Machine learning for radio galaxy morphology analysis

Mostert, R.I.J.

1 | Introduction

Homo sapiens are but one of the millions of species that exist on Earth today. Some species only existed in the past and many will cease to exist in the nearby future. A curious mind might wonder how Homo sapiens and other species evolved, how they function, and explore what the future might hold for them. Scientists observed that a Homo sapiens does not pop into static existence: a human is created, born, evolves as a person within a lifetime, and as a species in hundred thousands of years. The evolution of species, the metamorphosis of life-forms within their lifetime and the complex interaction between life-forms spark our curiosity, they imply that many causal relationships are at play and that many are left, waiting to be discovered.

In contrast, the night sky might seem static, non-evolving and devoid of interaction. Even so, it is hard for a curious mind not to look up at this glittering firmament which triggers so many questions about the past, present and future, the origin of life, the Universe, and everything. This curiosity leads us from Earth to knowledge of the Sun around which we orbit, the other planets within our solar system, to the other stars in the night sky, around which other planetary systems regularly appear to orbit(!). On a larger scale, stars around us appear to be gravitationally bound in a (spiral) structure known to us as the Milky Way galaxy. Countless other galaxies appear to exist, some of them grouped in dense galaxy clusters, the largest known gravitationally bound objects in the Universe.

Studies in the past ten decades determined that the Universe is not so static after all: Earth, other planets, the Sun, other stars, the Milky Way, other galaxies, are formed, and evolve within their lifetime. They also evolve as classes of objects over the measured 13.8 billion years that our universe exists. The planets, stars and galaxies close to us in place and time are different from those that existed in the past, as even the abundance of elements and energetic state of the entire universe changed over time. Thus it is that our curious minds try to uncover the intriguing nature and the hidden causal relationships that are at play for galaxies and galaxy clusters — the objects that form the relatively rare, matter-dense spots in a Universe that is overwhelmingly void. This thesis, originating from a curious mind privileged to contemplate out-of-this-Earth-matters amongst other curious minds, is my microscopic contribution to our understanding of galaxies.

The motion of celestial objects has always influenced our lives and routines. The rise and set of the sun divide the day, the phases of the moon divide the month, and the sun's trajectory and the location of constellations divide the year.

For centuries, the naked eye was used to study the relative motion of celestial objects. Astronomers tried to quantify and capture this movement early on, with the orientation of buildings, scratch marks in stones and rudimentary instruments, and later in exhaustive written catalogues (e.g. North, 2008). The position of the stars marked the beginning of festive events and the shared sense of time allowed for the planning of events.

With the invention of the optical telescope at the beginning of the 17th century, not just the motion, but also the properties of celestial objects could be studied (e.g. King, 2003). Optical telescopes are instruments that use lenses or curved mirrors to gather a larger amount of light than normally enters our eye thereby allowing one to detect fainter objects with a higher angular resolution. Angular resolution indicates the smallest angle between two point sources of light for which the two can still be distinguished from one another.

In the 1860s Gustav Kirchhoff observes that hot objects, like the sun, fire, (or the later invented wire filament of an incandescent light bulb) emit visible light in a process known as black body radiation (Kirchhoff, 1860). Around the same time, James Maxwell postulated the theory that light can be described as electromagnetic waves (Maxwell, 1865). In this theory, the light that we see with our eyes is electromagnetic radiation with very short wavelengths. Still, light with lower (and higher) wavelengths should also exist and be detectable using instruments other than our eyes. Indeed, our eyes are only sensitive to a narrow range of wavelengths between 380 nm and 750 nm (referred to as visible light). The rest of the electromagnetic spectrum remained yet to be explored.

In the 1880s, Heinrich Hertz developed antennas — an antenna is principally an (array of) metal conductors — and experimentally confirmed that electromagnetic radiation was indeed wave-like. Hertz's experiments (described by D'Agostino, 2000) showcased the emission and detection of electromagnetic radiation with long wavelengths using antennas. Electromagnetic radiation with (long) wavelengths above 1 mm (or below 300 GHz) are now referred to as radio waves.

In the 1920s Heber Curtis, Ernst Öpik and Edwin Hubble provided proof that the Andromeda nebula which was earlier believed to be a cluster of stars within the Milky Way was in fact located far outside the confines of the Milky Way (e.g. Curtis, 1988; Opik, 1922; Hubble, 1929). This object, the first known extra-galactic object, was correctly believed to be its very own galaxy full of stars.

The visible light that reaches us from the night sky is dominated by the black-body radiation from the stars within the Milky Way. Observing the sky at different wavelengths must reveal black-body radiation from objects of different temperatures. However, as our eyes evolved to be most sensitive around the frequency at which the sun's black-body radiation is brightest, it was initially thought that electromagnetic radiation far outside the bandwidth of visible light (radio waves for example) would not reveal any astronomical objects. So it came to be that not an astronomer, but an engineer called Karl Jansky was the first to discover extra-terrestrial radio emission (Jansky, 1933). In 1932, Jansky worked at a telecom company on a project to identify sources of radio noise that might interfere with trans-Atlantic radio communications. By serendipity, he discovered that one of the sources of noise was radio emission that had to be of extraterrestrial origin. In hindsight, this was the moment that the field of radio astronomy was born. In follow-up observations, the electrical engineer and amateur astronomer Grote Reber, discovered that there were more extra-terrestrial radio waves than extra-terrestrial visible light (as recounted by Tyson, 2003). Thereby allowing astronomers to speculate on non-stellar origin of the radiation and radiation mechanisms other than black-body radiation.

Just like ever larger optical telescopes advanced astronomy in the visible light regime, larger radio dishes could advance radio astronomy. For a diffraction-limited telescope[1] with a circular aperture, the Rayleigh criterion indicates that its angular resolution $\theta \approx 1.22\lambda/D$, with λ the wavelength of the observed light, and D the diameter of the telescope aperture. As expected, this means that we can resolve smaller details with a larger aperture D. However, the dependence of λ means that observing radio waves with a wavelength of ~ 2 m (or a frequency of 150 MHz) requires a telescope aperture of four thousand km to attain the same angular resolution as an optical telescope with an aperture of 1 m. Clearly, a single radio dish would never match the angular resolution of an optical

[1]Where diffraction-limited means the bottle-neck for a high angular resolution is not due to lens-aberrations or atmospheric effects.

telescope with a 1 m aperture.

Around 1946, Martin Ryle, Joseph Lade Pawsey and Ruby Payne-Scott came up with the technique that enabled high-resolution radio astronomy. In a process known as aperture synthesis or radio interferometry, they were able to combine the received signal of multiple separate radio antennas that were a known number of radio waves apart (Ryle & Vonberg, 1946; McCready et al., 1947). Aperture synthesis is a complex process that makes use of the van Cittert-Zernike theorem (van Cittert, 1934; Zernike, 1938) to turn the correlation of the radio signal received by pairs of radio antennas into a measured brightness distribution of the observed source. See Thompson et al. (2017) for an introduction to radio interferometry, or the introduction of the PhD thesis of Oei 2023 (subm.) for a step-by-step derivation of the theorem. Crucially, radio interferometry allows for a diffraction-limited angular resolution $\theta \approx 1.22\lambda/B$, where B is not the diameter of radio dishes, but the longest distance or baseline between the antennas of a radio telescope. This break-through influences the design of radio telescopes up to this day. As can be imagined, two three-metre radio dishes four hundred metres apart lead to a lower sensitivity than one dish with a four-hundred metres diameter. Furthermore, a single baseline is only sensitive to structures of a specific angular size. As such, multiple antennas are needed to increase the absolute sensitivity, and they need to be placed at different distances with respect to each other to increase the sensitivity to structures of different angular sizes.

