Unsupervised classification of simulated black hole shadows

By

MICAELA MENEGALDO

Department of Physics
UNIVERSITY OF PRETORIA

Submitted in fulfilment of the requirements for the degree of MASTER OF SCIENCE (MSC) IN PHYSICS in the Faculty of Natural and Agricultural Sciences.

February 24, 2022

Under the supervision of:
Prof. Roger DEANE

With co-supervisors
Dr. Jordy Davelaar
Prof. Heino Falcke

UNIVERSITY OF PRETORIA

Abstract

Faculty of Natural and Agricultural Sciences
Department of Physics

Master of Science (MSc) in Physics

Unsupervised classification of simulated black hole shadows

by MICAELA MENEGALDO

Keywords: Black hole shadows; M87; GRMHD models; unsupervised machine learning methods; self-organising maps, PINK

In April 2019, the Event Horizon Telescope (EHT) Collaboration released the first image of a black hole (BH) shadow. Theoretical models that aim to describe the environments of BHs are complex and highly-dimensional numerical simulations are often needed to outline the problem. While previous work has employed the use of machine learning (ML) algorithms to predict BH shadow model parameters from image data, in this thesis, we assess the suitability of a particular class of ML algorithms, namely self-organising maps (SOM), as a tool to classify simulated BH shadow images. We employ the SOM network PINK, which spatially compares visual input using a flip and rotation invariant similarity measure, to generate a set of representative BH shadow prototypes for a library of simulated images. Using this and the clustered input data parameter distributions, we find that the shadow ring size, which is related to BH mass in the model, is the dominant class determining factor of the images. Other model parameters, especially those that influence the orientation of the shadow on the image plane, were less influential on the clustering given PINK 's flip/rotation invariance. Despite this, PINK may be useful in determining persistent image-plane features of BH shadows for other model parameters, given a constant BH mass, to curate a subset of meaningfully different models that can then be used in more advanced analyses reducing the volume of data needing further consideration.

Declaration of Authorship

I, MICAELA MENEGALDO, declare that the thesis, which I hereby submit for the degree of MSc in Physics at the University of Pretoria, is my own work and has not previously been submitted by me for a degree at this or any other tertiary institution.

Signature: 

Date: 21/03/2022

Acknowledgements

This work would not have been possible without the financial support from the Inter-University Institute for Data Intensive Astronomy (IDIA). In addition to this assistance, this work was carried out using the data processing pipelines developed at IDIA and available at https://idia-pipelines.github.io. IDIA is a partnership of the University of Cape Town, the University of Pretoria, the University of the Western Cape and the South African Radio astronomy Observatory.

We acknowledge the use of the ilifu cloud computing facility – www.ilifu.ac.za, a partnership between the University of Cape Town, the University of the Western Cape, the University of Stellenbosch, Sol Plaatje University, the Cape Peninsula University of Technology and the South African Radio Astronomy Observatory (SARAO). The Ilifu facility is supported by contributions from IDIA, the Computational Biology division at UCT and the Data Intensive Research Initiative of South Africa (DIRISA).

This project was partly supported by an Erasmus+ Traineeship Scholarship which enabled a 6-month traineeship at Radboud University in The Netherlands. The European Commission support for the production of this publication does not constitute an endorsement of the contents which reflects the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.

This publication was partially supported as part of the joint Dutch Research Council (NWO) and the (South African) National Research Foundation (NRF) project (NWO-NRF2) which supplemented my stay in The Netherlands. In particular, I would like to acknowledge Prof Mark Bentum of the Netherlands Institute for Radio Astronomy (ASTRON) and Prof Oleg Smirnov, head of the Radio Astronomy Research Group at SARAO, for awarding me this funding.

I would like to thank the entire Event Horizon Telescope Collaboration members at Radboud University for providing me with knowledge and guidance during my time in The Netherlands. Spending time with this group was a privilege that I am tremendously appreciative of. In particular, I would like to thank Dr Jordy Davelaar and Prof Heino Falcke, the co-supervisors of this thesis, for your involvement and invaluable support during my traineeship.

Most importantly, this project would not be possible without the support and encouragement from my supervisor, Prof Roger Deane. Having reignited my

love for astrophysics in my Honours studies, I am immensely grateful for how you welcomed me back into academia after my time away. Your insight and expertise were invaluable not only to this work but to my personal growth.

Finally, I would not have been able to complete this work (especially during the COVID pandemic) without the motivation provided by my family, partner and friends. Thank you for being by my side throughout, even though you did not fully understand what I was doing.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In April 2019, The Event Horizon Telescope Collaboration released the first image of a black hole (BH; The Event Horizon Telescope Collaboration, 2019a, henceforth Paper I). M87*, pronounced 'M87-star', is the radio source associated with the supermassive black hole (SMBH) found at the centre of elliptical galaxy M87, located 16.8 ± 0.8 Mpc from Earth (Blakeslee et al., 2009; Bird et al., 2010; Cantiello et al., 2018). The image was captured using the technique of Very Long Baseline Interferometry (VLBI), synthesising telescopes from across the globe to form the apparent Earth-sized telescope, the Event Horizon Telescope (EHT; The Event Horizon Telescope Collaboration, 2019b, henceforth Paper II). At 1.3 mm, the EHT was able to image M87* on event horizon scales due to the instrument's superior angular resolution. Images of M87* were reconstructed from four separate days of observations completed in April 2017 (The Event Horizon Telescope Collaboration, 2019c, henceforth Paper IV). The images, shown in Figure 1.1, depict a ring-like structure with a dark central deficit of light (The Event Horizon Telescope Collaboration, 2019e, henceforth Paper VI).

BHs environments are some of the most extreme in the Universe. Theoretical models that aim to describe BHs are complex and typically span a large parameter space, especially those generated by modern numerical simulations. Exploring these models is key to improving our understanding of BHs (The Event Horizon Telescope Collaboration, 2019d, henceforth Paper V).

1.1 Black holes

The existence of BHs is a fundamental prediction of general relativity (GR), a gravitational theory developed by Einstein (1915). In GR, gravity is not a force

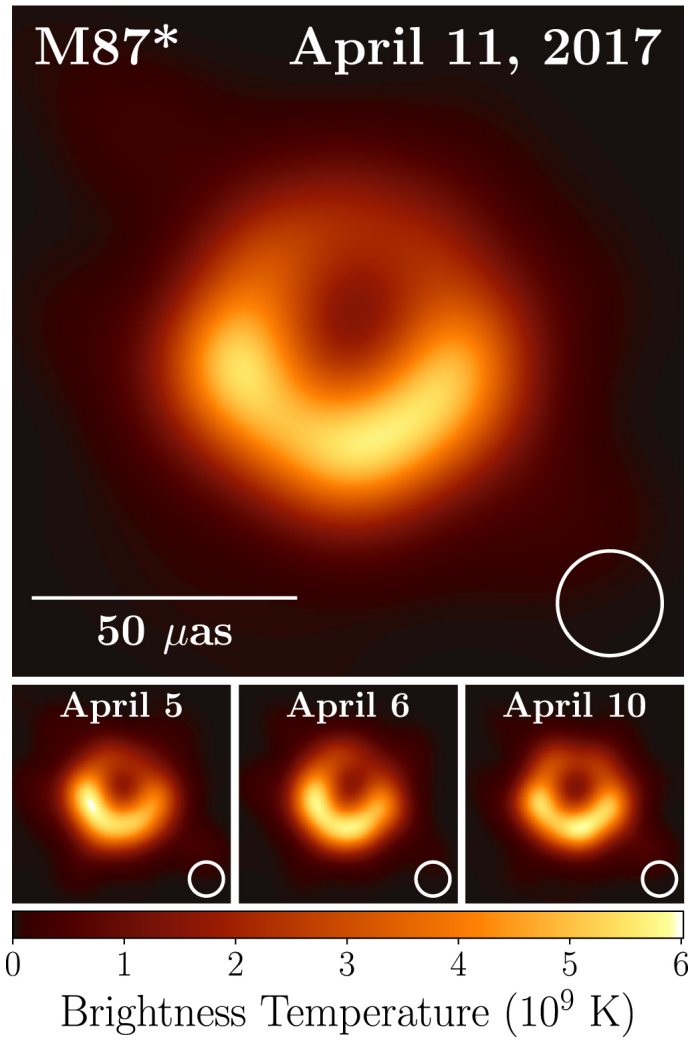

FIGURE 1.1: Reconstructed images from the EHT 2017 observation run. Each image is the average reconstruction from three different imaging methods for the given date. The reconstructions are convolved with Gaussian kernels to achieve the same resolution across the methods, an FWHM of 20 μas, the size of which is shown in the bottom right corner of the top image. Credit: EHT Paper I

but rather describes the interaction between mass and space-time geometry. An object curves space-time and this distortion is seen as gravitational potential. The curvature of space-time is proportional to an object's mass, hence, high-mass objects cause a stronger gravitational field.

