGN, since it is the only known mechanism that could account for the tremendous power necessary to produce the observed emission in all bands of the electromagnetic spectrum, including a significant portion of emission at radio frequencies. The radio emission from AGN comes from either powerful relativistic jets (see [Figure 4.2](#))

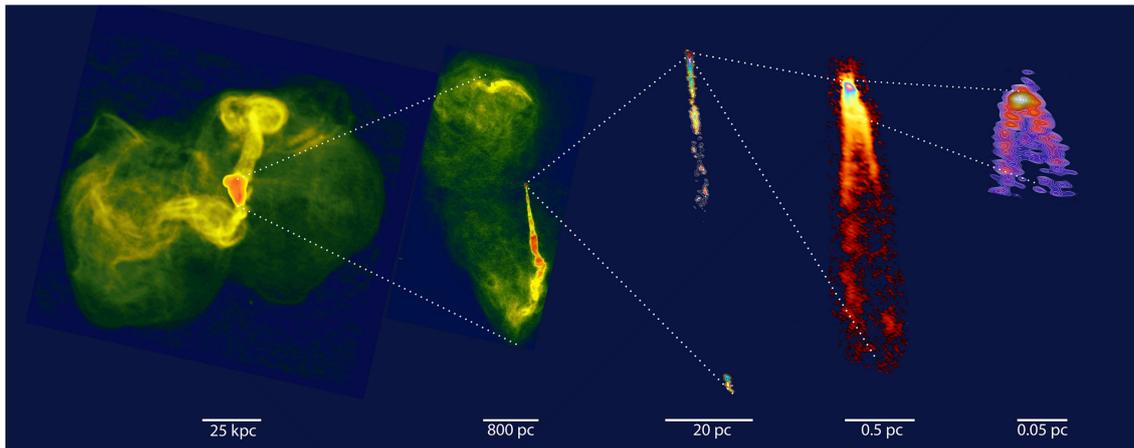


Figure 4.2: A progressively zoomed image of the radio galaxy M87, showing the structure of the jets at different scales. The instruments used for imaging are, from left to right: VLA 90 cm, VLA 20 cm, VLBA 20 cm, VLBA 7 mm, Global VLBI 3 mm. Image credit: [Blandford et al. \(2019\)](#)

that can extend well beyond the galactic scale up to \sim Mpc ([Blandford et al., 2019](#)), or from thermal radio emission originating from the black hole accretion disk. AGN and their associated phenomena form one of the two main components of the radio sky (the other being star-forming galaxies, see [section 4.4](#)). This, in addition to the advantages of observing radio waves that are undisturbed by dust, means that much of radio astronomy focuses on the observation of active galaxies.

The reasons to study AGN are numerous, though I will give some motivations here. Firstly, it is now understood that the evolution of the central SMBH and its host galaxy are inextricably linked ([Kormendy & Ho, 2013](#); [Heckman & Best, 2014](#)), which means that a good understanding of the mechanics of galaxy formation and evolution must come packaged with a similarly good understanding of AGN. Secondly, the underlying physics of the production of relativistic jets is still an area of heavy research, and while our understanding has very recently increased due to the expansion in observing capability, there still remains many questions that need clarifying such as the exact mechanism (or mechanisms) that generates the collimated jet, or the factors governing the occurrence of duty cycles that, in turn, play a major role in AGN feedback ([Fabian, 2012](#)). Pursuing a satisfactory governing theory of relativistic jets may reveal new physics, and all the implications associated with that. Similarly, it is well known that the singularity itself stretches the limit of even the most advanced astrophysical theory, and any attempt to resolve this long standing problem would be aided by a high-quality repository of knowledge surrounding active black holes and their environments.

4.1.1 The Unified Model of AGN

In the many decades since Maarten Schmidt ([Schmidt, 1963](#)) linked together the radio and optical components of 3C273 — discovering the first ‘quasi-stellar’ object or *quasar* — there have been numerous identifications of galaxies with a powerful radio component; what we now call AGN. These discoveries have often been accompanied by more questions than answers, and as such the literature surrounding AGN is filled with a variety of labels and examples of different, and sometimes extreme, cases of active black hole systems. The confusing taxonomy of AGN is an unfortunate byproduct of our knowledge of AGN often significantly preceding our understanding as, in an attempt to achieve the latter, various classification schemes have been formed to highlight similarities between the many known cases of AGN. Rather than exhaustively detail all these different categories and their observational characteristics, I instead refer the reader to the work of [Padovani et al. \(2017\)](#), which seeks to simplify the ‘zoo’ of AGN. This section will instead focus on a jargon-free, simpler introduction to the now widely accepted model of unification for AGN, as shown in [Figure 4.3](#).

It is generally assumed that at the center of every galaxy lies a SMBH, defined by [Netzer \(2015\)](#) as a black hole with mass roughly exceeding $10^5 M_{\odot}$. At some point during the life of that SMBH, mass accretion on to the black hole can be triggered by dynamical events, leading to the formation of an accretion disk $\lesssim 1\text{pc}$ wide. The system is highly energetic, causing a significant increase in luminosity across the electromagnetic spectrum that can outshine the entire host galaxy. When an SMBH is in this state, it is ‘active’ and is therefore referred to as an Active Galactic Nucleus, or AGN. There are thought to be several key components of an AGN, which the following sections will describe.

The Singularity

Black holes, as predicted and described by general relativity, are produced when the escape velocity of some gravitational body exceeds the speed of light. Within the radius known as the *Schwarzschild radius*, photons cannot even escape the hole’s influence, and so the object appears ‘black’. The existence of black holes remains a constant focal point of all theories of the universe, since they are one of, if not *the* most extreme example of physics in nature. Most black holes are formed at the end of the lifecycle of a very massive ($\gtrsim 60\text{-}80 M_{\odot}$) star, but that mechanism alone is insufficient to explain the presence of the populations of SMBHs observable today. Therefore, there are a number of

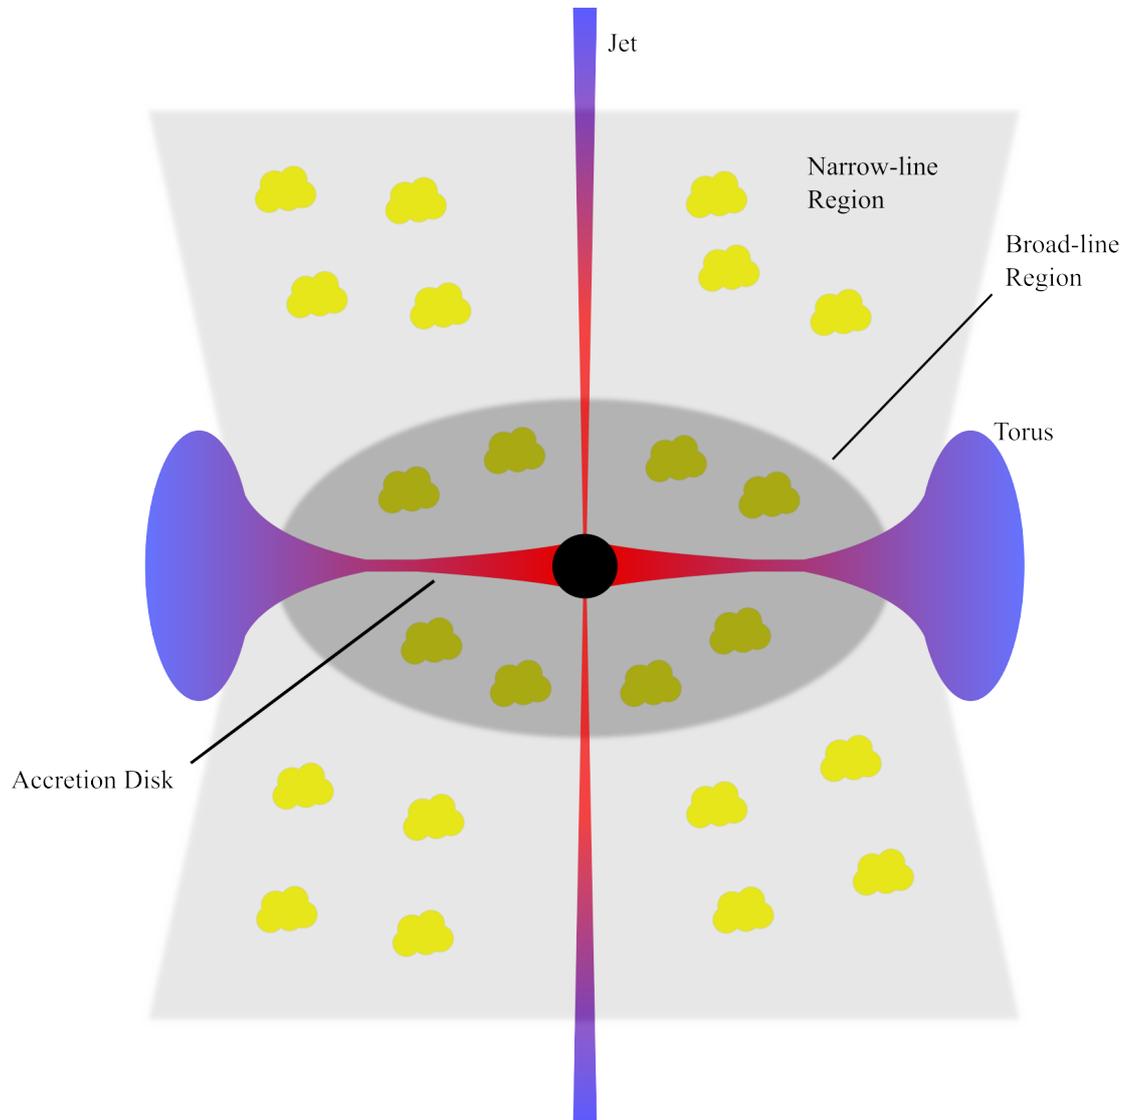


Figure 4.3: An approximate schematic of the anatomy of a typical AGN viewed side-on (not to scale). The black hole lies at the center, surrounded by an accretion disk that flows out into the dusty torus. In a region close to the black hole itself lies the broad-line region, and further out in a cone extending from the core of the nucleus lies the narrow-line region. Relativistic jets are ejected out of the poles extending out well-beyond the central structure.

explanations for the presence of SMBHs (Inayoshi et al., 2020), such as a now (likely) extinct population of massive ($\gtrsim 100 M_{\odot}$) stars known as Pop III stars: the first generation of stars to form after the big bang (Madau & Rees, 2001). Through a combination of mass accretion and major mergers, it is possible that these early massive stars seeded the SMBHs of today. Very recent JWST observations may show signs of Pop III stars in the halo of a very high redshift galaxy hosting an SMBH at $z = 10.6$ (Maiolino et al., 2023; Maiolino et al., 2024).

The Accretion Disk

Inevitably mass will be drawn in towards the black hole from a number of sources (Rees, 1984), which collects in the form of a disk surrounding it. The loss of angular momentum that drives the accretion is caused by a combination of disk viscosity, turbulence, and electromagnetic interaction with the black hole itself (Urry & Padovani, 1995). The accretion disk is also extremely hot, and therefore emits strongly in the X-ray and UV bands. The accretion of matter on to the black hole is one of the main sources of black hole growth, and the characteristics of the accretion disk is one of the primary factors in determining the characteristics of (or the existence of) a jet (Blandford et al., 2019). Accretion disks are a common occurrence not just for black holes, so for a comprehensive review see Pringle (1981) and Abramowicz & Fragile (2013).

The Torus

Some AGN will have a circularly symmetric donut-like cloud of colder dust and gas surrounding them, known as a ‘torus’ or sometimes a ‘dusty torus’. An exploration of tori and their formation is given by Rees (1984), but in general they occur when conditions change with increasing distance from the black hole such that internal forces in the disk begin to compete with gravity, and the mass of the disk balloons out slightly to form a thicker structure. How exactly the torus then connects to the host galaxy itself is still a subject of ongoing research (Netzer, 2015)

The presence of a dusty torus is one of the key factors of the unified model, as the torus obscures the inner region of the AGN when it is viewed side-on. It is now believed that at least some classes of AGN are a product of viewing angle, specifically when viewing directly through tori which obscure the Broad-line Region (BLR), or when viewing at small angles to the jet axis which exposes the BLR. The presence of a torus is another

reason why radio astronomy is a powerful tool for exploring AGN, as radio observations are not as susceptible to the dust in the torus as other modes of observation, providing a less biased view of the inner regions of AGN.

The Jets

With modern VLBI systems we now know that relativistic jets exist among many AGN, but the mechanism that produces them is still not well understood. The accelerated matter clearly originates from the accretion disk, but it is unknown at what point the matter is ‘confined’ to such a small area and ejected in a strongly collimated flow. It has been demonstrated that strongly magnetized outflows self-collimate under certain conditions (Heyvaerts & Norman, 1989) but this mechanism alone cannot explain the observed degree of collimation. A promising explanation involves magnetic interaction with the accretion disk itself to generate the initial jet which then self-regulates at large distances (Globus & Levinson, 2016). In any case, testing these theories through observation is difficult because of the extremely small spatial scales, so simulations of jets remain the best tool for testing theoretical models (e.g. Komissarov et al. 2007; Tchekhovskoy et al. 2011; McKinney et al. 2012)

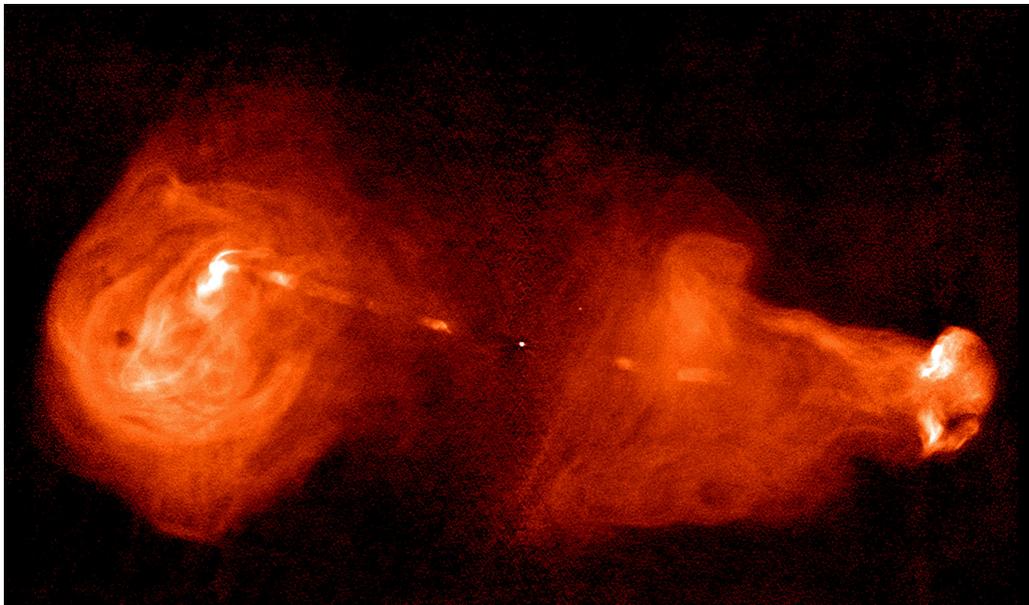


Figure 4.4: Dual radio lobes in the radio galaxy 3C353. The galaxy is viewed from the side, and the jets originate from a small point at the center of the image which is the AGN itself. Image credit: NRAO

The strongest jets will propagate to well-beyond the host galaxy itself where they may dissipate into the lobe-like structure seen in many radio galaxies, as in Figure 4.4. As mentioned previously, the role of jets in the evolution of host galaxy — the AGN feedback

mechanism — appears to be a critical factor in the development of the host galaxy (Fabian, 2012; Kormendy & Ho, 2013), and perhaps even its surroundings (Heckman et al., 1990; McNamara & Nulsen, 2012).

The Broad and Narrow-line Regions

The broad spectral-lines found in some AGN spectra are thought to originate from a region inside the boundary of the torus close to the black hole, specifically from clouds of hot, fast-moving gas emitting at optical and ultraviolet wavelengths (Urry & Padovani, 1995). This emission is obscured when the AGN is viewed through the torus but observable when viewed at inclinations that allow peering inside the inner region. Further out, slower, colder clouds of gas exist within the ‘ionisation cones’ of the AGN. These clouds are photoionised by radiation from the nucleus itself and produce narrow emission lines that are generally observable at all inclination angles. The seemingly random presence or lack of broad emission lines drove much of the confusion surrounding different classes of AGN, and is one of the main problems solved by the unified model.

4.2 AGN at Radio Frequencies

Because of their ability to observe extremely small spatial-scales, radio observatories (VLBI observatories in particular) are an attractive method of investigating AGN and their environments. Conveniently, the primary mechanism that dictates the dominant non-thermal component of continuum radio emission below ~ 30 GHz in AGN is the same mechanism that produces radio emission in star-forming galaxies: synchrotron radiation.

4.2.1 Synchrotron Radiation

The emission mechanism of synchrotron radiation, also known as magnetic bremsstrahlung radiation, is well understood, and neatly explains much of the emission seen from extragalactic radio sources. In many of these sources, ultra-relativistic particles travel parallel to a magnetic field, and are therefore accelerated into a helical motion about the field lines. In this state, they will emit radiation over a wide range of frequencies.

To explore this in more detail, consider particles moving at relativistic speeds (i.e. $\beta = v/c \approx 1$) in the presence of a magnetic field \mathbf{B} as shown in Figure 4.5. Particles in this state will make one full rotation every $T = 2\pi/\omega_B$, where the gyration frequency ω_B

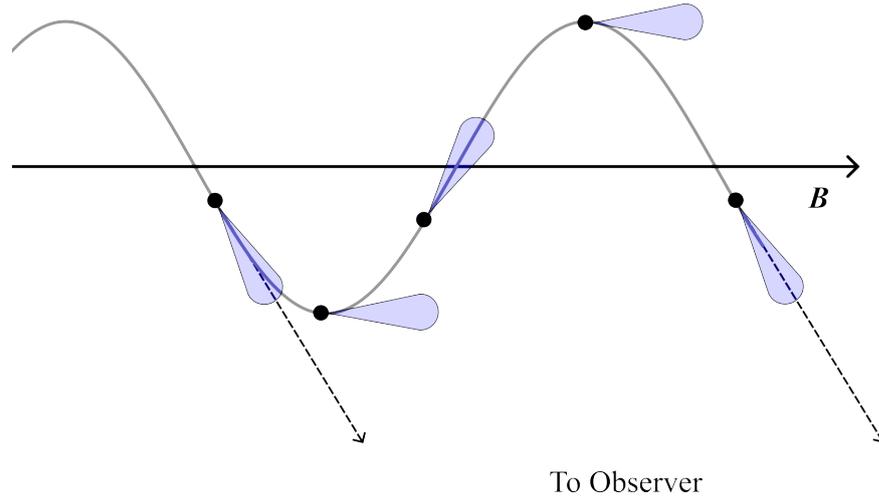


Figure 4.5: A diagram of ultra-relativistic electrons orbiting about a magnetic field line indicated by B . Blue cones indicate the beaming effect introduced by their speed. At a point on each orbit round the field line, the cone will be pointing in the direction of the observer, causing the observer to see pulses of light at regular intervals. The time between these pulses is governed by the properties of the particles.

is defined by the following from Rybicki & Lightman (1979):

$$\omega_B = \frac{qB}{\gamma mc} \quad (4.1)$$

for a particle of charge q , gyrating about a magnetic field of strength B , with mass m . Here, γ is the Lorentz factor $\gamma = 1/\sqrt{1-\beta^2}$, and c is the speed of light. Because these particles are traveling at highly relativistic speeds, a *beaming* effect is introduced, such that the particles radiate light in the forward direction with emission profiles indicated by the blue cones in Figure 4.5. Therefore, an observer would see pulses of light with a finite duration in regular intervals derived from the gyration frequency ω_B . Fourier transforming this signal yields the synchrotron spectrum of a single particle.

In astronomy, the source of synchrotron radiation is mostly from relativistic electrons, though not all electrons will have the exact same properties within an overall system (such as an AGN jet). It is common for studies of synchrotron emission to assume a power law distribution of electron energies over a small range such that:

$$N(E)dE \propto E^{-\alpha} dE \quad (4.2)$$

Considering the integrated light from such a distribution of electrons gyrating about a magnetic field gives a wide ranging spectrum with characteristics tightly bound to the

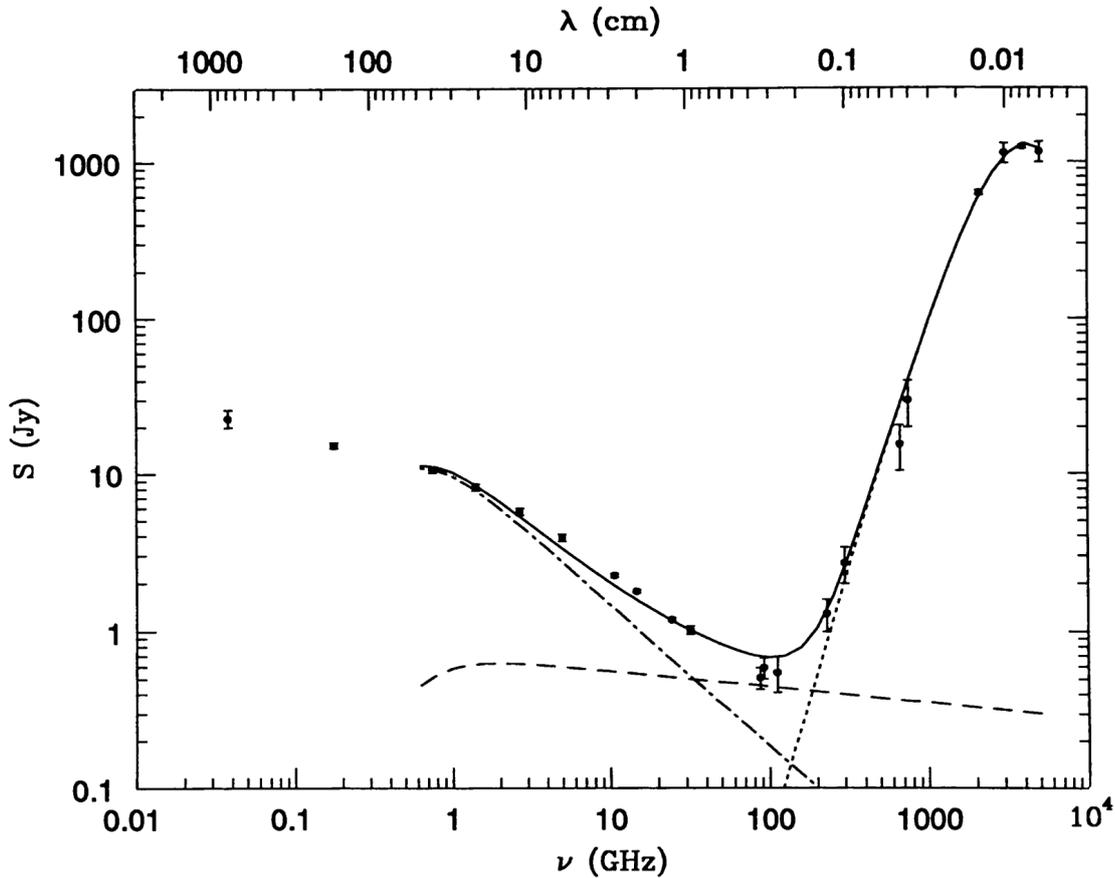


Figure 4.6: The radio spectrum of M82. At ~ 100 GHz the thermal dust emission (dotted line) decreases and the synchrotron spectrum (dash-dotted line) takes over with a smaller contribution from a free-free component (dashed line). Below this limit, the spectrum can be approximated by the spectral index Γ . Image credit: [Condon \(1992\)](#)

properties of underlying system. At frequencies above ~ 100 GHz, thermal dust emission is usually strong enough to outshine synchrotron radiation, but that component decreases at low frequencies. As it decreases, nonthermal processes — assumed to be mostly sourced from the synchrotron mechanism — become relatively stronger, until taking over at frequencies in the \sim GHz range and below. A prime example of this is the radio spectrum of M82, shown in [Figure 4.6](#)

4.3 Radio-loud and Radio-quiet AGN

Soon after the discovery of the first quasars, objects with a similar optical appearance were found to have no associated radio emission. We now know that there is in fact an associated radio source with these objects, but it is approximately one-thousand time less powerful and was below the detection limit at the time. This led to a distinction between ‘radio-loud’ AGN and ‘radio-quiet’ AGN. The separation is traditionally made in terms

of radio power and a radio-to-optical flux density ratio, R . Radio-quiet AGN are typically thought to have radio powers $L_{1.4\text{GHz}} \lesssim 10^{23} \text{WHz}^{-1}$ and radio-to-optical ratios $R \lesssim 10$. Less than 10% of all known AGN are radio-loud, with the vast majority being radio-quiet. A comprehensive review of the two populations is given by [Padovani \(2016\)](#) and [Padovani \(2017\)](#), both of which I summarise here.

The differences between the radio-loud and radio-quiet AGN is not just restricted to radio brightness; the dominant source of emission across the spectrum is different. Radio-loud AGN attribute the majority of their bolometric luminosity to strong jets emitting non-thermally, while radio-quiet AGN derive most of their energy output from thermal (free-free) emission associated with the accretion disk. Naturally, this means that radio-loud AGN tend to emit over much larger spatial scales than the quiet AGN, since the jets can extend far beyond the host galaxy. There also appears to be a difference in the host galaxy between the two: radio-quiet AGN tend to be located in late-type galaxies, while radio-louds tend to be present in early-type galaxies. The origin of *radio* emission in particular in the radio-quiet AGN is still being researched, but recent studies support the hypothesis that the majority of radio emission is attributable to star-formation ([Padovani et al., 2011](#); [Bonzini et al., 2013](#); [Panessa et al., 2019](#)). This correlates with the observation that many radio-quiet AGN are hosted in late-type galaxies which display more star-formation.

The recently developed ability of radio observatories to probe the $\sim \mu\text{Jy}$ radio population has shown instances of radio-loud AGN that violate the radio flux density limits described earlier ([Bonzini et al., 2013](#)). In fact, the available evidence indicates that a more accurate distinction can be made by the presence or lack of a strong, relativistic jet, as this has major implications for the source brightness across the spectrum. A more accurate and helpful terminology therefore, would be to follow that of ([Padovani, 2017](#)) which instead describes a population of ‘jetted’ and ‘non-jetted’ AGN. One of the strongest ways of distinguishing between these two classes is by examination of the famous FIR-Radio correlation ([van der Kruit, 1973](#); [Condon, 1992](#)). The FIR-Radio correlation is a trend indicative of star formation, and as jetted AGN are not dominated by star formation they do not follow it, while non-jetted AGN do. Now that the period of confusion around AGN classes appears to be passing, other questions now need answering such as: Why are there so few AGN with strong jets? What is the distribution of jetted and non-jetted AGN at fainter radio fluxes, and what can that tell us about their evolution? The latter experiment is only possible with modern radio telescopes, and is partly what this thesis aims to explore.

4.4 Star Forming Galaxies

All galaxies undergo some level of star formation, but there are a subset of galaxies with significantly greater levels of star formation, usually as a consequence of a merger or tidal interaction. Galaxies with a greater rate of star formation will naturally exhibit an above average rate of core collapse supernovae. Stars with mass $M \gtrsim 8M_{\odot}$ will eventually produce a supernova as part of its lifecycle (Smartt, 2009), and the supernova phenomenon is well-documented at almost all points on the spectrum by now, including at radio frequencies.

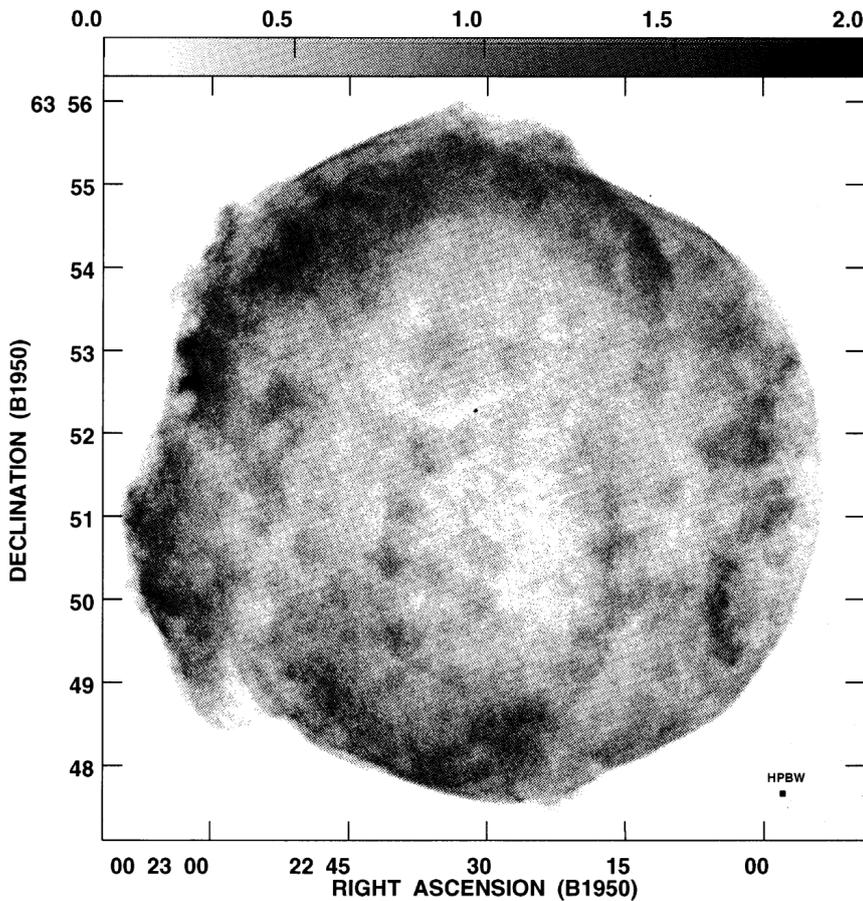


Figure 4.7: An image of the Tycho supernova remnant observed by the VLA at 1.4 GHz. At the centre, a point of emission indicates the remnant core of the star. Image Credit: Reynoso et al. (1997).

The classic view of a supernova is a singular powerful explosion that produces a spike in emission across the electromagnetic spectrum, falling away over a timescale $< 1\text{yr}$ (Woosley & Weaver, 1986). Naturally, this brief spike in energy output applies to radio emission as well, but radio telescopes are able to view a secondary effect related to the material ejected from the supernova remnant, for many *thousands of years* after the supernova occurred. The supernova will eject material from the star into the interstellar

medium and the interstellar magnetic field. There, shock fronts will form, and the particles will emit synchrotron radiation in the same way as described in [subsection 4.2.1](#), by interacting with the local interstellar magnetic field. An example of a supernova remnant seen at radio frequencies is given in [Figure 4.7](#). The flux density of the synchrotron emission decays over time as the particles diffuse through the galaxy, but the decay is slow enough such that a local $z \sim 0$ remnant is resolved and observable for much longer than the initial flash.

In galaxies teeming with star formation and the subsequent supernovae, there is a significantly increased density of relativistic electrons propagating through the interstellar medium, which means that these star forming galaxies are observable with radio telescopes even at high z , and make up one of the two primary components of the faint radio population. Tracking the strength and number of these star-forming galaxies as the universe evolves is a direct indication of the cosmic star formation rate.

4.5 A Dual Population

As a result of the mechanisms described in [section 4.2](#) and [section 4.4](#), the radio sky below the \sim mJy limit is dominated by SFGs and AGN. Distinguishing between these two at $z \sim 0$, is not so much of an observational challenge since it is trivial to determine the source of the radio emission by simply locating it within a galactic disk; if it resides in the center of a galaxy it is an AGN, if it resides in the disk it is likely a supernova remnant (or at least, *not* associated with an AGN).

However, at higher z where many objects are unresolved to most instruments this is no longer possible. Most studies of the faint radio sky therefore utilise multi-wavelength imaging and spectroscopy to classify objects, in particular at infrared wavelengths to search for signs of star formation, and X-rays to search for signs of black hole activity. Introducing multi-wavelength information is not easy, and complications can arise with the mixing of resolving power, systematic uncertainties, source matching, SED fitting, selection effects, and so on. A few notable selection effects include:

- As discussed in [subsection 4.1.1](#), radio powers from AGN can be affected by the orientation of the BH with respect to the observer. This also true for other observing bands as the torus can obscure the central region when viewed side-on
- VLBI arrays are predisposed towards detecting compact radio emission (AGN) due

to their lower surface brightness sensitivity

- SFGs are harder to detect at higher z as the more diffuse emission falls below the sensitivity limit of radio telescopes more easily than compact AGN

Many studies of the faint radio population have been conducted over the last two decades or so such that its properties are well observed above $\sim 100\mu\text{Jy}$. Beneath this limit information is more scarce. The source counts are in agreement from a broad range of surveys (Smolcic et al., 2015), but the relative distribution of AGN to SFGs is explored, but not certain. Figure 4.8 provides an estimate of the relative number of each class within the Extended Chandra Deep Field South (E-CDFS) study conducted by Padovani et al. (2015).

The picture at the μJy limit is just beginning to come together, as new surveys with upgraded instruments are being conducted. Vernstrom et al. (2016a) conduct a deep survey of the Lockman Hole with the VLA, reaching a detection limit of $5\mu\text{Jy}/\text{beam}$. From 558 sources they find that 10% are jetted AGN, 28% are non-jetted and 58% are SFGs. Muxlow et al. (2020) have conducted a deep survey of the GOODS-N field down to $\sim 1\mu\text{Jy}$, but have not yet determined classifications of sources. They intend to investigate the AGN and SFG populations through an analysis of morphology at different spatial scales. They combine data from both the VLA and e-MERLIN to achieve both high sensitivity and excellent resolving power. This also means that by weighting the baselines differently, they can explore progressively higher resolution images of sources, to see if the radio emission remains compact or becomes more diffuse, indicating AGN or star-formation respectively. If successful, this would sidestep the difficulties introduced by relying on multi-wavelength data for classifications.