More extra-terrestrial radio sources were detected in the 1950s, but it took some time before these measurements were of high enough resolution to pair them with an optical counterpart. When optical counterparts were found for three radio sources in 1960 (e.g. Matthews & Sandage, 1963), the optical sources turned out to be faint, point-like objects and were thus referred to as quasi-stellar objects (QSOs) or quasi-stellar radio sources ('quasars'). The spectra of these objects in the optical turned out to have emission lines that were redshifted to such an extent that they had to be extragalactic objects instead of Milky Way stars (e.g. Schmidt, 1963). Meanwhile, the radio emission of some quasars was shown to be highly variable (e.g. Pauliny-Toth & Kellermann, 1966), which put an upper limit on the sizes of the emitting sources that was way smaller than the extent of a galaxy and, given the enormous distances at which they were now known to be located, all of these quasars had to be extremely luminous. We now know that quasar emission originates from the region around a galaxy's central supermassive black hole — such a region is known as an active galactic nucleus or AGN (e.g. Netzer, 2013). Different types of AGN can be detected at different wavelengths, and we will refer to AGN from which we observe strong radio emission as radio-loud AGN or RLAGN.

Single-dish radio telescopes still exist, the largest being the Five-hundred-meter Aperture Spherical Telescope (FAST; Nan et al., 2011), but most modern radio telescopes are arrays of radio dishes. Instead of arrays of steerable dishes, a radio telescope can also be built from an array of radio antennas with no moving parts. Arrays of static radio antennas can change their observing direction in a process known as digital beamforming, which involves introducing digital delays between the signals of antenna pairs before correlating the signals. Digital beamforming allows for near-instantaneous changes to the observing direction of a radio telescope and saves materials and maintenance costs that come with the chassis and tracking motors of steerable radio dishes.

A key future pillar of radio astronomy is the Square Kilometre Array (SKA; Braun et al., 2015)[2], which is an intergovernmental radio telescope with both static antennas and steerable dish antennas that will be built in South-Africa and Australia, and is projected to reach unprecedented sensitivity.

[2] Another notable future radio telescope, not related to the SKA, is the Deep Synoptic Array (DSA-2000; Hallinan et al., 2022) which would be located in Nevada.

To experiment with the not-yet-existing hardware and software of the SKA, three pathfinder telescopes were envisioned. The Meerkat telescope in South-Africa (Jonas & MeerKAT Team, 2016), the Australian Square Kilometre Array Pathfinder (ASKAP Johnston et al., 2008) and the Low Frequency Array (LOFAR; van Haarlem et al., 2013) in the Netherlands.

For the SKA and all of its pathfinders, large-scale sky surveys are a key objective. In a sky survey, a telescope observes a large (contiguous) area of the sky. Sky surveys aim to discover unknown objects and phenomena, which might lurk in places where we usually do not point our telescopes at. Sky surveys also aim to gather large populations of different astronomical objects — for example, the most distant objects that we can detect and for which we do not know their location a priori due to their low apparent brightness. The LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al., 2017) alone will discover millions of radio sources.

The start of the research for this thesis, in Fall 2018, coincided with the internal availability of the calibrated data from the LoTSS data release one (LoTSS DR1; Shimwell et al., 2019). Classifying all of the radio sources in LoTSS according to their morphology, optical host galaxy, their position in the Cosmic Web and their distance/age would allow us to learn more about the evolution of galaxies. In this thesis we focus on the classification of all the radio sources in LoTSS based on their morphology. This provides a multitude of technical challenges, as manual visual inspection of the images containing millions of sources is infeasible. We need an automatic approach to inspecting images and there is a scientific field, called computer vision, which offers frameworks to break down this kind of problem. Computer vision is a field focusing on pattern recognition for digital images or video, with the goal of recovering the (three-dimensional) shape, position and condition of objects (Szeliski, 2022). Many of the techniques employed in various stages of a computer vision pattern recognition system use a form of machine learning (ML). ML is a subset of artificial intelligence that models relationships in one dataset to make automatic predictions or decisions based on patterns in another dataset (Murphy, 2012).

In this chapter, we provide more details about LOFAR (Sect. 1.1), the emission processes that create the radiation detected by LOFAR (Sect. 1.2), RLAGN (Sect. 1.3), the morphology of RLAGN (Sect. 1.3), and open questions concerning RLAGN (Sect. 1.3.2). Furthermore, we introduce the field of computer vision in Sect. 1.4, give an overview of ML in Sect. 1.5, introduce artificial neural networks which are a class of ML techniques in Sect. 1.5.1. We also introduce the specific ML techniques used in this thesis, namely Fast R-CNNs in Sect. 1.5.2, Self-Organising Maps in Sect. 1.5.3, and random forests in Sect. 1.5.4. Finally, we specify the exact questions we address in our thesis (Sect. 1.6) and suggest fruitful avenues for future research (Sect. 1.7).

1.1 Radio interferometry with LOFAR

The Low Frequency Array (LOFAR; van Haarlem et al., 2013) is a radio telescope, officially opened in 2010, that is capable of operating in the 10-90 MHz range using its dual linear droop dipole antennas (low-band antennas hereafter) and in the 120-240 MHz range using its bow-tie shaped dual dipole antennas (high-band antennas hereafter). The antennas are grouped in stations, 38 stations in the Netherlands and 14 international stations with locations ranging from Ireland to Latvia and from Sweden to France.[3] Each of the Dutch stations comprise 96 low-band antennas and 48 high-

[3]The number of international stations is still in flux, an up-to-date map can be retrieved at https://www.astron.nl/lofartools/lofarmap.html

band antennas, and each of the international stations comprise 96 low-band and 96 high-band antennas. The large field of view of the station beam of the Dutch LOFAR-stations (van Haarlem et al., 2013) makes them suitable for quickly observing large patches of the sky.

The LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al., 2017, 2019, 2022) uses the Dutch high-band antennas with the aim to observe the entire Northern sky with a resolution of 6 arcsec for at least 8 hours per observed pointing which has resulted in a mean sensitivity of 83 μJy beam^{-1}, which is twice as sensitive as the previous best radio survey with similar resolution (FIRST; Becker et al., 1995). The latest data release (LoTSS-DR2; Shimwell et al., 2022) covers 27% of the northern sky and detected more than 4 million radio sources. An estimated 8% of these 4 million radio sources are resolved and it is these sources that we want to classify based on their morphology.

Additionally, the LoTSS Deep fields (Tasse et al., 2021; Sabater et al., 2021; Kondapally et al., 2021) focus on three sky regions of around 20 square degrees: Boötes, Elais-N1, and Lockman Hole, which were observed for more than 100 hours to get to mean sensitivities of 25, 19, and 36 μJy beam^{-1} respectively.

Ground-based optical telescopes often do not reach their diffraction-limited angular resolution due to visible light scattering in the atmosphere — air bubbles or air streams with different temperatures or humidity levels have different diffraction indices and limit the resolution to the seeing (typically 1 arcsec). At radio frequencies, the observed extra-terrestrial electromagnetic radiation is not so much affected by air, but it is affected by the ionosphere. The ionosphere is a layer of the atmosphere which appears at heights between 50 to 1000 km above sea-level and contains a lot of free electrons and ionized nuclei. At lower frequencies, ionospheric scattering of the radio signal becomes stronger and increasingly harder to correct for. Even so, with the LOfar Low-band antenna Sky Survey (LoLSS; de Gasperin et al., 2021, 2023), the LOFAR consortium is conducting a survey at effectively 54 MHz, with a resolution of 15 arcsec, achieving a median sensitivity of 1.6 mJy beam^{-1}.

With the ongoing development of its hardware and software (e.g. Edler et al., 2021; Krüger et al., 2022), LOFAR will be able to simultaneously observe with the high-band and low-band antennas, allowing the higher frequency observations to be used to constrain the calibration of the lower frequency observations. LOFAR observations will enable cutting-edge science even when the SKA (with antennas operating from $50 - 350$ MHz) comes online, if the LOFAR consortium keeps pushing the boundaries of imaging at frequencies close to the atmospheric cutoff ($\sim$ 10 MHz), imaging at sub-arcsecond resolution using the high-band antennas of the international stations (e.g. Morabito et al., 2022; Sweijen et al., 2022), or at imaging at sub-arcminute resolution at frequencies close to the atmospheric cut-off (Groeneveld et al., 2022).