Schwarzschild (1916) solved Einstein's GR field equations for a spherically symmetric, non-rotating object. In his solution, an object is characterised by mass,

M, and described by the Schwarzschild Radius

$$R_S = \frac{2GM}{c^2} \tag{1.1}$$

where G is the gravitational constant and c is the speed of light (Schwarzschild, 1916). R_S gives the position of the event horizon, a defining boundary of BHs. Inside the event horizon, the escape velocity becomes $> c$ so nothing, not even light, can escape the BH's gravitational field within this boundary (Schwarzschild, 1916).

The Kerr (1963) Metric describes a rotating BH without charge in GR spacetime. A BH is characterised by mass, M, and angular momentum, J. The event horizon of a Kerr BH is smaller than for a Schwarzschild BH, making energy extraction more efficient as matter can orbit closer to a Kerr BH.

BHs are found to exist in two main mass populations: stellar-mass BHs, which are the result of high mass ($M > 8\,M_\odot$[1]) star core collapses (e.g. Neugebauer, 2003); and SMBHs, which have masses in the range $M \sim 10^6 - 10^{10}\,M_\odot$. SMBHs are thought to exist at the centre of every galaxy (e.g. Kormendy and Richstone, 1995; Miyoshi et al., 1995).

As BHs cannot be directly observed, evidence for their existence has been provided in alternative ways. We describe two here.

Gillessen et al. (2009) tracked the orbits of stars in the neighbourhood of Sgr A*, the SMBH at the centre of the Milky Way, from 1992 to 2008 to determine the gravitational field within which they orbit. From the stellar dynamics, they calculated that a mass of approximately $4.31 \times 10^6\,M_\odot$ exists at the galactic centre.

More recently, gravitational waves, originating from two coalescing solar mass BHs, were detected. Abbott et al. (2016) reported the two BHs to have initial masses of $\sim 36\,M_\odot$ and $\sim 29\,M_\odot$. The resultant BH mass, post-merger, is stated as $\sim 62\,M_\odot$. About $3\,M_\odot$ worth of energy was radiated away as gravitational energy during the merger. The gravitational waves produced were detected by the Laser Interferometer Gravitational-Wave Observatory (LIGO). This result was the first detection of its kind and confirms the existence of gravitational waves from GR.

[1] $M_\odot$ is a solar mass and is the mass of the Sun = 2×10^{30} kg

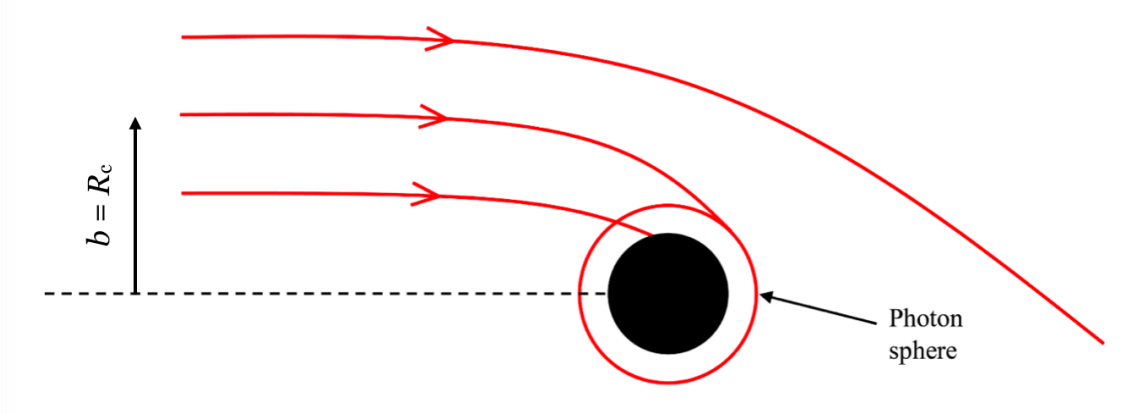

FIGURE 1.2: Incident photons from infinity approach a non-rotating black hole with impact parameter b. This determines the worldline of a photon based on the distance at which it approaches a BH. $b = R_c$ is the photon ring orbit radius. Credit: Bisnovatyi-Kogan and Tsupko (2017)

1.1.1 The Photon Ring Around a Black Hole

Another phenomenon resulting from the curvature of space-time by high-mass objects is gravitational lensing (Einstein, 1915). Light follows the curvature of space-time, and so, in 3D it appears to bend where there is curvature. In regions near BHs, the curvature is so large that photons can bend and orbit around the BH, creating the so-called photon ring (Hilbert, 1917).

Figure 1.2 depicts the path of a photon approaching a non-rotating BH. The BH has a photon ring, indicated as the photon sphere in the figure, at

$$R_c = \sqrt{27}\, r_g, \tag{1.2}$$

where r_g is the characteristic length scale of a BH given by

$$r_g \equiv \frac{GM}{c^2}. \tag{1.3}$$

Equation 1.2 gives the radius of the only circular orbit around a Schwarzschild BH. From Equation 1.1, it follows that $R_S = 2r_g$, indicating the photon ring is at a larger radius than the event horizon (Teo, 2003).

R_c determines the worldline of a photon approaching a BH with impact parameter b. Photons are captured by the BH if $b < R_c$; escape to infinity if $b > R_c$; and, in principle, orbit the BH if $b = R_c$ (Hilbert, 1917).

A Schwarzschild BH has only one photon ring for a given observer, however,

a Kerr BH has an infinite number of photon rings. Photons may be captured on one of these unstable orbits and exit at a later stage, making it possible to observe the light that originates from this region (Johnson et al., 2020; Bronzwaer and Falcke, 2021).

1.1.2 The Local Black Hole Environment

A schematic of the local environment of a BH is shown in Figure 1.3. Here, the BH is viewed edge-on, so the accretion disk is perpendicular to the viewing plane. This figure provides a composite view of the features of a BH.

The accretion flow consists of plasma moving near c. Its emission is lensed from behind the BH, leading it to appear above and below the central depression. The depression is known as the BH shadow (Falcke, Melia, and Agol, 2000; Narayan, Johnson, and Gammie, 2019; Bronzwaer and Falcke, 2021).

Doppler beaming is a viewing effect observed when matter is moving at relativistic velocities as in the accretion disk. As depicted in Figure 1.3, the matter moving away from the viewer is Doppler diminished and appears dimmer, while the matter moving towards the viewer is Doppler boosted and appears brighter. This is argued to be the cause of the apparent asymmetrical brightness of the ring seen in Figure 1.1.

The photon ring is made up of multiple photon orbits which become thinner and fainter moving towards the BH. Photons can orbit the BH on these rings several times before escaping and reaching the observer.

For M87* at 1.3 mm, the diameter of the emission ring was found to be $42 \pm 3\,\mu as$ (Paper I). Previous VLBI studies at 1.3 mm similarly found that, for M87*, the emission region would have a diameter $\sim 40\,\mu as$ (Doeleman et al., 2012; Akiyama et al., 2015). This diameter is consistent with the predicted shadow diameter and the BH mass estimate from stellar dynamics (Macchetto et al., 1997; Gebhardt and Thomas, 2009; Gebhardt et al., 2011; Walsh et al., 2013).