Reaching such deep sensitivities is not easy with the instruments of today, as they require large amounts of observing and processing time. Even if a survey can reach those deep flux bins there may be issues of source confusion. This is made even worse by the RFI environments of many modern observatories which introduce numerous sensitivity and calibration problems which must be overcome. Conducting deep radio surveys is one of the key motivations for the RFI mitigation and identification measures discussed at length in chapter 3, which will become even more important as future facilities such as the SKA, which are significantly larger and more advanced, will be able to explore much deeper.

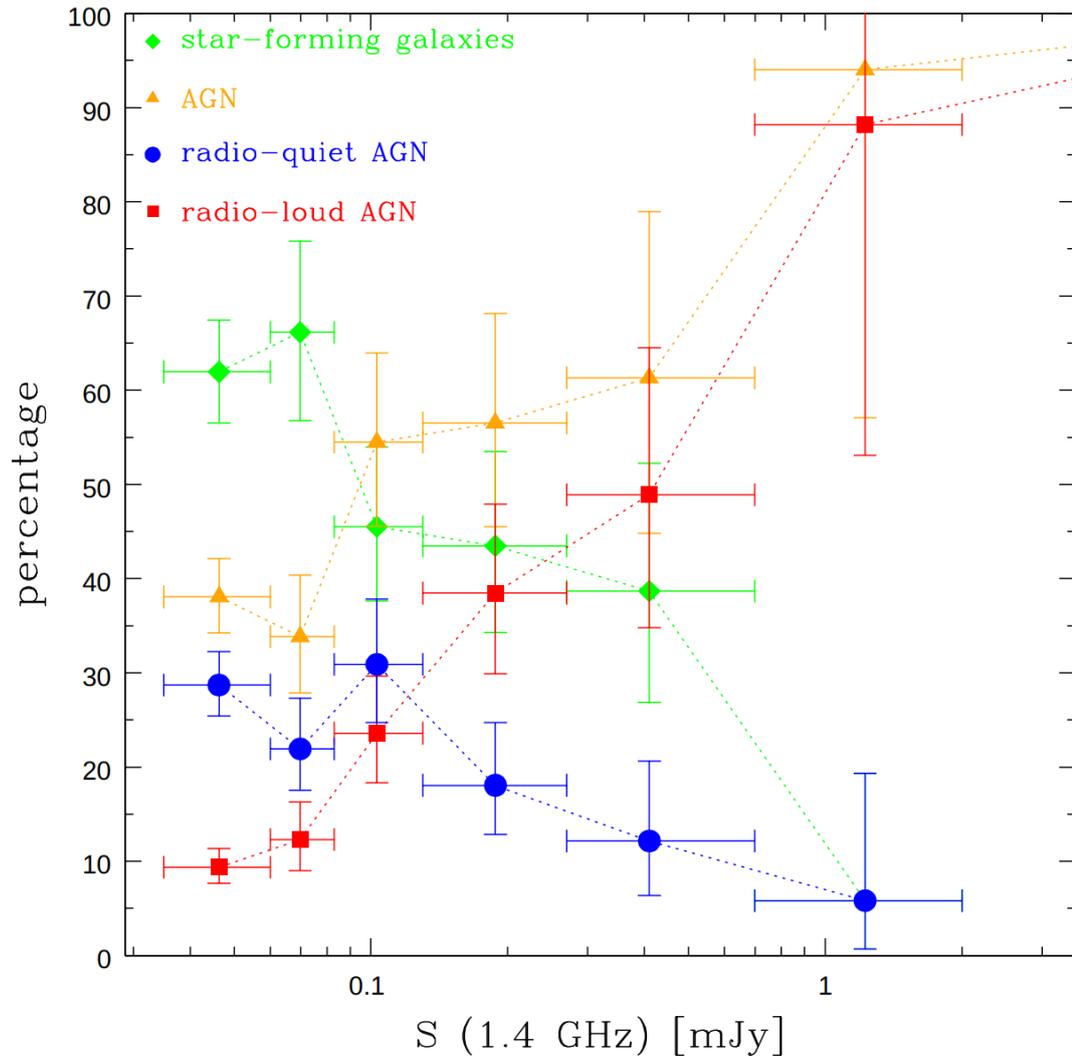


Figure 4.8: The relative distribution of the different classes of radio sources in the E-CDFS field. Jetted AGN dominate the sky at $\gtrsim 200 \mu\text{Jy}$, whereas below this limit SFGs and non-jetted AGN take over. Specifically, SFGs seem to dominate over AGN at these faint flux densities, an observation which is reinforced by the deep study conducted by [Vernstrom et al. \(2016a\)](#). Image Credit: [Padovani \(2016\)](#).

Chapter 5

Investigating a Binary AGN Candidate

The following chapter presents an updated version of a paper published in *Monthly Notices of the Royal Astronomical Society: Letters* in May 2022, about the potential discovery of a bound pair of AGN with \sim kpc scale separation. The paper title is ‘Revealing dual radio sources in a sub-kpc-scale binary active galactic nucleus candidate’. The authors are myself (J. E. Brooks) and M. K. Argo from the Jeremiah Horrocks Institute, Preston, UK; Hojin Cho and Jong-Hak Woo from the Department of Physics & Astronomy at Seoul National University, Republic of Korea; Taehyun Jung from the Korea Astronomy and Space Science Institute and the Department of Astronomy and Space Science at the University of Science and Technology, Korea; and N. Wrigley from the Jodrell Bank Centre for Astrophysics at The University of Manchester, UK. Then, [section 5.2](#) presents the results of a follow up study of the object by the EVN. All of the following material in this chapter was written by myself and is solely my own work. Furthermore, all of the data processing and analysis related to the e-MERLIN and EVN radio data presented in this chapter was performed by myself and is solely my own work. The valuable contributions from the co-authors listed above were in the form of suggestions and verbal discussions.

5.1 Revealing dual radio sources in a sub-kpc-scale binary active galactic nucleus candidate

5.1.1 Abstract

We present new VLA and e-MERLIN imaging of a sub-kpc-scale binary active galactic nuclei (AGN) candidate. Two unresolved radio sources of similar luminosity around 10^{22} WHz^{-1} are identified in ~ 35 hrs of e-MERLIN 6 cm imaging. These radio sources have an angular separation and position angle of $0.19 \pm 0.06''$ and PA $22 \pm 10^\circ$, corresponding to a projected separation of 0.95 ± 0.29 kpc at the epoch of the source. Our results suggest the presence of a kpc-scale active black hole pair hosted by two galaxies in the late stage of a merger at $z = 0.35$. This work follows (Woo et al., 2014) which presented two optical sources with a similar separation and position angle, and a velocity separation of 200 km s^{-1} . Our target adds to the currently limited sample of close-separation binary AGN that will aid in understanding the frequency of mergers and the stochastic gravitational wave background.

5.1.2 Introduction

The exact role of the central supermassive black hole (SMBH) in the evolution of a galaxy is still unclear, but it is clear that mergers between galaxies can have a major impact on the life cycle of a galaxy (Kormendy & Ho, 2013). During a merger of two massive gas-rich galaxies, accretion onto one or both of the central SMBHs can be triggered producing a bound pair of active galactic nuclei (AGN). This merger proceeds in a series of stages before coalescence (Begelman et al., 1980; Colpi, 2014; Komossa & Zensus, 2016). Firstly, the two galaxies begin their interaction via momentum losses to dynamical friction. This continues until the binary separation is approximately 1 pc, then the pair begins to secure itself and the binary ‘hardens’. It is here where our understanding becomes slightly more murky, as the processes that take over dynamical friction to drive the merger may not operate on timescales that agree with observation. Early models even suggested that the binary can stall for longer than a Hubble time; commonly referred to as the ‘final-parsec’ problem (Khan et al., 2011, 2013; Vasiliev et al., 2015; Holley-Bockelmann & Khan, 2015). Our understanding picks back up again as the binary approaches coalescence and efficient gravitational wave emission takes over to drive the merger to its conclusion. Here, a mas-

sive burst of gravitational energy is emitted that will be clearly detected by observatories like the Laser Interferometer Space Antenna (*LISA*, [Danzmann et al. 2017](#)).

Studying these mergers is therefore critical to the future of the burgeoning field of gravitational wave astronomy, which can in turn deliver important insights to the entire field of astrophysics. In addition to being a strong source of emission for even our early gravitational observatories, mergers are thought to contribute to a stochastic background signal known as the gravitational wave background (GWB, [Goulding et al. 2019](#)). As Pulsar Timing Arrays continue to operate at lower and lower frequencies, so the interest in this field grows ([Manchester et al., 2013](#); [Verbiest et al., 2016](#)), and decomposing the relative contributions of various mechanisms to the GWB becomes an active area of research. Quantifying the contribution of binary AGN to the GWB necessarily requires a large sample of mergers at varying stages to approximate the rate of gravitational wave events. Furthermore, this sample could also be used in studying the galaxy merger process beyond $z \sim 0$ where major mergers seem to have at least partly driven changes in galaxy morphology ([Conselice, 2014](#)). Therefore, compiling a sample of close separation binaries provides an important tool for testing future theories and models of galaxy evolution and cosmology. Currently, this sample is quite limited (see e.g. [Rubinur et al. 2018](#); [De Rosa et al. 2019](#), for a review), and the sample of \sim kpc scale binaries is especially small (e.g. [Komossa et al. 2003](#); [Rodriguez et al. 2006](#); [Fu et al. 2015](#); [Kharb et al. 2017](#); [Goulding et al. 2019](#)).

The origin of the radio emission in radio faint AGN, like those presented in this article, is an ongoing topic of research (see [Padovani \(2016\)](#) for a review). It has been shown that there is a loose correlation of jet power with radio power, such that radio-loud AGN contain strong jets and radio-weak AGN contain weak or non-existent jets, and the origin of emission is likely dominated by some other mechanism operating in the host galaxy or the local environment of the central engine ([Bonzini et al. 2015](#); [Panessa et al. 2019](#)). This other mechanism is usually star formation, therefore radio faint populations are mostly composed of AGN and star formation dominated sources ([Padovani et al., 2015](#)). As is often the case, this is not a global rule. An example of one such exception is the strongly lensed system studied by [Hartley et al. \(2019\)](#), which shows dominant jet emission in a sub-mJy source on scales below a parsec. Identifying the dominant origin of emission in radio faint AGN often requires multi-wavelength observations and a combination of long and short baseline interferometers in the radio.

This article presents new evidence for the binary AGN nature of the object SDSS

J132323.33-015941.9 using radio imaging from both the Karl G. Jansky Very Large Array (VLA) and the Multi-Element Radio Linked Interferometer Network (e-MERLIN). We will build off the findings of Woo et al. (2014) who first identified the object as a candidate sub-kpc-scale binary AGN through Hubble Space Telescope (HST) imaging and Very Large Telescope Integral Field Unit (VLT IFU) spectra. They find two stellar cores in the HST imaging with an angular separation of 0.20 ± 0.01 arcsec and a position angle (PA) of $12.9 \pm 4.0^\circ$. These stellar cores overlap with two velocity components found in the IFU spectra that share very similar spatial separation and PA measurements, along with a velocity separation of $\sim 200 \text{ km s}^{-1}$; typical of a late stage galaxy merger (Liu et al., 2010). Furthermore, the $[\text{OIII}]/\text{H}\beta$ flux ratio of both components is much larger than 3, indicating that the ionization source of each component is an AGN. Further details of the optical properties of the target can be found in Woo et al. (2014).

The object, hereafter referred to as J132323, is located at $\alpha = 13^{\text{h}}23^{\text{m}}23.33^{\text{s}}$ and $\delta = -01^\circ59'41.9''$, with a redshift $z = 0.350280 \pm 0.000014$ as measured in DR13 of the Sloan Digital Sky Survey (Albareti et al., 2017). Assuming $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_\Lambda = 0.7$, and $\Omega_m = 0.3$, one arcsecond corresponds to 4.94 kpc at the local epoch of J132323. In this paper, we first present new radio observations of J132323 in section two and our initial analyses in section three. The implications of our findings are discussed in section four and we present our conclusions in section five.

5.1.3 Observations

VLA Observations

The target was observed by the VLA under project code 15A-072 (PI: J. Woo) in both *C*-band and *L*-band in July and August 2015 respectively, in A configuration for 10 and 18 minutes (total on-source time) respectively. The observations were conducted with all available antennas and covered 2.048 GHz around 5.5 GHz for *C*-band and 1.024 GHz around 1.5 GHz for *L*-band. Both observations split up their respective bandwidths into 16 spectral windows of 64 channels. The data products were retrieved from the archive in early 2021 and passed through the VLA pipeline packaged with CASA¹. No extra flags were applied after the pipeline run.

The VLA observations were successful in detecting a single unresolved source of ra-

¹see the webpage
<https://science.nrao.edu/facilities/vla/data-processing/pipeline>

radio emission in both bands (see [Figure 5.1](#)). These sources show peaks of 0.236 ± 0.01 mJy/beam at *C*-band and 1.39 ± 0.03 mJy/beam at *L*-band. The VLA synthesised beam is 1.3 arcsec and 0.33 arcsec at *L* and *C*-band respectively; both are around 7 times larger than e-MERLIN beam sizes. Therefore, we expect that the VLA will measure greater peak emission values than equivalent measurements from e-MERLIN.

e-MERLIN Observations

Radio imaging of the target was conducted by e-MERLIN over seventeen *C* and *L*-band observations between November 2016 and September 2020 (see [Table 5.1](#)). This totals to ~ 35 hours on-source time at *C*-band (5 GHz/6 cm) and ~ 8 hours on-source time at *L*-band (1.5 GHz/21 cm), with flags applied. All observations utilised all available antennas excluding the Lovell telescope, giving a maximum baseline of 217 km (Knockin - Cambridge) and a minimum baseline of 11 km (Mk II - Pickmere). The *L*-band spectral range covered 512 MHz centred at 1.51 GHz, broken up into 8 spectral windows of 64 MHz each. Each spectral window was averaged to 128 channels of 0.5 MHz. The *C*-band spectral range covered 512 MHz centered at 5.07 GHz, broken up into four spectral windows of 128 MHz each. Each spectral window was also averaged to 128 channels, giving a slightly larger channel width of 1 MHz.

Observations conducted in 2016 were generally of a high quality, though the first *C*-band observation on the 25th November is significantly shorter (only 7 hrs including overheads) making calibration and imaging more difficult. One of the two follow-up observations conducted in 2017 (8th March) suffered critical data loss for unknown reasons and was omitted from study. In the first two *L*-band observations conducted in 2020, quite a few medium length baselines were lost due to hardware issues. The *C*-band observations conducted on 21st and 22nd August suffered critical data loss due to a combination of very poor weather conditions and hardware problems. These two datasets were also omitted from study.

After being passed through the e-MERLIN CASA Pipeline², any problematic data such as spurious signals or data from faulty antennas was manually identified and flagged to reinforce the quality of each observation. Image fidelity and point-source response is a particular issue with this target when observing with e-MERLIN, as the maximum elevation of target is only $\sim 35^\circ$ above the horizon producing an extremely elongated point

²see the GitHub repository:
https://github.com/e-merlin/eMERLIN_CASA_pipeline

| Epoch (dd/mm/yy) | Band | On-source time (hours) | RMS Noise (μ Jy/beam) | Source/s Peak (μ Jy/beam) |
|---------------------|----------|---------------------------|-------------------------------|-----------------------------------|
| 2016 Nov 25 | <i>C</i> | 1.6 | 48 | - |
| 2016 Nov 26 | <i>C</i> | 3.4 | 42 | - |
| 2016 Nov 27 | <i>C</i> | 3.6 | 42 | - |
| 2016 Dec 13 | <i>C</i> | 3.9 | 40 | - |
| 2016 Dec 14 | <i>C</i> | 3.9 | 29 | - |
| 2017 Apr 6 | <i>C</i> | 4.8 | 25 | - |
| 2020 Aug 18 | <i>C</i> | 3.5 | 23 | - |
| 2020 Aug 20 | <i>C</i> | 2.5 | 23 | - |
| 2020 Aug 23 | <i>C</i> | 3.2 | 28 | - |
| 2020 Aug 29 | <i>C</i> | 4.5 | 24 | - |
| 2016 Dec 21 | <i>L</i> | 2.9 | 32 | 404 |
| 2020 Sept 8 | <i>L</i> | 1.3 | 56 | 359 |
| 2020 Sept 10 | <i>L</i> | 1.0 | 66 | 373 |
| 2020 Sept 11 | <i>L</i> | 2.9 | 34 | 338 |
| All | <i>C</i> | 34.9 | 8.7 | 56/87 (NE/SW) |
| All | <i>L</i> | 8.1 | 22 | 357 |

Table 5.1: A table of all observations of the target conducted by e-MERLIN in the 2016-2020 period. Observations conducted in 2016-2017 are under the project code CY4205 (PI: J. Woo), while observations conducted in 2020 are under CY10204 (PI: M. Argo). The statistics of the total combined visibility data is presented at the bottom. The RMS values quoted for the individual *C*-band epochs are measured from the dirty images as no source is detected in these, therefore no cleaning is appropriate. The on-source time presented is the total integration time on the source with flags applied.

spread function (PSF). To correct the elongated image artefacts produced by this PSF, we restored the images in Figure 2 using a circular restoring beam with a full width half maximum (FWHM) appropriate to the observing frequency. After flagging, the data was self-calibrated until the RMS noise in the outer regions of the image field reached a minimum.

The image produced by each observation was inspected for sources, manually masked, and then cleaned using the standard Högbom algorithm implemented in `CASA` (Högbom, 1974) to a 3σ limit — three times the predicted RMS noise. The individual L -band observations show a single, clear, unresolved source (see Figure 5.2). Note that for some of the L -band images the source appears to be marginally resolved, though the difference between the total flux from the source and the peak flux density is minimal. Furthermore, the large-scale noise structures produced by sparser coverage on medium length baselines could distort the central region such that the target falsely appears marginally resolved. Each C -band image for a single epoch does not immediately show any signs of a dual source, however after combining and imaging the visibility data of all C -band observations two distinct radio sources appear above the noise (see Figure 5.2).

5.1.4 Methods & Results

Positional measurements of the sources apparent in the combined e-MERLIN C -band image were performed by first fitting a pair of two-dimensional elliptical Gaussian models. It is important to note here that the fitting was performed on the map which was restored with the circular restoring beam and not the highly elongated clean beam produced by the Högbom algorithm. In the map restored with the clean beam the two sources blur into one another and make fitting realistic models very difficult, we therefore decided a more accurate model of the sources could be obtained by fitting to the map created by the idealised circular restoring beam. The angular separation of the centroids of the two models is $0.19''$, with a position angle of 22° . We take the fitted sigma value within the Gaussian model as our uncertainty of the centroid position, giving a separation of $0.19 \pm 0.06''$ and a PA of $22 \pm 10^\circ$. This corresponds to a projected separation of 0.95 ± 0.29 kpc at the epoch of the source. These measurements are similar to the positional measurements of the stellar cores obtained by Woo et al. (2014), suggesting that the unresolved radio emission from sources A and B originates from the centre of the stellar cores. The radio sources are faint when compared to a typical AGN, with 5 GHz luminosities around $\sim 10^{22}$ WHz^{-1}

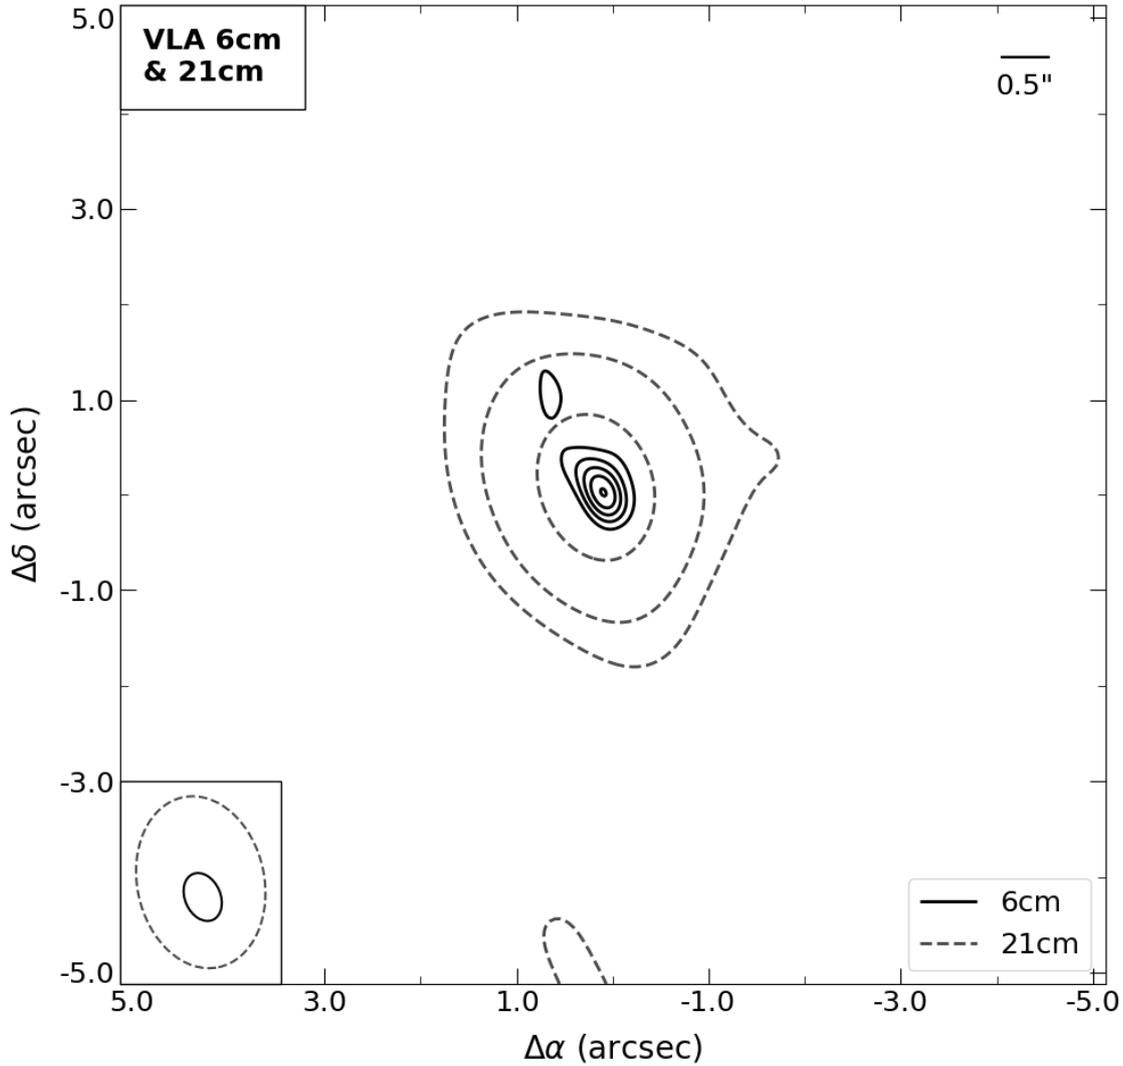


Figure 5.1: Contours of both the 21 cm and 6 cm VLA images. Contours for the 21 cm data are given on a logarithmic scale; -3σ , 3σ , 9σ , 27σ , ... where σ is the 21 cm RMS noise. Contours for the 6 cm data are given on a linear scale; -5σ , 5σ , 9σ , 13σ , ... where σ is the 6 cm RMS noise. The clean beams used to restore each of the images are given superimposed and to-scale in the bottom left. The image shows the initial detection of an unresolved radio source within J132323.

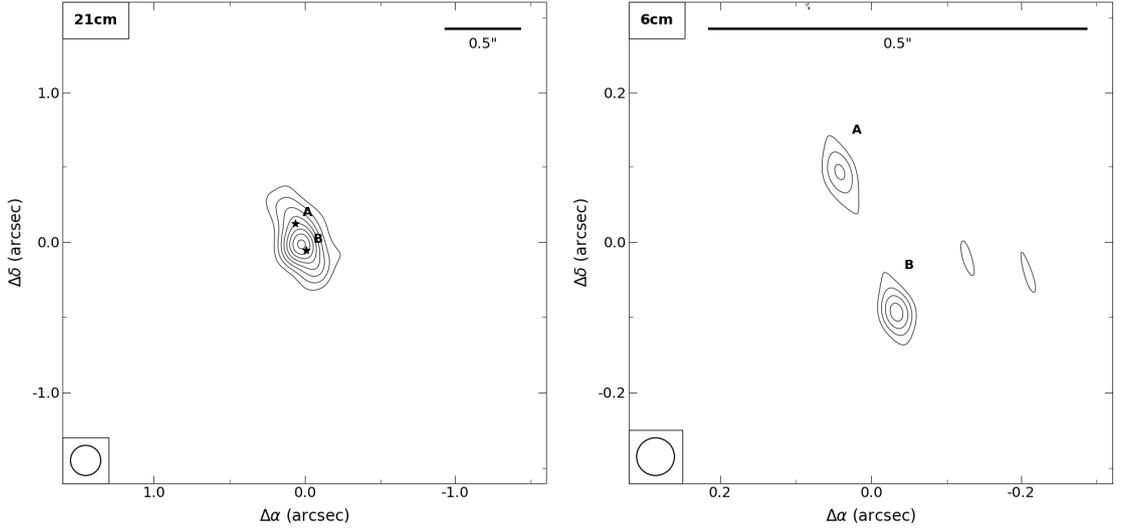


Figure 5.2: Contour images of the combined L -band and C -band e-MERLIN imaging. Contours for 6 cm map are given at -3σ , 3σ , 5σ , 7σ , ... where σ is the RMS noise in the map. Contours for the 21 cm map are instead given at -5σ , 5σ , 7σ , 9σ , ... where σ is the RMS noise in the map. The σ values are $22.6\mu\text{Jy}/\text{beam}$ and $8.7\mu\text{Jy}/\text{beam}$ for the 21 cm and 6 cm map respectively. A to-scale version of the circular *restoring* beam used to create each map is given in the bottom-left; the FWHM sizes are $0.2''$ and $0.05''$ for the 21 cm and 6 cm map respectively. The actual *clean* beam FWHM dimensions are $0.37''$ and $0.13''$ for the major and minor axes at L -band. The same dimensions at C -band are $0.16''$ and $0.03''$. The target is detected quite strongly in the 21 cm imaging. A weak detection of two unresolved radio sources is made in the 6 cm imaging with angular separation $0.19\pm 0.06''$ and PA $22\pm 10^\circ$. Both images have been cleaned to three times the theoretical noise limit.

measured with both the VLA and e-MERLIN. This increases slightly to around $\sim 10^{23}$ WHz^{-1} at 21 cm. More specifically, we measure $P_L = 3.1 \pm 0.1 \times 10^{23}$ WHz^{-1} for the VLA and slightly less from e-MERLIN with $P_L = 8.6 \pm 0.5 \times 10^{22}$ WHz^{-1} , both at a distance of 1.38 Gpc at the present epoch.

The temporal sampling of the target is very irregular so detecting any variability that might give clues as to the spatial scale over which the emission originates is not feasible. We measure a 10% decline in total 21 cm emission in the 4-year gap between observations, though it is currently impossible to determine if this is due to instrumental errors, if it is an absolute change, or if we are simply sampling an instant of a periodic fluctuation. It remains a possibility that if the target was detectable in the C -band imaging at single epochs some variability might be accounted for given the frequency of observation. Put simply, these e-MERLIN observations cannot, on their own, reveal any information about source variability. Further radio observations are therefore necessary to investigate this aspect of J132323.

Given that e-MERLIN cannot resolve two sources of 21 cm emission in addition to the 6 cm sources, a spectral index value of an assumed power law is limited in its inter-

pretation. Despite this we calculate an index of $\alpha = -0.76 \pm 0.14$ (where $S_\nu \propto \nu^\alpha$) by taking the peak flux density of the single *L*-band source and summing the peak flux densities of sources A and B. Note that this approach underestimates the total 6 cm flux that is comparable to the total 21 cm flux, so the quoted index is a slight overestimate. In addition, this method cannot account for any temporal variance of the source so the index will be an average over the observation epochs. Assuming that sources A and B are similar in nature, the value of α suggests the presence of two approximately steep-spectrum radio sources — commonly associated with extended synchrotron emission. Furthermore, the VLA imaging presents a single radio source with a spectral index of $\alpha = -1.38 \pm 0.06$, suggesting a steeper radio spectrum characteristic of stronger non-thermal emission. We stress though, that these indices are based on unsuitable measurements as we know there is two sources of radio emission within J132323, thus any reliable measurement or comparison of spectral index (or indeed any other attribute) must measure source A and B separately. A spectral index image of the e-MERLIN 21 cm and 6 cm data does not show any features of interest. A brightness temperature limit for each radio source is simpler, though again limited, in its interpretation. We calculate $T_{NE} > 1246$ K and $T_{SW} > 1713$ K using the $0.05''$ circular restoring beam, which only serves to rule out the presence of a region of cold thermal emission in both sources.

5.1.5 Discussion

One of the conclusions of [Woo et al. \(2014\)](#) was that high resolution radio observations were necessary to look for attached radio components of the dual narrow line region (NLR), thus indicating a pair of AGN. The discovery of these two unresolved sources of radio emission within J132323 is major evidence in favour of a binary AGN model. [Figure 5.3](#) shows the dual radio source overlaid on the F550M Hubble image. Note that the F550M data has an absolute positional error of $\sim 0.1''$, determined through foreground star matching to the GAIA catalogue ([Gaia Collaboration et al., 2021](#)). The relative positions remain accurate. Given the overlap of source A and B with the optical cores (accurate to within this absolute error), in addition to their similarity in separation and PA, it is very likely that a binary AGN is hosted by two galaxies in the late stages of a merger. Moving forward, the issues that warrant further investigation are: long-term variability in the radio emission, as the temporal sampling of all radio observations so far is quite poor; resolving the possibility of a jet-cloud interaction in one or both of source A and B rather than an

AGN; and further observations with an alternative radio observatory for deeper imaging and to overcome the severe impact of the e-MERLIN synthesised beam orientation. We discuss these issues in order.

Variability is a common indicator of an AGN (e.g. [Mooley et al. 2016](#)), though it is not possible to conduct a useful investigation of variability here due to the irregular spacing of epochs and their poor signal-to-noise ratio. No individual *C*-band epoch can be used for this reason. It is possible to detect A and B in the combined 2016/17 data and then again in the combined 2020 data, though the image fidelity for each epoch is significantly worse as a result of approximately halving the integration time. The *C*-band peak flux measurements for source A and B in the 2016/17 epoch are $73 \pm 14 \mu\text{Jy}/\text{beam}$ and $101 \pm 14 \mu\text{Jy}/\text{beam}$, respectively. Similar measurements for the 2020 epoch are $64 \pm 10 \mu\text{Jy}/\text{beam}$ and $80 \pm 10 \mu\text{Jy}/\text{beam}$. These peak flux measurements are consistent within the uncertainties, so drawing any conclusions is difficult. Unfortunately, the *L*-band epochs also have their own set of problems that limit their usefulness in looking for source variability. Of the four *L*-band observations, one was conducted in late 2016, while three are conducted four years later over the space of four days. Ideally, more observations would be regularly spaced to sample any changes in accretion flow.

If, at higher resolution, only one of the two radio sources shows a core-jet morphology, the orientation of the jet could give clues as to the origin of the other source. For example, if a jet is detected in source B and it is oriented randomly and away from source A, then it is more likely that the binary AGN only displays one jet structure strong enough to be detected, and that there is in fact a BH pair. If, however, the jet of source B is oriented towards source A, then a case might be made, depending on the structure of the emission in source A, for a jet-cloud interaction producing the dual radio structure. A binary AGN may still present in this way, it is just unlikely given the range of possible directions of the jet. This scenario would certainly require further investigation, most likely seeking deeper optical observations to look for clearer signs of a dual NLR. We note that a similar conclusion was reached in [Woo et al. \(2014\)](#).