1.2 Radiation mechanisms

Blackbody radiation is the most familiar form of radiation. It is a form of thermal emission, whereby the intensity of the emission of an object peaks around a characteristic frequency that is solely determined by the temperature of the object. The higher the temperature of the black body object, the shorter the wavelength (the higher the frequency) of the emitted light. The emission spectrum corresponding to black body radiation is continuous and its shape is described by Planck's law. Pure black body objects do not exist in nature, but stars are pretty good approximations and as such reddish stars are colder than blueish stars. The peak in a star's emission lies in the visible wavelengths and consequently stars are extremely faint in the radio regime.

Indeed, the radio signal detected by Jansky was not thermal radiation from stars. The radiation that he detected had a different origin. Charged particles emit electromagnetic radiation when accelerated under the influence of a magnetic field. This emission mechanism is generally known as gyromagnetic radiation. For ultra relativistic charged particles under the influence of a magnetic field, this emission process is referred to as synchrotron emission (Condon & Ransom, 2016). Synchrotron emission is the dominant emission mechanism of the radio waves detected by Jansky, and similarly so for the radio waves detected by LoTSS.

Extra-galactic synchrotron emission mostly originates from electrons that were accelerated by the magnetic fields that are generated by AGN (e.g. Condon & Ransom, 2016) or by supernova remnants in star-forming galaxies (e.g. Condon, 1992). Radio synchrotron sources persist for extended periods of time following the acceleration of their electrons to relativistic energies. As a result, synchrotron sources serve as archival records, preserving evidence of past energetic events over significant ($> 10^8$ years) durations (e.g. Morganti, 2017).

1.3 Radio-loud active galactic nuclei

Some galaxies are more luminous than would be expected from stellar activity alone. Furthermore, the associated line emission for certain high luminosity galaxies and or the variability in their emission is different than could be expected from stellar activity (e.g. Netzer, 2013). High resolution observations show that this unexpectedly luminous emission originates from the central regions of (usually massive) galaxies. These luminous regions are referred to as active galactic nuclei or AGN. Most galaxies have a (spinning) supermassive black hole at their centre, with a gravitationally bound toroidal accretion disk of matter around it. To account for the luminosity of AGN, the supermassive black hole is thought to convert potential energy (of the matter in the accretion disk) to radiation (and thus luminosity) via heat or kinetic energy (Rees, 1984).

The accretion process of matter onto a supermassive black hole needs to be efficient enough to account for the luminosities that we observe. The accretion efficiency of an AGN can be expressed as

$$\epsilon = \frac{L}{m_{acc}c^2}, \tag{1.1}$$

where L is the AGN luminosity, m_{acc} the mass accretion rate onto the black hole, and c the speed of light (for an introduction see e.g. Schneider, 2006). An optically thick, thin accretion disk allows accreting matter to be converted to radiation with an efficiency ϵ between $\sim$ 6% and $\sim$ 29% for a non-rotating black hole and a rotating black hole respectively. (Non-rotating and rotating black holes are known as 'Schwarzschild' and 'Kerr' black holes respectively.) AGN with optically thick, geometrically thin accretion disks where potential energy is mostly radiated away by the accretion disk are referred to as radiative-mode (or quasar-mode) AGN.

For lower accretion rates of mass onto the black hole, the (lower density) accretion disk of an AGN may not be optically thick (e.g. Netzer, 2013). In those cases, the potential energy is thought to result in a higher entropy of the gas in the accretion disk instead of converting to heat and then radiating away (Narayan & Yi, 1994). Such entropy-increasing accretion is referred to as advection-dominated accretion and the resulting accretion efficiency ϵ for its radiative component is low. The radio luminosity originating from these low accretion AGN is thought to be the result of significant particle outflows along AGN 'jets'. Therefore, these AGN with low radiation efficiencies, where potential energy is mostly converted to outflows (kinetic energy) are referred to as jet-mode (or radio-mode) AGN.

Orthogonal to the spin-direction of a supermassive black hole in the centre of a galaxy, strong magnetic field lines are generated, along which charged particles can be accelerated. These co-linear jets of charged particles, simply called 'jets' here-after, can pierce into the circum-galactic medium and extent to scales from pc[4] to Mpc (see Hardcastle & Croston, 2020, for a review on jetted AGN). RLAGN emission can thus extend far beyond the optical radius (star-related emission) of a galaxy. Once the particles accelerated along the jets lose enough energy, they spread out along a bow shock and/or dissipate adiabatically into plumes that we both refer to as the 'lobes' of a radio galaxy. The location where the jets deposit their energy in the lobes can sometimes be observed as bright spots (referred to as 'hot spots') within the radio lobes.

Most AGN show some radio emission, but a large variety exists in the radio to optical luminosity. This radio loudness factor R is quantified as

$$R = \frac{L_\nu(5\mathrm{GHz})}{L_\nu(4400\text{Å})}, \tag{1.2}$$

with L_ν the AGN luminosity at frequency ν. $R > 10$ is usually set as the arbitrary value above which an AGN is considered to be radio-loud.

RLAGN are among the highest energy phenomena in the universe. To produce 100 kpc radio jets alone, the total energy needed is on the order of 10^{54} J or more, which is equivalent to the instantaneous conversion of millions of solar masses of matter to energy (Hardcastle & Croston, 2020). As such, it is understandable that RLAGN have an impact on their host galaxy and the circum-galactic medium, although the detailed consequences are not yet clear. It is clear that jetted AGN magnetize their surroundings. It is also likely that there are multiple ways in which AGN influence both the star formation rate of their host galaxy and the availability of AGN accretion material — both causal chains are referred to as AGN feedback.

Stars are formed when pockets of cold and dense molecular gas (in the interstellar medium) of a galaxy collapse under their own gravity. The radiation provided by an AGN (possibly) triggers star formation in the AGN's host galaxy (e.g. Mannering et al., 2011; Ishibashi & Fabian, 2012), which would be an example of positive AGN feedback. In this scenario, the AGN radiation is absorbed by (dust embedded in the) molecular gas and causes a pressure-wave to propagate from the core of the galaxy to the outer shell of a galaxy. This pressure-wave allows pockets of molecular gas to pass the critical density threshold beyond which their self-gravity makes the gas pockets collapse into stars.

Simultaneously, negative AGN feedback can occur if the kinetic energy provided by the AGN jet activity — or the radiation from a radiative-mode AGN — heats the gas in the AGN host galaxy. The resulting thermal expansion causes the density of the molecular gas to decrease and thus decreases the star formation rate (e.g. Couto & Storchi-Bergmann, 2023). Negative AGN feedback can also occur when the kinetic energy of the AGN jettisons molecular gas out of the AGN host galaxy, thereby depleting the gas available for star formation and the gas available for AGN accretion and thus slowing down star formation, decreasing the growth of the supermassive black hole and hence lowering the activity of the AGN. AGN jets are usually oriented perpendicular to their host galaxy's galactic disk, but supermassive black hole merging events might change the jet inclination

[4] One parsec (pc) equals 3.262 light years (ly).

angle (Krause et al., 2018). Inclined AGN jets will pierce through a denser part of their host galaxy and are potentially able to remove $\sim 20\%$ of the galactic disc gas within 20 Myr, thereby decreasing the galaxy's star formation rate (Cielo et al., 2018). It is possible that the ejected gas precipitates back onto the galaxy from which it was jettisoned in its cooled state ('cooling flows') allowing star-formation rates to increase again (Oosterloo et al., 2023). As such, negative AGN feedback can either quench or simply delay star formation.