1.1.3 Active Galactic Nuclei

Active Galactic Nuclei (AGN) are active SMBHs that produce high luminosities. The most luminous objects in the universe are a class of AGN known as quasars

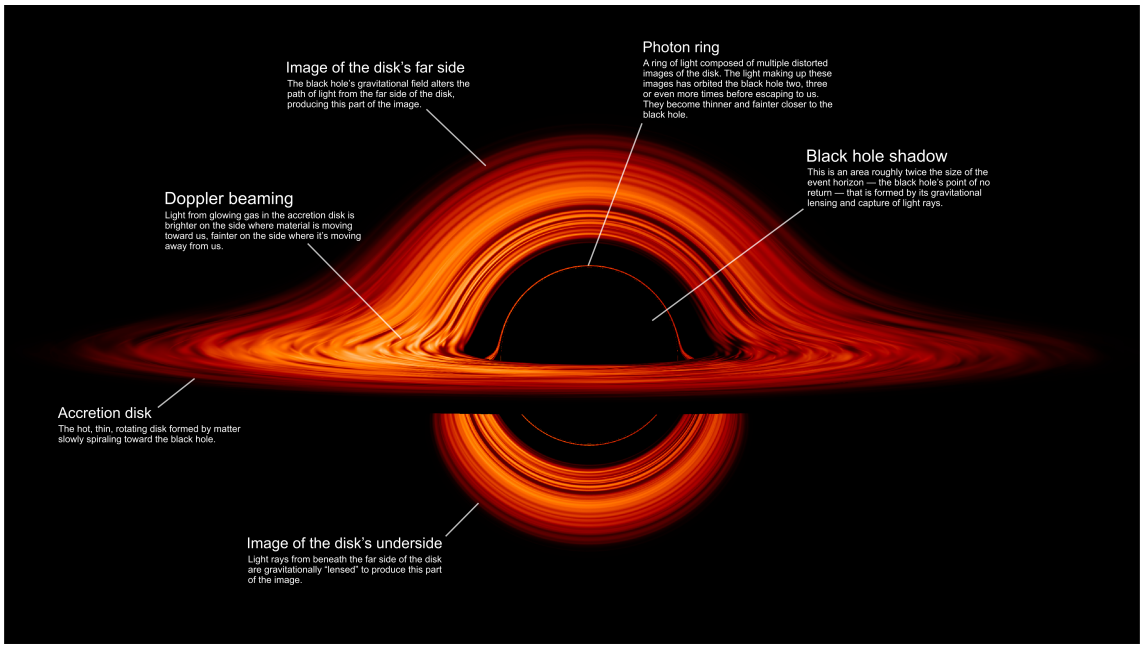

FIGURE 1.3: Schematic of a BH viewed edge-on. Light from the accretion disk is lensed from behind the BH to form a ring shape around the central depression. The photon ring, which consists of multiple unstable orbits from which photons can escape, surrounds the BH shadow. The relativistic plasma in the accretion disk is dimmed (brightened) when moving away (towards) the viewer due to Doppler beaming, causing an asymmetric ring brightness. Credits: NASA's Goddard Space Flight Center/Jeremy Schnittman

(Schmidt, 1963; Sanders et al., 1989). Their luminosity is due to the SMBH accreting material at a high rate from a geometrically thin and optically thick circumnuclear accretion disk (Shakura and Sunyaev, 1973; Sun and Malkan, 1989). AGN found in the local Universe, including Sgr A* and M87*, are low luminosity AGN (LLAGN) and their accretion is associated with a hot, tenuous accretion flow at a much lower rate (Ichimaru, 1977; Blandford and Begelman, 1999).

AGN are often seen to have jets of relativistic plasma emanating from their cores. These can extend from parsec (pc) to megaparsec (Mpc) scales. The jets are thought to be powered either by accretion (Blandford and Payne, 1982) or by the magnetic field extracting rotational energy from the BH (Blandford and Znajek, 1977).

In Figure 1.4, five radio images of M87* at different wavelengths and scales are depicted. From left to right, the figures show: (a) the extended structure of M87 which spans ~ 80 kpc (Owen, Eilek, and Kassim, 2000); (b) the M87* jet and radio lobes from within the red region of the previous image, extending $\sim$ 6 kpc (Livio,

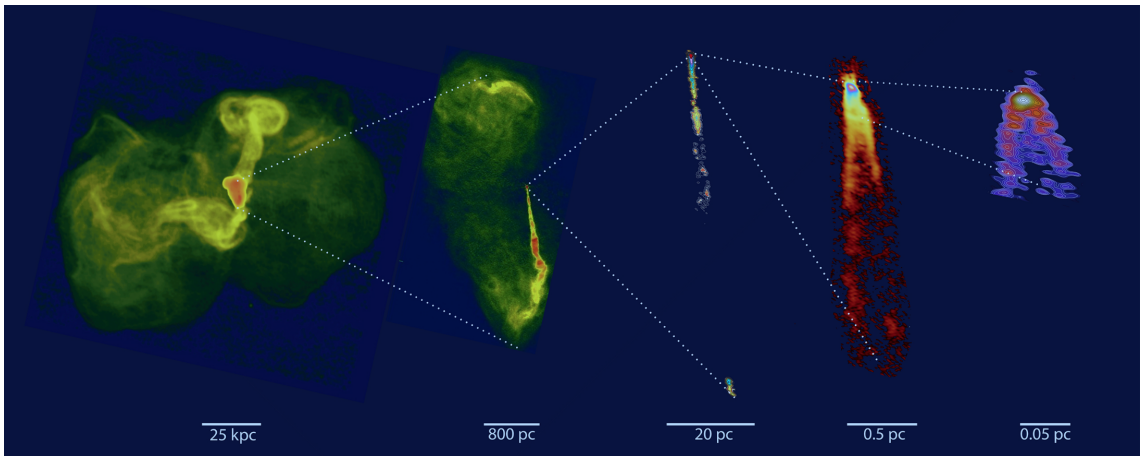

FIGURE 1.4: Radio images of M87 from the outer radio lobes to the inner jet launching region. From left to right: (a) the large radio lobes, formed by jet interactions with the intergalactic medium, spanning $\sim$ 80kpc; (b) a highly collimated jet and radio lobes within the red region of (a) spanning $\sim$ 6kpc; (c) the core and inner jet spanning $\sim$ 35pc; (d) the innermost jet spanning $\sim$ 5.5pc; (e) the jet launching region near the central SMBH spanning $\sim$ 0.08pc. Figure reproduced from Blandford, Meier, and Readhead (2019). Image credits: (a) NRAO, 90 cm VLA (Owen, Eilek, and Kassim, 2000); (b) NRAO, 20 cm VLA (Livio, 1997); (c) NRAO, 20 cm VLBA (Cheung, Harris, and Stawarz, 2007); (d) NRAO, 7 mm VLBA (Walker et al., 2018); (e) 3 mm global VLBI network (Kim et al., 2018).

1997); (c) the core and inner jet of M87*, which span $\sim$ 35 pc (Cheung, Harris, and Stawarz, 2007); (d) the innermost jet of M87*, extending across $\sim$ 5.5 pc, and dimmed counter-jet (Walker et al., 2018); and (e) the jet, with base diameter $\sim 4.0 - 5.5\, R_S$ and length $\sim 1.0\, \mu as \simeq 0.08$ pc $\simeq 140\, R_S$, launching in the region of the central SMBH (Kim et al., 2018). These images illustrate that jet launching occurs on horizon scales near the central SMBH in M87.