Another unlikely scenario is that J132323 is actually a compact symmetric object (CSO): a very young radio source usually presenting with dual radio lobes separated by < 1 kpc ([Conway, 2002](#)). These dual radio lobes often accompany a third compact radio source between them representing the central AGN. We do not expect our dual radio source to be a CSO for a few reasons, primary amongst these is the larger separation and weaker luminosity of our sources. In a young radio source, the luminosity of each lobe

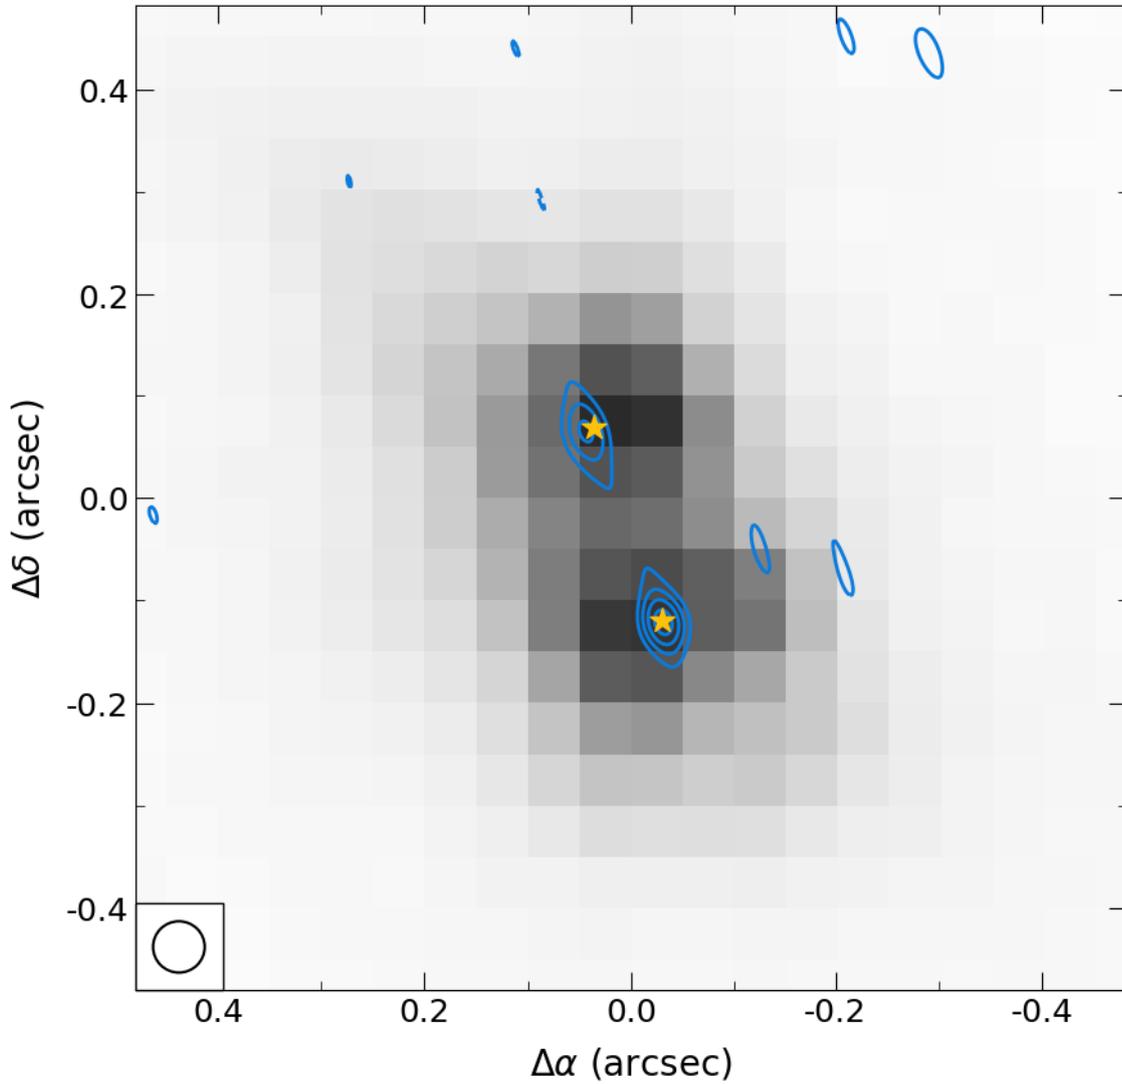


Figure 5.3: A plot of the C-band (6 cm) contours (in colour) overlaid on the F550M image (in grayscale) from the HST. A very small shift of \sim pixel (corresponding to ~ 0.05 arcsec) has been applied to the F550M image to bring radio source B into alignment with the brightest pixel in the south-west F550M component. Yellow markers indicating the centroids of the HST cores are based on the values in [Woo et al. \(2014\)](#). As the absolute positional information in the HST data is not reliable (see [subsection 5.1.5](#)), we have placed the south-west marker directly over radio source B to illustrate the similarity in separation and position angle.

has a loose positive scaling with projected distance — as the jet propagates through the surrounding medium. So one would expect that a projected separation of 1 kpc (as in the case of J132323) would be on the brighter side for a CSO. However, in a comparison with [An & Baan \(2012\)](#), J132323 appears to have a 1.4 GHz luminosity multiple orders of magnitude below even the average case of a 1 kpc CSO. In addition to this, our dual radio source is accompanied by dual stellar cores presumably associated with two distinct NLRs, whereas a CSO is by definition a single active black hole system with two unbeamed jets. Therefore, a dual NLR system is not expected in a CSO.

Changing the weighting scheme of the e-MERLIN data does not provide any extra information. We have tested Briggs weighting with robust parameters $R \geq 0$, after which the dual radio source detection becomes lost in the noise. More information can be inferred from the relative flux density differences between the VLA and e-MERLIN data. The 6 cm flux density remains similar between the two, suggesting that the majority of this emission is compact in nature. In contrast, the 21 cm emission sees a ~ 1 mJy increase in the VLA data, presumably because the larger VLA beam is detecting extended emission distributed throughout the disks and tails of the galaxies.

5.1.6 Conclusions

We present new e-MERLIN *C* (6 cm) and *L*-band (21 cm) imaging of a close-separation binary AGN candidate at $z \sim 0.35$, identified by [Woo et al. \(2014\)](#) using HST imaging and VLT IFU spectra. We detect a single unresolved *L*-band source with a peak flux density of $357 \pm 22 \mu\text{Jy}/\text{beam}$. We also detect two unresolved *C*-band sources of similar flux density with an angular separation and position angle of $0.19 \pm 0.06''$ and $22 \pm 10^\circ$ which is very similar to previous measurements of optical data representing dual cores of stellar emission. We conclude, therefore, that the target very likely harbours an active BH pair, hosted by two galaxies in the late stages of a merger. This pair is separated by 0.95 ± 0.29 kpc at the epoch of the source. To further investigate the nature of these sources, we are pursuing deeper, higher resolution observations to look for a core-jet morphology in one or both of the radio sources, to better constrain the separation and position angle, and to determine a tighter constraint on the brightness temperature of the two sources. Further work may also include deeper observations in the optical and infrared band, to investigate the dynamics of the merger and the properties of the stellar populations within each galaxy.

5.2 Follow up EVN Observations

During the investigation of the binary black hole system, a further observation proposal (PI: J. E. Brooks) was written and accepted by the EVN Programme Committee, to more deeply investigate the object for signs of AGN activity. Two C-band observations were conducted on the 12th and 13th March 2022, under the project codes EB086A and EB086B. Each observation lasted approximately eight hours. The target is difficult to observe at such a low declination as for many of the EVN stations it appears at low altitudes for most of the day. Observing close to the horizon is usually a problem for radio telescopes because of the significantly increased volume of terrestrial radio emission entering the receiver. Using two observations spread over two days, therefore, allows the target sensitivity to be reached, without observing close to the horizon for long periods.

Several antennas were lost or excluded from observation for various reasons: Russian stations were excluded from the EVN as a result of the Russian invasion of Ukraine in February 2022³, some antennas reported no data for unknown reasons, and a few antennas could not observe due to receiver issues. Consequently, only eleven stations reported any data at all for the first observation, and only ten for the second. Both observations used the same frequency setup: a 256 MHz bandwidth covering the range 4.8-5 GHz. The data was automatically processed through the EVN pipeline, which averaged the data to 8 channels of 32 MHz. Both observations observed the same flux and phase calibrators: 3C273 for the flux calibrator, and the same phase calibrator used by e-MERLIN to observe the binary AGN. This source was also the subject of the GAN training set, and is described at the end of [subsection 3.5.1](#).

No further processing is applied to the data received from the EVN pipeline apart from generating self calibration solutions from the phase calibrator and applying them to the target field. The final image reaches a modest (compared to e-MERLIN) sensitivity of $\sigma = 23\mu\text{Jy}/\text{beam}$. In theory, this should be sensitive enough to observe the less bright northern source at a significance of $2-3\sigma$, but unfortunately *no* radio emission was observed from the object at all, even from the brighter southern source. The final image is presented in [Figure 5.4](#).

I hypothesize the lack of a detection is largely due to the problems encountered during the EVN observation, combined with uncertainties in both the e-MERLIN and EVN observations. It cannot be assumed that a perfect calibration was reached given the number

³See the EVN statement: <https://www.jive.eu/statement-evn-ukraine>

of observational problems that occurred across a significant number of antennas; there may be chunks of the observations that lack proper phase information, for example. So even though the required sensitivity was technically (and just barely) reached, a good calibration may not be possible. Moreover, there are other factors to consider such as the spatial sampling of the much longer EVN baselines, which may have filtered out just enough extended emission (which may have been present) from the source for it to dip below the noise. The e-MERLIN radio study could not rule out source variability, which also may have caused the source to dip below the noise.

Taking the evidence at face value, the lack of a detection at EVN spatial scales does suggest that the origin of the radio emission is not from an AGN, since the spectral index must be steeper than the e-MERLIN measurement presented in [subsection 5.1.4](#), and the brightness temperature must be similar or less. Furthermore, a convincing argument can be made that since the radio emission in J132323 is extended, it is weakly observed at e-MERLIN scales and completely resolved out at EVN scales. However, the optical evidence presented by [Woo et al. \(2014\)](#) and the subsequent detection of two radio sources directly coincident with the two stellar cores is still strong evidence in favour of a binary AGN. Therefore, further radio observations with the EVN are needed to clarify the situation. At the very least, further EVN observations would be useful to place a deeper sensitivity constraint on the object (or objects), and perhaps reveal variability. Observations at X-ray wavelengths cannot resolve the closely bound sources, but may be useful to investigate the existence of X-ray emission at all.

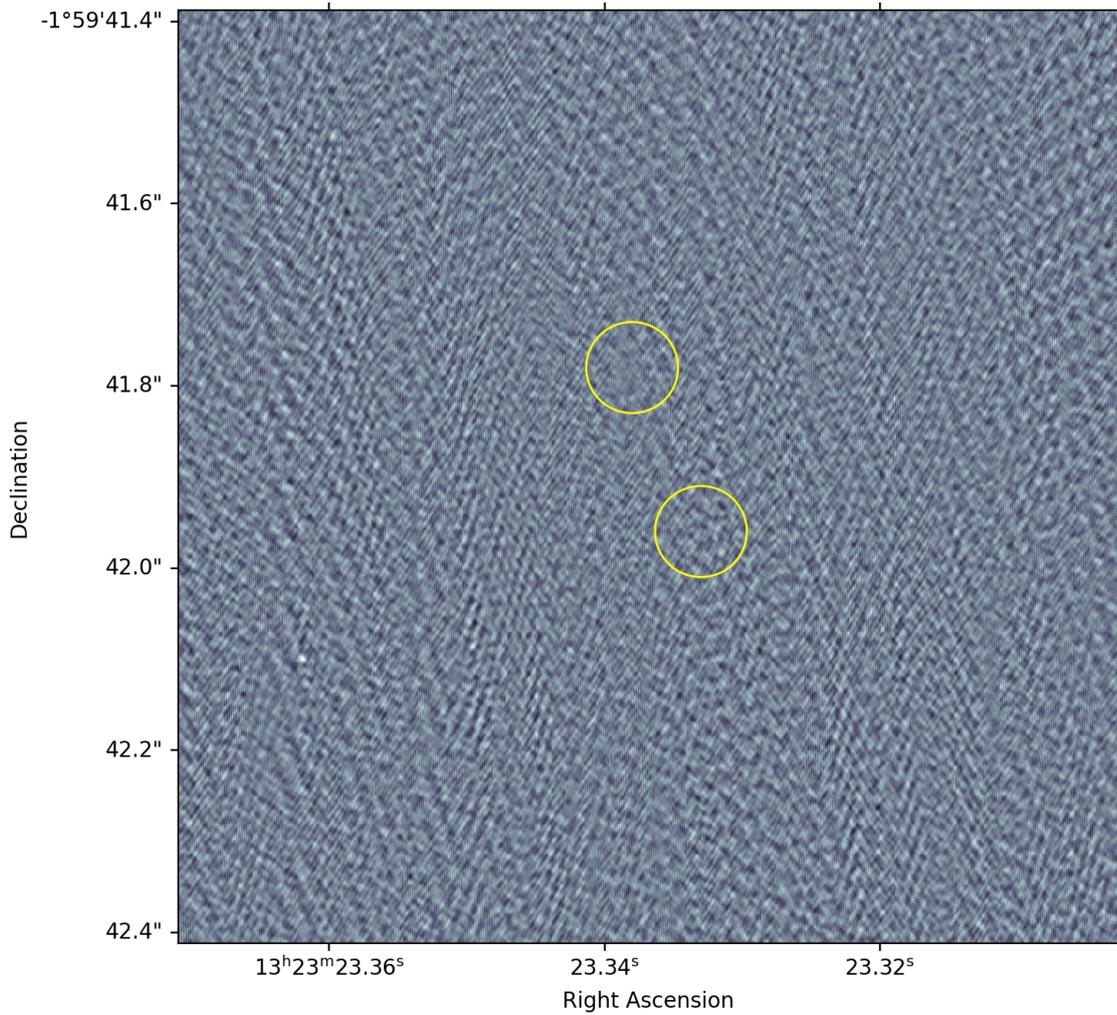


Figure 5.4: The EVN image of the location of the binary source detected by e-MERLIN. Yellow circles indicate the source locations measured in the e-MERLIN image, with a radius equal to the 6 cm e-MERLIN restoring beam size of $0.05''$. The RMS noise across the image is $\sigma = 23\mu\text{Jy}/\text{beam}$.

Chapter 6

Studying the Lockman Hole

This chapter will present data of the Lockman Hole produced by e-MERLIN, the process of calibration and imaging, the source catalogue, and some initial analysis of the field. The objective of this experiment is to serve as a complimentary study to the one conducted by [Muxlow et al. \(2020\)](#), which combined e-MERLIN and VLA data to create a high resolution image of the similarly well-studied GOODS-N field. By comparing these two images, any effects related to cosmic variance can be accounted for. Furthermore, producing these images and cataloguing the sources within is a necessary precursor to a detailed classification of sources in the field. The catalogue can be used to properly address the problem of source counts introduced at the start of [chapter 4](#), and a detailed classification can be used to explore the evolution of the population through cosmic history. The choice to observe the Lockman Hole is due to the properties of this particular patch of sky. The Earth itself sits in a sea of gas and dust belonging to the Milky Way ISM, and the gas contains varying amounts of neutral hydrogen (or HI) which can absorb photons with energies > 13.6 eV. This absorption obscures light passing through it on its way to the Earth, making observations at UV and X-ray wavelengths more difficult. In 1986, [Lockman et al.](#) published a study of HI column densities across the northern sky, and found a shallow patch around $\alpha = 10^{\text{h}}48^{\text{m}}00^{\text{s}}$, $\delta = +57^{\text{d}}00^{\text{m}}00^{\text{s}}$, which came to be known as the ‘Lockman Hole’. As a result of the low column density of HI, it has become one of the most heavily studied patches of sky, and observational data is available from a wide array of sources across the electromagnetic spectrum. This data can be utilised towards more reliable classifications of sources within the field.

| Observation Date (dd/mm/yyyy) | t_{obs} (hours) | σ_{LH} ($\mu\text{Jy}/\text{beam}$) | Flagged % |
|----------------------------------|----------------------|-------------------------------------------------|-----------|
| 02/01/2017 | 6 | 67 | 44 |
| 22/03/2017 | 14 | 53 | 54 |
| 24/03/2017 | 3 | 105 | 45 |
| 01/04/2017 | 8 | 101 | 57 |
| 16/04/2017 | 14 | 60 | 57 |
| 27/04/2017 | 26 | 36 | 50 |
| 16/06/2017 | 12 | 38 | 42 |
| 01/07/2017 | 3 | 67 | 39 |
| 09/08/2017 | 10 | 74 | 54 |
| 01/01/2018 | 7 | 70 | 53 |
| 18/03/2019 | 4 | 96 | 32 |
| 22/05/2019 | 11 | 65 | 57 |
| 30/05/2019 | 9 | 53 | 50 |
| 15/11/2021 | 5 | 45 | 47 |
| 10/09/2022 | 11 | 38 | 40 |
| 04/11/2022 | 6 | 41 | 44 |
| 08/01/2023 | 6 | 66 | 48 |
| 15/01/2023 | 6 | 76 | 48 |

Table 6.1: Details of each observation included in the final image of the Lockman Hole. Quoted observing times, given by t_{obs} are an estimate of the effective on-target observing time, with phase reference cycles removed. Image sensitivities, given by σ_{LH} are an approximate value obtained by measuring the RMS fluctuation in a cutout of size 256x256 located at the centre of the e-MERLIN primary beam. The flagged percentage is a percentage of the total time observing the target.

6.1 The e-MERLIN Observations

A total of 18 separate L-band (21 cm) observations of the Lockman Hole were included in the final image. This data is drawn from 24 L-band observations conducted by e-MERLIN, where the missing 6 were excluded due to irrecoverably bad or missing data. Some details of each remaining observation can be found in [Table 6.1](#). All observations were conducted using the same frequency configuration: e-MERLIN L-band averaged to 1024 channels arranged into eight spectral windows covering 512 MHz between 1.25 GHz - 1.75 GHz. The width of the averaged channels in this setup is 0.5 MHz. All observations utilised all e-MERLIN antennas *excluding* the Lovell telescope.

All observations used the phase-referencing technique, so a phase calibrator was periodically observed to maintain knowledge of phase disruptions throughout the observations of the target; the Lockman Hole does not contain any sources that appear bright enough to e-MERLIN that it can use in self-calibration. This phase-referencing source is a strong unresolved source located nearby to the Lockman Hole field. The sources observed remained consistent across all the L-band observations, so they are presented in

| Code | Right Ascension (hms) | Declination (dms) (dms) | Role |
|-----------|--------------------------|----------------------------|----------------------|
| 1407+2827 | 14h 07m 00.4s | +28d 27m 14.7s | Bandpass Calibration |
| 1331+3030 | 13h 31m 08.3s | +30d 30m 32.3s | Flux Calibration |
| 0319+4130 | 03h 19m 48.2s | +41d 30m 42.1s | Point Calibration |
| 1058+5628 | 10h 58m 37.7s | +56d 28m 11.2s | Phase Calibration |
| 1046+5900 | 10h 46m 00.0s | +59d 01m 00.0s | Target |

Table 6.2: Details of each source observed across all 18 observations.

[Table 6.2](#). The pointing observed in the Lockman Hole was specifically chosen to overlap with the observations presented in [Vernstrom et al. \(2016b\)](#), which also overlap with the observations presented in [Owen & Morrison \(2008\)](#).

6.2 Calibration & Imaging

Each observation was first processed using a customised AOFLAGGER ([Offringa et al., 2012b](#)) strategy that was designed to flag conservatively. As the analysis of the performance of the GAN did not show strong advantages over AOFLAGGER (see [section 3.5](#)), this approach should produce better results. Since the calibration was done manually with regular visual inspection of the solutions throughout the process, a more conservative flagging strategy was possible to cut down on unnecessary data loss. The e-MERLIN calibration pipeline ([Moldon, 2021](#)) was not used as it is now several versions behind the latest release of the Common Astronomy Software Applications (CASA) package on which it is built. To avoid the presence of errors caused by bugs in the software, a small calibration script was created that mimics the e-MERLIN pipeline in functionality which was then used for calibrating all the observations. The results of the script were regularly checked against the initial e-MERLIN pipeline run to ensure a similar output. Specifically, the script performed the following calibration strategy:

- Calculate and apply delay calibration solutions for the bandpass calibration source
- Calculate and apply phase and amplitude solutions for the bandpass calibration source
- Calculate an initial bandpass correction for the full observation and apply to all fields
- Calculate and apply delay calibration solutions for all fields

- Calculate and apply phase corrections for all fields
- Calculate and apply scaled amplitude solutions for all fields using the model of the flux calibrator
- Recalculate flux-corrected bandpass solutions
- Recalculate and apply delay, phase, and amplitude solutions for all sources

At each of these stages, the generated solutions were visually inspected and any outliers or erroneous solutions were flagged using the `PLOTMS` tool in `CASA`. `IFLAG` was not used since it is not designed to work on the calibration table format in `CASA`. After the strategy was completed, self-calibration was performed on the phase calibrator source with the intent of transferring the corrections over to the target. Though, this correction made little, if any, difference to the sensitivity in the target field.

Self-calibration of the target data using sources in the target field was attempted, but did not produce useable corrections. When strong point-like sources are available in the target field it is normally possible to self-calibrate. However, for e-MERLIN there are no bright enough sources for this to be possible; there simply isn't enough information in the field to calculate accurate solutions. Therefore, the data used to generate the sky image in [Figure 6.1](#) are **not** self-calibrated.

After calibration, the target data was separated from the original datasets and imaged using `WSCLEAN` (see [section 2.8](#), or [Offringa et al. 2014](#)). This and all other images in the Lockman Hole investigation used natural weighting to maximise sensitivity. Some structure was present in the image produced by `WSCLEAN` at this stage, and it was traced to a single observation. The problematic data in this observation was flagged — it is unknown whether it was caused by residual RFI or an incorrect calibration solution. A version of the final primary beam corrected image is presented in [Figure 6.1](#).

6.2.1 Noise Structure

One inherent problem of e-MERLIN is that it has relatively few antennas. Since it is also covering a large range in spatial frequency, it can be difficult to properly sample the full area of the uv plane covered in a given observation. In the specific case of the observations presented in this thesis, there are two large gaps in the spatial frequency coverage, shown in [Figure 6.2](#), which shows up in the sky image as a ‘spotty’ noise

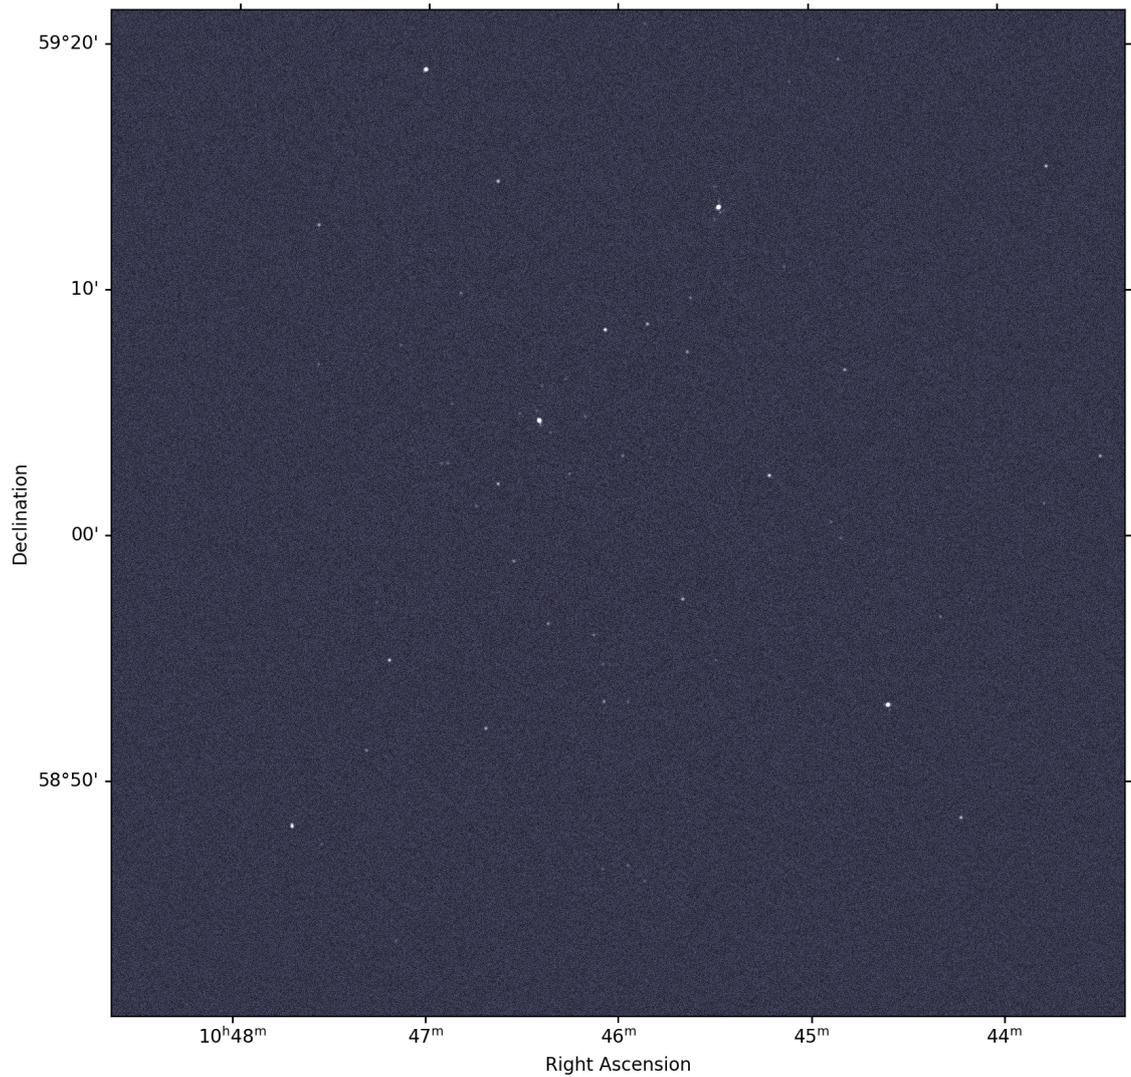


Figure 6.1: The final image of the Lockman Hole using all observations. The image resolution is twice the synthesised beam width (undersampled), and has a restoring beam of size $5''$. This is to allow the detected sources to be seen in a small image like the one above.

structure, shown in [Figure 6.3](#). This section will provide a brief analysis of this noise structure, and its potential impact.

It is widely, and often implicitly assumed that the background noise in a radio image is accurately described by a gaussian probability distribution centered about zero, and a spread equal to the root mean square (RMS) of the post-CLEAN residuals or a large area containing no significant emission. Consequently, as a measure of the depth of a particular image, wide-field surveys will often simply quote the RMS value. The RMS is also commonly used as a stopping threshold in CLEAN so that no sources are cleaned below a specified limit. Common choices for this limit are 3σ and 5σ , where σ is the calculated RMS value. Selection of this limit must be done carefully, since it is the primary factor in determining how many sources are detected, and where.

To ensure the ‘spotty’ noise structure does not significantly impact the selection of a CLEAN threshold, and by extension the final source catalogue, a brief analysis of its properties is presented in [Figure 6.4](#). The experiment is conducted by sampling a patch at the centre of the full-sky image to eliminate any direction dependent effects. The patch is progressively made smaller to sample smaller spatial scales, to test the assumption of gaussian noise in the presence of apparent small scale noise structure. At small scales the gaussian assumption still holds, with only a small deviation of $\sim 1\mu\text{Jy}$ from a mean of zero in the smallest cutout. The deviation from a perfect gaussian distribution does increase as the image patch is reduced in size, but this is to be expected for smaller sample sizes.

This analysis suggests that selecting a normal threshold of 3σ or 5σ is possible, though there is another factor to consider. The size of the Lockman Hole image is extremely large as it covers the full e-MERLIN primary beam at full resolution. This results in an image of size 65536×65536 pixels; a total of approximately 4.3 billion pixels. Even assuming perfectly gaussian noise across an empty field of that size, one expects to find approximately 11.6 million pixels exceeding a 3σ RMS threshold, compared to just ~ 2000 pixels exceeding 5σ . Because of this, coupled with the noise structure shown in [Figure 6.3](#), a selection of a 5σ CLEAN threshold is most appropriate. This greatly reduces the number of erroneous CLEAN components introduced from simple random variations, and potentially from ‘spots’ in the small scale noise structure.



Figure 6.2: The e-MERLIN spatial frequency coverage when observing the Lockman Hole, obtained by fourier transforming the CLEANed sky image. There are two noticeable gaps in the coverage at the very centre, and slightly out from the centre.

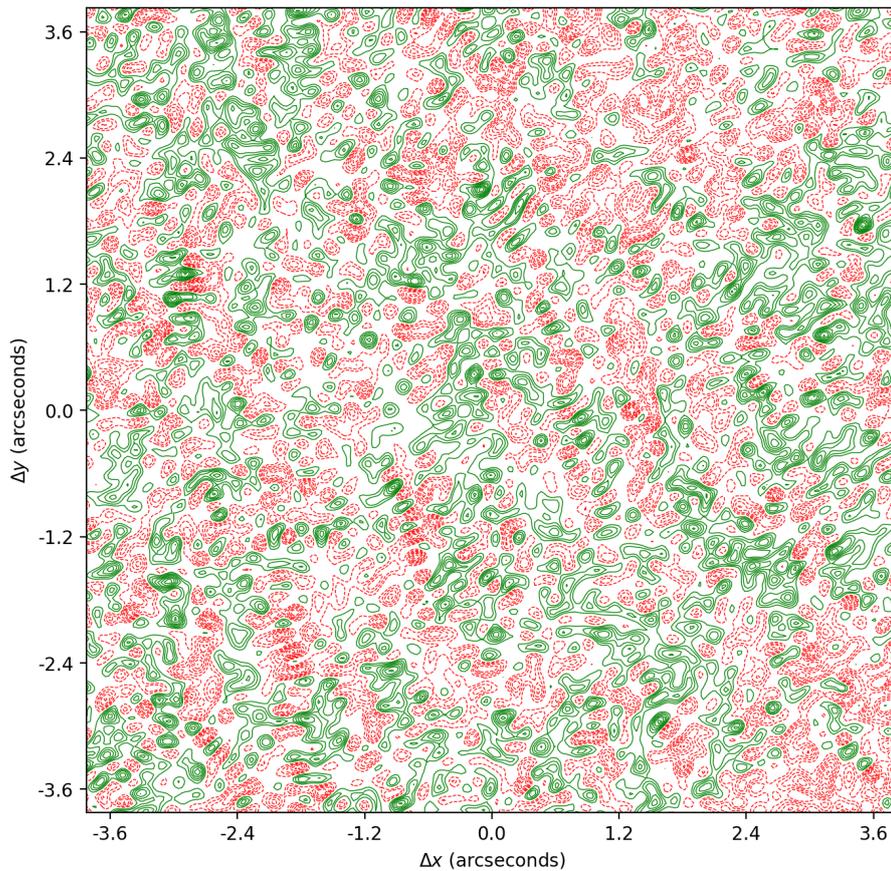


Figure 6.3: A sample cutout of the noise structure in the Lockman Hole image. Contours are given in increments of $\frac{\sigma}{2}$, where $\sigma = 12\mu\text{Jy}/\text{beam}$ is the RMS in the small field shown. Green contours indicate positive values, whereas red dashed contours indicate negative values. There is a clear structure of peaks and troughs on small scales that is not real emission.

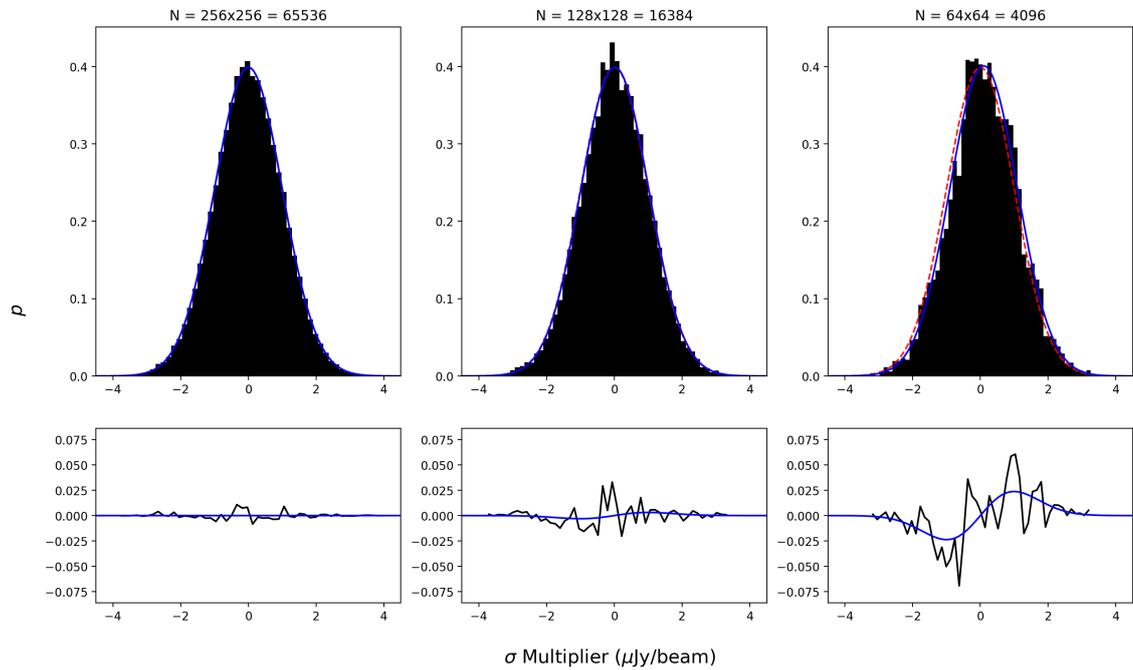


Figure 6.4: The probability distribution of pixel values in progressively smaller cutouts, measured relative the RMS, σ , in the cutout. The perfect gaussian model is drawn in a red-dashed line (though only visible in the rightmost panel). The panels along the bottom show the magnitude of the deviation from a perfect gaussian distribution for the actual pixel values (black) and the fitted gaussian model (blue). There is clearly a small divergence from a gaussian distribution with smaller sample sizes, but this is somewhat expected. There also appears to be a small trend towards positive valued noise pixels, presumably due to confusion from extremely faint undetected sources.

6.3 Source Catalogue

Source extraction is performed using the `PYBDSF` tool (Mohan & Rafferty, 2015). The image given to `PYBDSF` is a primary beam corrected image of size 32768x32768 pixels, or 2^{15} pixels in width and height, for a more efficient FFT calculation. Note that this image has a slightly lower resolution than the full resolution image used in subsection 6.2.1 and to make the cutouts shown in section 6.4, but the synthesised beam is still Nyquist sampled. The image is generated using `WSCLEAN`, enforcing no negative clean components and using a clean threshold of 5σ . The pixel scale is 70 milliarcseconds, so the total image area is just over 38 arcminutes in width and height, centred on the coordinate 10h46m00s +59d01m00s. The RMS at the centre of the field is $\sigma = 12\mu\text{Jy}/\text{beam}$, and the primary beam correction is applied such that pixel values drift towards zero as in Figure 2.9. The size of the synthesised beam is $0.19 \times 0.14''$, which covers approximately 3×2 pixels in the image used for producing the catalogue. A total of 82 sources are extracted from the field using `PYBDSF`, but 4 are judged to be image artefacts and are removed, leaving a total of 78 distinct sources. Three of the artefacts belonged to the powerful source studied in subsection 6.4.2, and the remaining artefact belonged to the source studied in subsection 6.4.4. Note that the *two* components of that source are still present in the catalogue; a third `PYBDSF` entry belonged to that source and was judged to be an artefact. Details of the sources are presented in Table 6.4, and cutouts of the sources are presented in Figure 6.5.