Another complicating factor in determining the net role of RLAGN feedback is the variation in AGN activity throughout the lifetime of the host galaxy (Morganti, 2017). Periods of quiescence following an active period in the galaxy nucleus can be detected in the radio either through the spectral index of the RLAGN emission or through the RLAGN's radio morphology (e.g. Brienza et al., 2017). Radio sources for which accretion onto the supermassive black hole and thus AGN activity has stopped, are known as AGN remnants. Consequently, for AGN remnants, the fresh supply of charged particles accelerated through the jets into the lobes has stopped. Charged particles with higher energies will radiate away their energy more quickly than charged particles with lower energies. As such, the bright parts of an RLAGN (the jets and the lobe hot spots) will quickly become fainter and the RLAGN turns into an AGN remnant. However, the low energetic particles will keep on radiating for years after the AGN has shut off. Furthermore, the lobes of an AGN remnant will not be pushing outward any more under the pressure of the jets. Instead, the lobes of an AGN remnant will adiabatically expand into the surrounding medium. These AGN remnant lobes can be recognized in radio observations as roundish, amorphous lobes with an even surface brightness.

A restart of accretion activity in the nucleus of a galaxy can also be observed in the radio. The clearest cut cases are double-double radio galaxies (DDRG; Schoenmakers et al., 2000; Mahatma et al., 2019), for which an inner (young) set of lobes are co-linearly straddled by an outer (older) more diffuse set of lobes with a lower surface-brightness, indicating two phases of activity. Some core-bright radio sources might, upon (higher-resolution) inspection, also indicate (very) recent restarted activity (Jurlin et al., 2020). Extended periods of quiescence and activity from a galaxy's central region are referred to as the AGN duty cycle. A completeness-corrected estimate of the percentage of RLAGN that are active, remnant or restarted can allow us to put bounds on the AGN duty cycle (e.g. Godfrey et al., 2017; Jurlin, 2022).

1.3.1 Morphology

The morphology of jetted AGN typically shows a radio core at the location of the AGN and two colinear jets ending in two more or less symmetric radio lobes. The jet that is oriented slightly more towards the observer is usually much brighter due to Doppler boosting or beaming (Garrington et al., 1988). Based on the brightness profile of the radio lobes, RLAGN are typically divided into class I or II Fanaroff-Riley objects (FRIs or FRIIs; Fanaroff & Riley, 1974), whereby FRIIs are edge-brightened and FRIs are centre-bright.[5] FRIIs are mainly thought to be produced by RLAGN with radiatively efficient accretion producing high-powered jets and higher luminosities, while FRIs are mainly thought to be produced through RLAGN with radiatively inefficient accretion producing low-powered jets and lower luminosities. Exceptions are abundant, high luminosity FRIs and low luminosity FRIIs exist (e.g. Mingo et al., 2019), AGN can likely switch between radiative and jet-mode accretion throughout their lifetime, and RLAGN can have variable jet-power and luminosity

[5]RLAGN that have one FRI-like and one FRII-like lobe are sometimes referred to as FR-hybrid (Mingo et al., 2019).

throughout their lifetimes. The (density of the) circum-galactic medium also plays a role, where a dense medium is thought to decelerate the relativistic particles in a jet faster, starting the adiabatic expansion at an earlier point in time, typically giving rise to more FRI-like objects (Bicknell, 1995).

An RLAGN's motion with respect to its surrounding medium can affect its lobe morphology (Miley et al., 1972). Especially if the medium is dense, linear motion of the RLAGN through the surrounding medium can lead to lobes that drag behind the RLAGN host galaxy, creating curved or bent morphologies. The 'bending sequence' (see Fig. 6 in Miley, 1980) explains a range of observed RLAGN lobe morphologies as the result of different amounts of bending of the lobes. RLAGN with a slight bending of the lobes are known as wide-angle tail (WAT) sources (e.g. van Breugel, 1980), and RLAGN with strong bending of the lobes are known as narrow-angle tail (NAT) sources (e.g. O'Dea & Owen, 1986). As RLAGN with strongly bent lobes are likely moving fast through their surrounding medium, these RLAGN are conjectured to have host galaxies that are relatively less massive (and are thus less luminous) than the host galaxies of other RLAGN (e.g. Owen & Rudnick, 1976). As bent RLAGN lobes require an RLAGN to move through a sufficiently dense medium, and we know that merging galaxy clusters provide the densest medium in the Universe[6], bent RLAGN lobes can indicate cluster merger activity (e.g. Sakelliou & Merrifield, 2000). In some cases, bulk gas motion from galaxy cluster merging activity — potentially combined with the sudden axial shift in the direction of RLAGN jets induced by a black hole merging event — results in RLAGN lobe-deformation that can be observed as an 'X-shaped' morphology (e.g. Hardcastle et al., 2019).

Whereas the charged particles in the jets and lobes of an RLAGN are accelerated by the magnetic field created by the RLAGN's supermassive black hole, in some cases an external magnetic field can further accelerate these charged particles. Magnetic fields in the shock-fronts of merging galaxy clusters can re-accelerate the charged particles that were originally accelerated along the RLAGN jet and thereby extend the emission coming from an RLAGN to intricate structures that can span several Mpc (e.g. van Weeren et al., 2017).

As mentioned in Sect. 1.3, the (jet) activity of an AGN can affect the jet and lobe morphology. As such, RLAGN morphology can be an indicator for AGN activity. However, care must be taken, as the diffuse amorphous radio emission from starforming galaxies or radio 'halo' emission from merging galaxy clusters[7] (e.g. van Weeren et al., 2019) can for example be mistaken for the diffuse amorphous radio emission from AGN remnants.

Observed RLAGN range in angular sizes from sub-arcsec (Sweijen et al., 2022) to degrees on the sky, and with proper sizes ranging from pcs for optical-galaxy–scale jets (Webster et al., 2021) to radio galaxies spanning several Mpc (e.g. Oei et al., 2022). RLAGN with a proper size above 0.7 Mpc are referred to as giant radio galaxies. Due to selection effects it is not trivial to deduce the average size of RLAGN based on the observed RLAGN size distribution (e.g. Oei et al., 2023a). As such, it is not yet clear whether giant radio galaxies are a special and rare type of RLAGN, or perfectly average old RLAGN that are only rarely observed because lobe surface-brightness drops sharply with increased RLAGN proper size.

Zooming out, we know that on large scales, the Cosmic Web is magnetised, but it is an open

[6]The expansion of the Universe combined with gravity's effect on matter makes that the Universe is mostly filled with low-density voids. Under the influence of gravity, matter aggregates in planes or 'sheets'. Over time, these 'sheets' tend to collapse into even higher density 'filaments' and these filaments can in turn collapse into even higher density 'clusters'. This high-level (sponge-like) structure of matter in the Universe is known as the Cosmic Web.

[7]Merging galaxy clusters create magnetic fields that can accelerate charged particles in the intra-cluster medium, thereby creating a roundish 'halo' of gyromagnetic emission.

question how this came to be (e.g. Subramanian, 2016a; Vazza et al., 2021). Primordial cosmic seed fields could have been generated during the inflationary stage of the early Universe (e.g. Widrow et al., 2012), originate from the Biermann-battery process, or originate from aperiodic turbulent fluctuations in the intergalactic plasma (e.g. Kulsrud et al., 1997). Several processes can have amplified these seed fields to the field strengths that we observe today, among which are star formation through supernova-induced turbulence (e.g. Kronberg et al., 1999) and RLAGN outflow (e.g. Furlanetto & Loeb, 2001). Quantification of the intrinsic length distribution and number density of (giant) radio galaxies leads to an estimate of the volume filling fraction of radio galaxy lobes and can thus shed light on what fraction of the cosmic magnetisation is due to RLAGN outflow (Oei et al., 2023a), Mostert et al. (in prep.).