1.1.4 Simulations of Emission Around a Black Hole

Comparing theoretical BH models with observations is required to extract information and meaning from real data. To model a Kerr BH, the 3D dynamics of magnetised plasma moving relativistically near a rotating BH need to be determined. This is a complex problem, and analytical solutions only exist for simplified cases, such as for Schwarzschild (non-rotating) BHs (Schwarzschild, 1916). General relativistic magnetohydrodynamic (GRMHD) simulations, the equations of which couple GR curved space-time and energy conservation laws, can numerically model this problem (e.g. Dexter, McKinney, and Agol, 2012; Mościbrodzka,

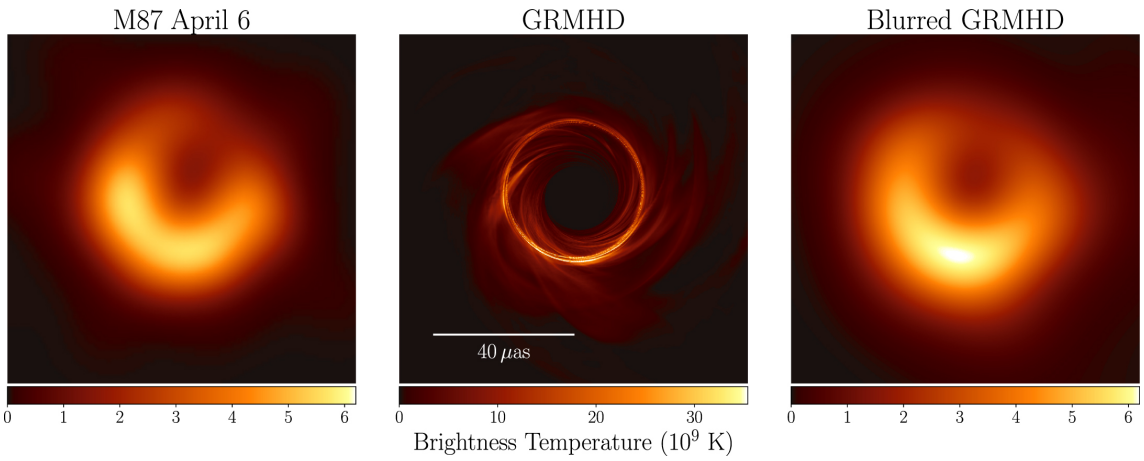

FIGURE 1.5: Left panel: M87* as imaged by the EHT in 2017. Middle panel: a GRMHD model of a BH shadow. Right panel: The GRMHD model blurred with a 20 μas FWHM Gaussian beam. The true M87* image and the blurred GRMHD image are comparable. Credit: EHT Paper I

Falcke, and Shiokawa, 2016; Mościbrodzka et al., 2017).

GRMHD simulations of M87* for different physical models typically produce images of an emission ring with asymmetric brightness around a dark central deficit (Paper V). Simulating a variety of models with varying input parameters gives insight into how the parameters impact the resultant ring morphology.

Figure 1.5 illustrates the importance of comparing theoretical and intrinsic data. In the figure, three depictions of a BH shadow are shown: the first is the recovered source structure from the 2017 EHT observation; the second, a GRMHD model image; and the third is the GRMHD model image blurred to the nominal angular resolution of the EHT. The intrinsic observation and blurred simulated image look similar even though the (unblurred) GRMHD BH ring shows a complex structure.

The GRMHD simulations in Paper V are characterised by two parameters; a_*, a dimensionless spin parameter, and ϕ, the dimensionless magnetic flux across the event horizon (Tchekhovskoy, Narayan, and McKinney, 2011; Porth et al., 2019). Both parameters describe the state of the disk accretion flow.

a_* is defined as

$$a_* \equiv \frac{Jc}{GM^2} \tag{1.4}$$

where J, c, G and M are as previously defined. a_* is restricted to $-1 < a_* < 1$. For a non-rotating BH, $a_* = 0$, and, for a maximally spinning BH, $|a_*| = 1$.

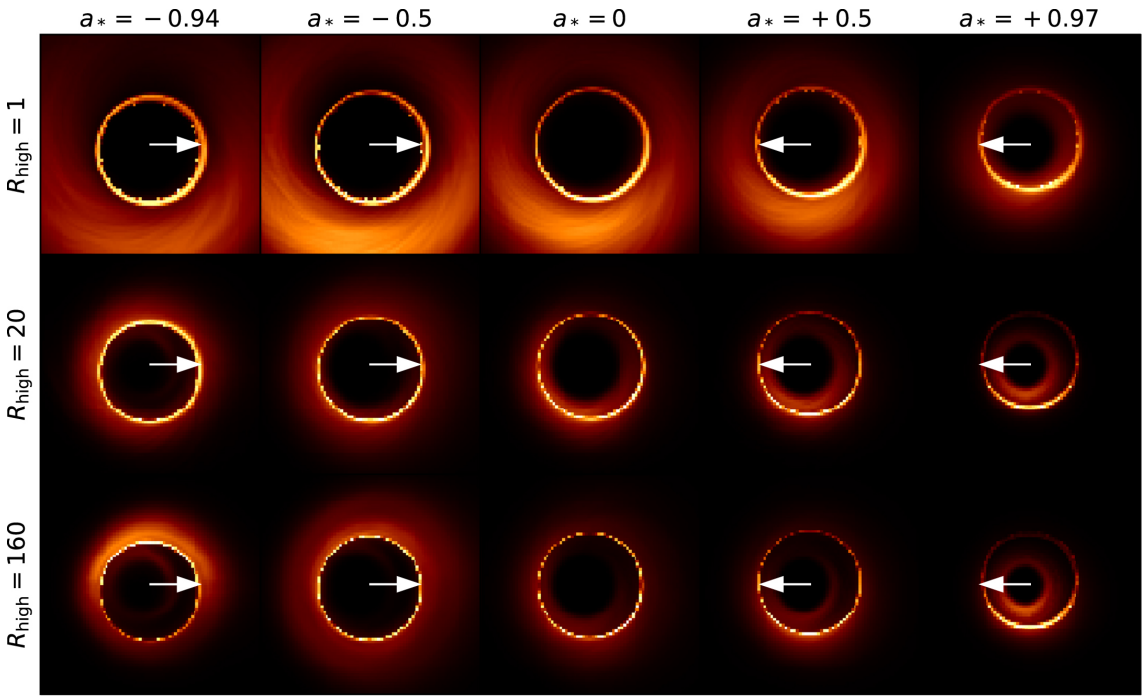

FIGURE 1.6: GRMHD simulations of a BH in a SANE magnetism environment. Each 1.3 mm image is time-averaged; a_* increases from left to right, $R_{\rm high}$ from top to bottom. The inclination is constant at $i = 163°$. The colour scale is linear. The jet approaching Earth is on the right in all images. The arrow indicates the direction of the BH spin vector projected onto the plane of the sky; pointing right indicates the BH is rotating in an anticlockwise direction, pointing left, a clockwise direction. The image resolution is 20 times finer than that of the EHT. Credit: The Event Horizon Telescope Collaboration Paper V

$a_* < 0$ indicates that the angular momentum of the accretion disk and BH are anti-aligned. For $a_* > 0$, the momentum of the accretion disk and BH are aligned.

The dimensionless magnetic flux, ϕ, is defined as

$$\phi = \frac{\Phi}{\sqrt{\dot{M} r_{\rm g}^2}} \tag{1.5}$$

where Φ is the magnetic flux across the event horizon, $\dot{M}$ the mass accretion rate and $r_{\rm g}$ is as defined in Equation 1.3.