Where possible, cross matches are made with the LOFAR 2 m Sky Survey (Shimwell et al. 2017, hereafter simply ‘the LOFAR survey/catalogue’) and the Vernstrom et al. (2016b) VLA survey (hereafter simply ‘the VLA survey/catalogue’), both of the Lockman Hole. An e-MERLIN source is considered to have a match if it falls within one synthesised beam FWHM of an entry in the associated catalogue. The LOFAR beam width is taken to be $6''$, which is taken directly from the header information in the Shimwell et al. (2017) image of the Lockman Hole. For the VLA survey, the synthesised beam width is set to $8''$, which is quoted by Vernstrom et al. (2016b). Note that they used a combination of data from C and BnA configurations of the VLA, so $8''$ is an upper limit for the VLA survey synthesised beam width. Using these criteria, exactly one match is found in the LOFAR catalogue for 46/78 (59%) sources, and exactly one match is found in the VLA catalogue for 33/78 (42%) sources. There are 28 sources in the e-MERLIN catalogue that do not have a corresponding source in either the VLA or LOFAR catalogues. Details of

the cross matches can be found in [Table 6.3](#).

The overwhelming majority of sources with no match in the VLA catalogue exist at distances $> 10'$ from the pointing centre, and therefore fall outside the sensitive region of the [Vernstrom et al. \(2016b\)](#) image. There are 8 remaining sources inside the $10'$ boundary with no corresponding source within the $8''$ matching area. However, one of these sources (S29, see [subsection 6.4.3](#)) can be matched to the complex source TV16 J104624+590522 in the VLA survey, also shown in Figure 6a of [Vernstrom et al. \(2016b\)](#).

The lack of matches in the LOFAR catalogue is likely due to a combination of factors. Firstly, the quoted RMS of the LOFAR field is approximately twice as large as the e-MERLIN field, introducing an inherent bias towards $\Gamma < 0$ sources; faint sources with $\Gamma \geq 0$ are unlikely to be cross matched between the two surveys. Secondly, there are some sources with no LOFAR match at the edge of the e-MERLIN field, so it is possible that direction dependent errors in the e-MERLIN field play a part. Finally, it may be that the 140MHz morphology of a particular e-MERLIN source is different enough that the source appears to be located further away than a single LOFAR synthesised beam width. Special attention would be required for these cases to determine if the sources are related or not.

Redshifts are gathered from the Spitzer Wide-area Infrared Extragalactic Survey (SWIRE) photometric redshift catalogue ([Rowan-Robinson et al., 2008](#)), which contains both spectroscopic and photometric redshifts. Where possible, spectroscopic redshifts are prioritised over photometric measurements, but spectroscopic redshifts are not available for all sources. Quantifying the uncertainty on the photometric redshift is complex (see [Rowan-Robinson et al. \(2008\)](#) for a full analysis), but in general it is assumed that spectroscopic redshifts, z_{sp} , are accurate, and that photometric redshifts z_{ph} obey $z_{\text{ph}} \approx z_{\text{sp}}$. Sources are matched with the redshift catalogue using the same method as above, with a FWHM of $2''$, which is the approximate size of the *Spitzer Infrared Array Camera* PSF¹. This method produces matches in the redshift catalogue for 35 out of 78 sources (45%).

¹See the handbook: <https://doi.org/10.26131/irsa486>

6. STUDYING THE LOCKMAN HOLE

| ID | α (hms) | δ (dms) | LOFAR ID | VLA ID | SWIRE ID |
|----|-------------------|-------------------|----------|---------------------------------|----------|
| 1 | 10h48m15.91s | 58d51m58.92s | - | - [†] | - |
| 2 | 10h48m5.8s | 59d13m7.59s | - | - [†] | - |
| 3 | 10h47m54.68s | 58d57m44.38s | - | - [†] | - |
| 4 | 10h47m55.44s | 59d12m30.22s | - | - [†] | - |
| 5 | 10h47m49.81s | 59d14m39.21s | 24383 | - [†] | - |
| 6 | 10h47m41.82s | 58d48m17.03s | 24617 | - [†] | - |
| 7 | 10h47m41.73s | 58d48m12.28s | 24617 | - [†] | - |
| 8 | 10h47m36.67s | 59d6m15.65s | - | - [†] | - |
| 9 | 10h47m35.26s | 58d47m41.49s | - | - [†] | - |
| 10 | 10h47m34.5s | 59d12m41.09s | 24925 | - [†] | - |
| 11 | 10h47m15.69s | 58d57m20.63s | - | TV16J104716+585719 [†] | 573520 |
| 12 | 10h47m11.59s | 58d54m59.57s | 25647 | - [†] | - |
| 13 | 10h47m9.18s | 58d43m34.02s | 25750 | - [†] | 411447 |
| 14 | 10h47m0.85s | 59d19m1.91s | 26019 | - [†] | - |
| 15 | 10h47m0.21s | 59d1m7.66s | 26075 | TV16J10470+590106 | 574709 |
| 16 | 10h46m55.52s | 59d3m0.86s | 26215 | TV16J104654+590259 | 575556 |
| 17 | 10h46m54.72s | 58d57m33.13s | - | TV16J104655+585747 | - |
| 18 | 10h46m53.56s | 59d2m59.93s | 26293 | TV16J104653+590316 | 575380 |
| 19 | 10h46m53.57s | 59d3m0.21s | 26293 | TV16J104653+590316 | 575380 |
| 20 | 10h46m49.64s | 59d9m56.22s | 26419 | TV16J104650+590955 [†] | - |
| 21 | 10h46m44.54s | 59d1m16.29s | - | TV16J104645+590115 | 573464 |
| 22 | 10h46m41.38s | 58d52m13.82s | 26669 | TV16J104641+585213 [†] | 567292 |
| 23 | 10h46m37.75s | 59d2m10.69s | 26775 | TV16J104638+590210 | - |
| 24 | 10h46m37.97s | 59d14m29.57s | 26747 | - [†] | - |
| 25 | 10h46m37.43s | 58d46m25.43s | - | - [†] | - |
| 26 | 10h46m36.31s | 59d4m43.33s | - | - | - |
| 27 | 10h46m32.77s | 58d59m2.23s | 26936 | TV16J104631+585907 | - |
| 28 | 10h46m26.89s | 59d15m27.15s | - | - [†] | - |
| 29 | 10h46m24.87s | 59d4m45.94s | 27139 | - | 574122 |
| 30 | 10h46m24.01s | 59d5m22.21s | - | TV16J104624+590522 | 574436 |
| 31 | 10h46m21.97s | 58d56m29.71s | 27248 | TV16J104622+585629 | 568388 |
| 32 | 10h46m16.06s | 58d55m58.58s | 27421 | TV16J104616+585558 | 567575 |
| 33 | 10h46m15.36s | 59d2m35.96s | 27432 | TV16J104615+590235 | 571806 |
| 34 | 10h46m11.45s | 58d49m11.77s | - | - [†] | - |
| 35 | 10h46m7.63s | 58d56m1.98s | 27672 | TV16J10468+585601 | - |
| 36 | 10h46m4.79s | 58d54m52.08s | - | TV16J10465+585452 | 566058 |
| 37 | 10h46m4.54s | 59d6m47.25s | 27697 | TV16J10464+590640 | 573713 |
| 38 | 10h46m4.41s | 58d53m19.46s | 27758 | TV16J10464+585319 | - |
| 39 | 10h46m4.11s | 59d8m27.69s | 27698 | TV16J10464+590827 | 574778 |
| 40 | 10h46m4.07s | 59d8m27.68s | 27698 | TV16J10464+590827 | 574778 |
| 41 | 10h46m3.6s | 59d4m23.45s | 27780 | TV16J10462+590425 | 572019 |
| 42 | 10h46m3.4s | 59d8m55.78s | - | - | - |
| 43 | 10h46m2.6s | 59d16m36.53s | - | - [†] | - |
| 44 | 10h45m58.66s | 59d3m19.54s | 27928 | TV16J104559+590319 | - |
| 45 | 10h45m57.56s | 58d56m35.66s | 27952 | TV16J104558+585636 | - |
| 46 | 10h45m52.06s | 59d3m59.81s | - | TV16J104551+590345 | 570755 |
| 47 | 10h45m51.85s | 59d3m11.26s | 28129 | TV16J104552+590311 | 570173 |
| 48 | 10h45m51.84s | 58d46m1.05s | 28155 | - [†] | 559748 |
| 49 | 10h45m50.82s | 59d8m41.43s | 28186 | TV16J104549+590830 | 573830 |
| 50 | 10h45m43.62s | 59d14m57.96s | - | - [†] | - |
| 51 | 10h45m39.81s | 58d57m30.18s | 28425 | TV16J104540+585730 | 565697 |
| 52 | 10h45m39.79s | 58d57m29.69s | 28425 | TV16J104540+585730 | 565697 |
| 53 | 10h45m38.3s | 59d7m33.34s | 28501 | TV16J104537+590732 | 571963 |
| 54 | 10h45m37.81s | 58d55m52.17s | - | - | - |
| 55 | 10h45m37.19s | 59d9m45.37s | 28536 | TV16J104537+590945 | 573353 |
| 56 | 10h45m28.32s | 59d13m26.52s | 28782 | - [†] | 575056 |
| 57 | 10h45m27.48s | 58d47m36.72s | - | - [†] | - |
| 58 | 10h45m19.48s | 59d5m2.5s | - | - | - |
| 59 | 10h45m14.78s | 58d43m14.23s | - | - [†] | - |

| ID | α (hms) | δ (dms) | LOFAR ID | VLA ID | SWIRE ID |
|----|-------------------|-------------------|----------|---------------------------------|----------|
| 60 | 10h45m12.52s | 59d2m31.18s | 29246 | TV16J104513+590231 | 566610 |
| 61 | 10h45m11.07s | 59d7m11.69s | 29292 | TV16J104511+590712 | 569359 |
| 62 | 10h45m10.58s | 59d14m41.26s | 29305 | - [†] | 574389 |
| 63 | 10h45m6.75s | 58d55m17.69s | - | - | - |
| 64 | 10h45m5.88s | 59d18m32.44s | - | - [†] | - |
| 65 | 10h44m50.41s | 59d19m26.38s | 29895 | - [†] | 575821 |
| 66 | 10h44m50.21s | 58d59m57.6s | 29957 | TV16J104450+585958 | 563382 |
| 67 | 10h44m48.74s | 59d6m48.57s | 29990 | TV16J104449+590649 [†] | 567344 |
| 68 | 10h44m46.84s | 58d56m56.08s | - | - [†] | - |
| 69 | 10h44m45.57s | 59d14m36.11s | - | - [†] | - |
| 70 | 10h44m44.34s | 59d14m37.44s | - | - [†] | - |
| 71 | 10h44m35.65s | 58d53m10.38s | 30384 | - [†] | - |
| 72 | 10h44m13.05s | 58d48m33.74s | - | - [†] | - |
| 73 | 10h43m50.5s | 58d50m21.78s | 31885 | - [†] | - |
| 74 | 10h43m46.54s | 59d18m25.91s | - | - [†] | - |
| 75 | 10h43m46.56s | 59d18m25.71s | - | - [†] | - |
| 76 | 10h43m46.22s | 59d1m19.03s | 32001 | - [†] | 559320 |
| 77 | 10h43m44.74s | 59d15m2.61s | 31978 | - [†] | 567430 |
| 78 | 10h43m38.86s | 59d17m58.02s | - | - [†] | - |

Table 6.3: Cross matches for the 78 sources extracted using `PYBDSF` from the Lockman Hole field cleaned to a 5σ limit. From left to right, columns are: e-MERLIN Source ID; Right Ascension, α ; Declination, δ ; [Shimwell et al. \(2017\)](#) LOFAR Source ID; [Vernstrom et al. \(2016b\)](#) VLA Source ID; [Rowan-Robinson et al. \(2008\)](#) SWIRE Source ID. Entries with a [†] in the VLA Source ID column are $>10'$ from the phase centre.

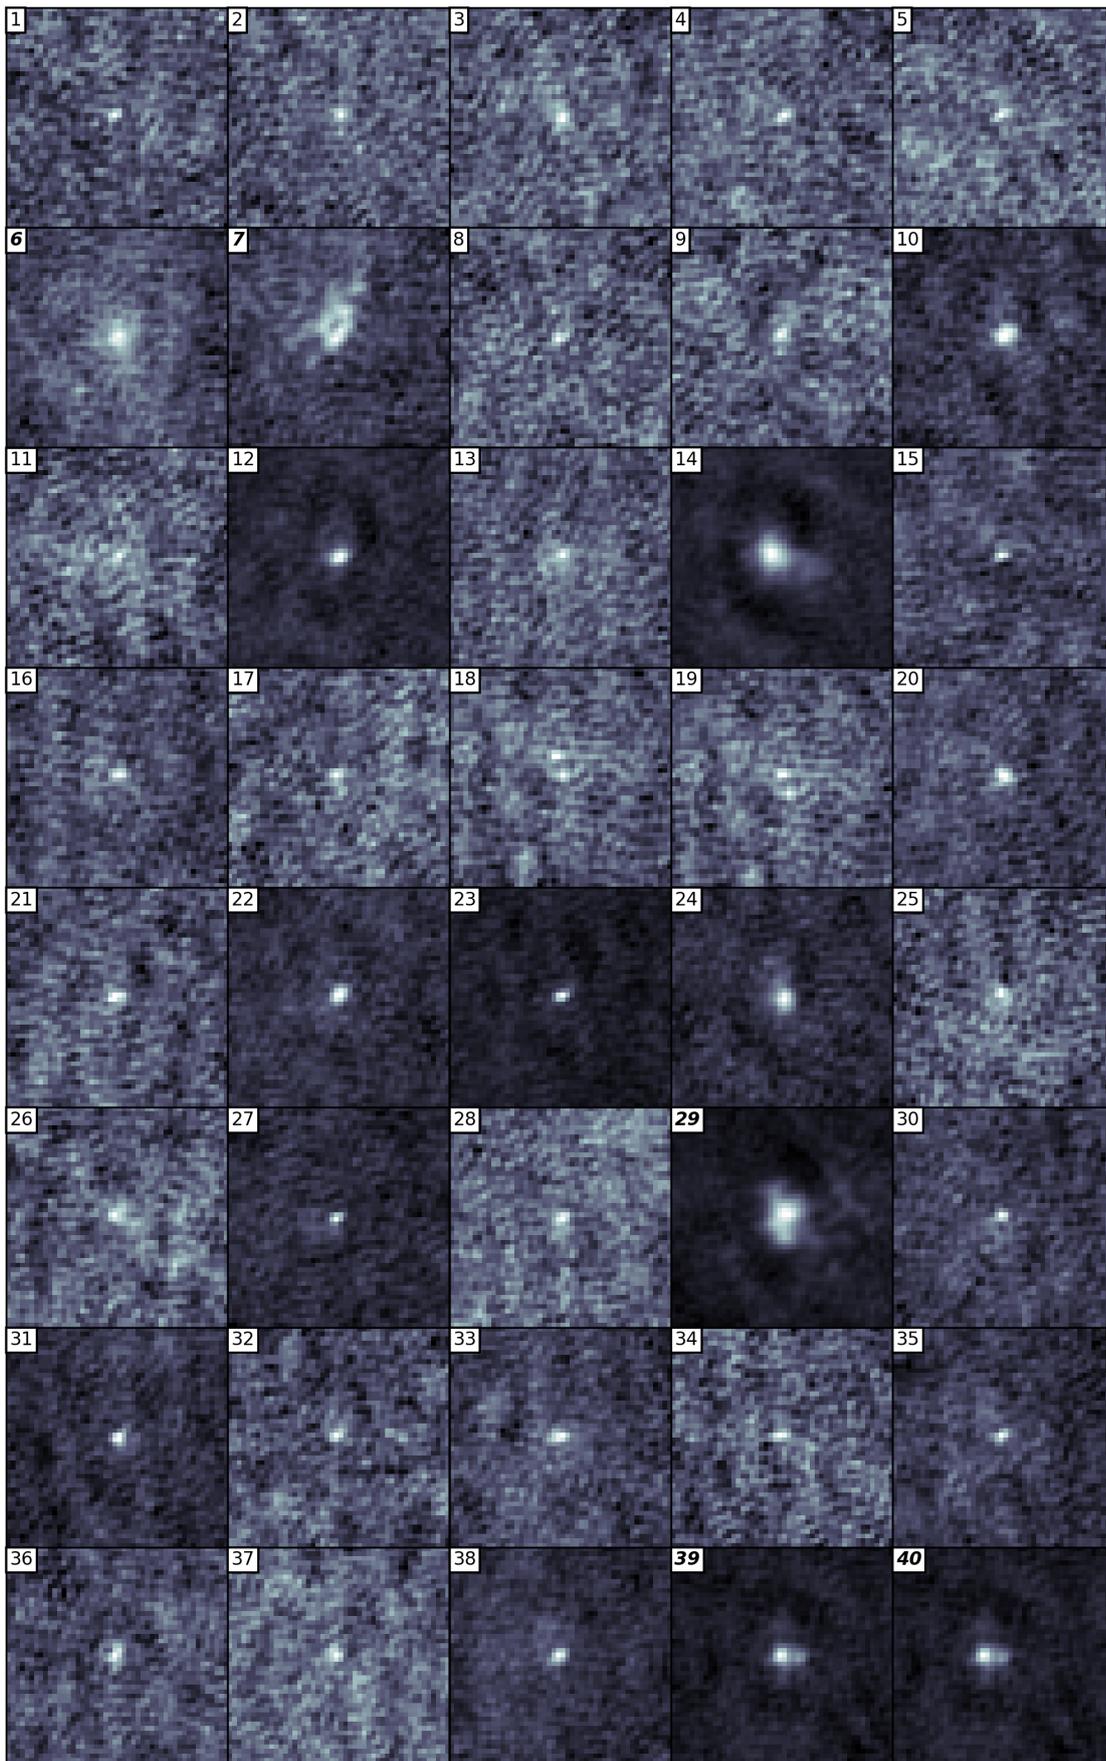
6. STUDYING THE LOCKMAN HOLE

| ID | $P_{140\text{MHz}}$ ($\mu\text{Jy}/\text{beam}$) | $S_{1.5\text{GHz}}$ (μJy) | $\Delta S_{1.5\text{GHz}}$ (μJy) | $P_{1.5\text{GHz}}$ ($\mu\text{Jy}/\text{beam}$) | $\Delta P_{1.5\text{GHz}}$ ($\mu\text{Jy}/\text{beam}$) | $P_{3\text{GHz}}$ ($\mu\text{Jy}/\text{beam}$) | z_{sp} | z_{ph} | Synthesised Beam Area Ratio |
|----|-------------------------------------------------------|-------------------------------------------|--------------------------------------------------|-------------------------------------------------------|--------------------------------------------------------------|-----------------------------------------------------|-----------------|-----------------|--------------------------------|
| 1 | - | 94.0 | 26.6 | 65.9 | 11.6 | - [†] | - | - | 1.43 |
| 2 | - | 87.5 | 25.9 | 69.5 | 12.7 | - [†] | - | - | 1.26 |
| 3 | - | 206.4 | 49.6 | 60.6 | 11.5 | - [†] | - | - | 3.41 |
| 4 | - | 112.0 | 30.7 | 67.8 | 12.2 | - [†] | - | - | 1.65 |
| 5 | 317.4 | 97.0 | 29.9 | 61.6 | 12.2 | - [†] | - | - | 1.57 |
| 6 | 73850.3 | 1250.9 | 145.0 | 75.1 | 8.2 | - [†] | - | - | 16.65 |
| 7 | 73850.3 | 150.2 | 39.6 | 69.7 | 12.7 | - [†] | - | - | 2.16 |
| 8 | - | 110.8 | 32.1 | 62.4 | 12.0 | - [†] | - | - | 1.78 |
| 9 | - | 136.5 | 39.0 | 60.8 | 12.5 | - [†] | - | - | 2.24 |
| 10 | 2676.2 | 610.1 | 64.9 | 133.5 | 11.8 | - [†] | - | - | 4.57 |
| 11 | - | 104.0 | 30.9 | 62.9 | 12.1 | 23.0 [†] | - | 1.24 | 1.65 |
| 12 | 6557.2 | 762.2 | 50.1 | 248.3 | 12.7 | - [†] | - | - | 3.07 |
| 13 | 2759.0 | 139.4 | 37.2 | 66.6 | 12.5 | - [†] | - | 3.9 | 2.09 |
| 14 | 101606.6 | 4687.7 | 223.2 | 371.9 | 15.1 | - [†] | - | - | 10.63 |
| 15 | 637.6 | 99.7 | 21.4 | 96.5 | 11.6 | 55.2 | 2.56 | 2.89 | 1.03 |
| 16 | 1125.9 | 177.4 | 37.6 | 84.1 | 12.6 | 112.5 | - | 0.21 | 2.11 |
| 17 | - | 94.6 | 29.6 | 62.7 | 12.6 | 11.6 | - | - | 1.51 |
| 18 | 390.0 | 138.8 | 40.8 | 59.6 | 12.7 | 42.2 | 0.12 | 0.19 | 2.33 |
| 19 | 390.0 | 142.8 | 35.1 | 76.3 | 12.8 | 42.2 | 0.12 | 0.19 | 1.87 |
| 20 | 260.0 | 374.1 | 57.9 | 103.2 | 12.8 | 97.1 [†] | - | - | 3.63 |
| 21 | - | 159.8 | 37.4 | 81.3 | 13.3 | 83.3 | 0.8 | 0.85 | 1.97 |
| 22 | 3230.3 | 442.7 | 43.5 | 179.4 | 13.1 | 72.3 [†] | - | 2.98 | 2.47 |
| 23 | 323.0 | 384.1 | 29.1 | 263.3 | 12.6 | 265.6 | - | - | 1.46 |
| 24 | 7972.5 | 911.7 | 75.2 | 155.1 | 11.0 | - [†] | - | - | 5.88 |
| 25 | - | 259.5 | 63.0 | 58.4 | 11.7 | - [†] | - | - | 4.44 |
| 26 | - | 226.2 | 59.6 | 56.0 | 12.0 | - | - | - | 4.04 |
| 27 | 267.0 | 202.6 | 23.2 | 186.7 | 12.5 | 15.0 | - | - | 1.09 |
| 28 | - | 125.5 | 35.0 | 65.1 | 12.5 | - [†] | - | - | 1.93 |
| 29 | 77150.1 | 5019.8 | 217.9 | 406.8 | 18.0 | - | - | 1.5 | 9.63 |
| 30 | - | 242.3 | 40.2 | 106.7 | 12.7 | 4911.2 | - | 1.59 | 2.27 |
| 31 | 1244.9 | 323.0 | 43.5 | 132.3 | 13.2 | 149.8 | 0.25 | 0.23 | 2.44 |
| 32 | 580.1 | 115.8 | 32.1 | 69.4 | 12.9 | 85.6 | 0.69 | 0.69 | 1.67 |
| 33 | 1421.2 | 259.4 | 46.1 | 105.3 | 13.8 | 156.3 | - | 0.44 | 2.46 |
| 34 | - | 89.4 | 26.2 | 66.8 | 11.8 | - [†] | - | - | 1.34 |
| 35 | 812.4 | 211.2 | 35.1 | 116.9 | 13.2 | 197.6 | - | - | 1.81 |
| 36 | - | 218.3 | 43.5 | 81.1 | 12.1 | 107.2 | 0.39 | 0.42 | 2.69 |
| 37 | 897.3 | 160.6 | 44.2 | 60.6 | 12.5 | 18.1 | 0.51 | 0.41 | 2.65 |
| 38 | 2536.1 | 277.1 | 33.3 | 161.0 | 13.1 | 75.6 | - | - | 1.72 |
| 39 | 11257.2 | 880.6 | 56.0 | 265.8 | 13.4 | 371.4 | - | 1.14 | 3.31 |
| 40 | 11257.2 | 326.1 | 43.5 | 139.2 | 13.6 | 371.4 | - | 1.14 | 2.34 |
| 41 | 864.1 | 167.6 | 38.1 | 81.3 | 13.0 | 10.6 | - | 2.3 | 2.06 |
| 42 | - | 236.5 | 57.9 | 60.1 | 11.9 | - | - | - | 3.93 |
| 43 | - | 124.1 | 36.3 | 60.2 | 12.4 | - [†] | - | - | 2.06 |
| 44 | 664.9 | 181.5 | 21.2 | 187.3 | 12.0 | 296.9 | - | - | 0.97 |
| 45 | 1014.9 | 216.3 | 54.4 | 62.9 | 12.6 | 47.2 | - | - | 3.44 |
| 46 | - | 103.9 | 32.3 | 62.6 | 12.8 | 7.7 | - | 0.53 | 1.66 |
| 47 | 831.6 | 128.5 | 35.4 | 71.3 | 13.3 | 98.9 | 0.43 | 0.37 | 1.8 |
| 48 | 6151.5 | 323.2 | 58.2 | 81.4 | 12.0 | - [†] | - | 1.67 | 3.97 |
| 49 | 254.1 | 456.4 | 39.5 | 208.8 | 13.0 | 6.0 | - | 0.73 | 2.19 |
| 50 | - | 78.5 | 24.2 | 66.3 | 12.4 | - [†] | - | - | 1.18 |
| 51 | 3025.7 | 435.6 | 83.0 | 74.8 | 12.3 | 804.9 | 0.39 | 0.41 | 5.82 |
| 52 | 3025.7 | 302.8 | 59.7 | 88.5 | 13.8 | 804.9 | 0.39 | 0.41 | 3.42 |
| 53 | 2966.8 | 751.8 | 97.0 | 84.7 | 9.9 | 7.0 | - | 1.51 | 8.87 |
| 54 | - | 137.7 | 40.4 | 60.0 | 12.7 | - | - | - | 2.29 |
| 55 | 3162.8 | 284.4 | 44.3 | 110.6 | 12.9 | 86.4 | - | 2.03 | 2.57 |
| 56 | 38863.3 | 5630.3 | 135.6 | 1554.5 | 19.7 | - [†] | - | 1.25 | 4.06 |
| 57 | - | 91.1 | 24.9 | 72.3 | 11.8 | - [†] | - | - | 1.26 |
| 58 | - | 130.1 | 37.9 | 61.1 | 12.8 | - | - | - | 2.13 |
| 59 | - | 104.2 | 30.6 | 63.8 | 12.2 | - [†] | - | - | 1.63 |

| ID | $P_{140\text{MHz}}$ ($\mu\text{Jy/beam}$) | $S_{1.5\text{GHz}}$ (μJy) | $\Delta S_{1.5\text{GHz}}$ (μJy) | $P_{1.5\text{GHz}}$ ($\mu\text{Jy/beam}$) | $\Delta P_{1.5\text{GHz}}$ ($\mu\text{Jy/beam}$) | $P_{3\text{GHz}}$ ($\mu\text{Jy/beam}$) | z_{sp} | z_{ph} | Synthesised Beam Area Ratio |
|----|------------------------------------------------|-------------------------------------------|--------------------------------------------------|------------------------------------------------|-------------------------------------------------------|----------------------------------------------|-----------------|-----------------|--------------------------------|
| 60 | 5692.1 | 601.9 | 32.2 | 371.7 | 13.2 | 323.8 | - | 0.89 | 1.62 |
| 61 | 977.4 | 151.4 | 37.7 | 73.1 | 12.8 | 49.1 | - | 1.04 | 2.07 |
| 62 | 205.3 | 244.2 | 58.8 | 62.0 | 12.1 | - [†] | - | 1.17 | 3.94 |
| 63 | - | 145.7 | 41.6 | 61.8 | 13.0 | - | - | - | 2.36 |
| 64 | - | 244.2 | 56.2 | 63.3 | 11.8 | - [†] | - | - | 3.86 |
| 65 | 30733.7 | 1256.6 | 157.4 | 71.0 | 8.4 | - [†] | - | 0.97 | 17.69 |
| 66 | 508.1 | 267.5 | 57.1 | 72.8 | 12.5 | 71.3 | 0.92 | 0.88 | 3.67 |
| 67 | 1045.5 | 531.8 | 49.7 | 170.5 | 12.4 | 106.0 [†] | 0.39 | 0.37 | 3.12 |
| 68 | - | 670.7 | 114.4 | 53.6 | 8.5 | - [†] | - | - | 12.52 |
| 69 | - | 113.0 | 33.3 | 64.4 | 12.9 | - [†] | - | - | 1.76 |
| 70 | - | 205.6 | 60.7 | 47.3 | 11.5 | - [†] | - | - | 4.35 |
| 71 | 13914.8 | 4446.9 | 145.8 | 868.3 | 15.9 | - [†] | - | - | 4.86 |
| 72 | - | 715.2 | 69.4 | 140.0 | 11.5 | - [†] | - | - | 5.11 |
| 73 | 143.0 | 461.9 | 95.7 | 48.6 | 9.2 | - [†] | - | - | 9.5 |
| 74 | - | 101.0 | 30.4 | 61.7 | 12.4 | - [†] | - | - | 1.64 |
| 75 | - | 44.9 | 20.1 | 44.9 | 11.4 | - [†] | - | - | 1.0 |
| 76 | 219.2 | 455.4 | 91.7 | 57.2 | 10.3 | - [†] | - | 0.56 | 7.96 |
| 77 | 8464.6 | 1637.2 | 162.0 | 70.2 | 6.7 | - [†] | - | 1.11 | 23.34 |
| 78 | - | 131.0 | 36.6 | 63.1 | 12.4 | - [†] | - | - | 2.08 |

Table 6.4: Details of the 78 sources extracted using PYBDSF from the Lockman Hole field cleaned to a 5σ limit. From left to right, columns are: e-MERLIN Source ID; LOFAR 140 MHz peak flux, $P_{140\text{MHz}}$; e-MERLIN 1.5 GHz total flux, $S_{1.5\text{GHz}}$; e-MERLIN 1.5 GHz total flux error, $\Delta S_{1.5\text{GHz}}$; e-MERLIN 1.5 GHz peak flux, $P_{1.5\text{GHz}}$; e-MERLIN 1.5 GHz peak flux error, $\Delta P_{1.5\text{GHz}}$; VLA 3 GHz peak flux, $S_{3\text{GHz}}$; spectroscopic redshift, z_{sp} ; photometric redshift, z_{ph} ; The ratio of the area of the gaussian fitted by PYBDSF, and the area of the synthesised beam. Peak and total fluxes for e-MERLIN, and their errors, are extracted from the PYBDSF output catalogue. As in Table 6.3, entries with a [†] in the VLA peak flux column are $>10'$ from the phase centre.

6. STUDYING THE LOCKMAN HOLE



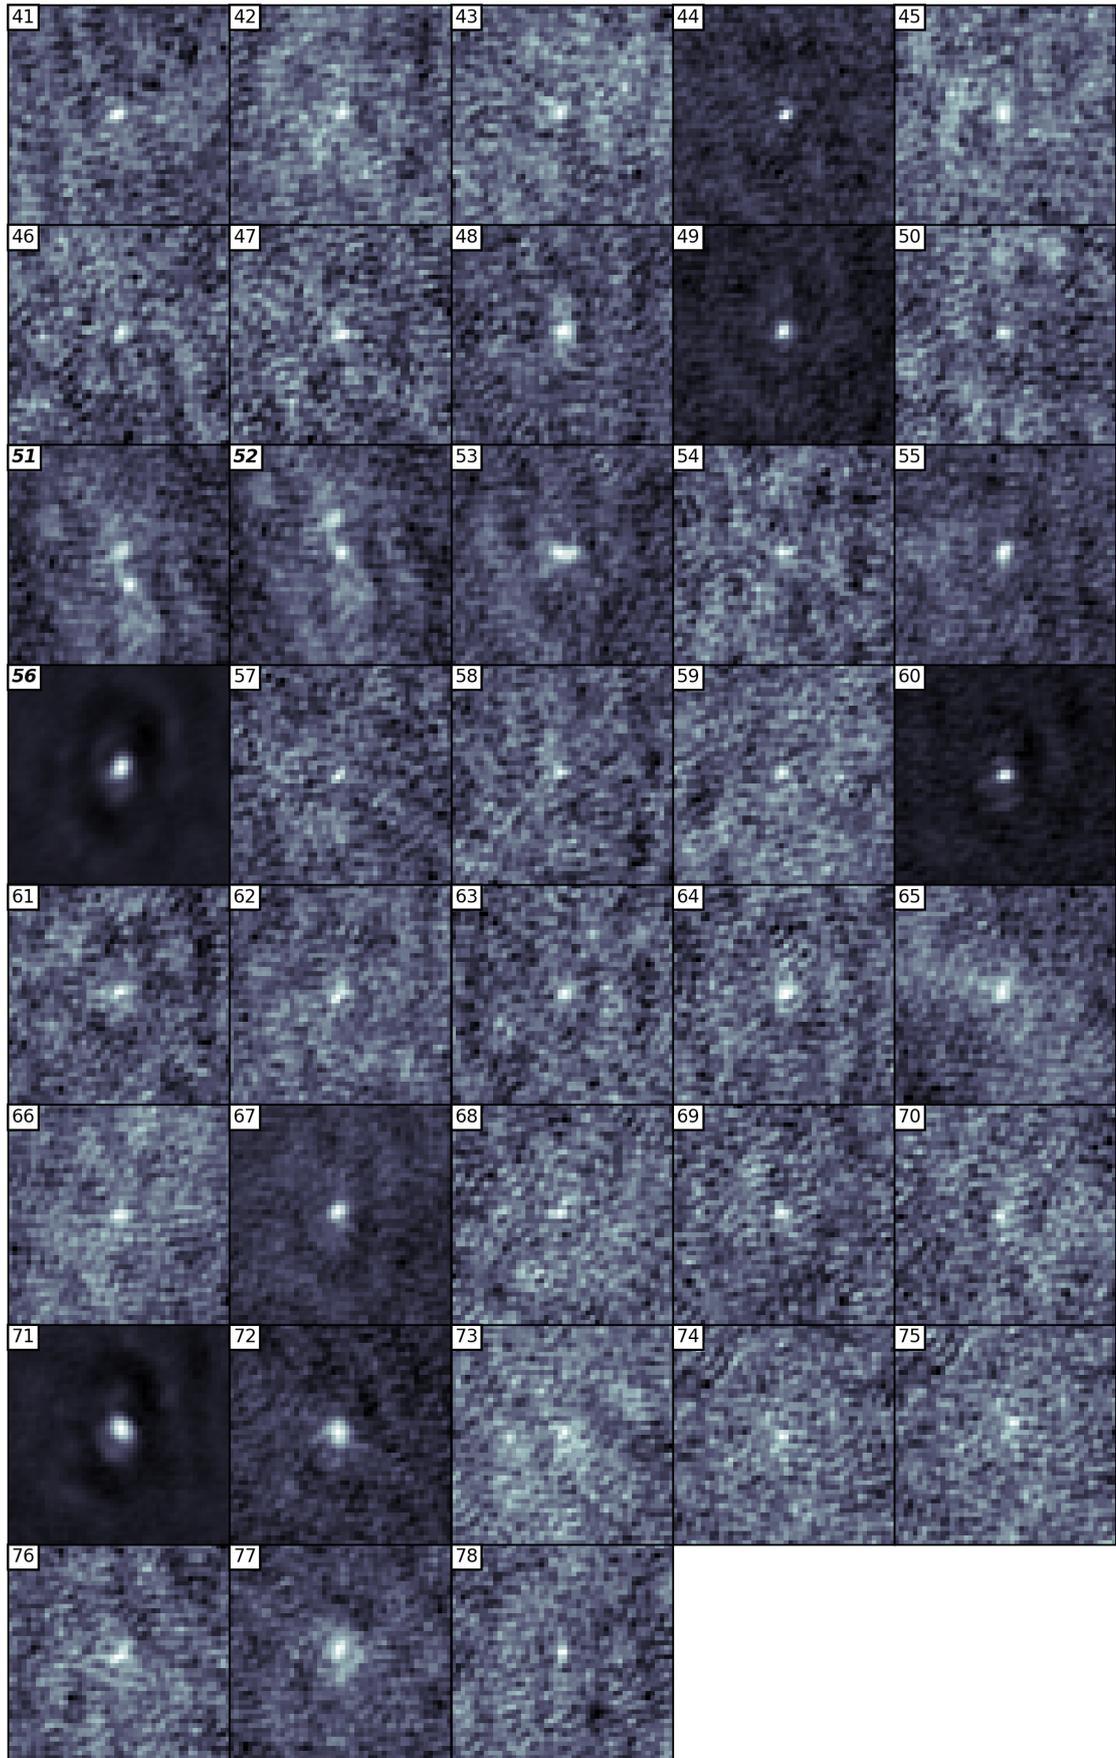


Figure 6.5: Cutouts of all 78 sources extracted from the Lockman Hole image. Labels given in the top left of each image correspond to the IDs in [Table 6.4](#), where labels given in bold italics also correspond with sources studied in detail in [section 6.4](#). Each cutout is 48x48 pixels in size, covering a patch of sky 3.36x3.36'' in size. The size of the synthesised beam is 0.19x0.14'', which is approximately 3x2 pixels in these images.