1.3.2 Open questions

There are many open questions related to the evolution of galaxies that can likely be answered at least partially by studying RLAGN. To what extent are RLAGN responsible for cosmic rays? What triggers an AGN to switch between its 'active' and 'dormant' phases? How often do AGN switch back and forth, how long do they typically spend in their different life phases and how does this impact the evolution of the host galaxy and its star formation rate? How do FRI/FRII type RLAGN evolve over time, and how does the relation between RLAGN morphology, host galaxy properties and the circum-galactic medium change over time? What role do RLAGN play in the magnetisation of the Cosmic Web?

Open scientific questions often require us to solve various technical problems first. For example, to derive the proper size of RLAGN from their angular size we need to estimate their distance. In the optical and infrared regime, the distance of an extra-galactic object is determined via photometric or spectroscopic redshift estimates. As we cannot determine redshifts from continuum observations around 150 MHz, sources in radio surveys are usually cross-matched to sources in infrared or optical catalogues to obtain distance estimates.[8] Cross-matching unresolved sources for surveys with similar resolution is a solved problem (e.g. Sutherland & Saunders, 1992), but cross-matching methods become less reliable when the relevant radio and optical surveys diverge in resolution. Matching a well-resolved, morphologically complex RLAGN that spans arcminutes on the plane of the sky, to an optical survey that has a source density that is orders of magnitudes higher than the radio survey can be hard, even through visual inspection by a human expert. Even so, for most resolved RLAGN in LoTSS, the jet and lobe morphology can provide compelling evidence for a corresponding host galaxy observed in the optical or infrared.[9] Combining the right radio morphology information (see Chapter 2 and Chapter 3) with the available optical/infrared information (e.g. colour and magnitude) to come to an automated cross-matching is a non-trivial problem that has only recently been tackled through a heuristic (Barkus et al., 2022), and lends itself to be improved in the future through computer vision (Sect. 1.4) and machine learning (Sect. 1.5).

Technical problems also arise in the morphological classification of RLAGN. Making statements about the evolution of RLAGN in different morphological classes requires that the radio sources are morphologically classified in the first place. FR classification for example, started out

[8] Given enough sensitivity and minimal attenuation, the spectrographic signal of the recombination of electrons and ions can sometimes be observed. Detecting these radio recombination lines is an active field of research and especially challenging for extra-galactic objects, see Emig (2021) for an overview.

[9] Or at least limit the set of likely corresponding host galaxies to a hand full.

and is still done manually through visual inspection (e.g. Fanaroff & Riley, 1974). However, with thousands of sources detected in LoTSS (and many more to be expected), visual inspection becomes too time consuming. Automated morphological classification of radio sources is thus an active field of research, with current research turning to computer vision approaches to solve the problem either with rule-based heuristics (Mingo et al., 2019, e.g.) or machine learning methods (e.g. Aniyan & Thorat, 2017; Mostert et al., 2021; Scaife & Porter, 2021; Bowles et al., 2021; Hossain et al., 2023).

1.4 Computer vision

Computer vision is a field focusing on the development and application of techniques and algorithms for automated visual pattern recognition, with the goal of recovering the (three-dimensional) shape, position or condition of objects in digital images or video (e.g. Szeliski, 2022). Computer vision techniques are relevant for a wide range of applications in a broad range of scientific disciplines, from diagnosis in medicine (Esteva et al., 2021), taxonomy in biology (Fernandes et al., 2020), to galaxy classification in astronomy (e.g. Dieleman et al., 2015).

To give an idea of the complexity of vision for computers, it is helpful to know that only in 2015 a computer algorithm was able to classify (not localise or describe, just classify) everyday objects in images with an error rate that surpassed that of a *regular* human (He et al., 2016), while the Deep Blue computer already beat the *world's best* human chess player at playing chess in 1997 (Campbell et al., 2002).[10] Indeed, to achieve a robust computer vision system, there are many challenges to overcome, of which we list six below:

- Most objects are anisotropic: observing them at different angles will lead to different apparent morphologies.
- Objects of the same class (or category) may have varying intrinsic sizes and different apparent sizes based on the observer's viewpoint. Scale variation inhibits size estimates without the aid of some reference object or frame.
- Object variation due to interaction with the environment can lead to deformation (changes in appearance), challenging classification.
- Within a class, the intrinsic appearance of objects may vary widely, to the extent that they are not distinguishable from other related classes by external appearance alone. Intra-class variation may also refer to objects (living, dead or inanimate) manifesting themselves in various forms, as a function of their life cycle.
- Objects can be (partly) occluded by other objects in a scene.[11]

[10] Deep Blue was not able to 'see' its opponent, the chess board or the pieces. A human used a keyboard to notify the algorithm of the moves of its opponent and moved the pieces on the board based on the textual output of the algorithm.

[11] Astronomers have a history of using occlusions to their advantage. Dyson et al. (1920) observed a solar eclipse to prove the prediction from Einstein's general relativity theory that a gravitational field could impact the path of a light wave. Occlusions are used to perform transit photometry, whereby the existence of exo-planets can be inferred from the periodic dimming of a star (e.g. Charbonneau et al., 2000; Deeg & Alonso, 2018). Transit spectroscopy even allows for the characterisation of the exoplanet's atmosphere (Brown, 2001), water vapour and several other molecules have since been detected (e.g. Tinetti et al., 2007; Snellen et al., 2010). Partial lunar occultations were used to achieve super-resolution, both in the optical (e.g. to measure the separation of binary stars; Evans et al., 1985) and in the radio regime (e.g. Elsmore & Whitfield, 1955; Link, 1956).

- The detected light emitted or reflected by objects of interest may be so faint that the object is barely distinguishable from the background or surrounding objects.

With these challenges, it is unsurprising that automated visual pattern recognition cannot be boiled down to a single algorithm and that solutions do not naturally generalise from one dataset to another. However, a high-level computer vision approach (or solution design) offers a useful process to follow. A general automated visual pattern recognition solution generally comprises: data acquisition through sensors, followed by data preprocessing, the assembly of a dataset with a homogeneous structure, a dimensionality reduction step to reduce the dataset to its most salient properties or features, model selection, model tuning and prediction, and finally assessment and application of these predictions. This process is a cycle, which requires multiple rounds of iteration to converge to the right sensors, preprocessing, dataset creation, dimensionality reduction, appropriate model, model tuning, and application.

Machine learning currently provides powerful solutions for most steps in an automated pattern recognition system, from data processing, dimensionality reduction and feature extraction, to model training, prediction and selection (Voulodimos et al., 2018).

1.5 Machine learning

Machine learning (ML) is a broad set of algorithms to make automatic predictions or decisions based on patterns detected in data (Murphy, 2012).

Traditionally, automated pattern matching involves comparing data to hand-crafted rules (or 'templates') that describe the sought-after pattern. For example, to detect a specific type of car in a set of photographs, we might come up with rules of thumb that describe the likely situation: we expect to see two round objects with a metallic colour surrounded by an annular ring of black colour (the wheels), furthermore, between the wheels, we expect to find the car-body, and assuming it is mostly of a single colour, we might trace the outline of this colour to find the outline of the car. Subsequently, we could check if the proportions of the outline of the car and the tyres match with the specific type of car that we are after. It is easy to see that to make this work for all edge-cases requires deep expert knowledge and still would result in a system that is quite brittle. It is hard to create hand-crafted rules that account for different viewing angles or one of the other computer vision challenges listed in Sect. 1.4.

Alternatively, in ML, (complex) patterns are not manually pre-scribed. Instead, in ML, patterns are modelled by using a machine learning algorithm (or self-learning algorithm) that processes numerous data samples that contain the sought-after pattern. Going back to our car example: by processing hundreds of photographs that contain the specific car that we are looking for, certain ML algorithms are capable of creating a model that captures the characteristic features of this car, generalising over the different lighting conditions, viewing angles and backgrounds that the algorithm encountered in the example photographs. In ML terms, the example photographs are referred to as the 'training data', the ML algorithm processing the photographs to improve its internal model is referred to as the 'learning process' or the 'training process' (or simply 'training').