The magnetism state of the accretion flow can be SANE (Standard and Normal Evolution; Narayan et al., 2012) with $\phi \sim 1$, or MAD (Magnetically Arrested Disk; Narayan, Igumenshchev, and Abramowicz, 2003) with $\phi \sim 15$. SANE accretion flows are a result of a smaller initial magnetic field model than MAD flows in the

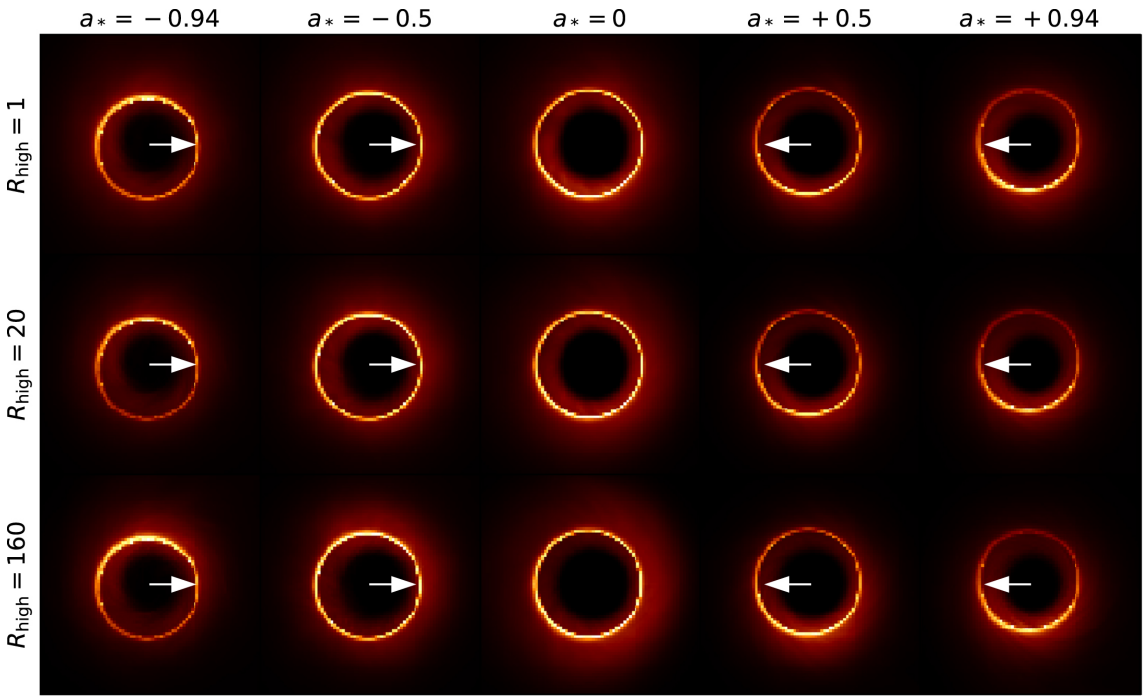

FIGURE 1.7: The same as 1.6 but realised in a MAD magnetism environment. Credit: The Event Horizon Telescope Collaboration Paper V

GRMHD BH simulations.

At 1.3 mm, the bulk of the emission observed is believed to be from synchrotron emission. Synchrotron emission is dependent on the electron distribution of the emitting plasma. Within the GRMHD BH simulations, the accretion material is considered to consist of non-relativistic ions, with temperature T_i, and relativistic electrons, with temperature T_e.

While T_i is given within the GRMHD simulations, T_e is dependent on the magnetisation of the plasma (Mościbrodzka, Falcke, and Shiokawa, 2016). The ratio of these temperatures is given by

$$R \equiv T_i/T_e = R_{\rm high}\frac{\beta_{\rm p}^2}{1+\beta_{\rm p}^2} + \frac{1}{1+\beta_{\rm p}^2}, \tag{1.6}$$

where $\beta_{\rm p} = P_{\rm gas}/P_{\rm mag}$ and $P_{\rm mag} = B^2/2$, where P is power and B is the magnetic field strength of the source (Mościbrodzka, Falcke, and Shiokawa, 2016; Paper I; Paper V). In the highly magnetised regions of the flow, T_i and T_e are assumed to be equal while these can differ in the gas dominated regions.

Equation 1.6 is characterised by one parameter, $R_{\rm high}$, which indicates the

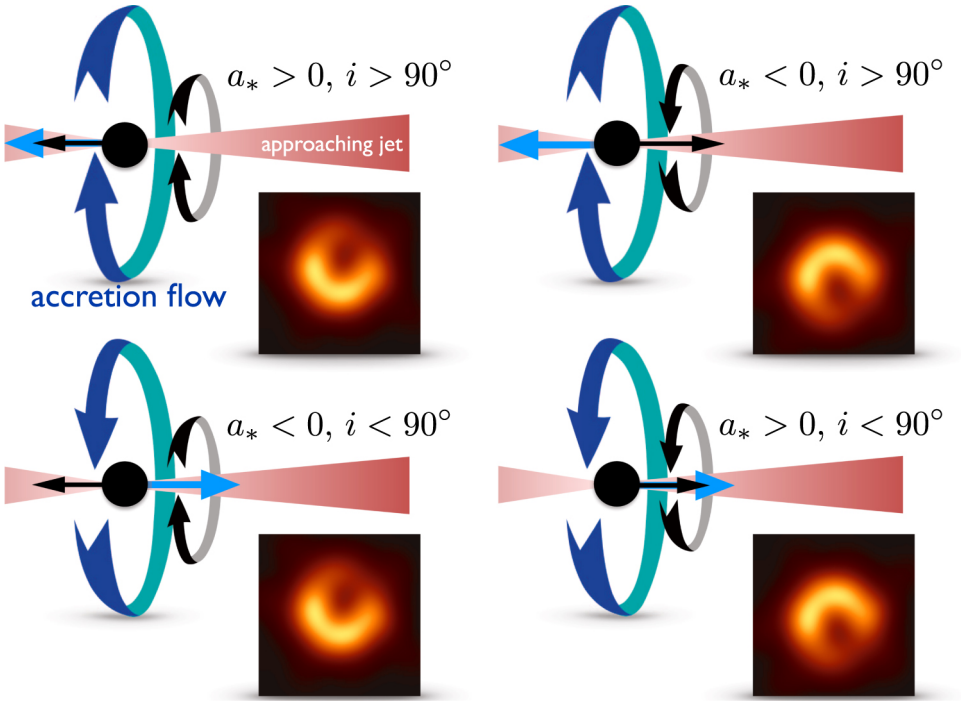

FIGURE 1.8: The orientation of the ring asymmetry on the sky depends on a_*, the BH spin, and i, the angle between the disk angular momentum vector and the line of sight. The bright flux is caused by plasma moving towards the observer so that it is Doppler boosted. In the figures, the blue ribbon represents the direction of the disk rotation and the black ribbon represents the BH spin direction. Across each column the rings look similar; this reflects the degeneracy with respect to the two geometric parameters a_* and i. Credit: The Event Horizon Telescope Collaboration Paper V.

region that produces the bulk of the observed emission (Mościbrodzka, Falcke, and Shiokawa, 2016; Mościbrodzka et al., 2017; Davelaar et al., 2018; Paper V) which is either from the weakly magnetised regions (small $R_{\rm high}$, $T_e \approx T_i/R_{\rm high}$) or strongly magnetised regions (large $R_{\rm high}$, $T_e \approx T_i$).

Other parameters included in the GRMHD simulations are inclination, i, and position angle, PA. i is the angle between the accretion flow angular momentum vector and line of sight; PA is the angle, East from North, of the accretion-flow momentum vector projected onto the sky. For an observer, this is counter-clockwise on the image plane (Paper V).

Figures 1.6 and 1.7 show images of GRMHD BH simulations for SANE and MAD flow, respectively. These figures establish the influence $R_{\rm high}$ and a_* have on the observed ring structure. In these figures, i is constant at 163°.

In SANE models (Figure 1.6), most of the emission comes from the weakly

magnetised region (the disk) when R_{high} is small, and from the strongly magnetised region (the jet) when R_{high} is large. This results in extended emission outside the photon ring for low R_{high} and inner ring emission when R_{high} is high. MAD models are unaffected by changes in R_{high} as all regions are strongly magnetised and the bulk of the emission comes from the disk (Figure 1.7; Paper V).

The bright region observed in the ring is caused by Doppler beaming. For both MAD and SANE accretion flows, the location of this region changes with a_*. In Figures 1.6 and 1.7 the bright section of the ring is at the top of the ring when $a_* > 0$, at the bottom when $a_* < 0$ and the rings look symmetric when $a_* = 0$.[2] This indicates that the peak flux location is influenced by the BH spin.

In addition to a_*, ring asymmetry is influenced by i. Ring asymmetry is insensitive to small changes in i but does depend on whether i is above or below 90°. This is illustrated in Figure 1.8 (Paper V).

1.1.5 Summary

In the previous section, we demonstrated how the GRMHD BH model parameters change a BH shadow's appearance. In the simulations, we can view these environments in great detail and can identify differences between the finer structure of models. In reality, however, achieving clarity to such a degree is not possible.