6.4 Source Studies

To outline the information contained in the e-MERLIN image of the Lockman Hole, a few sources are selected from the catalogue to be studied in detail. These sources cover a range of radio flux densities and morphologies, and are analyzed using data from the LOFAR and VLA surveys. Redshifts are primarily taken from the SWIRE photometric redshift catalogue. Spectral indices are defined by $S_\nu \propto \nu^\Gamma$, and specific measurements are quoted using a lower and upper format in MHz. For example, Γ_{140}^{1500} is a spectral index calculated using a 140 MHz and a 1.5 GHz measurement. Since most of the sources are partially or fully resolved by e-MERLIN, total fluxes are generally used over peak fluxes. For distance calculations I assume a flat universe with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_\Lambda = 0.7$, and $\Omega_m = 0.3$.

6.4.1 S51 + S52

The sources referred to here as S51 and S52 (hereafter simply S51), are collectively identified in the NRAO VLA Sky Survey (NVSS) as NVSS J104540+585725, which appears to be an AGN with a resolved jet structure (see [Figure 6.6](#)). The brightest component is located at $\alpha = 10^{\text{h}}45^{\text{m}}39.8^{\text{s}}$, $\delta = +58^{\text{d}}57^{\text{m}}29.7^{\text{s}}$. The spectroscopic redshift is $z_{\text{sp}} = 0.39$. It is also detected in both the LOFAR and VLA surveys. It is known as TV16 J104539+585730 in the VLA catalogue and is identified as a complex source by [Vernstrom et al. \(2016b\)](#) (see [Figure 6b](#) of that paper). They identify their complex source as belonging to one object: likely an AGN with extended jets. [Figure 6.6](#) shows the resolved central component of their complex object.

The total flux of the southern source is $303 \mu\text{Jy}$, and $436 \mu\text{Jy}$ for the northern source which appears to be slightly more extended. It is unknown which of the two sources is the AGN, and since resolution-matched studies are not available it is difficult to determine this from L-band morphology alone. The peak flux of the source measured in the VLA survey is $805 \mu\text{Jy}/\text{beam}$, and $3026 \mu\text{Jy}/\text{beam}$ in the LOFAR survey. There may be a faint third source located about halfway along the bridge of emission connecting the two sources, though this possible third source is not detected in the source catalogue and could be a noise fluctuation (see [subsection 6.2.1](#)); it is therefore ignored.

The 3 GHz in-band spectral index $\Gamma_{3000} = -0.6$ quoted in the [Vernstrom et al. \(2016b\)](#) catalogue suggests optically thin synchrotron emission characteristic of extended emission in radio galaxies. The spectral index $\Gamma_{140}^{3000} = -0.43$ that can be calculated by inter-

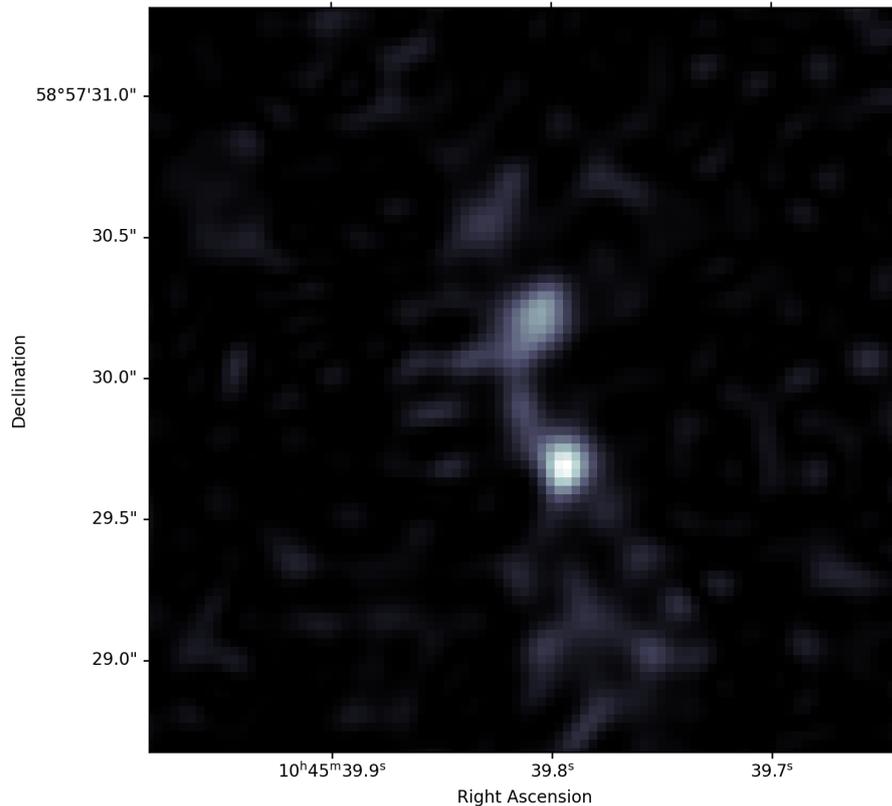


Figure 6.6: The two sources collectively referred to as S51. They are separated by an angular distance of $0.5''$, which corresponds to 2.6 kpc at the epoch of the source. Values in this image are squared to emphasise the pair of sources.

polating between the LOFAR and VLA peak fluxes suggests a flatter spectrum, though summing together the two e-MERLIN measurements of the total flux, and using the LOFAR peak flux, gives a similar spectral index to Γ_{3000} of $\Gamma_{140}^{1500} = -0.59$. It remains possible that some contributions from the jet are being resolved out by e-MERLIN so the real slope may be flatter. Assuming the radio bridge accurately tracks the trajectory of the jet, the AGN must have some degree of precession to produce the observed morphology, which agrees with the structure of emission seen with the VLA. The angular distance between the two sources is approximately $0.5''$, which corresponds to 2.6 kpc at the epoch of the source. In a low inclination galaxy, this would likely place the two sources well within the galactic disk. The bridge of emission connecting the two sources strongly suggests a radio galaxy that may have only entered its duty cycle relatively recently. S51 is an excellent candidate for follow-up, focused observation to investigate the properties of the jet structure, precession, and origin. At this angular scale, optical observations that could collocate the galactic component with one of the radio sources are difficult to acquire, so follow-up observations should account for this.

6.4.2 S56

One of the brightest sources in the e-MERLIN image is S56, shown in [Figure 6.7](#). It is partially resolved with a peak flux of $P_{1.5\text{GHz}} = 1.57\text{mJy/beam}$ and a total flux of $S_{1.5\text{GHz}} = 5.63\text{mJy}$. The source is located at $\alpha = 10^{\text{h}}45^{\text{m}}28.3^{\text{s}}$, $\delta = +59^{\text{d}}13^{\text{m}}26.6^{\text{s}}$ and is studied in detail by [Weedman et al. \(2006\)](#) as source A8. They quote a spectroscopic redshift of $z_{\text{sp}} = 2.31 \pm 0.2^2$, and classify the source as an AGN as it passes their luminosity criteria at both X-rays and in the mid-infrared, suggesting that whatever is powering the source is still X-ray luminous even after dust obscuration. The presence of a strong, unresolved radio component in the e-MERLIN image supports the AGN classification.

Unfortunately, S56 is 13' from the phase centre which places it right at the edge of the VLA field, so no match is made. However, the source does have a very strong component in the LOFAR catalogue with a 140 MHz peak flux of 38.9 mJy/beam. This makes S56 the second strongest 140 MHz source within the e-MERLIN Lockman Hole field.

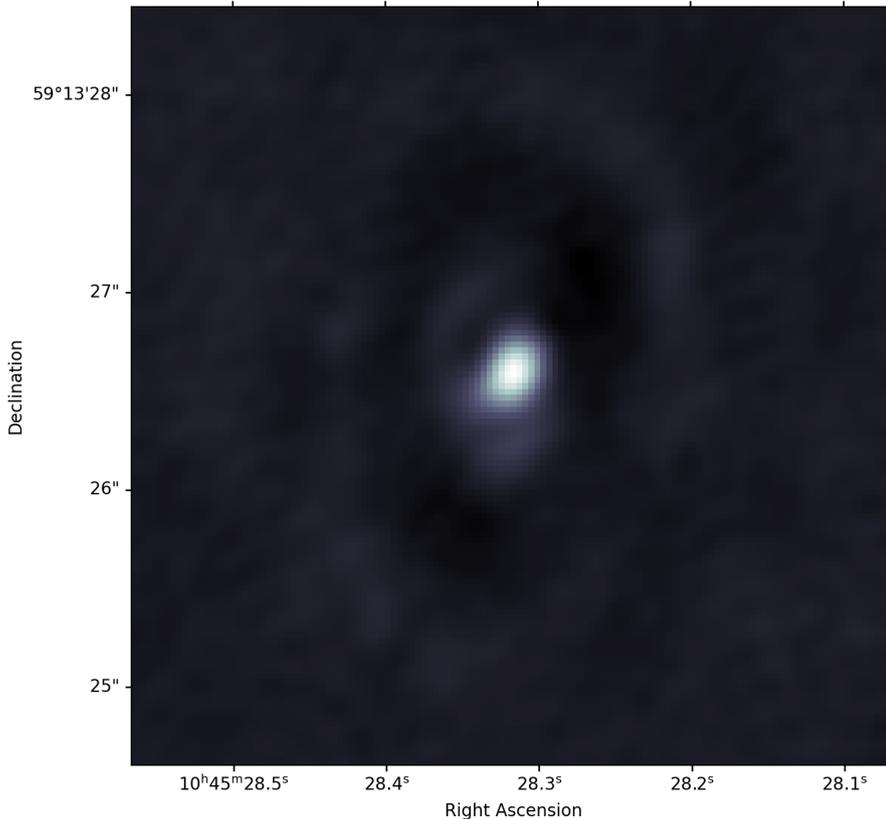


Figure 6.7: The source referred to as S56. It is a strong source that demonstrates the cleaning artefacts present around strong sources of emission in the e-MERLIN Lockman Hole image.

At such a high redshift, it is unlikely any morphology would be detectable even by e-MERLIN, so the structure around the source is very likely an improper subtraction

²Note that this spectroscopic redshift significantly differs from the photometric redshift given in [Table 6.4](#)

of the PSF from a strong source during the cleaning process, given the ‘negative bowl’ surrounding it. This conclusion is also reinforced by the presence of a similar pattern around other strong sources such as S71 and S60. Alternatively, the artefacts surrounding this source and others may be resolved hotspots in extended components, but given the pattern similarity between multiple strong sources this is difficult to determine.

The 1.5 GHz radio power and steep-spectrum are characteristic of extended synchrotron emission in a radio galaxy, which agrees with prior classification of AGN by [Weedman et al. \(2006\)](#) and [Vernstrom et al. \(2016b\)](#).

6.4.3 S29

The source referred to here as S29 is a resolved source with a total flux of $S_{1.5\text{GHz}} = 5\text{mJy}$ (see [Figure 6.8](#)). The source is located at $\alpha = 10^{\text{h}}46^{\text{m}}24.8^{\text{s}}$, $\delta = +59^{\text{d}}04^{\text{m}}46.4^{\text{s}}$. As with S56, this source has an extremely strong unresolved 140 MHz companion with peak flux $S_{140\text{MHz}} = 68.8\text{mJy}$. While missing the matching criteria described in [section 6.3](#), this source can in fact be related to the complex source TV16J104624+590522 from the VLA catalogue, which is identified as an FR II quasar with a spectroscopic redshift of $z = 3.63$ ([Véron-Cetty & Véron, 2010](#); [Vernstrom et al., 2016b](#)). The reason for the criteria being missed is because e-MERLIN is viewing the radio lobe approximately $\frac{3}{4}'$ south of the VLA catalogue source, shown in [Figure 6a](#) of [Vernstrom et al. \(2016b\)](#). There is also a clear detection of this source in the Faint Images of the Radio Sky at Twenty centimetres (FIRST) survey ([Becker et al., 1995](#)) at 1.4 GHz. That source is unresolved with a peak flux $P_{1.4\text{GHz}} = 11.2\text{mJy}$. It is possible to view both the southern radio lobe (S29) and the northern radio lobe in the 140MHz image. The two components are separated by $\sim 80''$ (see [Figure 6.9](#)).

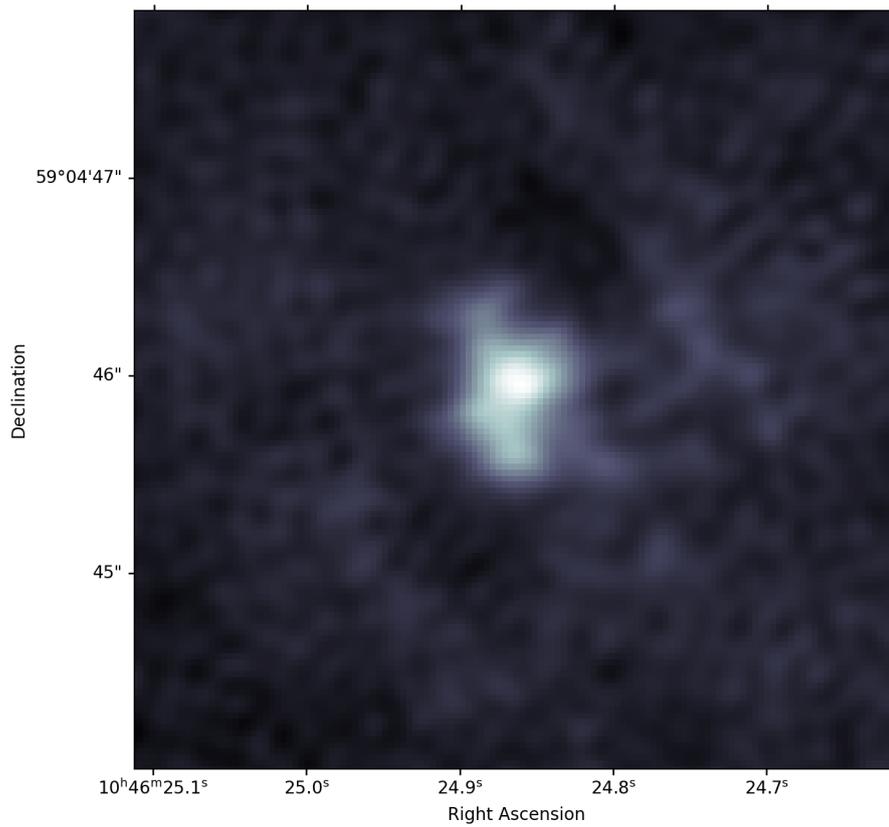


Figure 6.8: The source referred to as S29. It is related to the southern radio lobe of source TV16J104624+590522 in the VLA survey. The entire source can be seen in Figure 6a of [Vernstrom et al. \(2016b\)](#), and in [Figure 6.9](#) at 140MHz.

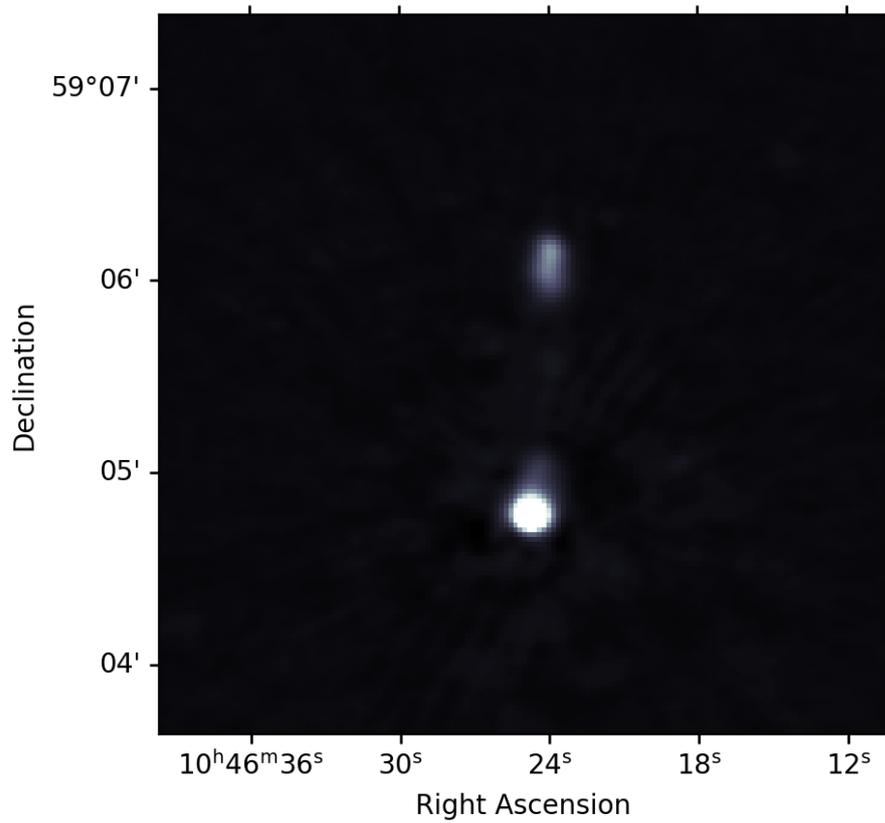


Figure 6.9: The southern source in this image is the 140 MHz match to source S29. The northern source is the northern radio lobe of source TV16J104624+590522, shown in Figure 6a of [Vernstrom et al. \(2016b\)](#). The actual location of TV16J104624+590522 is between the two sources in this image. The image is a cutout of the LOFAR 2 m Sky Survey ([Shimwell et al., 2017](#)).

6.4.4 S39 + S40

This section studies the sources S39 and S40, which are collectively referred to as just S39. They are two partially resolved sources with $< 0.3''$ separation (see [Figure 6.10](#)). The eastern source is brighter with a total flux of $881\mu\text{Jy}$, and the western source has a total flux of $326\mu\text{Jy}$. The brightest source is located at $\alpha = 10^{\text{h}}46^{\text{m}}04.1^{\text{s}}$, $\delta = +59^{\text{d}}08^{\text{m}}27.7^{\text{s}}$. The source has a photometric redshift of $z_{\text{ph}} = 1.14$, which corresponds to a projected separation of ~ 2.5 kpc. Assuming the redshift measurement is accurate, it's possible that e-MERLIN is viewing radio hotspots within the galactic disk.

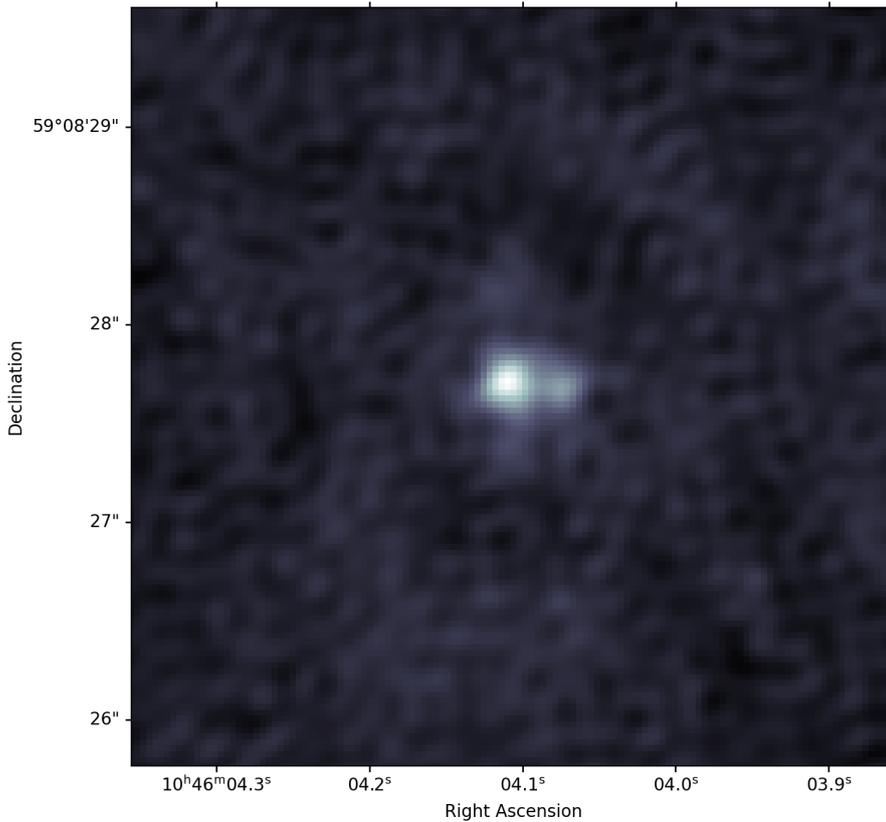


Figure 6.10: The sources S39 and S40, collectively referred to as S39. They have a small projected separation of < 2.5 kpc, demonstrating the resolving power of e-MERLIN.

The source has a 140 MHz peak flux of 9.8 mJy/beam and a 3 GHz peak flux of $371\mu\text{Jy/beam}$, giving spectral indices $\Gamma_{140}^{1500} = -0.94$ by summing the total fluxes, and $\Gamma_{140}^{3000} = -1.11$. These are in broad agreement both with each other, and with the value of $\Gamma = -1.3$ calculated solely from the S-band data presented by [Vernstrom et al. \(2016b\)](#). These values suggest extended synchrotron emission, and since e-MERLIN is able to observe two radio components on galactic scales it is possible it is caused by star formation. Moreover, this particular source has a similarity to the source studied in [chapter 5](#), so a bound pair of black holes is also possible, though this conclusion would contest with

the spectral indices calculated earlier. The source does not have an associated entry in the Chandra Source Catalog (Evans et al., 2020), reducing (though not eliminating) the likelihood of an AGN. Overall, more evidence is needed to classify this particular object. Follow up observations of this source on longer baseline VLBI scales would test the compactness of these two sources, and potentially reveal further SF components, a bound system, or clearer jet activity. The source is too distant for meaningful morphological information to be gained from further optical observation, but similar optical evidence could be gathered as in Woo et al. (2014) to find support for or against a star formation origin. This source is an excellent indicator of the ability of e-MERLIN to resolve even the most closely separated radio sources.

6.4.5 S6 + 7

The objects S6 and S7 are located quite close together, and the morphology of S7 suggests it may be a jet component with multiple hotspots. It is unknown whether the parent AGN is S6, S7, or remains undetected (see Figure 6.11). The northern radio source, S6, is located at $\alpha = 10^{\text{h}}47^{\text{m}}41.8^{\text{s}}$, $\delta = +58^{\text{d}}48^{\text{m}}17.0^{\text{s}}$. The two sources are separated by approximately $4.6''$, though unfortunately, neither of these sources has a match in the Rowan-Robinson et al. (2008) catalogue so calculating a distance is not possible with the available information. A match can also not be found in the VLA survey since it falls outside the survey area. The southern source, does not appear to have been properly fit by PYBDSF as the ratio of the fitted gaussian area to the area of the synthesised beam is too small given its structure. Though this is understandable given its complexity. It is likely then, that the total flux of S7 has been underestimated. The northern source clearly has some resolved emission. The total flux from the northern source is $S_{1.5\text{GHz}} = 1.25\text{mJy}$. Both sources are matched to a single very strong 140 MHz source with peak flux $P_{140\text{MHz}} = 73.9\text{mJy/beam}$, and are close enough together to be encompassed by a single synthesised beam in the LOFAR survey. Assuming the 140 MHz flux is divided evenly between the two e-MERLIN sources, and using the total fluxes calculated by PYBDSF, spectral indices can be approximated as $\Gamma_{140}^{1500} = -1.72$ for the northern source S6, and $\Gamma_{140}^{1500} = -2.61$ for the southern source S7 (though this is very likely a bad estimation due to the incorrect fit). This evidence points to the presence of a radio galaxy, with S7 being a likely jet component given its structure.

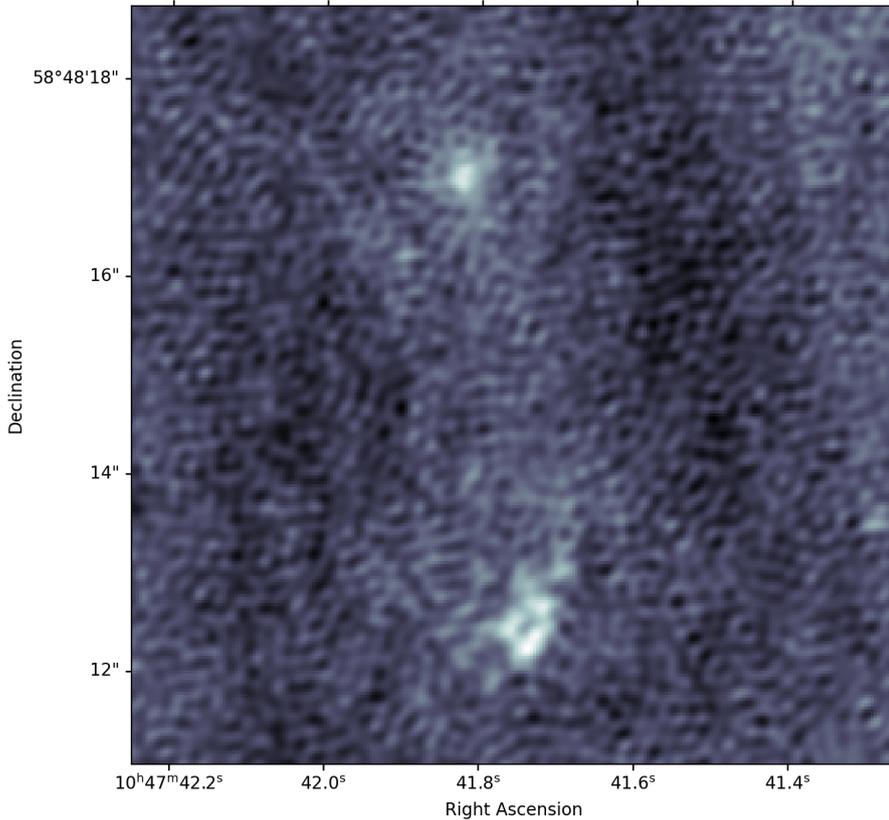


Figure 6.11: The sources S6 (northern) and S7 (southern). The two sources separated by less than the LOFAR synthesised beam width, and are therefore encompassed by a single strong $P_{140\text{MHz}} = 73.9\text{mJy/beam}$ source in the LOFAR survey.

6.5 Differential Source Counts

For radio interferometers, counting the number of sources detected within a survey must come with a recognition of the biases introduced by the instrumentation; primary beam width, baseline length, and surface brightness sensitivity are all factors that can significantly impact the number of sources detected. Because of these factors, certain instruments are better suited to producing source catalogues that more accurately represent the quantity of sources in a certain patch of sky at a certain frequency. LOFAR, for example, is a powerful 140 MHz sky surveyor because of its wide field of view, and strong surface brightness sensitivity. In contrast, e-MERLIN is less capable of producing sky catalogues because it is, first and foremost, a VLBI facility that is less capable of detecting fainter, low surface brightness phenomena. It is expected that e-MERLIN is biased in its detections towards compact radio emitters like jetted and non-jetted AGN, since SFGs are more likely to resolve into more and fainter sources, increasing the likelihood that they slip below the noise limit, whereas AGN will generally remain as at most a few compact sources at higher resolutions.

Despite these limitations, it is useful to compare the source counts to other studies, and radio astronomers have settled on a shared metric for tracking the radio source counts at fainter flux densities: the Differential Source Counts (DSCs). A good introduction to the DSC is provided in section 4 of [Vernstrom et al. \(2016b\)](#) which, for completeness, I summarise here with similar notation. The DSC fundamentally measures the *number* of sources *per* Jansky, *per* steradian. It is, therefore, most simply written as:

$$\frac{dN}{dS_a} = \frac{n_a}{\Omega_a \Delta S_a} \quad (6.1)$$

where n_a is the number of sources in bin number a , Ω_a is the area over which sources are detectable, and ΔS_a is the width of bin a . The values of n_a and ΔS_a are trivial, but determining the value of Ω_a requires the consideration of more complex factors such as the shape of the primary beam. [Vernstrom et al. \(2016b\)](#) define the angle over which a source can be detected as:

$$\Theta_a = \sqrt{\frac{\ln P_a}{\ln 2}} \frac{\theta_{\text{FWHM}}}{2} \quad (6.2)$$

where the quantity P_a is the signal to detection limit ratio and θ_{FWHM} is the FWHM of the primary beam. As most of the sources in the e-MERLIN field are partially or fully resolved, P_a is defined as $P_a = \langle S \rangle_a / 5\sigma$ in this work, where $\langle S \rangle_a$ is the mean of the total fluxes within bin a . For e-MERLIN the primary beam width is taken as $\theta_{\text{FWHM}} = 30'$. From Θ_a , an area can be easily computed as $\Omega_a = \pi \Theta_a^2$, giving the final equation for the DSCs:

$$\frac{dN}{dS_a} = \pi \left(\frac{\ln 2}{\ln P_a} \right)^{3/2} \frac{4n_a}{\theta_{\text{FWHM}}^2 \Delta S_a} \quad (6.3)$$

Note that [Vernstrom et al. \(2016b\)](#) introduce weightings to the DSCs that are not applied here, but these weightings only seem to make a significant difference in lower flux bins $< 50\mu\text{Jy}$. A final correction of $S^{5/2}$ is commonly applied to [Equation 6.3](#) to correct for an inherent increase in the number of sources with increasing distance from the observer. The result is known as the Euclidean-normalized DSCs, which are plotted in [Figure 6.12](#) for the e-MERLIN catalogue. The counts are plotted using a range of number of bins A , to illustrate the spread in making different choices for the bins. In every case, the upper and lower limits for the bins are $40\mu\text{Jy}$ and 6mJy .

As expected, e-MERLIN underestimates the source count in all flux bins compared to other instruments due to the reasons discussed above. There is a flattening of the

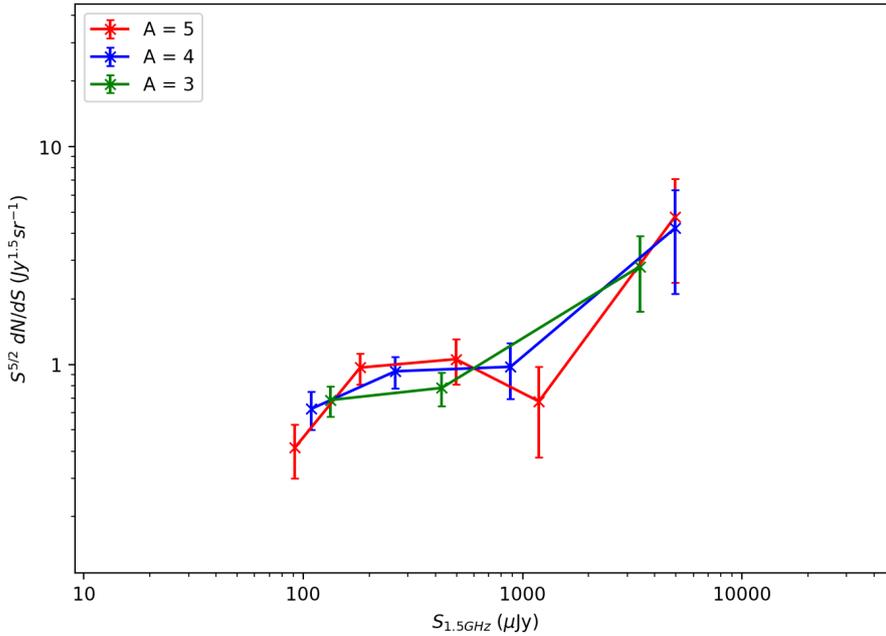


Figure 6.12: The DSCs for the e-MERLIN Lockman Hole catalogue. Measurements are repeated for number of bins $A = 3, 4,$ and 5 , to illustrate the spread from different choices for the bins. The upper and lower flux density limits are the same in each case. Errors on the source counts are Poisson errors (i.e. $1/\sqrt{N}$).

counts between $\sim 100 - 1000 \mu\text{Jy}$, followed by an increase at the highest flux densities. The overall trend in this range is largely in alignment with the predictions from many other studies (see Figure 4.1). At $\lesssim 100 \mu\text{Jy}$, these results show a decrease in the source count at higher fluxes compared to other studies, which is again expected since the fainter region is dominated by non-jetted AGN and SFGs; e-MERLIN preferentially detects only one of these classes. Therefore, it is possible that the number of SFGs in this flux density range could be estimated by examining the gap between the e-MERLIN DSCs and the Vernstrom et al. (2016b) DSCs obtained using the VLA. This gap is most apparent in Figure 6.13, which shows the DSCs calculated for the full VLA catalogue. These counts are calculated using the same method as above, with $\sigma = 1 \mu\text{Jy}$, $\theta_{\text{FWHM}} = 15'$, and using integrated fluxes scaled to 1.5 GHz using a spectral index $\Gamma = -0.7^3$. The comparison shows clearly the missing SFG population at $\lesssim 1000 \mu\text{Jy}$.