The learning process in ML can take on different forms and can be expressed through many algorithms, which are suited for different types of data, and different tasks. In the current ML paradigm, the type of learning process is categorised on a spectrum between supervised and unsupervised (Murphy, 2012). These learning types are not based on the modality of the data[12], nor are they based on the complexity of the algorithm, instead they are based on the availability of explicit external annotations (or 'labels') for data that indicate the desired outcome of a data input given a prediction task. For example, for the predictive task "classify all RLAGN in the northern sky as either FRI or FRII", a useful dataset with explicit external annotations (or simply a 'labelled dataset') would be a dataset (of a small part of the northern sky) with images of RLAGN that are already classified (preferably through visual inspection by an expert) as either FRI or FRII.

In general, supervised learning (see e.g. Murphy, 2012; Singh et al., 2016, for a review) covers all ML techniques where 1) for each input x in dataset X, we want to learn the mapping $g(x) = y$, where g is an unknown function and y is an unknown output, on the condition that 2) we also have a different dataset $\tilde{X}$ at our disposal for which its inputs $\tilde{x}$ are drawn from the same distribution as those in dataset X and for which $g(\tilde{x}) = \tilde{y}$, where we *do* explicitly know $\tilde{y}$. The function g can either map to a continuous output, in which case we talk about a 'regression' problem, or to a discrete output, in which case we talk about a 'classification' problem. In ML jargon, x and $\tilde{x}$ are 'input', y and $\tilde{y}$ are 'output' or 'labels', X is an 'unlabelled' dataset and $\tilde{X}$ is a 'labelled dataset', the function g (or more accurately a function that approximates g) is referred to as the 'model', and to learn g a model is 'trained' using 'labelled' data. If X and $\tilde{X}$ do not draw from the same but still from a comparable distribution, the (supervised) learning approach is referred to as 'transfer learning' (see e.g. Zhuang et al., 2020, for a review). One could for example start training an FRI/II classifier for LOFAR by tweaking a model that already works on FIRST, or take a model trained on LOFAR data as the starting point for a classifier aimed at SKA data.

Unsupervised learning covers all ML techniques where we want to learn about structures in dataset X for which we have no explicit labels at our disposal or when we do not want to use the labels that we have — we might for example not trust our labels, or we might want to look for more general emergent structures or patterns in the data without applying pre-conceived theory ('human bias') straight away. In unsupervised learning, the learning goal can be achieved by setting up an algorithm in which a form of $\tilde{y}$ is derived from aspects of the data itself (without human intervention).[13] Unsupervised learning is especially useful for the data-exploration stage of a project. Clustering and outlier detection are tasks that are well-suited for unsupervised learning. Some pattern matching problems can benefit from multiple sequentially applied algorithms. The output of an unsupervised learning algorithm might well be used as (one of the) inputs of a supervised learning algorithm.

1.5.1 Artificial neural networks

The animal brain is the most generally capable example of a biological self-learning system. As such, it is not very surprising that artificial neural networks — the most powerful set of algorithms within ML — are loosely inspired on the biological neural networks in the animal brain. Artificial neural networks (referred to as 'neural networks' hereafter) are the main driving force behind breakthroughs within the field of ML and computer vision over the last decade (e.g. Kingma & Welling, 2013; Ronneberger et al., 2015; He et al., 2016; Silver et al., 2016; Vaswani et al., 2017; Redmon & Farhadi, 2018; Goodfellow et al., 2020; Dosovitskiy et al., 2020; Jumper et al., 2021).

[12] Examples of different data modalities are: text, image, audio, video, and graphs.

[13] ML algorithms are not always supervised or unsupervised, various in-between forms of learning exist, like self-supervised learning (e.g. Shwartz-Ziv & LeCun, 2023) and semi-supervised learning (Chapelle et al., 2009; Zhou & Zhou, 2021).

1

Neural networks exist of neurons, which are interconnected processing units, usually grouped in layers through which computation usually moves unidirectionally. Neural networks are built to predict a vector as output, given a vector as input. Depending on the layout of the network, the input and output vector can have arbitrary but fixed dimensions, as the connections between neurons are usually static.[14] Artificial neurons and the computation executed in them can take on various forms. The first proposed artificial neuron is the Perceptron neuron (McCulloch & Pitts, 1943; Rosenblatt, 1958). A Perceptron neuron stores an adjustable scalar value, called a 'weight', for each of its incoming connections. Perceptron neurons function by multiplying the scalar values of an incoming signal with the corresponding weight for the connection, they do so for all of the incoming connections, sum these products, feed the sum plus an extra weight called the 'bias' through a non-linear function referred to as the 'activation function'[15] and send the (same) result to all of their outgoing connections. In the case of the Perceptron neuron, the heavy-side function was used as its activation function. A single layer of Perceptrons with different weights can model different logic functions, and multiple layers of Perceptrons can be combined to model more complex logic functions. Currently, many more complex neuron designs and neural network configurations exist, but the basic idea of unidirectional signal flow through neurons arranged in layers, a set of adjustable (or 'learnable') weights per neuron and a non-linear activation function persists. The advantage of multi-layered neural networks (even when composed of basic Perceptron neurons) is that, with the right weights, they can in theory approximate any continuous function on a compact domain (Hornik et al., 1989; Csáji et al., 2001).

Although a neural network can in principle approximate very complex functions, adjusting the weights of the model to achieve this is non-trivial. The first requirement in optimal weight adjustment is a 'loss function' or 'cost function' to evaluate how well the current state of a neural network approximates the sought-after function. The second requirement for optimal weight adjustment is a method to indicate how to tweak (or 'update') the weights such that the loss function goes down. Backpropagation (Rumelhart et al., 1988) is an algorithm that computes the gradient of a loss function with respect to the weights of the network for a specific input and corresponding desired output (label), allowing one to update the weights of the network to lower the loss (and thus improve the model). Popular implementations of the backpropagation algorithm are stochastic gradient descent (as described in Goodfellow et al., 2016) and Adam (Kingma & Ba, 2014; Loshchilov & Hutter, 2017).

Convolutional Neural Networks (CNNs; LeCun et al., 1989) are neural networks composed of convolutional layers, which are especially well suited to process high-dimensional data like images.[16] A convolutional layer has kernels (or 'filters') instead of neurons, where a filter is a two-dimensional matrix of adjustable weights. A convolutional operation is performed by sliding each filter over the entire input volume and calculating the dot product between the filter weights and the input at each position of the sliding window. This process generates a two-dimensional activation or feature map for each filter as output, allowing the network to identify spatially located features of interest in the input. In a regular fully connected (or 'feedforward') neural network layer, each neuron in one layer is connected to and has a weight for each neuron in the previous layer. For high-dimensional data where many dimensions are correlated (such as neighbouring pixels in an image), having and

[14] Pre- and/or post-processing are usually required to navigate the fixed in and output dimensions that a network requires.

[15] *tanh*(), the logistic function, and (variants of) the ReLu function (Nair & Hinton, 2010) are commonly used activation functions.

[16] Each pixel in an image is its own dimension, or more than one if an image as multiple channels (or wavelength 'bands'), for example red, green, blue in the optical regime.

updating weights for each of these pixels is computationally inefficient and prone to overfitting. By comparison, the number of weights in each filter of a convolutional layer can be way lower than the number of pixels of the input image. By increasing the stride of the convolutional sliding window, the output of a convolutional layer can be sparser, requiring a smaller number of weights in the next layer of the network. CNNs are great tools for feature extraction (also called 'featurising' or creating 'feature maps') from high dimensional, correlated data like images, while fully connected layers of a neural network are great at boiling these features down to a much smaller number of output values for classification or regression.[17]

Convolutional operations are naturally equivariant to a geometric translation of their input, which is beneficial in object classification — we like to be able to classify an RLAGN as FRI or FRII, regardless of small offsets to its location within an image. Augmentation of input data (e.g. Dieleman et al., 2015; Mostert et al., 2022) or group-equivariant convolutions (e.g. Scaife & Porter, 2021) can make CNNs (more) equivariant to rotations — essential for astronomical image applications — and approaches like feature pyramid networks can make CNNs (more) equivariant to scale (e.g. Lin et al., 2017). To further reduce the number of parameters of a CNN, sets of neighbouring weights within a layer of a CNN might be condensed down to the maximum weight or to the average weight in each set. Such down-sampling is known as a 'maximum' or 'average pooling' operation respectively, and apart from reducing the number of parameters (and thus computational cost), it also introduces invariance to slight translations of the input (Boureau et al., 2010). Pooling operations are generally used in conjunction with convolutional operations (e.g. Krizhevsky et al., 2012).