This challenge is exemplified in Figure 1.9; three BH GRMHD models, unblurred (top row) and blurred to the nominal resolution of the EHT (bottom row), are shown. The structure of emission inside and outside of the ring can be differentiated between the top models, but these differences are not evident in the blurred images. The bottom row of rings are morphologically similar as the finer ring structure is lost.

The loss of detail, which is a result of limited observing capabilities, introduces uncertainty into the deductions that can be made from EHT images. This is a motivating factor for exploring an extensive model dataset, even if some models are unphysical.

In the following section, we give a brief overview of radio astronomy and the techniques used to improve images' resolutions.

[2] This pattern is true for all except for SANE models with $R_{\text{high}} = 1$ where the bright region will always be at the bottom of the ring when $i > 90°$ (Paper V)

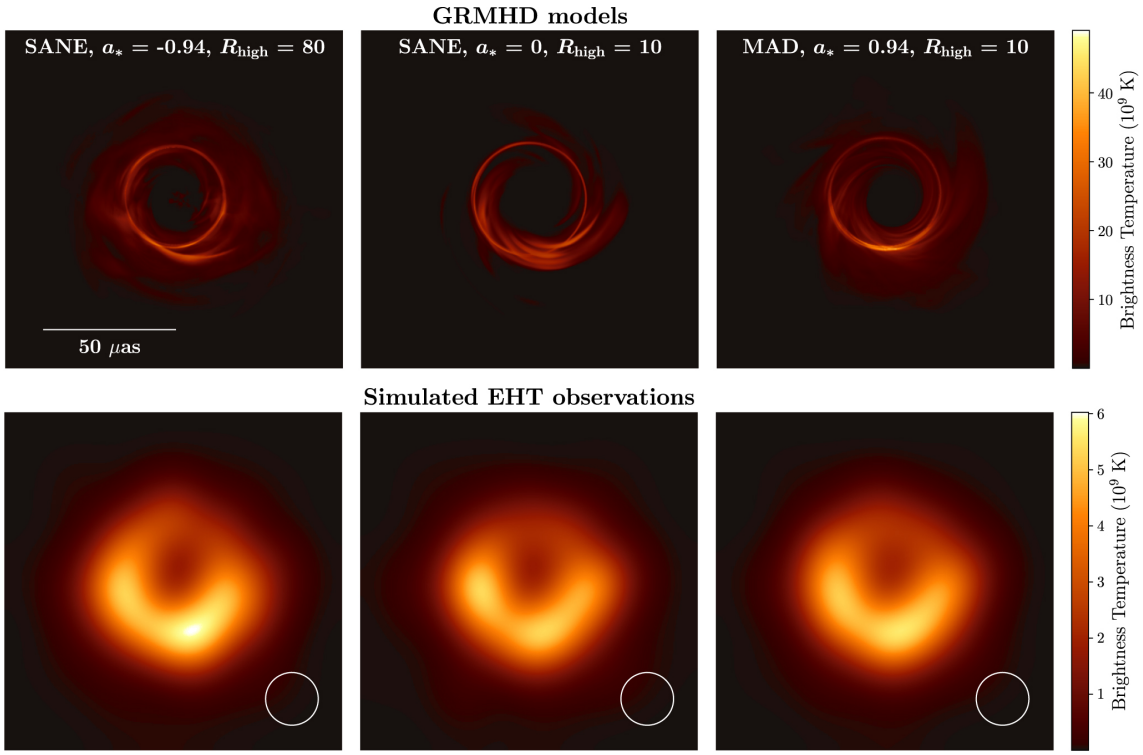

FIGURE 1.9: Top: GRMHD simulation snapshots for three different BH models. Three parameters have been varied; the accretion flow magnetism state, a_* and $R_{\rm high}$. Bottom: the top row images simulated through a VLBI pipeline set to mimic the EHT 2017 observation conditions. Visible morphological differences are present between the model snapshots in the top row. Once these are blurred to the resolution of the EHT, the models look similar. This emphasises the need to explore a variety of models as multiple fit the 2017 result. Credit: The Event Horizon Telescope Collaboration Paper I

1.2 Radio Astronomy

Radio waves span a wide frequency range, from 10 MHz to 1 THz. This corresponds to wavelengths of 0.2 mm to hundreds of meters, making radio waves mostly unaffected by cosmic or atmospheric dust extinction. However, in the mm-regime, radio waves are affected by absorption and scattering that occurs in the Earth's atmosphere. These effects are not explored further in this thesis but are an important consideration for future analyses of this sort.

The specific intensity of a source at frequency v is given by I_v. I_v is independent of distance to the source and has units of $\mathrm{W\,m^{-2}\,Hz^{-1}\,sr^{-1}}$. I_v is the measure of the emission per unit area, per unit time, per unit solid angle at v as seen by the observer. I_v can also be expressed in units of Jansky (Jy) as $\mathrm{Jy\,sr^{-1}}$ where

$1\,\mathrm{Jy} = 10^{-26}\,\mathrm{W\,m^{-2}\,Hz^{-1}}$.

A source's radio strength is measured in terms of flux density, S_v, and is the integral of I_v over all directions. Hence, S_v is given by

$$S_v = \int I_v(\theta, \phi) \cos\theta d\Omega \tag{1.7}$$

which has units of Jy and is proportional to the inverse of the square of the distance to the source. S_v is therefore the power of a source radiated at frequency v per unit area of a detector. A typical radio source will have S_v of the order of 1 Jy.

Integrating S_v over a sphere gives the specific luminosity, L_v, which is the power of the source at frequency v. This value is intrinsic to a source. L_v is defined as

$$L_v = \int S_v(R, \theta, \phi) R^2 d\Omega \tag{1.8}$$

where R is the radius of the sphere.

I_v is defined by the Planck function. When $hv << k_\mathrm{B}T$, we use the Rayleigh-Jeans approximation of the Planck function, given by

$$I_v = B_v \approx \frac{2\,k_\mathrm{B}\,T_B\,v^2}{c^2}, \tag{1.9}$$

where k_B is the Boltzmann constant, T_B is the brightness temperature of the source and h and T are the Planck constant and physical temperature of the source respectively (Rybicki and Lightman, 1979). Rearranging for T_B gives

$$T_B(v) = \frac{I_v\,c^2}{2\,k_\mathrm{B}\,v^2}, \tag{1.10}$$

which is independent of distance.

T_B is not necessarily a measure of a physical temperature. This follows from the approximation made in Equation 1.9 for a thermal emitter. Since radio sources' radiation is often from non-thermal emissions, such as synchrotron radiation, T_B is the temperature required for a source to produce the specific intensity from thermal radiation processes (i.e. if it were a blackbody.)

At 1.3 mm, M87* has a high brightness temperature, in the range $\sim 1 \times 10^{10}$ K (Akiyama et al., 2015; Paper IV). This is the peak brightness temperature of the ring in the left panel of Figure 1.5.

1.2.1 Interferometry

In this section, we give an overview of interferometry. As the EHT is an interferometric telescope, it is important to highlight the methods of this technique. However, in the proceeding chapters, interferometric effects are not considered as they are beyond the scope of this thesis as we aim solely to explore an already highly-dimensional BH dataset with a novel technique. Including the interferometric effects will be of interest for future work.

The maximum angular resolution a telescope can achieve, θ, is defined as

$$\theta \sim \frac{\lambda}{D}, \tag{1.11}$$

where λ and D are the wavelength of the incoming signal and diameter of the receiving antenna, respectively. Improving θ is desirable as a telescope with a superior angular resolution captures the finer structure of a source.

As radio waves are the longest waves on the electromagnetic spectrum, creating dishes that are sufficiently large to achieve high angular resolutions present insurmountable engineering and mechanical problems. To overcome these limitations, telescopes are synthesised to form radio interferometers that have apparent diameters needed to achieve desired angular resolutions.

An interferometer consists of multiple receiver antennas at various distances from each other. The distance between any two antennas in an array is known as a baseline. An interferometer has an apparent dish diameter that is the length of the longest baseline. Hence, θ in Equation 1.11 is no longer limited by the size of an individual dish but rather is dependent on the distance between the two furthest antennas (*Synthesis Imaging in Radio Astronomy II* 1999), to the first order. Some interferometers have baselines spanning entire continents, or even across oceans. These use a technique called Very Long Baseline Interferometry (VLBI), with some having angular resolutions of milli-arcseconds and even μas.