6.6 Discussion

The e-MERLIN Lockman Hole image provides a clear census of the 1.5 GHz radio population at brightnesses $\gtrsim 50 \mu\text{Jy}$, at VLBI scales, and over a redshift range $0.1 < z < 3$

³This is the approach taken by Vernstrom et al. (2016b) to compare their results to 1.4 GHz source counts.

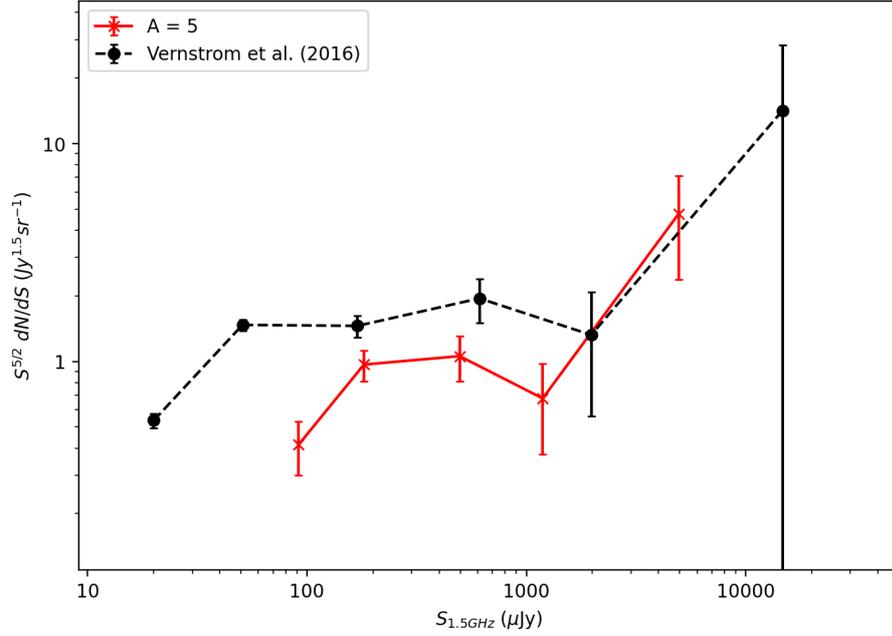


Figure 6.13: The DSCs for the e-MERLIN field with $A = 5$ bins, with the DSCs from the [Vernstrom et al. \(2016b\)](#) catalogue plotted as a dashed black line. The method used to calculate these counts are the same in both cases. Errors on the source counts are Poisson errors (i.e. $1/\sqrt{N}$).

using estimates from [Rowan-Robinson et al. \(2008\)](#). The primary distinction between this study and other studies that must be taken into account is the significantly higher resolution which increases information about source morphology at the cost of a decreased sensitivity to low surface brightness sources, and the population of SFGs expected to dominate in the lower flux bins. As expected, the e-MERLIN DSCs are smaller than equivalent measurements of the DSCs using the VLA. It may be possible to estimate the number of SFGs by examining the gap in the DSCs between these two measurements, but this should be done carefully. Accurately characterising the degree to which e-MERLIN samples AGN over SFGs requires the study of an already classified population; preferably a larger one than is presented in this catalogue. However, the classification of sources in this catalogue that may allow a study such as this is the subject of future research. Regardless of whether this study is conducted, it should be a priority for future work to obtain more e-MERLIN observations of the same target to further lower the detection limit, ideally to seek parity with the sensitivity of the [Vernstrom et al. \(2016b\)](#) VLA observations of the same area. The observations presented here were crippled by large amounts of data lost to RFI, so the final sensitivity was worse than desired. As a result, the e-MERLIN study of the Lockman Hole is a perfect example of a situation where more effective RFI mitigation/excision methods are critically important. Succeeding in doing this, and obtaining more observations, would enable a richer comparison with the 3 GHz catalogue in

particular, but also other radio studies more generally. It may also be desirable to expand the search area to include more sources in the $100 < S_{1.5\text{GHz}} < 1000\mu\text{Jy}$ flux bins, to more accurately measure the gap in source counts observed by e-MERLIN.

The biggest step towards classifying sources in the Lockman Hole would come from combining the e-MERLIN uv data presented here with the [Owen & Morrison \(2008\)](#) uv data, using the same method of [Muxlow et al. \(2020\)](#) in the GOODS-N field. The resulting dataset would contain the resolving power of e-MERLIN without the downside of reduced sensitivity to SFGs, allowing a more complete catalogue of sources to be classified either through the morphological decomposition technique demonstrated by [Muxlow et al. \(2020\)](#), or some other method. This study would be invaluable in mapping the faint radio population within the Lockman Hole, and would highlight any discrepancies due to cosmic variance in both fields. To go further, more and deeper observations at low frequencies $< 1\text{GHz}$ would provide further context to the 1.5-3 GHz population, and specifically target sources with flatter ($\Gamma \sim 0$) spectral indices that are missed due to the sensitivity constraints of the current 140 MHz LOFAR observations. In general, forming a stronger pool of evidence that can be used in cosmological studies requires more and higher quality radio observations of deep fields such as the Lockman Hole, GOODS-N, and the Hubble Deep Fields to settle the disagreement in the DSCs and to accurately describe the number of non-jetted AGN and SFGs as a function of redshift. New facilities like the SKA are a critical step towards forming this pool of evidence, by probing further down to the $\sim \mu\text{Jy}$ limit and perhaps even lower into the nJy range. This range of extremely faint flux densities representing the higher redshift population is something the next generation of radio telescopes will explore, and is currently the largest gap in our knowledge of the faint radio population.

Chapter 7

Conclusions

This work has explored a few topics: experimenting with machine learning to flag interference in astronomical radio observations, investigating a rare sub-kpc scale binary black hole candidate, and conducting a deep survey of the Lockman Hole with the e-MERLIN radio telescope. The first experiment was conducted with the aim of improving the outcome of the last, as interference is a growing problem within radio astronomy and new, more advanced techniques are needed.

7.1 The RFI Experiment

The RFI experiment used manually flagged data from the e-MERLIN telescope to train a GAN to automatically produce flag masks. Using a test set of similarly manually flagged data, the trained GAN achieved a True Positive Rate of 36.4%, a True Negative Rate of 94.5%, and finally an F_1 -score of 40%. The model was also tested using a range of simulated inputs, where perfect knowledge of RFI is achievable. On simulated images containing RFI, the GAN achieves a mean F_1 -score of 64.3% before the performance sharply decreases at high dynamic ranges. The GANs performance is found not to be advantageous over currently existing statistical methods for a number of possible reasons including: lack of precision in the training set, an inefficient network architecture, lack of variation in the training set, and the common problem (not just to this technique) of flagging the most faint instances of RFI. It is therefore concluded that, while manually identified RFI should be included in the training process of any machine learning method in some form, exclusive use does not appear to give desirable results for this particular network architecture. Furthermore, while the adversarial model may be a desirable frame-

work, it is not clear that the U-Net/CNN combination is optimal for the task of identifying instances of RFI.

Further experiments may improve the pre-processing pipeline to better exaggerate instances of RFI, or they may explore the possibility of modifying the generator side of the GAN to be better suited to producing RFI masks, using knowledge about the range of RFI morphologies to improve performance. It may also be beneficial to broaden the problem to more than two classes; instead of a binary classification problem between RFI and ‘not RFI’, including more categories such as ‘unsure’ or ‘bordering RFI’ may reduce the tension in the adversarial model and produce better results.

Machine learning techniques more generally are uniquely suited to the problem of RFI flagging, and deserve further testing. Radio observatories necessarily occupy a unique RFI environment that will change with infrastructure developments surrounding it; machine learning methods trained on locally sourced visibilities are able to learn about that particular RFI environment, and use this knowledge to improve performance. This is a clear advantage of machine learning models over statistical tools. Furthermore, machine learning techniques do not necessarily require the same amount of parameter optimisation as statistical methods to achieve the best result during inference. This can greatly simplify the problem of flagging radio visibilities for the astronomer by outsourcing the task of constructing an effective training set to the radio observatory, where knowledge of the local RFI environment is much greater. Machine learning techniques are also a promising alternative to meet the challenge of a significantly increased volume of satellite RFI; a problem that will affect the operation of the next generation of radio observatories.

7.2 The Binary Black Hole Study

Following on from an optical study by [Woo et al. \(2014\)](#), a radio investigation of a sub-kpc scale binary black hole candidate was conducted. The investigation successfully identified two distinct, but faint sources of emission within the object that directly align with two stellar cores identified from imaging and integral field unit spectra. The overlap between the two optical sources and the two radio sources is highly indicative of a bound pair of AGN. The redshift of the source is $z = 0.35$, and the two radio sources are separated by $0.19 \pm 0.06''$, which translates to 0.95 ± 0.29 kpc at the epoch of the source. Unfortunately, the PSF of the e-MERLIN observations skews the radio emission directly along the position angle of the binary source, making separation measurements extremely difficult.

To investigate this object further, an observing proposal was submitted to the European VLBI Network (EVN) which has a much greater resolving power than e-MERLIN and may have revealed jet activity, confirming the dual AGN hypothesis. However, no radio emission was detected from the source at all by the EVN. Given how faint the object is, the null detection by the EVN does not preclude the presence of radio emission as the object may have simply slipped under the noise limit. Several problems occurred during observation which lowered the sensitivity of the final image, raising this probability. Further EVN observations may detect the two sources observed by e-MERLIN and reveal the origin of the radio emission, though the available evidence suggests that the original radio detection was the product of two points of extended synchrotron emission: not AGN. At such small angular scales finding useful ancillary data is difficult, but X-ray observations in particular may at least confirm the existence of a single AGN.

7.3 The Lockman Hole Study

The final part of this work involved processing 18 e-MERLIN 1.5GHz observations of the patch of sky known as the Lockman Hole. This area was chosen because of the wealth of multi-wavelength data available not just at radio frequencies. The data was calibrated and combined to produce a high resolution wide-field image of the Lockman Hole $\sim 30'$ in size, with a mean sensitivity of $\sigma = 12\mu\text{Jy}/\text{beam}$. From this image, a catalogue of 78 sources above a 5σ limit was produced, and $\sim 50\%$ of these sources were cross-matched with two other studies of the same area at lower (140MHz) and higher (3GHz) radio frequencies. Redshifts were also identified for 35 sources. Using this information, 5 sources were investigated in detail, to exercise the capabilities of the e-MERLIN image. Differential Source Counts (DSCs) are calculated using the e-MERLIN catalogue, and compared to the DSCs obtained from the [Vernstrom et al. \(2016b\)](#) VLA catalogue of the Lockman Hole. Naturally, because of bias towards detecting compact, high surface brightness sources (AGN in particular) the counts are underestimated in all but the highest flux bin. Though, the findings are consistent with an observed faint radio population with a suppressed SFG component, and are consistent with the overall trend observed by the VLA in the $\lesssim 1\text{mJy}$ range.

Further work on the e-MERLIN survey of the Lockman Hole should prioritise combining the data with the VLA survey of the Lockman Hole by [Owen & Morrison \(2008\)](#), to conduct a similar study to that of [Muxlow et al. \(2020\)](#). This would allow a strong

comparison to be made between the populations of the GOODS-N field and the Lockman Hole, and highlight any differences in the source count due to cosmic variance. In addition to this, further observations by e-MERLIN would continue to bring down the noise limit, and allow an investigation of the very faint $\lesssim 10\mu\text{Jy}$ population where most of the disagreement in source count lies. However reaching this flux density limit would require significantly more observing time, because of the impact of RFI. Further e-MERLIN observations would simultaneously allow a deeper investigation of the selected source S51, and potentially detect radio emission from the host galaxy. The Lockman Hole is also a strong candidate for observation by the next generation of radio telescopes. These observations would allow the $\sim \mu\text{Jy}$ range to be explored with a high degree of accuracy, and should settle the disagreement in source counts. Though, the growing impact of satellite RFI may impede the next generation of observations as well; further highlighting the need for better methods of RFI removal.

Overall, the objective should be to further constrain the properties of the faint radio population within the Lockman Hole over a range of radio frequencies, with the end goal of using the properties of the SFG/AGN population to explore the history of star formation and black hole formation in this field, which in turn will inform all theories of cosmology — past, present, and future.

Acronyms

CASA Common Astronomy Software Applications. 68, 117, 118

AGN Active Galactic Nuclei. 14, 15, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 95, 96, 97, 99, 112, 113, 132, 133, 134, 140, 142, 143, 144, 146, 147, 148

AIPS Astronomical Image Processing System. 67

ARQZWA Australian Radio Quiet Zone WA. 60

BLR Broad-line Region. 88

CCD Charge-Coupled Device. 17

CGAN Conditional Generative Adversarial Network. 54, 55

CMB Cosmic Microwave Background. 9

CNN Convolutional Neural Network. 48, 49, 50, 51, 52, 66

DCGAN Deep Convolutional Generative Adversarial Network. 53, 54

DSC Differential Source Count. 141, 142, 143, 144, 147

E-CDFS Extended Chandra Deep Field South. 96, 97

e-MERLIN *enhanced* Multi Element Remotely Linked Interferometer Network. 15, 17, 20, 21, 42, 66, 72, 96, 99, 112, 113, 114, 115, 116, 118, 120, 121, 124, 125, 127, 132, 133, 134, 135, 138, 139, 140, 141, 142, 143, 144, 145, 146, 147, 148

EHT Event Horizon Telescope. 11, 84

EVN European VLBI Network. 66, 99, 112, 113, 114, 147

FAST Five-Hundred Aperture Spherical radio Telescope. 10, 56, 60, 66

FFT Fast Fourier Transform. 30, 124

FIRST Faint Images of the Radio Sky at Twenty centimetres. 135

FRB Fast Radio Burst. 66

FWHM Full Width at Half Maximum. 42, 124, 125, 141

GAN Generative Adversarial Network. 51, 52, 53, 54, 55, 56, 58, 66, 68, 69, 75, 76, 77, 78, 79, 80, 112, 117, 145, 146

GOODS-N Great Observatories Origins Deep Survey - North. 15, 96, 115, 144, 148

HST Hubble Space Telescope. 10, 11

ISM Interstellar Medium. 22, 115

ITU International Telecommunications Union. 58, 59, 81

JWST James Webb Space Telescope. 11, 88

LOFAR LOw Frequency ARray. 19, 29, 62, 63, 124, 125, 132, 133, 134, 137, 140, 144

MS Measurement Set. 68

NRAO National Radio Astronomy Observatory. 10, 18, 89

NVSS NRAO VLA Sky Survey. 132

PPV Positive Predictive Value. 75

PSF Point Spread Function. 38, 39, 40, 125, 135, 146

RFI Radio Frequency Interference. 31, 32, 57, 58, 60, 61, 62, 63, 64, 65, 66, 68, 72, 73, 74, 75, 76, 77, 78, 79, 80, 82, 96, 118, 143, 145, 146, 148

RQZ Radio Quiet Zone. 60

SED Spectral Energy Distribution. 57

SETI Search for Extraterrestrial Intelligence. 44

SFG Star-forming Galaxy. 14, 95, 96, 97, 140, 142, 143, 144, 147, 148

SKA Square-Kilometer Array. 11, 44, 60, 62, 96, 144

SMBH Supermassive Black Hole. 84, 85, 86, 88

SWIRE Spitzer Wide-area Infrared Extragalactic Survey. 125, 132

TNR True Negative Rate. 75, 76, 78

TPR True Positive Rate. 75, 76, 78

VLA Very Large Array. 10, 11, 15, 18, 38, 61, 94, 96, 115, 124, 125, 132, 133, 134, 135, 142, 143, 147

VLBI Very-Long-Baseline Interferometry. 18, 57, 62, 66, 67, 89, 90, 95, 139, 140, 142

WSRT Westerbork Synthesis Radio Telescope. 41

Bibliography

- Abramowicz, M. A. & Fragile, P. C., 2013. Foundations of Black Hole Accretion Disk Theory, *Living Reviews in Relativity*, **16**(1), 1.
- Agarwal, D., Aggarwal, K., Burke-Spolaor, S., Lorimer, D. R., & Garver-Daniels, N., 2020. FETCH: A deep-learning based classifier for fast transient classification, *MNRAS*, **497**(2), 1661–1674.
- Akeret, J., Seehars, S., Chang, C., Monstein, C., Amara, A., & Refregier, A., 2017. HIDE & SEEK: End-to-end packages to simulate and process radio survey data, *Astronomy and Computing*, **18**, 8–17.
- Albaret, F. D., Allende Prieto, C., Almeida, A., Anders, F., Anderson, S., Andrews, B. H., Aragón-Salamanca, A., Argudo-Fernández, M., Armengaud, E., Aubourg, E., Avila-Reese, V., Badenes, C., Bailey, S., Barbuy, B., Barger, K., Barrera-Ballesteros, J., Bartosz, C., Basu, S., Bates, D., Battaglia, G., Baumgarten, F., Baur, J., Bautista, J., Beers, T. C., Belfiore, F., Bershad, M., Bertran de Lis, S., Bird, J. C., Bizyaev, D., Blanc, G. A., Blanton, M., Blomqvist, M., Bolton, A. S., Borissova, J., Bovy, J., Brandt, W. N., Brinkmann, J., Brownstein, J. R., Bundy, K., Burtin, E., Busca, N. G., Camacho Chavez, H. O., Cano Díaz, M., Cappellari, M., Carrera, R., Chen, Y., Cherinka, B., Cheung, E., Chiappini, C., Chojnowski, D., Chuang, C.-H., Chung, H., Cirolini, R. F., Clerc, N., Cohen, R. E., Comerford, J. M., Comparat, J., Correa do Nascimento, J., Cousinou, M.-C., Covey, K., Crane, J. D., Croft, R., Cunha, K., Darling, J., Davidson, James W., J., Dawson, K., Da Costa, L., Da Silva Ilha, G., Deconto Machado, A., Delubac, T., De Lee, N., De la Macorra, A., De la Torre, S., Diamond-Stanic, A. M., Donor, J., Downes, J. J., Drory, N., Du, C., Du Mas des Bourboux, H., Dwelly, T., Ebelke, G., Eigenbrot, A., Eisenstein, D. J., Elsworth, Y. P., Emsellem, E., Eracleous, M., Escoffier, S., Evans, M. L., Falcón-Barroso, J., Fan, X., Favole, G., Fernandez-Alvar, E., Fernandez-Trincado, J. G., Feuillet, D., Fleming, S. W., Font-Ribera, A., Freischlad, G., Frinchaboy, P., Fu, H., Gao, Y., Garcia, R. A., Garcia-Dias, R., Garcia-Hernández, D. A., Garcia Pérez, A. E., Gaulme, P., Ge, J., Geisler, D., Gillespie, B., Gil Marin, H., Girardi, L., Goddard, D., Gomez Maqueo Chew, Y., Gonzalez-Perez, V., Grabowski, K., Green, P., Grier, C. J., Grier, T., Guo, H., Guy, J., Hagen, A., Hall, M., Harding, P., Harley, R. E., Hasselquist, S., Hawley, S., Hayes, C. R., Hearty, F., Hekker, S., Hernandez Toledo, H., Ho, S., Hogg, D. W., Holley-Bockelmann, K., Holtzman, J. A., Holzer, P. H., Hu, J., Huber, D., Hutchinson, T. A., Hwang, H. S., Ibarra-Medel, H. J., Ivans, I. I., Ivory, K., Jaehnig, K., Jensen, T. W., Johnson, J. A., Jones, A., Jullo, E., Kallinger, T., Kinemuchi, K., Kirkby, D., Klaene, M., Kneib, J.-P., Kollmeier, J. A., Lacerna, I., Lane, R. R., Lang, D., Laurent, P., Law, D. R., Leauthaud, A., Le Goff, J.-M., Li, C., Li, C., Li, N., Li, R., Liang, F.-H., Liang, Y., Lima, M., Lin, L., Lin, L., Lin, Y.-T., Liu, C., Long, D., Lucatello, S., MacDonald, N., MacLeod, C. L., Mackereth, J. T., Mahadevan, S., Maia, M. A. G., Maiolino, R., Majewski, S. R., Malanushenko, O., Malanushenko, V., Mallmann, N. D., Manchado, A., Maraston,

C., Marques-Chaves, R., Martinez Valpuesta, I., Masters, K. L., Mathur, S., McGreer, I. D., Merloni, A., Merrifield, M. R., Mészáros, S., Meza, A., Miglio, A., Minchev, I., Molaverdikhani, K., Montero-Dorta, A. D., Mosser, B., Muna, D., Myers, A., Nair, P., Nandra, K., Ness, M., Newman, J. A., Nichol, R. C., Nidever, D. L., Nitschelm, C., O’Connell, J., Oravetz, A., Oravetz, D. J., Pace, Z., Padilla, N., Palanque-Delabrouille, N., Pan, K., Parejko, J., Paris, I., Park, C., Peacock, J. A., Peirani, S., Pellejero-Ibanez, M., Penny, S., Percival, W. J., Percival, J. W., Perez-Fournon, I., Petitjean, P., Pieri, M., Pinsonneault, M. H., Pisani, A., Prada, F., Prakash, A., Price-Jones, N., Raddick, M. J., Rahman, M., Raichoor, A., Barboza Rembold, S., Reyna, A. M., Rich, J., Richstein, H., Ridl, J., Riffel, R. A., Riffel, R., Rix, H.-W., Robin, A. C., Rockosi, C. M., Rodríguez-Torres, S., Rodrigues, T. S., Roe, N., Roman Lopes, A., Román-Zúñiga, C., Ross, A. J., Rossi, G., Ruan, J., Ruggeri, R., Runnoe, J. C., Salazar-Albornoz, S., Salvato, M., Sanchez, S. F., Sanchez, A. G., Sanchez-Gallego, J. R., Santiago, B. X., Schiavon, R., Schimoia, J. S., Schlafly, E., Schlegel, D. J., Schneider, D. P., Schönrich, R., Schultheis, M., Schwobe, A., Seo, H.-J., Serenelli, A., Sesar, B., Shao, Z., Shetrone, M., Shull, M., Silva Aguirre, V., Skrutskie, M. F., Slosar, A., Smith, M., Smith, V. V., Sobek, J., Somers, G., Souto, D., Stark, D. V., Stassun, K. G., Steinmetz, M., Stello, D., Storch Bergmann, T., Strauss, M. A., Streblyanska, A., Stringfellow, G. S., Suarez, G., Sun, J., Taghizadeh-Popp, M., Tang, B., Tao, C., Tayar, J., Tembe, M., Thomas, D., Tinker, J., Tojeiro, R., Tremonti, C., Troup, N., Trump, J. R., Unda-Sanzana, E., Valenzuela, O., Van den Bosch, R., Vargas-Magaña, M., Vazquez, J. A., Villanova, S., Vivek, M., Vogt, N., Wake, D., Walterbos, R., Wang, Y., Wang, E., Weaver, B. A., Weijmans, A.-M., Weinberg, D. H., Westfall, K. B., Whelan, D. G., Wilcots, E., Wild, V., Williams, R. A., Wilson, J., Wood-Vasey, W. M., Wylezalek, D., Xiao, T., Yan, R., Yang, M., Ybarra, J. E., Yeche, C., Yuan, F.-T., Zakamska, N., Zamora, O., Zasowski, G., Zhang, K., Zhao, C., Zhao, G.-B., Zheng, Z., Zheng, Z., Zhou, Z.-M., Zhu, G., Zinn, J. C., & Zou, H., 2017. The 13th Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDSS-IV Survey Mapping Nearby Galaxies at Apache Point Observatory, *Astrophysical Journal Supplement Series*, **233**(2), 25.

An, T. & Baan, W. A., 2012. The Dynamic Evolution of Young Extragalactic Radio Sources, *ApJ*, **760**(1), 77.

Aniyan, A. K. & Thorat, K., 2017. Classifying Radio Galaxies with the Convolutional Neural Network, *ApJ Supplement*, **230**(2), 20.

Argo, M., 2012. Sir Bernard Lovell (1913 - 2012), *WGN, Journal of the International Meteor Organization*, **40**(4), 114–116.

Baan, W. A., 2019. Implementing RFI Mitigation in Radio Science, *Journal of Astronomical Instrumentation*, **8**(1), 1940010.

Baars, J. W. M. & Hooghoudt, B. G., 1974. The Synthesis Radio Telescope at Westerbork. General Lay-out and Mechanical Aspects, *A&A*, **31**, 323.

Becker, R. H., White, R. L., & Helfand, D. J., 1995. The FIRST Survey: Faint Images of the Radio Sky at Twenty Centimeters, *ApJ*, **450**, 559.

Begelman, M. C., Blandford, R. D., & Rees, M. J., 1980. Massive black hole binaries in active galactic nuclei, *Nature*, **287**(5780), 307–309.

Bethapudi, S. & Desai, S., 2018. Separation of pulsar signals from noise using supervised machine learning algorithms, *Astronomy and Computing*, **23**, 15.

- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M., 2008. Correcting direction-dependent gains in the deconvolution of radio interferometric images, *A&A*, **487**(1), 419–429.
- Blandford, R., Meier, D., & Readhead, A., 2019. Relativistic Jets from Active Galactic Nuclei, *ARA&A*, **57**, 467–509.
- Bonaldi, A., Bonato, M., Galluzzi, V., Harrison, I., Massardi, M., Kay, S., De Zotti, G., & Brown, M. L., 2019. The Tiered Radio Extragalactic Continuum Simulation (T-RECS), *MNRAS*, **482**(1), 2–19.
- Bonzini, M., Padovani, P., Mainieri, V., Kellermann, K. I., Miller, N., Rosati, P., Tozzi, P., & Vattakunnel, S., 2013. The sub-mJy radio sky in the Extended Chandra Deep Field-South: source population, *MNRAS*, **436**(4), 3759–3771.
- Bonzini, M., Mainieri, V., Padovani, P., Andreani, P., Berta, S., Bethermin, M., Lutz, D., Rodighiero, G., Rosario, D., Tozzi, P., & Vattakunnel, S., 2015. Star formation properties of sub-mJy radio sources, *MNRAS*, **453**(1), 1079–1094.
- Boonstra, A. J. & van der Tol, S., 2005. Spatial filtering of interfering signals at the initial Low Frequency Array (LOFAR) phased array test station, *Radio Science*, **40**(5), RS5S09.
- Bracewell, R. N., 1956. Strip Integration in Radio Astronomy, *Australian Journal of Physics*, **9**, 198.
- Braun, R., Bonaldi, A., Bourke, T., Keane, E., & Wagg, J., 2019. Anticipated Performance of the Square Kilometre Array – Phase 1 (SKA1), *arXiv e-prints*, p. arXiv:1912.12699.
- Burke, B. F., Graham-Smith, F., & Wilkinson, P. N., 2019. *An Introduction to Radio Astronomy*, Cambridge University Press, 4th edn.
- Carozzi, T. D. & Woan, G., 2009. A generalized measurement equation and van Cittert-Zernike theorem for wide-field radio astronomical interferometry, *MNRAS*, **395**(3), 1558–1568.
- CASA Team, Bean, B., Bhatnagar, S., Castro, S., Donovan Meyer, J., Emonts, B., Garcia, E., Garwood, R., Golap, K., Gonzalez Villalba, J., Harris, P., Hayashi, Y., Hoskins, J., Hsieh, M., Jagannathan, P., Kawasaki, W., Keimpema, A., Kettenis, M., Lopez, J., Marvil, J., Masters, J., McNichols, A., Mehringer, D., Miel, R., Moellenbrock, G., Montesino, F., Nakazato, T., Ott, J., Petry, D., Pokorny, M., Raba, R., Rau, U., Schiebel, D., Schweighart, N., Sekhar, S., Shimada, K., Small, D., Steeb, J.-W., Sugimoto, K., Suoranta, V., Tsutsumi, T., van Bemmelen, I. M., Verkouter, M., Wells, A., Xiong, W., Szomoru, A., Griffith, M., Glendenning, B., & Kern, J., 2022. CASA, the Common Astronomy Software Applications for Radio Astronomy, *Publications of the Astronomical Society of the Pacific*, **134**(1041), 114501.
- Colpi, M., 2014. Massive binary black holes in galactic nuclei and their path to coalescence, *Space Science Reviews*, **183**(1-4), 189–221.
- Condon, J. J., 1992. Radio emission from normal galaxies., *ARA&A*, **30**, 575–611.
- Conselice, C. J., 2014. The evolution of galaxy structure over cosmic time, *Annu. Rev. Astron. Astrophys.*, **52**, 291–337.

- Conway, J. E., 2002. Compact symmetric objects—newborn radio galaxies?, *New Astronomy Reviews*, **46**(2-7), 263–271.
- Cooley, J. W. & Tukey, J. W., 1965. An algorithm for the machine calculation of complex Fourier series, *Mathematics of Computation*, **19**, 297–301.
- Cornwell, T. J., 2008. Multiscale CLEAN Deconvolution of Radio Synthesis Images, *IEEE Journal of Selected Topics in Signal Processing*, **2**(5), 793–801.
- Czech, D., Mishra, A., & Inggis, M., 2018. A CNN and LSTM-based approach to classifying transient radio frequency interference, *Astronomy and Computing*, **25**, 52–57.
- Danzmann, K., Amaro-Seoane, P., Audley, H., Babak, S., Baker, J., Barausse, E., Bender, P., Berti, E., Binetruy, P., Born, M., Bortoluzzi, D., Camp, J., Caprini, C., Cardoso, V., Colpi, M., Conklin, J., Cornish, N., Cutler, C., Dolesi, R., Ferraioli, L., Ferroni, V., Fitzsimons, E., Gair, J., Bote, L. G., Giardini, D., Gibert, F., Grmani, C., Halloin, H., Heinzel, G., Hertog, T., Hewitson, M., Holley-Bockelmann, K., Hollington, D., Hueller, M., Inchauspe, H., Jetzer, P., Karnesis, N., Killow, C., Klein, A., Klipstein, B., Korsakova, N., Larson, S. L., Livas, J., Lloro, I., Man, N., Mance, D., Martino, J., Mateos, I., McKenzie, K., McWilliams, S. T., Miller, C., Mueller, G., Nardini, G., Nelemans, G., Nofrarias, M., Petiteau, A., Pivato, P., Plagnol, E., Porter, E., Reiche, J., Robertson, D., Robertson, N., Rossi, E., Russano, G., Schutz, B., Sesana, A., Shoemaker, D., Slutsky, J., Sopuerta, C. F., Sumner, T., Tamanini, N., Thorpe, I., Troebis, M., Vallisneri, M., Vecchio, A., Vetrugno, D., Vitale, S., Volonteri, M., Wanner, G., Ward, H., Wass, P., Weber, W., Ziemer, J., & Zweifel, P., 2017. LISA laser interferometer space antenna.
- De Rosa, A., Vignali, C., Bogdanović, T., Capelo, P. R., Charisi, M., Dotti, M., Husemann, B., Lusso, E., Mayer, L., Paragi, Z., Runnoe, J., Sesana, A., Steinborn, L., Bianchi, S., Colpi, M., del Valle, L., Frey, S., Gabányi, K. É., Giustini, M., Guainazzi, M., Haiman, Z., Herrera Ruiz, N., Herrero-Illana, R., Iwasawa, K., Komossa, S., Lena, D., Loiseau, N., Perez-Torres, M., Piconcelli, E., & Volonteri, M., 2019. The quest for dual and binary supermassive black holes: A multi-messenger view, *New Astronomy Reviews*, **86**, 101525.
- DeBoer, D. R., Gough, R. G., Bunton, J. D., Cornwell, T. J., Beresford, R. J., Johnston, S., Feain, I. J., Schinckel, A. E., Jackson, C. A., Kesteven, M. J., Chippendale, A., Hampson, G. A., O’Sullivan, J. D., Hay, S. G., Jacka, C. E., Sweetnam, T. W., Storey, M. C., Ball, L., & Boyle, B. J., 2009. Australian SKA Pathfinder: A High-Dynamic Range Wide-Field of View Survey Telescope, *IEEE Proceedings*, **97**(8), 1507–1521.
- Di Vruno, F., Winkel, B., Bassa, C. G., Józsa, G. I. G., Brentjens, M. A., Jessner, A., & Garrington, S., 2023. Unintended electromagnetic radiation from Starlink satellites detected with LOFAR between 110 and 188 MHz, *A&A*, **676**, A75.
- Do, T., Tran, H., Tjiputra, E., Tran, Q. D., & Nguyen, A., 2022. Fine-Grained Visual Classification using Self Assessment Classifier, *arXiv e-prints*, p. arXiv:2205.10529.
- Doeleman, S., Agol, E., Backer, D., Baganoff, F., Bower, G. C., Broderick, A., Fabian, A., Fish, V., Gammie, C., Ho, P., Honman, M., Krichbaum, T., Loeb, A., Marrone, D., Reid, M., Rogers, A., Shapiro, I., Strittmatter, P., Tilanus, R., Weintroub, J., Whitney, A., Wright, M., & Ziurys, L., 2009. Imaging an Event Horizon: submm-VLBI of a Super Massive Black Hole, in *astro2010: The Astronomy and Astrophysics Decadal Survey*, vol. 2010, p. 68.