With one or two CNN layers, textures and simple features like lines can be extracted from an image. With a CNN of five consecutive layers, Krizhevsky et al. (2012) showed that higher level features within an image can be featurised, like a human head, or a car-wheel. However, adding even more layers usually means adding more weights, and with more parameters, a model is prone to be unstable during training and prone to overfitting. As such, neural networks with many layers (referred to as 'deep neural networks') have long been technically unattainable. It took a number of regularisation innovations in the 2010s, like skip-connections (He et al., 2016), drop-out (Hinton et al., 2012) and thoughtful initialisation of the model parameters (e.g. Mishkin & Matas, 2015; He et al., 2015), to enable the training of deep neural networks.

1.5.2 Neural networks for object detection

For large datasets with hundreds of labelled examples, CNNs excel at image classification (e.g. Krizhevsky et al., 2012; He et al., 2015). Image classification entails predicting which of a predetermined number of classes best describes an image. A typical CNN usually consists of consecutive layers of convolutional and pooling layers, with a fully connected layer at the end. Object detection, which entails predicting what kind of objects reside in which exact location within an image, is a slightly more complex task than image classification. Moving from the task of image classification to object detection requires an adaptation of the typical CNN that is used for image classification. First, an object detection architecture has to be able to evaluate different (separate) regions of the input image. Second, an object detection architecture has to be able to provide a variable number of output predictions. Third, the dimensions of the to-be-detected objects will need to be variable.

[17] A feedforward network can, for example, be designed to produce one scalar output for each of the expected number of class objects.

1

The region-based CNN (R-CNN Girshick et al., 2014) solved these problems by starting out with a separate algorithm that proposes many (rectangular) regions with varying dimensions within an image that might delineate an object. Girshick et al. (2014) choose the Selective Search algorithm (Uijlings et al., 2013) for region-proposals[18], but in essence, the proposed R-CNN architecture was agnostic of the region-proposal algorithm.

In the R-CNN, the proposed image regions suggested by Selective Search are then warped to the fixed image dimensions dictated by a CNN. This CNN architecture is adapted such that its final fully connected layers predict the class of the image region together with a confidence score that expresses the network's relative certainty in the prediction. On top of that, the fully connected layers are also trained to predict an offset from the proposed region[19] such that the predicted region better encapsulates the to-be-detected objects. Above a user-set confidence threshold, objects can be regarded as detected by the R-CNN.

The R-CNN approach is slow since the CNN features must be calculated from scratch for each of the hundreds of considered regions of interest. The Fast Region-based CNN (Girshick, 2015, Fast R-CNN) improved upon the algorithm design by calculating the feature maps of each image only once. The algorithm starts by 1) getting region proposals for an image using Selective Search, 2) then feeding the entire image to a CNN, 3) then warping the region proposals to the dimensions that they take up in the final CNN feature maps (before the fully connected layers), 4) then rescaling this region of interest of the feature map to a fixed size, 5) then using fully connected layers to get for each of the fixed size feature map regions a class prediction with a confidence score and predicted offsets to the location and dimension of the proposed region.

In Chapter 3 we discuss our application of a Fast R-CNN for the purpose of radio source component association. With current and upcoming large sky-surveys uncovering many resolved radio sources with complex morphologies, the relevancy of neural networks for the purpose of detecting radio objects is growing. The first SKA data challenge (Bonaldi & Braun, 2018), in which detection of radio sources in simulated continuum observations was a key part, saw two of the nine competing teams using neural networks in their approach, and both of them ending in the bottom half of the brackets (Bonaldi et al., 2021). The results of the second SKA data challenge saw six of the twelve competing teams, and notably both the winner and the runner-up, using neural networks in their approach (Hartley et al., 2023).

1.5.3 Self-Organising Maps

Self-organising maps (SOM; Kohonen, 2001) are a form of unsupervised learning. They can be seen as a single-layer neural network, where the neurons are organised on a lattice which is usually two-dimensional for visualisation purposes. The weights of each neuron of the SOM can be directly inspected and, over the course of training, will come to resemble a group (or cluster) of inputs in the training set. SOMs have been used in astronomy to assist in the search for eclipsing binaries in GAIA data (Süveges et al., 2017) and to estimate photometric redshifts (Speagle & Eisenstein, 2017a,b).

[18] The Selective Search algorithm proposes rectangular bounding boxes around all objects within an image using a rule-based approach based on grouping neighbouring pixels for example if they share similar colour attributes. Further hierarchical grouping of these initial proposals leads to proposals with diversity in size.

[19] This offset is framed as a regression problem: the networks need to predict the values for a four-dimensional vector that represents an offset for the position (x and y) and an offset for the width and height of the proposed region.

The implementation by Polsterer et al. (2015) adapted the original SOM algorithm to be equivariant to the rotation of image inputs, enabling its use for dimensionality reduction, clustering and outlier detection for astronomical images which in general have no preferred orientation on the sky-plane and thus greatly benefit from rotational equivariance. The adapted SOM algorithm was used to cluster radio galaxies (Polsterer et al., 2015; Galvin et al., 2019; Ralph et al., 2019; Galvin et al., 2020; Mostert et al., 2021) and is further discussed in Chapter 2 and Chapter 4.

1.5.4 Decision trees for explainable predictions

Although flexible and powerful, (deep) neural networks are not always the best solution for every pattern matching problem. First, deep neural networks are more expensive to train and run in terms of computational and energy costs than many traditional ML algorithms. Second, Grinsztajn et al. (2022) show that tree-based models still out-compete deep neural networks for structured datasets with up to $\sim$ 10k labelled data points. Third, the predictions of deep neural networks are relatively opaque: it is not imminently clear why and which features of an image triggered a convolutional neural network to come up with a certain label. Getting to explainable features in a computationally tractable way is an ongoing research topic (e.g. by propagating Shapley values; Chen et al., 2022). The random forest (Breiman, 2001) algorithm, a common ensemble of decision trees, is a supervised learning technique that scores favourably on all three points above and will be further introduced and applied in Chapter 4 to create an AGN remnant classifier.

1.6 This thesis

In this thesis, we investigate how to morphologically classify the well-resolved jetted RLAGN population in the LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al., 2018) using machine learning.

Many outstanding questions related to jetted RLAGN require large samples of specific RLAGN subpopulations to answer. Therefore, the contributions in this thesis are all related to the crucial question: how do we extract specific RLAGN subpopulations from radio images at scale? Specifically, the main questions addressed in this thesis are:

- What can morphology in total radio intensity maps tell us about the observed radio sources in a sky survey, without complementary wavelength information and with limited visual inspection (Chapter 2)?
- How can we scale radio source-component association (Chapter 3)?
- How can we reduce visual inspection for identifying RLAGN remnant candidates based on radio morphology (Chapter 4)?
- How can we gather the largest sample of observed giant radio galaxies to date? How does this sample constrain the properties of the intrinsic giant radio galaxy population (Chapter 5)?

The chapters in this thesis are arranged by their submission dates, Chapter 4 builds on Chapter 2, while Chapter 5 builds on the work in Chapter 3.

Chapter 2 Unveiling the rarest morphologies of the LoTSS radio source population with self-organised maps.