Antennas convert electromagnetic wavefronts into voltages. A pair of spatially separated antennas p and q measure voltages $\mathbf{v}_p$ and $\mathbf{v}_q$. Each voltage, $\mathbf{v}$, is a vector with two components, a and b, which correspond to the polarisation orientation of the antenna.

Voltages are amplified before being sent to a correlator for processing. Pairwise signals are combined and interfered by multiplying voltages together and

integrating over time. Through this process, the correlator produces four correlations that form the visibility matrix, V_{pq},

$$\mathrm{V}_{pq} = 2 \begin{pmatrix} \left\langle v_{pa} v_{qa}^{*} \right\rangle & \left\langle v_{pa} v_{qb}^{*} \right\rangle \\ \left\langle v_{pb} v_{qa}^{*} \right\rangle & \left\langle v_{pb} v_{qb}^{*} \right\rangle \end{pmatrix} = 2 \left\langle \mathbf{v}_p \mathbf{v}_q^{H} \right\rangle \tag{1.12}$$

where angled brackets indicate averaging over some time and frequency bin, x^* is the complex conjugate of x and H is the Hermitian transpose operator.

For an N element interferometer, there are $\frac{N(N-1)}{2}$ independent baselines. Each baseline can be described by vector $\mathbf{b} = (u, v, w)$ where u and v are the projected length of the baseline as seen at the source in the West (u) and South (v) directions and w is in the direction of the source which is perpendicular to the u, v-plane.

The visibility of a baseline, $V_{u,v}$, is a complex representation of the antennas' response to the incoming signal. It is given by the Van Cittert-Zenike Theorem (van Cittert, 1934; Zernike, 1938) as

$$V_{u,v} = \int \int I(l, m) e^{-i2\pi(ul+vm)} \, dl \, dm. \tag{1.13}$$

An interferometer records visibilities, or correlations, which, in the absence of other instrumental effects, represents spacial coherency. If $I(l, m)$ is the sky brightness distribution in the (l, m) or image-plane, by measuring the spatial correlation of the incoming signal in an array, we can sample, and in principle recover, the source brightness distribution by taking the Fourier Transform of Equation 1.13 (Bracewell, 1965).

The above is a simplistic description of correlation. This process actually requires highly complex calculations and is computationally expensive but has been included for completeness.

The Event Horizon Telescope

The image of M87* in Figure 1.1 was captured with the EHT, a VLBI array consisting of mm- and sub-mm-telescopes from across the globe. The 2017 EHT array included 8 stations. The locations of these are shown in Figure 1.10 with baselines indicated.

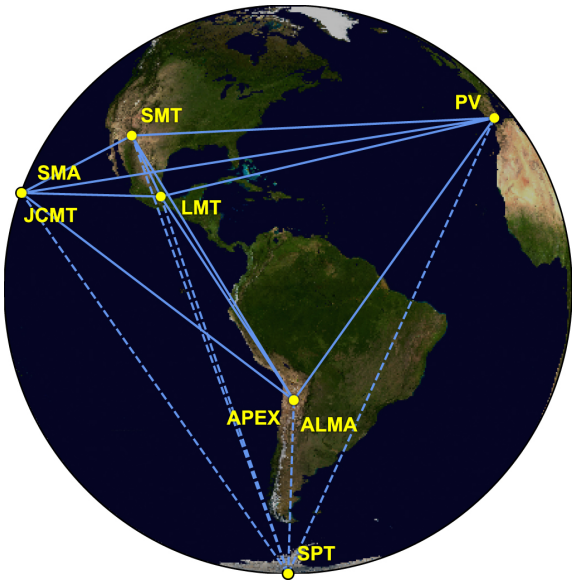

FIGURE 1.10: Geographical locations of the eight telescopes included in the EHT 2017 array. The solid lines are the baselines of pairwise telescopes; the dotted lines are the baselines of pairwise telescopes of the calibration source, 3C279. Credit: The Event Horizon Telescope Collaboration Paper II

The EHT synthesised an Earth-sized telescope that, at the observational wavelength of 1.3 mm, enabled it to achieve an angular resolution of 25 μas. This is the resolution needed to view M87* on event horizon scales (Paper II). However, the limited instrumental resolution resulted in the blurred ring structure seen in Figure 1.1 (Paper II).

1.2.2 Summary

The overview of radio astronomy and interferometry we have conveyed in this section serves to provide the reader with a broad understanding of the field and emphasise the complex nature of recovering the shadow image of M87* in Figure 1.1. These observational and technological complexities, combined with theoretical data and the products thereof, generate an extensive parameter dataspace requiring analysis. Using tools to constrain the scope of this space and identifying data of significance may prove invaluable given the increasing rate at which data is being produced, and hence, needing effective and efficient processing.

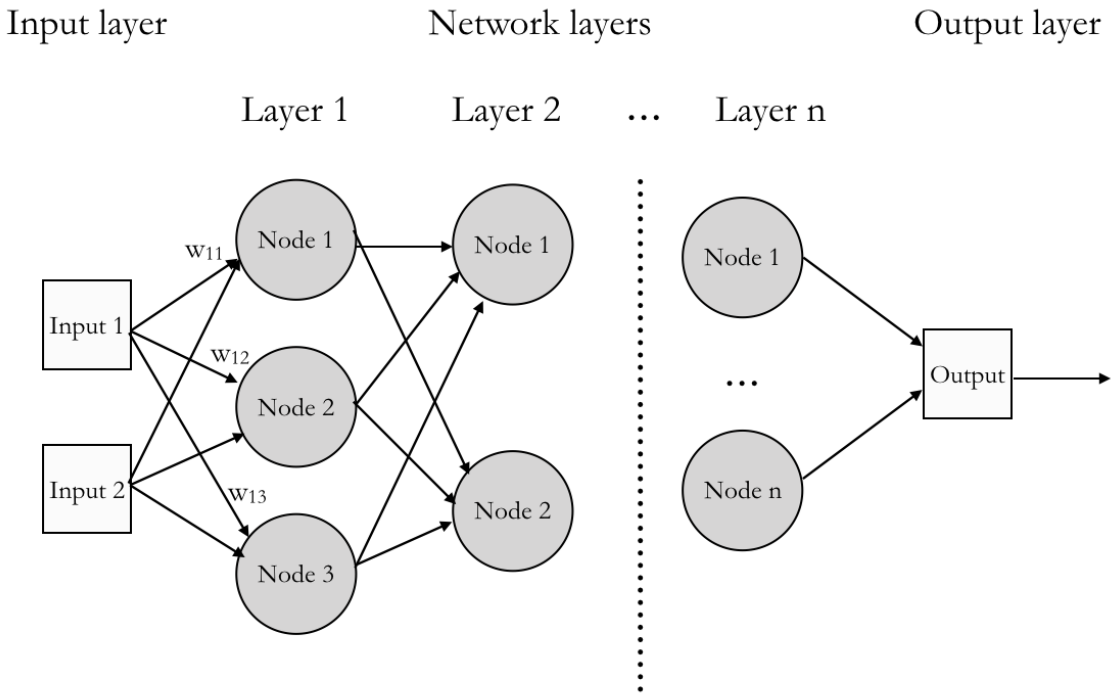

FIGURE 1.11: A schematic of an artificial neural network (ANN). The network receives input(s) which are passed to the nodes in the first network layer. Each connection has a certain weight. w_{11} is the weight of the connection between Input 1 and Node 1. Each node then calculates a result from the inputs based on the weight of each connection. This result is then passed on to the node(s) in the following network layer and the process continues. After moving through all the layers, the network will produce a result consisting of any number of outputs.

1.3 Unsupervised Machine Learning

Machine learning (ML) is a branch of artificial intelligence that utilises computational analysis to build a model based on some observed data. By analysing the given data through a number of iterations, the model improves, learning through practice. The model is both a conjecture regarding the world of the data and can be used to solve problems (Russell and Norvig, 2021).