- Dunlop, J. S. & Peacock, J. A., 1990. The redshift cut-off in the luminosity function of radio galaxies and quasars., *MNRAS*, **247**, 19.
- Ekers, R., 2012. The History of the Square Kilometre Array (SKA) - Born Global, *arXiv e-prints*, p. arXiv:1212.3497.
- Enriquez, J. E., Siemion, A., Foster, G., Gajjar, V., Hellbourg, G., Hickish, J., Isaacson, H., Price, D. C., Croft, S., DeBoer, D., Lebofsky, M., MacMahon, D. H. E., & Werthimer, D., 2017. The Breakthrough Listen Search for Intelligent Life: 1.1-1.9 GHz Observations of 692 Nearby Stars, *ApJ*, **849**(2), 104.
- Evans, I. N., Primini, F. A., Miller, J. B., Evans, J. D., Allen, C. E., Anderson, C. S., Becker, G., Budynkiewicz, J. A., Burke, D., Chen, J. C., Civano, F., D'Abrusco, R., Doe, S. M., Fabbiano, G., Martinez Galarza, J., Gibbs, D. G., I., Glotfelty, K. J., Graessle, D. E., Grier, J. D., J., Hain, R. M., Hall, D. M., Harbo, P. N., Houck, J. C., Lauer, J. L., Laurino, O., Lee, N. P., McCollough, M. L., McDowell, J. C., McLaughlin, W., Morgan, D. L., Mossman, A. E., Nguyen, D. T., Nichols, J. S., Nowak, M. A., Paxson, C., Perdikeas, M., Plummer, D. A., Rots, A. H., Siemiginowska, A. L., Sundheim, B. A., Thong, S., Tibbetts, M. S., Van Stone, D. W., Winkelman, S. L., & Zografou, P., 2020. The Chandra Source Catalog — A Billion X-ray Photons, in *American Astronomical Society Meeting Abstracts #235*, vol. 235 of **American Astronomical Society Meeting Abstracts**, p. 154.05.
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., Alef, W., Asada, K., Azulay, R., Baczko, A.-K., Ball, D., Baloković, M., Barrett, J., Bintley, D., Blackburn, L., Boland, W., Bouman, K. L., Bower, G. C., Bremer, M., Brinkerink, C. D., Brisenenden, R., Britzen, S., Broderick, A. E., Brogiere, D., Bronzwaer, T., Byun, D.-Y., Carlstrom, J. E., Chael, A., Chan, C.-k., Chatterjee, S., Chatterjee, K., Chen, M.-T., Chen, Y., Cho, I., Christian, P., Conway, J. E., Cordes, J. M., Crew, G. B., Cui, Y., Davelaar, J., De Laurentis, M., Deane, R., Dempsey, J., Desvignes, G., Dexter, J., Doeleman, S. S., Eatough, R. P., Falcke, H., Fish, V. L., Fomalont, E., Fraga-Encinas, R., Freeman, W. T., Friberg, P., Fromm, C. M., Gómez, J. L., Galison, P., Gammie, C. F., García, R., Gentaz, O., Georgiev, B., Goddi, C., Gold, R., Gu, M., Gurwell, M., Hada, K., Hecht, M. H., Hesper, R., Ho, L. C., Ho, P., Honma, M., Huang, C.-W. L., Huang, L., Hughes, D. H., Ikeda, S., Inoue, M., Issaoun, S., James, D. J., Jannuzi, B. T., Janssen, M., Jeter, B., Jiang, W., Johnson, M. D., Jorstad, S., Jung, T., Karami, M., Karuppusamy, R., Kawashima, T., Keating, G. K., Kettenis, M., Kim, J.-Y., Kim, J., Kim, J., Kino, M., Koay, J. Y., Koch, P. M., Koyama, S., Kramer, M., Kramer, C., Krichbaum, T. P., Kuo, C.-Y., Lauer, T. R., Lee, S.-S., Li, Y.-R., Li, Z., Lindqvist, M., Liu, K., Liuzzo, E., Lo, W.-P., Lobanov, A. P., Loinard, L., Lonsdale, C., Lu, R.-S., MacDonald, N. R., Mao, J., Markoff, S., Marrone, D. P., Marscher, A. P., Martí-Vidal, I., Matsushita, S., Matthews, L. D., Medeiros, L., Menten, K. M., Mizuno, Y., Mizuno, I., Moran, J. M., Moriyama, K., Moscibrodzka, M., Müller, C., Nagai, H., Nagar, N. M., Nakamura, M., Narayan, R., Narayanan, G., Natarajan, I., Neri, R., Ni, C., Noutsos, A., Okino, H., Olivares, H., Ortiz-León, G. N., Oyama, T., Özel, F., Palumbo, D. C. M., Patel, N., Pen, U.-L., Pesce, D. W., Piétu, V., Plambeck, R., PopStefanija, A., Porth, O., Prather, B., Preciado-López, J. A., Psaltis, D., Pu, H.-Y., Ramakrishnan, V., Rao, R., Rawlings, M. G., Raymond, A. W., Rezzolla, L., Ripperda, B., Roelofs, F., Rogers, A., Ros, E., Rose, M., Roshanineshat, A., Rottmann, H., Roy, A. L., Ruszczyk, C., Ryan, B. R., Rygl, K. L. J., Sánchez, S., Sánchez-Arguelles, D., Sasada, M., Savolainen, T., Schloerb, F. P., Schuster, K.-F., Shao, L., Shen, Z., Small, D., Sohn, B. W., SooHoo, J., Tazaki, F., Tiede, P., Tilanus, R. P. J., Titus, M., Toma, K., Torne, P., Trent, T., Trippe,

S., Tsuda, S., van Bemmelen, I., van Langevelde, H. J., van Rossum, D. R., Wagner, J., Wardle, J., Weintraub, J., Wex, N., Wharton, R., Wielgus, M., Wong, G. N., Wu, Q., Young, K., Young, A., Younsi, Z., Yuan, F., Yuan, Y.-F., Zensus, J. A., Zhao, G., Zhao, S.-S., Zhu, Z., Algaba, J.-C., Allardi, A., Amestica, R., Anczarski, J., Bach, U., Baganoff, F. K., Beaudoin, C., Benson, B. A., Berthold, R., Blanchard, J. M., Blundell, R., Bustamente, S., Cappallo, R., Castillo-Domínguez, E., Chang, C.-C., Chang, S.-H., Chang, S.-C., Chen, C.-C., Chilson, R., Chuter, T. C., Córdova Rosado, R., Coulson, I. M., Crawford, T. M., Crowley, J., David, J., Derome, M., Dexter, M., Dornbusch, S., Dudevoir, K. A., Dzib, S. A., Eckart, A., Eckert, C., Erickson, N. R., Everett, W. B., Faber, A., Farah, J. R., Fath, V., Folkers, T. W., Forbes, D. C., Freund, R., Gómez-Ruiz, A. I., Gale, D. M., Gao, F., Geertsema, G., Graham, D. A., Greer, C. H., Grosslein, R., Gueth, F., Haggard, D., Halverson, N. W., Han, C.-C., Han, K.-C., Hao, J., Hasegawa, Y., Henning, J. W., Hernández-Gómez, A., Herrero-Illana, R., Heyminck, S., Hirota, A., Hoge, J., Huang, Y.-D., Impellizzeri, C. M. V., Jiang, H., Kamble, A., Keisler, R., Kimura, K., Kono, Y., Kubo, D., Kuroda, J., Lacasse, R., Laing, R. A., Leitch, E. M., Li, C.-T., Lin, L. C. C., Liu, C.-T., Liu, K.-Y., Lu, L.-M., Marson, R. G., Martin-Cocher, P. L., Massingill, K. D., Matulonis, C., McColl, M. P., McWhirter, S. R., Messias, H., Meyer-Zhao, Z., Michalik, D., Montaña, A., Montgomerie, W., Mora-Klein, M., Muders, D., Nadolski, A., Navarro, S., Neilsen, J., Nguyen, C. H., Nishioka, H., Norton, T., Nowak, M. A., Nystrom, G., Ogawa, H., Oshiro, P., Oyama, T., Parsons, H., Paine, S. N., Peñalver, J., Phillips, N. M., Poirier, M., Pradel, N., Primiani, R. A., Raffin, P. A., Rahlin, A. S., Reiland, G., Risacher, C., Ruiz, I., Sáez-Madaín, A. F., Sassella, R., Schellart, P., Shaw, P., Silva, K. M., Shiokawa, H., Smith, D. R., Snow, W., Souccar, K., Sousa, D., Sridharan, T. K., Srinivasan, R., Stahm, W., Stark, A. A., Story, K., Timmer, S. T., Vertatschitsch, L., Walther, C., Wei, T.-S., Whitehorn, N., Whitney, A. R., Woody, D. P., Wouterloot, J. G. A., Wright, M., Yamaguchi, P., Yu, C.-Y., Zeballos, M., Zhang, S., & Ziurys, L., 2019. First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole, *ApJ Letters*, **875**(1), L1.

Fabian, A. C., 2012. Observational Evidence of Active Galactic Nuclei Feedback, *ARA&A*, **50**, 455–489.

Findlay, J. W., Hvatum, H., & Waltman, W. B., 1965. An Absolute Flux-Density Measurement of Cassiopeia a at 1440 MHz., *ApJ*, **141**, 873.

Fridman, P. A. & Baan, W. A., 2001. RFI mitigation methods in radio astronomy, *A&A*, **378**, 327–344.

Fu, H., Myers, A. D., Djorgovski, S. G., Yan, L., Wrobel, J. M., & Stockton, A., 2015. Radio-selected Binary Active Galactic Nuclei from the Very Large Array Stripe 82 Survey, *ApJ*, **799**(1), 72.

Fukushima, K., 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position, *Biological Cybernetics*, **36**(4), 193–202.

Gaia Collaboration, Brown, A. G. A., Vallenari, A., Prusti, T., de Bruijne, J. H. J., Babusaix, C., Biermann, M., Creevey, O. L., Evans, D. W., Eyer, L., Hutton, A., Jansen, F., Jordi, C., Klioner, S. A., Lammers, U., Lindegren, L., Luri, X., Mignard, F., Panem, C., Pourbaix, D., Randich, S., Sartoretti, P., Soubiran, C., Walton, N. A., Arenou, F., Bailer-Jones, C. A. L., Bastian, U., Cropper, M., Drimmel, R., Katz, D., Lattanzi, M. G., van Leeuwen, F., Bakker, J., Cacciari, C., Castañeda, J., De Angeli, F., Ducourant, C.,

Fabricius, C., Fouesneau, M., Frémat, Y., Guerra, R., Guerrier, A., Guiraud, J., Jean-Antoine Piccolo, A., Masana, E., Messineo, R., Mowlavi, N., Nicolas, C., Nienartowicz, K., Pailler, F., Panuzzo, P., Riclet, F., Roux, W., Seabroke, G. M., Sordo, R., Tanga, P., Thévenin, F., Gracia-Abril, G., Portell, J., Teyssier, D., Altmann, M., Andrae, R., Bellas-Velidis, I., Benson, K., Berthier, J., Blomme, R., Brugaletta, E., Burgess, P. W., Busso, G., Carry, B., Cellino, A., Cheek, N., Clementini, G., Damerджи, Y., Davidson, M., Delchambre, L., Dell’Oro, A., Fernández-Hernández, J., Galluccio, L., García-Lario, P., Garcia-Reinaldos, M., González-Núñez, J., Gosset, E., Haigron, R., Halbwachs, J. L., Hambly, N. C., Harrison, D. L., Hatzidimitriou, D., Heiter, U., Hernández, J., Hestroffer, D., Hodgkin, S. T., Holl, B., Janßen, K., Jevardat de Fombelle, G., Jordan, S., Krone-Martins, A., Lanzafame, A. C., Löffler, W., Lorca, A., Manteiga, M., Marchal, O., Marrese, P. M., Moitinho, A., Mora, A., Muinonen, K., Osborne, P., Pancino, E., Pauwels, T., Petit, J. M., Recio-Blanco, A., Richards, P. J., Riello, M., Rimoldini, L., Robin, A. C., Roegiers, T., Rybizki, J., Sarro, L. M., Siopis, C., Smith, M., Sozzetti, A., Ulla, A., Utrilla, E., van Leeuwen, M., van Reeve, W., Abbas, U., Abreu Aramburu, A., Accart, S., Aerts, C., Aguado, J. J., Ajaj, M., Altavilla, G., Álvarez, M. A., Álvarez Cid-Fuentes, J., Alves, J., Anderson, R. I., Anglada Varela, E., Antoja, T., Audard, M., Baines, D., Baker, S. G., Balaguer-Núñez, L., Balbinot, E., Balog, Z., Barache, C., Barbato, D., Barros, M., Barstow, M. A., Bartolomé, S., Bassilana, J. L., Bauchet, N., Baudesson-Stella, A., Becciani, U., Bellazzini, M., Bernet, M., Bertone, S., Bianchi, L., Blanco-Cuaresma, S., Boch, T., Bombrun, A., Bossini, D., Bouquillon, S., Bragaglia, A., Bramante, L., Breedt, E., Bressan, A., Brouillet, N., Bucciarelli, B., Burlacu, A., Busonero, D., Butkevich, A. G., Buzzi, R., Caffau, E., Cancelliere, R., Cánovas, H., Cantat-Gaudin, T., Carballo, R., Carlucci, T., Carnerero, M. I., Carrasco, J. M., Casamiquela, L., Castellani, M., Castro-Ginard, A., Castro Sampol, P., Chaoul, L., Charlot, P., Chemin, L., Chiavassa, A., Cioni, M. R. L., Comoretto, G., Cooper, W. J., Cornez, T., Cowell, S., Crifo, F., Crosta, M., Crowley, C., Dafonte, C., Dapergolas, A., David, M., David, P., de Laverny, P., De Luise, F., De March, R., De Ridder, J., de Souza, R., de Teodoro, P., de Torres, A., del Peloso, E. F., del Pozo, E., Delbo, M., Delgado, A., Delgado, H. E., Delisle, J. B., Di Matteo, P., Diakite, S., Diener, C., Distefano, E., Dolding, C., Eappachen, D., Edvardsson, B., Enke, H., Esquej, P., Fabre, C., Fabrizio, M., Faigler, S., Fedorets, G., Fernique, P., Fienga, A., Figueras, F., Fouron, C., Fragkoudi, F., Fraile, E., Franke, F., Gai, M., Garabato, D., Garcia-Gutierrez, A., García-Torres, M., Garofalo, A., Gavras, P., Gerlach, E., Geyer, R., Giacobbe, P., Gilmore, G., Girona, S., Giuffrida, G., Gomel, R., Gomez, A., Gonzalez-Santamaria, I., González-Vidal, J. J., Granvik, M., Gutiérrez-Sánchez, R., Guy, L. P., Hauser, M., Haywood, M., Helmi, A., Hidalgo, S. L., Hilger, T., Hładczuk, N., Hobbs, D., Holland, G., Huckle, H. E., Jasniewicz, G., Jonker, P. G., Juaristi Campillo, J., Julbe, F., Karbevská, L., Kervella, P., Khanna, S., Kochoska, A., Kontizas, M., Kordopatis, G., Korn, A. J., Kostrzewa-Rutkowska, Z., Kruszyńska, K., Lambert, S., Lanza, A. F., Lasne, Y., Le Campion, J. F., Le Fustec, Y., Lebreton, Y., Lebzelter, T., Leccia, S., Leclerc, N., Lecoœur-Taïbi, I., Liao, S., Licata, E., Lindstrøm, E. P., Lister, T. A., Livanou, E., Lobel, A., Madrero Pardo, P., Managau, S., Mann, R. G., Marchant, J. M., Marconi, M., Marcos Santos, M. M. S., Marinoni, S., Marocco, F., Marshall, D. J., Martin Polo, L., Martín-Fleitas, J. M., Masip, A., Massari, D., Mastrobuono-Battisti, A., Mazeh, T., McMillan, P. J., Messina, S., Michalik, D., Millar, N. R., Mints, A., Molina, D., Molinaro, R., Molnár, L., Montegriffo, P., Mor, R., Morbidelli, R., Morel, T., Morris, D., Mulone, A. F., Muñoz, D., Muraveva, T., Murphy, C. P., Musella, I., Noval, L., Ordénovic, C., Orrù, G., Osinde, J., Pagani, C., Pagano, I., Palaversa, L., Palicio, P. A., Panahi, A., Pawlak, M., Peñalosa Esteller, X., Penttilä, A., Piersimoni, A. M., Pineau,

- F. X., Plachy, E., Plum, G., Poggio, E., Poretti, E., Poujoulet, E., Prša, A., Pulone, L., Racero, E., Ragaini, S., Rainer, M., Raiteri, C. M., Rambaux, N., Ramos, P., Ramos-Lerate, M., Re Fiorentin, P., Regibo, S., Reylé, C., Ripepi, V., Riva, A., Rixon, G., Robichon, N., Robin, C., Roelens, M., Rohrbasser, L., Romero-Gómez, M., Rowell, N., Royer, F., Rybicki, K. A., Sadowski, G., Sagristà Sellés, A., Sahlmann, J., Salgado, J., Salguero, E., Samaras, N., Sanchez Gimenez, V., Sanna, N., Santoveña, R., Sarasso, M., Schultheis, M., Sciacca, E., Segol, M., Segovia, J. C., Ségransan, D., Semeux, D., Shahaf, S., Siddiqui, H. I., Siebert, A., Siltala, L., Slezak, E., Smart, R. L., Solano, E., Solitro, F., Souami, D., Souchay, J., Spagna, A., Spoto, F., Steele, I. A., Steidelmüller, H., Stephenson, C. A., Süveges, M., Szabados, L., Szegedi-Elek, E., Taris, F., Tauran, G., Taylor, M. B., Teixeira, R., Thuillot, W., Tonello, N., Torra, F., Torra, J., Turon, C., Unger, N., Vaillant, M., van Dillen, E., Vanel, O., Vecchiato, A., Viala, Y., Vicente, D., Voutsinas, S., Weiler, M., Wevers, T., Wyrzykowski, Ł., Yoldas, A., Yvard, P., Zhao, H., Zorec, J., Zucker, S., Zurbach, C., & Zwitter, T., 2021. Gaia Early Data Release 3. Summary of the contents and survey properties, *A&A*, **649**, A1.
- Gardner, J. P., Mather, J. C., Clampin, M., Doyon, R., Greenhouse, M. A., Hammel, H. B., Hutchings, J. B., Jakobsen, P., Lilly, S. J., Long, K. S., Lunine, J. I., McCaughrean, M. J., Mountain, M., Nella, J., Rieke, G. H., Rieke, M. J., Rix, H.-W., Smith, E. P., Sonneborn, G., Stiavelli, M., Stockman, H. S., Windhorst, R. A., & Wright, G. S., 2006. The James Webb Space Telescope, *Space Science Reviews*, **123**(4), 485–606.
- Garrington, S. T., Anderson, B., Baines, C., Battilana, J. A., Bentley, M. N., Brown, D., Burgess, P., Diamond, P. J., Kitching, G. J., McCool, R., Muxlow, T. W., Noble, R. G., Roddis, N., Spencer, R. E., & Thomasson, P., 2004. e-MERLIN, in *Ground-based Telescopes*, vol. 5489 of **Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series**, pp. 332–343.
- Globus, N. & Levinson, A., 2016. The collimation of magnetic jets by disc winds, *MNRAS*, **461**(3), 2605–2615.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y., 2014. Generative adversarial nets, in *Advances in Neural Information Processing Systems*, vol. 27, Curran Associates, Inc.
- Goulding, A. D., Pardo, K., Greene, J. E., Mingarelli, C. M., Nyland, K., & Strauss, M. A., 2019. Discovery of a close-separation binary quasar at the heart of a $z \approx 0.2$ merging galaxy and its implications for low-frequency gravitational waves, *ApJL*, **879**(2), L21.
- Greisen, E. W., 2003. AIPS, the VLA, and the VLBA, in *Information Handling in Astronomy - Historical Vistas*, vol. 285 of **Astrophysics and Space Science Library**, p. 109.
- Grigg, D., Tingay, S. J., Sokolowski, M., Wayth, R. B., Indermuehle, B., & Prabu, S., 2023. Detection of intended and unintended emissions from Starlink satellites in the SKA-Low frequency range, at the SKA-Low site, with an SKA-Low station analogue, *A&A*, **678**, L6.
- Hamaker, J. P., Bregman, J. D., & Sault, R. J., 1996. Understanding radio polarimetry. I. Mathematical foundations., *A&AS*, **117**, 137–147.
- Hannun, A., Case, C., Casper, J., Catanzaro, B., Diamos, G., Elsen, E., Prenger, R., Satheesh, S., Sengupta, S., Coates, A., & Ng, A. Y., 2014. Deep Speech: Scaling up end-to-end speech recognition, *arXiv e-prints*, p. arXiv:1412.5567.

- Hartley, P., Jackson, N., Sluse, D., Stacey, H. R., & Vives-Arias, H., 2019. Strong lensing reveals jets in a sub-microJy radio-quiet quasar, *MNRAS*, **485**(3), 3009–3023.
- Heckman, T. M. & Best, P. N., 2014. The Coevolution of Galaxies and Supermassive Black Holes: Insights from Surveys of the Contemporary Universe, *ARA&A*, **52**, 589–660.
- Heckman, T. M., Armus, L., & Miley, G. K., 1990. On the Nature and Implications of Starburst-driven Galactic Superwinds, *ApJ Supplement*, **74**, 833.
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A., 1968. Observation of a Rapidly Pulsating Radio Source, *Nature*, **217**(5130), 709–713.
- Heyvaerts, J. & Norman, C., 1989. The Collimation of Magnetized Winds, *ApJ*, **347**, 1055.
- Högbom, J. A., 1974. Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines, *A&AS*, **15**, 417.
- Högbom, J. A. & Brouw, W. N., 1974. The Synthesis Radio Telescope at Westerbork. Principles of Operation, Performance and Data Reduction, *A&A*, **33**, 289.
- Holley-Bockelmann, K. & Khan, F. M., 2015. Galaxy Rotation and Rapid Supermassive Black Hole Binary Coalescence, *ApJ*, **810**(2), 139.
- Hopkins, A., Windhorst, R., Cram, L., & Ekers, R., 2000. What Will the Next Generation Radio Telescope Detect at 1.4 GHz?, *Experimental Astronomy*, **10**(4), 419–437.
- Hopkins, A. M. & Beacom, J. F., 2006. On the Normalization of the Cosmic Star Formation History, *ApJ*, **651**(1), 142–154.
- Hunter, T. R., Schwab, F. R., White, S. D., Ford, J. M., Ghigo, F. D., Maddalena, R. J., Mason, B. S., Nelson, J. D., Prestage, R. M., Ray, J., Ries, P., Simon, R., Srikanth, S., & Whiteis, P., 2011. Holographic Measurement and Improvement of the Green Bank Telescope Surface, *Publications of the Astronomical Society of the Pacific*, **123**(907), 1087–1099.
- Inayoshi, K., Visbal, E., & Haiman, Z., 2020. The Assembly of the First Massive Black Holes, *ARA&A*, **58**, 27–97.
- Isola, P., Zhu, J., Zhou, T., & Efros, A. A., 2016. Image-to-image translation with conditional adversarial networks, *CoRR*, **abs/1611.07004**.
- ITU-R, 2019. *Radio Regulations*, Sharm el-Sheik WRC-19.
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., Angel, J. R. P., Angeli, G. Z., Ansari, R., Antilogus, P., Araujo, C., Armstrong, R., Arndt, K. T., Astier, P., Aubourg, É., Auza, N., Axelrod, T. S., Bard, D. J., Barr, J. D., Barrau, A., Bartlett, J. G., Bauer, A. E., Bauman, B. J., Baumont, S., Bechtol, E., Bechtol, K., Becker, A. C., Becla, J., Beldica, C., Bellavia, S., Bianco, F. B., Biswas, R., Blanc, G., Blazek, J., Blandford, R. D., Bloom, J. S., Bogart, J., Bond, T. W., Booth, M. T., Borgland, A. W., Borne, K., Bosch, J. F., Boutigny, D., Brackett, C. A., Bradshaw, A., Brandt, W. N., Brown, M. E., Bullock, J. S., Burchat, P., Burke, D. L., Cagnoli, G., Calabrese, D., Callahan, S., Callen, A. L., Carlin, J. L., Carlson, E. L., Chandrasekharan, S., Charles-Emerson, G.,

Chesley, S., Cheu, E. C., Chiang, H.-F., Chiang, J., Chirino, C., Chow, D., Ciardi, D. R., Claver, C. F., Cohen-Tanugi, J., Cockrum, J. J., Coles, R., Connolly, A. J., Cook, K. H., Cooray, A., Covey, K. R., Cribbs, C., Cui, W., Cutri, R., Daly, P. N., Daniel, S. F., Daruich, F., Daubard, G., Daues, G., Dawson, W., Delgado, F., Dellapenna, A., de Peyster, R., de Val-Borro, M., Digel, S. W., Doherty, P., Dubois, R., Dubois-Felsmann, G. P., Durech, J., Economou, F., Eifler, T., Eracleous, M., Emmons, B. L., Fausti Neto, A., Ferguson, H., Figueroa, E., Fisher-Levine, M., Focke, W., Foss, M. D., Frank, J., Freemon, M. D., Gangler, E., Gawiser, E., Geary, J. C., Gee, P., Geha, M., Gessner, C. J. B., Gibson, R. R., Gilmore, D. K., Glanzman, T., Glick, W., Goldina, T., Goldstein, D. A., Goodenow, I., Graham, M. L., Gressler, W. J., Gris, P., Guy, L. P., Guyonnet, A., Haller, G., Harris, R., Hascall, P. A., Haupt, J., Hernandez, F., Herrmann, S., Hileman, E., Hoblitt, J., Hodgson, J. A., Hogan, C., Howard, J. D., Huang, D., Huffer, M. E., Ingraham, P., Innes, W. R., Jacoby, S. H., Jain, B., Jammes, F., Jee, M. J., Jenness, T., Jernigan, G., Jevremović, D., Johns, K., Johnson, A. S., Johnson, M. W. G., Jones, R. L., Juramy-Gilles, C., Jurić, M., Kalirai, J. S., Kallivayalil, N. J., Kalmbach, B., Kantor, J. P., Karst, P., Kasliwal, M. M., Kelly, H., Kessler, R., Kinnison, V., Kirkby, D., Knox, L., Kotov, I. V., Krabbendam, V. L., Krughoff, K. S., Kubánek, P., Kuczewski, J., Kulkarni, S., Ku, J., Kurita, N. R., Lage, C. S., Lambert, R., Lange, T., Langton, J. B., Le Guillou, L., Levine, D., Liang, M., Lim, K.-T., Lintott, C. J., Long, K. E., Lopez, M., Lotz, P. J., Lupton, R. H., Lust, N. B., MacArthur, L. A., Mahabal, A., Mandelbaum, R., Markiewicz, T. W., Marsh, D. S., Marshall, P. J., Marshall, S., May, M., McKercher, R., McQueen, M., Meyers, J., Migliore, M., Miller, M., Mills, D. J., Miraval, C., Moeyens, J., Moolekamp, F. E., Monet, D. G., Moniez, M., Monkewitz, S., Montgomery, C., Morrison, C. B., Mueller, F., Muller, G. P., Muñoz Arancibia, F., Neill, D. R., Newbry, S. P., Nief, J.-Y., Nomerotski, A., Nordby, M., O'Connor, P., Oliver, J., Olivier, S. S., Olsen, K., O'Mullane, W., Ortiz, S., Osier, S., Owen, R. E., Pain, R., Palecek, P. E., Parejko, J. K., Parsons, J. B., Pease, N. M., Peterson, J. M., Peterson, J. R., Petravick, D. L., Libby Petrick, M. E., Petry, C. E., Pierfederici, F., Pietrowicz, S., Pike, R., Pinto, P. A., Plante, R., Plate, S., Plutchak, J. P., Price, P. A., Prouza, M., Radeka, V., Rajagopal, J., Rasmussen, A. P., Regnault, N., Reil, K. A., Reiss, D. J., Reuter, M. A., Ridgway, S. T., Riot, V. J., Ritz, S., Robinson, S., Roby, W., Roodman, A., Rosing, W., Roucelle, C., Rumore, M. R., Russo, S., Saha, A., Sassolas, B., Schalk, T. L., Schellart, P., Schindler, R. H., Schmidt, S., Schneider, D. P., Schneider, M. D., Schoening, W., Schumacher, G., Schwamb, M. E., Sebag, J., Selvy, B., Sembroski, G. H., Seppala, L. G., Serio, A., Serrano, E., Shaw, R. A., Shipsey, I., Sick, J., Silvestri, N., Slater, C. T., Smith, J. A., Smith, R. C., Sobhani, S., Soldahl, C., Storrie-Lombardi, L., Stover, E., Strauss, M. A., Street, R. A., Stubbs, C. W., Sullivan, I. S., Sweeney, D., Swinbank, J. D., Szalay, A., Takacs, P., Tether, S. A., Thaler, J. J., Thayer, J. G., Thomas, S., Thornton, A. J., Thukral, V., Tice, J., Trilling, D. E., Turri, M., Van Berg, R., Vanden Berk, D., Vetter, K., Virieux, F., Vucina, T., Wahl, W., Walkowicz, L., Walsh, B., Walter, C. W., Wang, D. L., Wang, S.-Y., Warner, M., Wiecha, O., Willman, B., Winters, S. E., Wittman, D., Wolff, S. C., Wood-Vasey, W. M., Wu, X., Xin, B., Yoachim, P., & Zhan, H., 2019. LSST: From Science Drivers to Reference Design and Anticipated Data Products, *ApJ*, **873**(2), 111.

Jarvis, M. J. & Rawlings, S., 2004. The accretion history of the universe with the SKA, *New Astronomy Reviews*, **48**(11-12), 1173–1185.

Jester, S. & Falcke, H., 2009. Science with a lunar low-frequency array: From the dark ages of the Universe to nearby exoplanets, *New Astronomy Reviews*, **53**(1-2), 1–26.

- Jones, R. C., 1941. A new calculus for the treatment of optical systems. description and discussion of the calculus, *J. Opt. Soc. Am.*, **31**(7), 488–493.
- Kessler, D. J., 1991. Collisional cascading: The limits of population growth in low earth orbit, *Advances in Space Research*, **11**(12), 63–66.
- Khan, F. M., Just, A., & Merritt, D., 2011. Efficient merger of binary supermassive black holes in merging galaxies, *ApJ*, **732**(2), 89.
- Khan, F. M., Holley-Bockelmann, K., Berczik, P., & Just, A., 2013. Supermassive black hole binary evolution in axisymmetric galaxies: The final parsec problem is not a problem, *ApJ*, **773**(2), 100.
- Kharb, P., Lal, D. V., & Merritt, D., 2017. A candidate sub-parsec binary black hole in the Seyfert galaxy NGC, *Nat. Astron.*, **1**(10), 727–733.
- Khosla, A., Jayadevaprakash, N., Yao, B., & Fei-Fei, L., 2011. Novel dataset for fine-grained image categorization, in *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO.
- Kingma, D. P. & Ba, J., 2014. Adam: A method for stochastic optimization, *arXiv preprint arXiv:1412.6980*.
- Komissarov, S. S., Barkov, M. V., Vlahakis, N., & Königl, A., 2007. Magnetic acceleration of relativistic active galactic nucleus jets, *MNRAS*, **380**(1), 51–70.
- Komossa, S. & Zensus, J. A., 2016. Compact object mergers: Observations of supermassive binary black holes and stellar tidal disruption events, in *Proc. Int. Astron. Union*, vol. 10, pp. 13–25, Cambridge University Press.
- Komossa, S., Burwitz, V., Hasinger, G., Predehl, P., Kaastra, J. S., & Ikebe, Y., 2003. Discovery of a Binary Active Galactic Nucleus in the Ultraluminous Infrared Galaxy NGC 6240 Using Chandra, *ApJ*, **582**(1), L15–L19.
- Kormendy, J. & Ho, L. C., 2013. Coevolution (Or Not) of Supermassive Black Holes and Host Galaxies, *ARA&A*, **51**(1), 511–653.
- Le Conte, Z. A., Elvis, M., & Gläser, P. A., 2023. Lunar far-side radio arrays: a preliminary site survey, *RAS Techniques and Instruments*, **2**(1), 360–377.
- LeCun, Y., Bengio, Y., & Hinton, G., 2015. Deep learning, *Nature*, **521**(7553), 436–444.
- Li, D. & Pan, Z., 2016. The Five-hundred-meter Aperture Spherical Radio Telescope Project, *Radio Science*, **51**, 1060–1064.
- Li, Z., Yu, C., Xiao, J., Long, M., & Cui, C., 2021. Detection of radio frequency interference using an improved generative adversarial network, *Astronomy and Computing*, **36**, 100482.
- Lieu, M., Conversi, L., Altieri, B., & Carry, B., 2019. Detecting Solar system objects with convolutional neural networks, *MNRAS*, **485**(4), 5831–5842.
- Lilly, S. J., Le Fevre, O., Crampton, D., Hammer, F., & Tresse, L., 1995. The Canada-France Redshift Survey. I. Introduction to the Survey, Photometric Catalogs, and Surface Brightness Selection Effects, *ApJ*, **455**, 50.