The observation, calibration, imaging and radio blob detection for LoTSS is fully automated, cranking out thousands of unresolved sources and hundreds of (parts of) resolved sources on a weekly base. Transforming these radio images and rough radio blob catalogues into a polished source catalogue is a time-consuming process. Image artefacts need to be discarded, multiple radio blobs need to be combined into single radio sources for correct source counts and flux estimates, the host galaxy corresponding to each radio source needs to be identified. In this chapter, we explore the possibilities of clustering the morphologies of radio sources using just the radio images and rough radio blob catalogues. Specifically, we explore how a flipping and rotation equivariant self-organising map (Sect. 1.5.3), a form of unsupervised machine learning, is able to model the most common radio morphologies. With a model of common radio morphologies and a metric for how much each radio source diverges from these common morphologies in hand, we can automatically reveal the rarest morphologies in LoTSS.

Chapter 3 Radio source-component association for the LoTSS with region-based convolutional neural networks.
In this chapter, we contribute to the automated creation of a radio source catalogue, by focusing on the challenge of automatic radio source-component association. Although existing rule-based radio source detection algorithms are excellent at detecting connected islands of significant radio emission, the grouping of these emission islands into unique radio sources for well-resolved multi-component radio sources is hard to automate and hence still performed manually, even for large-scale surveys like LoTSS. We aim to automate the radio component association of large ($>$ 15 arcsec) radio components using ML techniques. Specifically, we turned the association problem into an object detection problem and trained an adapted Fast R-CNN (Sect. 1.5.2) to mimic the grouping of source components into unique sources as performed by astronomers for LoTSS DR1.

Chapter 4 Finding RLAGN remnant candidates based on radio morphology with machine learning.
Large samples of remnant radio galaxies enable quantification of the radio galaxy life cycle (Sect. 1.3). The remnants of radio-loud RLAGN can be identified in sky surveys through visual inspection based on their radio morphology (Sect. 1.3.1). However, this process is extremely time-consuming when applied to large-scale sky surveys like LoTSS. In this chapter, we aim to reduce the amount of visual inspection required to find RLAGN remnants based on their morphology, through supervised ML. First, we automated the extraction of a wide range of morphological features for all radio sources from their corresponding Stokes-I images. Second, we trained a random forest classifier (Sect. 1.5.4) to separate an existing sample of RLAGN remnant candidates from yet-to-be inspected radio sources.

Chapter 5 Constraining the giant radio galaxy population with machine learning–accelerated detection and Bayesian inference.
In this chapter, we combine our work on automated radio source-component association (Chapter 3), with recent work on automated host galaxy identification to create a pipeline that can automatically create a radio-optical catalogue for any observed pointing in LoTSS. We use this pipeline to identify giant radio galaxies (GRG) candidates[20], visually inspect these GRG candidates and combine our sample with existing GRG samples to create a sample that contains more than 10^4 GRG.

[20] Where GRG is defined as radio sources with a projected proper length $>$ 0.7 Mpc.

Using a subset of this sample and a forward model, we quantify the intrinsic GRG proper length distribution, the comoving GRG number density, and a current-day GRG lobe volume-filling fraction in clusters and filaments of the Cosmic Web.

1.7 Future directions

Following the line of research explored in this thesis, it is fruitful to further improve upon automated radio component association and optical host galaxy identification for well-resolved radio sources. The two are interdependent: The radio component association should take possible optical hosts in mind and optical host identification should consider the probability of different radio component combinations. The current independent approaches lead to a loss of information and compounding errors. Future work should focus on an algorithm that combines the entire process into a single step. The resulting algorithm can be used to assess the completeness of this method across different angular sizes and apparent surface brightness levels. This can be achieved by repeatedly inserting (fainter) copies of segmented, manually verified, well-resolved radio sources in an observed field and retrieving them using the algorithm. This will accurately assess the retrieval completeness of this algorithm without having to rely on comparisons with non-perfect manual search campaigns, or on assumptions of how radio source morphology impacts their retrieval.

Taking a higher-level view on future research directions, it seems sensible to accelerate scientific endeavours, lower the barrier of entry into astronomical research and lower the barriers in between the frequency-based sub-domains, by investing in the capabilities of virtual observatories (VOs). Already, the added value of general VO tools is clear, with the widely used bibliographic retrieval tools like the SAO/NASA Astrophysics Data System (ADS)[21], interactive sky atlas viewers like Aladin (Bonnarel et al., 2000; Boch & Fernique, 2014; Baumann et al., 2022), catalogue repositories like VizieR (Ochsenbein et al., 2000), and catalogue query systems like SIMBAD (Wenger et al., 2000). Three avenues of VO improvements should be considered:

A first avenue of VO improvements lies in not only publishing datasets but also publishing standardized metrics relevant to the community at large. This practice has accelerated advancements in machine learning, and computer vision in particular, by allowing researchers to easily and directly compare their results (e.g. Everingham et al., 2006; Deng et al., 2009; Lin et al., 2014; Cordts et al., 2016). In astronomy, the SKA data challenges set a good example (Bonaldi & Braun, 2018; Bonaldi et al., 2021; Hartley et al., 2023) and the Radio Camera Initiative (Hallinan et al., 2022) seems to be heading in that direction as well.[22]

The second avenue of VO improvements lies in the accessibility of image-derived features. Imaging is key to the field of astronomy and the current cutout services provided by VOs only scratch the surface of what is possible if specific subsets of the data are more readily accessible. There is more work to do on extracting image-derived features, like the morphology of astronomical objects from images (e.g. Dagli, 2023) as we explore in this thesis. There is more work to do in coupling well-calibrated uncertainty measures to these derived features (e.g. Gal & Ghahramani, 2016; Killestein et al., 2021; Scaife & Porter, 2021; Mohan et al., 2022), and in making these uncertainties explicit and queryable. Finally, this avenue could benefit from a natural language interface

[21] https://ui.adsabs.harvard.edu
[22] https://www.radiocamera.io/image-contest

to image features (e.g. Bowles et al., 2023) or a combined text- and visual-based query interface.

This leads us to the third avenue of VO improvements. Machine learning research in astronomy could help to integrate various (existing) VO offerings while allowing prompting using natural language. Since 2009 many natural science students reap the benefits of natural language processing (NLP) in their education by using WolframAlpha[23] to answer basic encyclopaedic and mathematical questions. In the astronomical domain, Schaaff et al. (2019) experimented with an astronomy-focussed chatbot that served as an interface to SIMBAD and VizieR. However, since then, with the invention of transformer-based neural networks (Vaswani et al., 2017) the NLP field has progressed dramatically.[24] With the introduction of ChatGPT (Radford et al., 2019), the potential of machine learning to assist in information retrieval in a broad domain is now palpable. For the astronomical community specifically, NLP could provide the interface of a general information retrieval system. Foundational large language models (LLMs) could be fine-tuned to the astronomical domain for this purpose, while the truthfulness of these systems can be supported through citation-graphs and other VO LLM-interoperable plug-ins[25], thereby making sure that answers refer to and are always grounded in existing publications, databases and catalogues.

[23] https://www.wolframalpha.com

[24] Recurrent neural networks (e.g. Hochreiter & Schmidhuber, 1997), the previous state-of-the-art technology in NLP, is sequential in nature. Predicting the next word of a sentence during training would always require the previous word (or the previous model 'state'). A key advantageous property of transformers is the non-sequential nature of their execution, allowing training to be highly parallelisable. The capabilities of transformer models are thus easier to expand by (pre-)training on ever-increasing text datasets. Tech giants like Microsoft, Google and Facebook now compete in training the best giant pre-trained transformers (GPTs) also referred to more generally as large language models (LLM). Increasing the dataset size (e.g. Sun et al., 2017) or model size (e.g. Tan & Le, 2019) do so consistently increase the performance of GPTs that it is possible to reliably extrapolate model performance based on these factors (referred to as 'scaling laws', e.g. Alabdulmohsin et al., 2022).

[25] WolframAlpha and the Wolfram Language is already inter-operable with ChatGPT4 via a plugin.