- Liu, X., Greene, J. E., Shen, Y., & Strauss, M. A., 2010. Discovery of Four kpc-Scale Binary AGNs, *ApJL*, **715**(1 PART 2), 30–34.
- Lockman, F. J., Jahoda, K., & McCammon, D., 1986. The Structure of Galactic H I in Directions of Low Total Column Density, *ApJ*, **302**, 432.
- Long, M., Yang, Z., Xiao, J., Yu, C., & Zhang, B., 2019. U-NetIM: An Improved U-Net for Automatic Recognition of RFIs, in *Astronomical Data Analysis Software and Systems XXVII*, vol. 523 of **Astronomical Society of the Pacific Conference Series**, p. 123.
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F., 2007. A Bright Millisecond Radio Burst of Extragalactic Origin, *Science*, **318**(5851), 777.
- Ma, P. X., Ng, C., Rizk, L., Croft, S., Siemion, A. P. V., Brzycki, B., Czech, D., Drew, J., Gajjar, V., Hoang, J., Isaacson, H., Lebofsky, M., MacMahon, D. H. E., de Pater, I., Price, D. C., Sheikh, S. Z., & Worden, S. P., 2023. A deep-learning search for technosignatures from 820 nearby stars, *Nature Astronomy*, **7**, 492–502.
- Madau, P. & Dickinson, M., 2014. Cosmic Star-Formation History, *ARA&A*, **52**, 415–486.
- Madau, P. & Rees, M. J., 2001. Massive Black Holes as Population III Remnants, *ApJ Letters*, **551**(1), L27–L30.
- Madau, P., Ferguson, H. C., Dickinson, M. E., Giavalisco, M., Steidel, C. C., & Fruchter, A., 1996. High-redshift galaxies in the Hubble Deep Field: colour selection and star formation history to $z \sim 4$, *MNRAS*, **283**(4), 1388–1404.
- Maiolino, R., Uebler, H., Perna, M., Scholtz, J., D’Eugenio, F., Witten, C., Laporte, N., Witstok, J., Carniani, S., Tacchella, S., Baker, W., Arribas, S., Nakajima, K., Eisenstein, D., Bunker, A., Charlot, S., Cresci, G., Curti, M., Curtis-Lake, E., de Graaff, A., Ji, Z., Johnson, B. D., Kumari, N., Looser, T. J., Maseda, M., Robertson, B., Rodríguez Del Pino, B., Sandles, L., Simmonds, C., Smit, R., Sun, F., Venturi, G., Williams, C., & Willmer, C., 2023. JWST-JADES. Possible Population III signatures at $z=10.6$ in the halo of GN-z11, *arXiv e-prints*, p. arXiv:2306.00953.
- Maiolino, R., Scholtz, J., Witstok, J., Carniani, S., D’Eugenio, F., de Graaff, A., Übler, H., Tacchella, S., Curtis-Lake, E., Arribas, S., Bunker, A., Charlot, S., Chevillard, J., Curti, M., Looser, T. J., Maseda, M. V., Rawle, T. D., Rodríguez del Pino, B., Willott, C. J., Egami, E., Eisenstein, D. J., Hainline, K. N., Robertson, B., Williams, C. C., Willmer, C. N. A., Baker, W. M., Boyett, K., DeCoursey, C., Fabian, A. C., Helton, J. M., Ji, Z., Jones, G. C., Kumari, N., Laporte, N., Nelson, E. J., Perna, M., Sandles, L., Shivaiei, I., & Sun, F., 2024. A small and vigorous black hole in the early universe, *Nature*, **627**(8002), 59–63.
- Manchester, R. N., Hobbs, G., Bailes, M., Coles, W. A., Van Straten, W., Keith, M. J., Shannon, R. M., Bhat, N. D., Brown, A., Burke-Spolaor, S. G., Champion, D. J., Chaudhary, A., Edwards, R. T., Hampson, G., Hotan, A. W., Jameson, A., Jenet, F. A., Kesteven, M. J., Khoo, J., Kocz, J., MacIesiak, K., Oslowski, S., Ravi, V., Reynolds, J. R., Sarkissian, J. M., Verbiest, J. P., Wen, Z. L., Wilson, W. E., Yardley, D., Yan, W. M., & You, X. P., 2013. The parkes pulsar timing array project, *Publ. Astron. Soc. Aust.*, **30**(1), e017.

- Mancuso, C., Lapi, A., Cai, Z.-Y., Negrello, M., De Zotti, G., Bressan, A., Bonato, M., Perrotta, F., & Danese, L., 2015. Predictions for Ultra-deep Radio Counts of Star-forming Galaxies, *ApJ*, **810**(1), 72.
- McDowell, J. C., 2020. The Low Earth Orbit Satellite Population and Impacts of the SpaceX Starlink Constellation, *ApJ Letters*, **892**(2), L36.
- McKinney, J. C., Tchekhovskoy, A., & Blandford, R. D., 2012. General relativistic magnetohydrodynamic simulations of magnetically choked accretion flows around black holes, *MNRAS*, **423**(4), 3083–3117.
- McNamara, B. R. & Nulsen, P. E. J., 2012. Mechanical feedback from active galactic nuclei in galaxies, groups and clusters, *New Journal of Physics*, **14**(5), 055023.
- Mirza, M. & Osindero, S., 2014. Conditional generative adversarial nets, *CoRR*, **abs/1411.1784**.
- Mitchell, D., Greenhill, L. J., Clark, M., Briggs, F., Bowman, J., Cappallo, R., Kaplan, D. L., Kasper, J., Kocz, J., Lonsdale, C. J., Stevens, J., Tingay, S., & Whitney, A., 2010. The Murchison Widefield Array, in *RFI Mitigation Workshop*, p. 16.
- Mohan, N. & Rafferty, D., 2015. PyBDSF: Python Blob Detection and Source Finder, Astrophysics Source Code Library, record ascl:1502.007.
- Moldon, J., 2021. eMCP: e-MERLIN CASA pipeline, Astrophysics Source Code Library, record ascl:2109.006.
- Mooley, K. P., Hallinan, G., Bourke, S., Horesh, A., Myers, S. T., Frail, D. A., Kulkarni, S. R., Levitan, D. B., Kasliwal, M. M., Cenko, S. B., Cao, Y., Bellm, E., & Laher, R. R., 2016. THE CALTECH-NRAO STRIPE 82 SURVEY (CNSS) PAPER. I. THE PILOT RADIO TRANSIENT SURVEY IN 50 DEG 2, *ApJ*, **818**(2), 105.
- Morris, D., Radhakrishnan, V., & Seielstad, G. A., 1964. On the Measurement of Polarization Distributions Over Radio Sources., *ApJ*, **139**, 551.
- Muxlow, T. W. B., Thomson, A. P., Radcliffe, J. F., Wrigley, N. H., Beswick, R. J., Smail, I., McHardy, I. M., Garrington, S. T., Ivison, R. J., Jarvis, M. J., Prandoni, I., Bondi, M., Guidetti, D., Argo, M. K., Bacon, D., Best, P. N., Biggs, A. D., Chapman, S. C., Coppin, K., Chen, H., Garratt, T. K., Garrett, M. A., Ibar, E., Kneib, J.-P., Knudsen, K. K., Koopmans, L. V. E., Morabito, L. K., Murphy, E. J., Njeri, A., Pearson, C., Pérez-Torres, M. A., Richards, A. M. S., Röttgering, H. J. A., Sargent, M. T., Serjeant, S., Simpson, C., Simpson, J. M., Swinbank, A. M., Varenus, E., & Venturi, T., 2020. The e-MERGE Survey (e-MERLIN Galaxy Evolution Survey): overview and survey description, *MNRAS*, **495**(1), 1188–1208.
- Netzer, H., 2015. Revisiting the Unified Model of Active Galactic Nuclei, *ARA&A*, **53**, 365–408.
- Offringa, A. R., de Bruyn, A. G., Biehl, M., Zaroubi, S., Bernardi, G., & Pandey, V. N., 2010a. Post-correlation radio frequency interference classification methods, *MNRAS*, **405**(1), 155–167.
- Offringa, A. R., de Bruyn, A. G., Zaroubi, S., & Biehl, M., 2010b. Post-correlation RFI detection, in *RFI Mitigation Workshop*, p. 36.

- Offringa, A. R., van de Gronde, J. J., & Roerdink, J. B. T. M., 2012a. A morphological algorithm for improving radio-frequency interference detection, *A&A*, **539**, A95.
- Offringa, A. R., van de Gronde, J. J., & Roerdink, J. B. T. M., 2012b. A morphological algorithm for improving radio-frequency interference detection, *A&A*, **539**, A95.
- Offringa, A. R., McKinley, B., Hurley-Walker, N., Briggs, F. H., Wayth, R. B., Kaplan, D. L., Bell, M. E., Feng, L., Neben, A. R., Hughes, J. D., Rhee, J., Murphy, T., Bhat, N. D. R., Bernardi, G., Bowman, J. D., Cappallo, R. J., Corey, B. E., Deshpande, A. A., Emrich, D., Ewall-Wice, A., Gaensler, B. M., Goeke, R., Greenhill, L. J., Hazelton, B. J., Hindson, L., Johnston-Hollitt, M., Jacobs, D. C., Kasper, J. C., Kratzenberg, E., Lenc, E., Lonsdale, C. J., Lynch, M. J., McWhirter, S. R., Mitchell, D. A., Morales, M. F., Morgan, E., Kudryavtseva, N., Oberoi, D., Ord, S. M., Pindor, B., Procopio, P., Prabu, T., Riding, J., Roshi, D. A., Shankar, N. U., Srivani, K. S., Subrahmanyam, R., Tingay, S. J., Waterson, M., Webster, R. L., Whitney, A. R., Williams, A., & Williams, C. L., 2014. WSCLEAN: an implementation of a fast, generic wide-field imager for radio astronomy, *MNRAS*, **444**(1), 606–619.
- OpenAI, :, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Leoni Aleman, F., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Baltescu, P., Bao, H., Bavarian, M., Belgum, J., Bello, I., Berdine, J., Bernadett-Shapiro, G., Berner, C., Bogdonoff, L., Boiko, O., Boyd, M., Brakman, A.-L., Brockman, G., Brooks, T., Brundage, M., Button, K., Cai, T., Campbell, R., Cann, A., Carey, B., Carlson, C., Carmichael, R., Chan, B., Chang, C., Chantzis, F., Chen, D., Chen, S., Chen, R., Chen, J., Chen, M., Chess, B., Cho, C., Chu, C., Chung, H. W., Cummings, D., Currier, J., Dai, Y., Decareaux, C., Degry, T., Deutsch, N., Deville, D., Dhar, A., Dohan, D., Dowling, S., Dunning, S., Ecoffet, A., Eleti, A., Eloundou, T., Farhi, D., Fedus, L., Felix, N., Posada Fishman, S., Forte, J., Fulford, I., Gao, L., Georges, E., Gibson, C., Goel, V., Gogineni, T., Goh, G., Gontijo-Lopes, R., Gordon, J., Grafstein, M., Gray, S., Greene, R., Gross, J., Gu, S. S., Guo, Y., Hallacy, C., Han, J., Harris, J., He, Y., Heaton, M., Heidecke, J., Hesse, C., Hickey, A., Hickey, W., Hoeschele, P., Houghton, B., Hsu, K., Hu, S., Hu, X., Huizinga, J., Jain, S., Jain, S., Jang, J., Jiang, A., Jiang, R., Jin, H., Jin, D., Jomoto, S., Jonn, B., Jun, H., Kaftan, T., Kaiser, L., Kamali, A., Kanitscheider, I., Shirish Keskar, N., Khan, T., Kilpatrick, L., Kim, J. W., Kim, C., Kim, Y., Kirchner, H., Kiros, J., Knight, M., Kokotajlo, D., Kondraciuk, Ł., Kondrich, A., Konstantinidis, A., Kosic, K., Krueger, G., Kuo, V., Lampe, M., Lan, I., Lee, T., Leike, J., Leung, J., Levy, D., Li, C. M., Lim, R., Lin, M., Lin, S., Litwin, M., Lopez, T., Lowe, R., Lue, P., Makanju, A., Malfacini, K., Manning, S., Markov, T., Markovski, Y., Martin, B., Mayer, K., Mayne, A., McGrew, B., McKinney, S. M., McLeavey, C., McMillan, P., McNeil, J., Medina, D., Mehta, A., Menick, J., Metz, L., Mishchenko, A., Mishkin, P., Monaco, V., Morikawa, E., Mossing, D., Mu, T., Murati, M., Murk, O., Mély, D., Nair, A., Nakano, R., Nayak, R., Neelakantan, A., Ngo, R., Noh, H., Ouyang, L., O’Keefe, C., Pachocki, J., Paino, A., Palermo, J., Pantuliano, A., Parascandolo, G., Parish, J., Parparita, E., Passos, A., Pavlov, M., Peng, A., Perelman, A., de Avila Belbute Peres, F., Petrov, M., Ponde de Oliveira Pinto, H., Michael, Pokorny, Pokrass, M., Pong, V., Powell, T., Power, A., Power, B., Proehl, E., Puri, R., Radford, A., Rae, J., Ramesh, A., Raymond, C., Real, F., Rimbach, K., Ross, C., Rotsted, B., Roussez, H., Ryder, N., Saltarelli, M., Sanders, T., Santurkar, S., Sasstry, G., Schmidt, H., Schnurr, D., Schulman, J., Selsam, D., Sheppard, K., Sherbakov, T., Shieh, J., Shoker, S., Shyam, P., Sidor, S., Sigler, E., Simens, M., Sitkin, J., Slama, K., Sohl, I., Sokolowsky, B., Song, Y., Staudacher, N., Petroski Such, F., Summers, N.,

- Sutskever, I., Tang, J., Tezak, N., Thompson, M., Tillet, P., Tootoonchian, A., Tseng, E., Tuggle, P., Turley, N., Tworek, J., Cerón Uribe, J. F., Vallone, A., Vijayvergiya, A., Voss, C., Wainwright, C., Wang, J. J., Wang, A., Wang, B., Ward, J., Wei, J., Weinmann, C., Welihinda, A., Welinder, P., Weng, J., Weng, L., Wiethoff, M., Willner, D., Winter, C., Wolrich, S., Wong, H., Workman, L., Wu, S., Wu, J., Wu, M., Xiao, K., Xu, T., Yoo, S., Yu, K., Yuan, Q., Zaremba, W., Zellers, R., Zhang, C., Zhang, M., Zhao, S., Zheng, T., Zhuang, J., Zhuk, W., & Zoph, B., 2023. GPT-4 Technical Report, *arXiv e-prints*, p. arXiv:2303.08774.
- Owen, F. N. & Morrison, G. E., 2008. The Deep Swire Field. I. 20 cm Continuum Radio Observations: A Crowded Sky, *Astronomical Journal*, **136**(5), 1889–1900.
- Padovani, P., 2016. The faint radio sky: radio astronomy becomes mainstream, *A&AR*, **24**(1), 13.
- Padovani, P., 2017. On the two main classes of active galactic nuclei, *Nature Astronomy*, **1**, 0194.
- Padovani, P., Miller, N., Kellermann, K. I., Mainieri, V., Rosati, P., & Tozzi, P., 2011. The VLA Survey of Chandra Deep Field South. V. Evolution and Luminosity Functions of Sub-millijansky Radio Sources and the Issue of Radio Emission in Radio-quiet Active Galactic Nuclei, *ApJ*, **740**(1), 20.
- Padovani, P., Bonzini, M., Kellermann, K. I., Miller, N., Mainieri, V., & Tozzi, P., 2015. Radio-faint AGN: a tale of two populations, *MNRAS*, **452**(2), 1263–1279.
- Padovani, P., Alexander, D. M., Assef, R. J., De Marco, B., Giommi, P., Hickox, R. C., Richards, G. T., Smolčić, V., Hatziminaoglou, E., Mainieri, V., & Salvato, M., 2017. Active galactic nuclei: what’s in a name?, *A&AR*, **25**(1), 2.
- Panessa, F., Baldi, R. D., Laor, A., Padovani, P., Behar, E., & McHardy, I., 2019. The origin of radio emission from radio-quiet active galactic nuclei, *Nature Astronomy*, **3**, 387–396.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., & Chintala, S., 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library, in *Advances in Neural Information Processing Systems 32*, pp. 8024–8035, Curran Associates, Inc.
- Perley, R. A., Chandler, C. J., Butler, B. J., & Wrobel, J. M., 2011. The Expanded Very Large Array: A New Telescope for New Science, *ApJ Letters*, **739**(1), L1.
- Petroff, E., Hessels, J. W. T., & Lorimer, D. R., 2019. Fast radio bursts, *A&AR*, **27**(1), 4.
- Pinchuk, P. & Margot, J.-L., 2022. A Machine Learning-based Direction-of-origin Filter for the Identification of Radio Frequency Interference in the Search for Technosignatures, *Astronomical Journal*, **163**(2), 76.
- Prestage, R. M., Constantikes, K. T., Hunter, T. R., King, L. J., Lacasse, R. J., Lockman, F. J., & Norrod, R. D., 2009. The Green Bank Telescope, *IEEE Proceedings*, **97**(8), 1382–1390.
- Pringle, J. E., 1981. Accretion discs in astrophysics, *ARA&A*, **19**, 137–162.

- Qahwaji, R. & Colak, T., 2007. Automatic Short-Term Solar Flare Prediction Using Machine Learning and Sunspot Associations, *Solar Physics*, **241**(1), 195–211.
- Rees, M. J., 1984. Black Hole Models for Active Galactic Nuclei, *ARA&A*, **22**, 471–506.
- Reynoso, E. M., Moffett, D. A., Goss, W. M., Dubner, G. M., Dickel, J. R., Reynolds, S. P., & Giacani, E. B., 1997. A VLA Study of the Expansion of Tycho’s Supernova Remnant, *ApJ*, **491**(2), 816–828.
- Robbins, H. & Monro, S., 1951. A Stochastic Approximation Method, *The Annals of Mathematical Statistics*, **22**(3), 400 – 407.
- Rodriguez, C., Taylor, G. B., Zavala, R. T., Peck, A. B., Pollack, L. K., & Romani, R. W., 2006. A Compact Supermassive Binary Black Hole System, *ApJ*, **646**(1), 49–60.
- Ronneberger, O., Fischer, P., & Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pp. 234–241, Springer International Publishing, Cham.
- Rowan-Robinson, M., Babbedge, T., Oliver, S., Trichas, M., Berta, S., Lonsdale, C., Smith, G., Shupe, D., Surace, J., Arnouts, S., Ilbert, O., Le Fèvre, O., Afonso-Luis, A., Perez-Fournon, I., Hatziminaoglou, E., Polletta, M., Farrah, D., & Vaccari, M., 2008. Photometric redshifts in the SWIRE Survey, *MNRAS*, **386**(2), 697–714.
- Rubinur, K., Das, M., & Kharb, P., 2018. Searching for dual active galactic nuclei, *J. Astrophys. Astron.*, **39**(1), 8.
- Rybicki, G. B. & Lightman, A. P., 1979. *Radiative processes in astrophysics*.
- Schmidt, M., 1963. 3C 273 : A Star-Like Object with Large Red-Shift, *Nature*, **197**(4872), 1040.
- Schwab, F. R., 1984. Relaxing the isoplanatism assumption in self-calibration; applications to low-frequency radio interferometry, *Astronomical Journal*, **89**, 1076–1081.
- Sekido, M. & Fukushima, T., 2006. A VLBI Delay Model for Radio Sources at a Finite Distance, *Journal of Geodesy*, **80**, 137–149.
- Shallue, C. J. & Vanderburg, A., 2018. Identifying Exoplanets with Deep Learning: A Five-planet Resonant Chain around Kepler-80 and an Eighth Planet around Kepler-90, *Astronomical Journal*, **155**(2), 94.
- Shimwell, T. W., Röttgering, H. J. A., Best, P. N., Williams, W. L., Dijkema, T. J., de Gasperin, F., Hardcastle, M. J., Heald, G. H., Hoang, D. N., Horneffer, A., Intema, H., Mahony, E. K., Mandal, S., Mechev, A. P., Morabito, L., Oonk, J. B. R., Rafferty, D., Retana-Montenegro, E., Sabater, J., Tasse, C., van Weeren, R. J., Brügger, M., Brunetti, G., Chyży, K. T., Conway, J. E., Haverkorn, M., Jackson, N., Jarvis, M. J., McKean, J. P., Miley, G. K., Morganti, R., White, G. J., Wise, M. W., van Bemmell, I. M., Beck, R., Brienza, M., Bonafede, A., Calistro Rivera, G., Cassano, R., Clarke, A. O., Cseh, D., Deller, A., Drabent, A., van Driel, W., Engels, D., Falcke, H., Ferrari, C., Fröhlich, S., Garrett, M. A., Harwood, J. J., Heesen, V., Hoeft, M., Horellou, C., Israel, F. P., Kapińska, A. D., Kunert-Bajraszewska, M., McKay, D. J., Mohan, N. R., Orrú, E., Pizzo, R. F., Prandoni, I., Schwarz, D. J., Shulevski, A., Sipiior, M., Smith, D. J. B., Sridhar, S. S., Steinmetz, M., Stroe, A., Varenus, E., van der Werf, P. P., Zensus, J. A., & Zwart, J. T. L., 2017. The LOFAR Two-metre Sky Survey. I. Survey description and preliminary data release, *A&A*, **598**, A104.

- Smartt, S. J., 2009. Progenitors of Core-Collapse Supernovae, *ARA&A*, **47**(1), 63–106.
- Smirnov, O. M., 2011. Revisiting the radio interferometer measurement equation. I. A full-sky Jones formalism, *A&A*, **527**, A106.
- Smirnov, O. M. & Tasse, C., 2015. Radio interferometric gain calibration as a complex optimization problem, *MNRAS*, **449**(3), 2668–2684.
- Smolcic, V., Padovani, P., Delhaize, J., Prandoni, I., Seymour, N., Jarvis, M., Afonso, J., Magliocchetti, M., Huynh, M., Vaccari, M., & Karim, A., 2015. Exploring AGN Activity over Cosmic Time with the SKA, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, p. 69.
- Tchekhovskoy, A., Narayan, R., & McKinney, J. C., 2011. Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole, *MNRAS*, **418**(1), L79–L83.
- Thomasson, P., 1986. MERLIN., *Quarterly Journal of the Royal Astronomical Society*, **27**, 413–431.
- Thompson, A. R., Moran, J. M., & Swenson, George W., J., 2017. *Interferometry and Synthesis in Radio Astronomy*, 3rd edn.
- Tyson, J. A., Ivezić, Ž., Bradshaw, A., Rawls, M. L., Xin, B., Yoachim, P., Parejko, J., Greene, J., Sholl, M., Abbott, T. M. C., & Polin, D., 2020. Mitigation of LEO Satellite Brightness and Trail Effects on the Rubin Observatory LSST, *Astronomical Journal*, **160**(5), 226.
- Urry, C. M. & Padovani, P., 1995. Unified Schemes for Radio-Loud Active Galactic Nuclei, *Publications of the Astronomical Society of the Pacific*, **107**, 803.
- Vafaei Sadr, A., Bassett, B. A., Oozeer, N., Fantaye, Y., & Finlay, C., 2020. Deep learning improves identification of Radio Frequency Interference, *MNRAS*, **499**(1), 379–390.
- van Cittert, P., 1934. Die wahrscheinliche schwingungsverteilung in einer von einer lichtquelle direkt oder mittels einer linse beleuchteten ebene, *Physica*, **1**(1), 201–210.
- van der Kruit, P. C., 1973. High-resolution Radio Continuum Observations of Bright Spiral Galaxies at 1415 MHz: A General Discussion, *A&A*, **29**, 263.
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., Heald, G., McKean, J. P., Hessels, J. W. T., de Bruyn, A. G., Nijboer, R., Swinbank, J., Fallows, R., Brentjens, M., Nelles, A., Beck, R., Falcke, H., Fender, R., Hörandel, J., Koopmans, L. V. E., Mann, G., Miley, G., Röttgering, H., Stappers, B. W., Wijers, R. A. M. J., Zaroubi, S., van den Akker, M., Alexov, A., Anderson, J., Anderson, K., van Ardenne, A., Arts, M., Asgekar, A., Avruch, I. M., Batejat, F., Bähren, L., Bell, M. E., Bell, M. R., van Bemmell, I., Bannema, P., Bentum, M. J., Bernardi, G., Best, P., Bîrzan, L., Bonafede, A., Boonstra, A. J., Braun, R., Bregman, J., Breitling, F., van de Brink, R. H., Broderick, J., Broekema, P. C., Brouw, W. N., Brüggem, M., Butcher, H. R., van Cappellen, W., Ciardi, B., Coenen, T., Conway, J., Coolen, A., Corstanje, A., Damstra, S., Davies, O., Deller, A. T., Dettmar, R. J., van Diepen, G., Dijkstra, K., Donker, P., Doorduyn, A., Dromer, J., Drost, M., van Duin, A., Eislöffel, J., van Enst, J., Ferrari, C., Frieswijk, W., Gankema, H., Garrett, M. A., de Gasperin, F., Gerbers, M., de Geus, E., Griebmeier, J. M., Grit, T., Gruppen, P., Hamaker, J. P., Hassall, T., Hoefl, M., Holties, H. A.,

Horneffer, A., van der Horst, A., van Houwelingen, A., Huijgen, A., Iacobelli, M., Intema, H., Jackson, N., Jelic, V., de Jong, A., Juette, E., Kant, D., Karastergiou, A., Koters, A., Kollen, H., Kondratiev, V. I., Kooistra, E., Koopman, Y., Koster, A., Kuniyoshi, M., Kramer, M., Kuper, G., Lambropoulos, P., Law, C., van Leeuwen, J., Lemaitre, J., Loose, M., Maat, P., Macario, G., Markoff, S., Masters, J., McFadden, R. A., McKay-Bukowski, D., Meijering, H., Meulman, H., Mevius, M., Middelberg, E., Millenaar, R., Miller-Jones, J. C. A., Mohan, R. N., Mol, J. D., Morawietz, J., Morganti, R., Mulcahy, D. D., Mulder, E., Munk, H., Nieuwenhuis, L., van Nieuwpoort, R., Noordam, J. E., Norden, M., Noutsos, A., Offringa, A. R., Olofsson, H., Omar, A., Orrú, E., Overeem, R., Paas, H., Pandey-Pommier, M., Pandey, V. N., Pizzo, R., Polatidis, A., Rafferty, D., Rawlings, S., Reich, W., de Reijer, J. P., Reitsma, J., Renting, G. A., Riemers, P., Rol, E., Romein, J. W., Roosjen, J., Ruiter, M., Scaife, A., van der Schaaf, K., Scheers, B., Schellart, P., Schoenmakers, A., Schoonderbeek, G., Serylak, M., Shulevski, A., Sluman, J., Smirnov, O., Sobey, C., Spreeuw, H., Steinmetz, M., Sterks, C. G. M., Stiepel, H. J., Stuurwold, K., Tagger, M., Tang, Y., Tasse, C., Thomas, I., Thoudam, S., Toribio, M. C., van der Tol, B., Usov, O., van Veelen, M., van der Veen, A. J., ter Veen, S., Verbiest, J. P. W., Vermeulen, R., Vermaas, N., Vocks, C., Vogt, C., de Vos, M., van der Wal, E., van Weeren, R., Weggemans, H., Weltevrede, P., White, S., Wijnholds, S. J., Wilhelmsson, T., Wucknitz, O., Yatawatta, S., Zarka, P., Zensus, A., & van Zwieten, J., 2013. LOFAR: The LOw-Frequency ARray, *A&A*, **556**, A2.

Vasiliev, E., Antonini, F., & Merritt, D., 2015. The Final-parsec Problem in the Collisionless Limit, *ApJ*, **810**(1), 49.

Venturi, T., Paragi, Z., Lindqvist, M., Bartkiewicz, A., Beswick, R., Bogdanović, T., Brisken, W., Charlot, P., Colomer, F., Conway, J., Frey, S., Guirado, J. C., Gurvits, L., van Langevelde, H., Lobanov, A., McKean, J., Morganti, R., Muxlow, T., Pérez-Torres, M., Rygl, K., Schulz, R., Szomoru, A., de Vicente, P., An, T., Anglada, G., Argo, M., Azulay, R., van Bemmelen, I., Bocanegra, T., Boccardi, B., Castangia, P., Chibueze, J., Cimò, G., Climent, J.-B., Deane, R., Deller, A., Dodson, R., Duev, D., Etoke, S., Fenech, D., Gabányi, K., Gabuzda, D., Garrett, M., Gawroński, M., Ghirlanda, G., Giroletti, M., Goddi, C., Gómez, J. L., Gray, M., Greaves, J., Hessels, J., van der Horst, A., Hunter, T., Laing, R., Vir Lal, D., Lambert, S., Loinard, L., Marcote, B., Merloni, A., Miller-Jones, J., Molera Calvés, G., Moscadelli, L., Olofsson, H., Petrov, L., Pizzo, R., Possenti, A., Quiroga-Núñez, L. H., Reynolds, C., Richards, A., Rioja, M., Sanna, A., Savolainen, T., Sbarrato, T., Spingola, C., Surcis, G., Trigilio, C., Varenus, E., Vlemmings, W., van Velzen, S., & van der Walt, J., 2020. VLBI20-30: a scientific roadmap for the next decade – The future of the European VLBI Network, *arXiv e-prints*, p. arXiv:2007.02347.

Verbiest, J. P., Lentati, L., Hobbs, G., Van Haasteren, R., Demorest, P. B., Janssen, G. H., Wang, J. B., Desvignes, G., Caballero, R. N., Keith, M. J., Champion, D. J., Arzoumanian, Z., Babak, S., Bassa, C. G., Bhat, N. D., Brazier, A., Brem, P., Burgay, M., Burke-Spolaor, S., Chamberlin, S. J., Chatterjee, S., Christy, B., Cognard, I., Cordes, J. M., Dai, S., Dolch, T., Ellis, J. A., Ferdman, R. D., Fonseca, E., Gair, J. R., Garver-Daniels, N. E., Gentile, P., Gonzalez, M. E., Graikou, E., Guillemot, L., Hessels, J. W., Jones, G., Karuppusamy, R., Kerr, M., Kramer, M., Lam, M. T., Lasky, P. D., Lassus, A., Lazarus, P., Lazio, T. J., Lee, K. J., Levin, L., Liu, K., Lynch, R. S., Lyne, A. G., Mckee, J., McLaughlin, M. A., McWilliams, S. T., Madison, D. R., Manchester, R. N., Mingarelli, C. M., Nice, D. J., Osłowski, S., Palliyaguru, N. T., Pennucci, T. T., Perera, B. B., Perrodin, D., Possenti, A., Petiteau, A., Ransom, S. M., Reardon, D., Rosado,

- P. A., Sanidas, S. A., Sesana, A., Shaifullah, G., Shannon, R. M., Siemens, X., Simon, J., Smits, R., Spiewak, R., Stairs, I. H., Stappers, B. W., Stinebring, D. R., Stovall, K., Swiggum, J. K., Taylor, S. R., Theureau, G., Tiburzi, C., Toomey, L., Vallisneri, M., Van Straten, W., Vecchio, A., Wang, Y., Wen, L., You, X. P., Zhu, W. W., & Zhu, X. J., 2016. The international pulsar timing array: First data release, *MNRAS*, **458**(2), 1267–1288.
- Vernstrom, T., Scott, D., Wall, J. V., Condon, J. J., Cotton, W. D., & Perley, R. A., 2016a. Deep 3-GHz observations of the Lockman Hole North with the Very Large Array - I. Source extraction and uncertainty analysis, *MNRAS*, **461**(3), 2879–2895.
- Vernstrom, T., Scott, D., Wall, J. V., Condon, J. J., Cotton, W. D., Kellermann, K. I., & Perley, R. A., 2016b. Deep 3-GHz observations of the Lockman Hole North with the Very Large Array - II. Catalogue and μ Jy source properties, *MNRAS*, **462**(3), 2934–2949.
- Véron-Cetty, M. P. & Véron, P., 2010. A catalogue of quasars and active nuclei: 13th edition, *A&A*, **518**, A10.
- Weedman, D., Polletta, M., Lonsdale, C. J., Wilkes, B. J., Siana, B., Houck, J. R., Surace, J., Shupe, D., Farrah, D., & Smith, H. E., 2006. Active Galactic Nucleus and Starburst Classification from Spitzer Mid-Infrared Spectra for High-Redshift SWIRE Sources, *ApJ*, **653**(1), 101–111.
- Wilman, R. J., Miller, L., Jarvis, M. J., Mauch, T., Levrier, F., Abdalla, F. B., Rawlings, S., Klöckner, H. R., Obreschkow, D., Olteanu, D., & Young, S., 2008. A semi-empirical simulation of the extragalactic radio continuum sky for next generation radio telescopes, *MNRAS*, **388**(3), 1335–1348.
- Wilman, R. J., Jarvis, M. J., Mauch, T., Rawlings, S., & Hickey, S., 2010. An infrared-radio simulation of the extragalactic sky: from the Square Kilometre Array to Herschel, *MNRAS*, **405**(1), 447–461.
- Wilson, T. L., Rohlfs, K., & Hüttemeister, S., 2013. *Tools of Radio Astronomy*, 6th edn.
- Woo, J.-H., Cho, H., Husemann, B., Komossa, S., Park, D., & Bennert, V., 2014. A sub-kpc-scale binary AGN with double narrow-line regions, *MNRAS*, **437**(1), 32–37.
- Woosley, S. E. & Weaver, T. A., 1986. The physics of supernova explosions., *ARA&A*, **24**, 205–253.
- Wootten, A. & Thompson, A. R., 2009. The Atacama Large Millimeter/Submillimeter Array, *IEEE Proceedings*, **97**(8), 1463–1471.
- Young, T., 1804. The bakerian lecture: Experiments and calculations relative to physical optics, *Philosophical Transactions of the Royal Society of London*, **94**, 1–16.
- Zernike, F., 1938. The concept of degree of coherence and its application to optical problems, *Physica*, **5**(8), 785–795.
- Zhang, H., Wu, M., Yue, Y., Gan, H., Hu, H., Huang, S., Sun, J., Song, J., Peng, P., & Nan, R., 2019. Rfi mitigation of fast: Challenge solution, in *2019 URSI Asia-Pacific Radio Science Conference (AP-RASC)*, pp. 1–3.

Frodo sat silent and motionless. Fear seemed to stretch out a vast hand, like a dark cloud rising in the East and looming up to engulf him. 'This ring!' he stammered. 'How, how on earth did it come to me?' 'Ah!' said Gandalf. 'That is a very long story. The beginnings lie back in the Black Years, which only the lore-masters now remember. If I were to tell you all that tale, we should still be sitting here when Spring had passed into Winter.'

'But last night I told you of Sauron the Great, the Dark Lord. The rumours that you have heard are true: he has indeed arisen again and left his hold in Mirkwood and returned to his ancient fastness in the Dark Tower of Mordor. That name even you hobbits have heard of, like a shadow on the borders of old stories. Always after a defeat and a respite, the Shadow takes another shape and grows again.'

'I wish it need not have happened in my time,' said Frodo.

'So do I,' said Gandalf, 'and so do all who live to see such times. But that is not for them to decide. All we have to decide is what to do with the time that is given us.'

– J. R. R. Tolkien
The Fellowship of the Ring