ML algorithms are employed to analyse and generate models for data because the designers do not necessarily know all the eventualities that the model may encounter and/or how to programme a possible solution to a problem (Russell and Norvig, 2021).

The algorithm used to produce a model learns from a set of training data.

Training can either be once-off, generating a fixed resultant model, or done continuously, which allows the model to change and adapt as it processes new information throughout its lifetime. While the resultant model can change, the algorithm used to build it cannot.

To produce a satisfactory model, the training data must fully represent the data population. If biases exist within the training data, the ML algorithm will build these into the model, which will, in turn, influence the solution(s) the model produces. So, even though the algorithm itself is unbiased, using biased data for training will skew the learning which leads to inaccurate outcomes from the trained model.

Broadly, ML algorithms can be classified as either supervised or unsupervised. Supervised algorithms determine a mapping function between a given set of inputs and known outputs, i.e. they determine the relationship between input and target data. An example would be finding the best fit function parameters for a given set of data points with $(x; y)$ coordinates.

Unsupervised algorithms only receive inputs, for example, x coordinates, without any output information. They use the inputs to describe the structure or features of the data. An example of this is clustering; data is divided into smaller groups based on similarity. The similarity metric is explicitly defined but how the clusters are partitioned out of the population as a whole is discovered through iterative analysis of the data.

An artificial neural network (ANN) is a type of ML algorithm structure. An ANN consists of a network of nodes that are interconnected. Each connection has a certain weight, or impact, on the resultant model built by the algorithm. ANNs often have multiple layers and information moves through them in one forward direction. The basic structure of such a network is shown in Figure 1.11.

In training ANNs, the initial weights of each connection are typically randomised. This will give a bad initial resultant model, but, as the algorithm iterates over the training data, the connection weights will be adjusted to produce a network with higher accuracy.

Essentially, an ANN node is a mathematical function that takes inputs from nodes in the preceding network layer and produces a result based on the weight of each connection. This is then forwarded to nodes in the following network layer. Hence, a neural network produces a result based on the impact of a varying number of connections from given inputs.

1.3.1 Self-Organising Maps

A self-organising map (SOM), also known as a Kohonen map, is a type of unsupervised ANN algorithm primarily employed as a dimension reduction tool (Kohonen, 1998). It produces low-dimensional, often 2D, representations of high-dimensional input spaces. In some cases, it can be a useful anomaly detection tool in conjunction with other techniques. A SOM typically consists of a single layer of nodes in either a square or hexagonal shape with each node representing a prototype of the data.

A prototype is a representative description of a class of objects from the data. This class will have similar attributes and the prototype portrays these, typifying the class.

SOMs use a competitive training technique to maintain the structure of the input data. This training style uses a similarity measure $\Delta(\mathbf{y}, \mathbf{w}_j)$ to assess the likeness between each input $\mathbf{y}_i$ and each node which is defined by its weight vector $\mathbf{w}_j$. The smaller the similarity measure, the higher the similarity between the input and node. The inverse is also true.

The node most like an input, i.e. $\mathbf{w}_j$ such that $\Delta(\mathbf{y}, \mathbf{w}_j)$ is minimised, is considered the 'winning' node or the best matching unit (BMU) for the given input and training iteration. After a training iteration, the weights of the nodes are updated to optimise (i.e. minimise) the similarity measures between inputs and their corresponding BMUs. These updates are not done in isolation; the weight vectors of other nodes in the map are also adjusted, the severity of which is dependent on its distance to the BMU for a training data point. The formalisation of this can be found in Section 2.2.

Post-training, each node represents a derived prototype of a subset of the data. These nodes exist in the latent space of the data. The latent space is a multi-dimensional space containing 'hidden' variables that are not directly observable in the data but meaningfully encompass the observable characteristics. These spaces can drastically reduce the number of variables to consider when faced with a problem (Russell and Norvig, 2021).

The nodes represent a compression of the data. Each node resides at a certain position, $\mathbf{c}_j$, in the map and will be most similar to the others in its locale. This provides a spatial ordering of the nodes based on similarity.

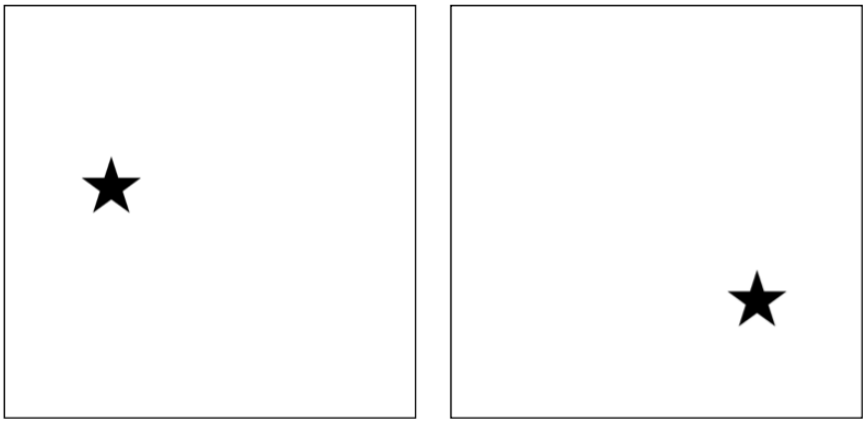

FIGURE 1.12: Pictures of a star placed at different locations in the image plane. A traditional ML algorithm will not be able to assess these images as being of the same object. PINK solves this by using a rotation and flip invariant similarity measure.

Finally, the data is mapped to the SOM. The similarity measures between inputs $\mathbf{y}_j$ and nodes $\mathbf{w}_j$ are calculated. $\mathbf{y}_j$ is mapped to node $\mathbf{w}_j$ for the minimum of $\Delta(\mathbf{y}, \mathbf{w}_j)$. The winning node best represents the data. As a result, similar data will be mapped to the same node, forming classes of similar objects (Teuvo, 2001).

The competitive training algorithm needs a large training set, often as large as the dataset to be analysed, and multiple training iterations to produce a satisfactory model. SOM training is affected by the ordering of the training data, the number of data patterns in each (unknown) prototype class, the similarity measure used and the number of training iterations completed. Finely tuning these parameters so that a satisfactory model is produced is imperative for this class of algorithms.

1.3.2 Parallelized rotation/flipping INvarient Kohonen maps

Polsterer et al. (2016) built the Parallelized rotation/flipping INvarient Kohonen maps algorithm, PINK, as a parameter reduction and classification tool for 2D images. Motivated by crowd-sourced classification projects, and the need for scale and objectivity for upcoming astronomical surveys with expected catalogue sizes in excess of 10^9, Polsterer, Gieseke, and Igel (2015) endeavoured to create an unsupervised ML algorithm with the visual pattern-recognition capabilities needed to classify such surveys.

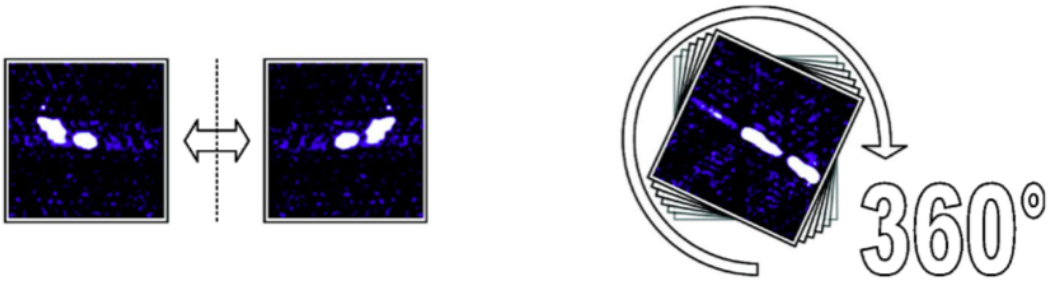

FIGURE 1.13: A flipping and rotation of images are conducted to calculate the most accurate similarity measure between images and nodes in PINK. On the left, `FIRSTJ075843.0+611936` is flipped; on the right, `FIRSTJ072529` is rotated through 360°. Images are from the Radio Galaxy Zoo project from the FIRST survey (Becker, White, and Helfand, 1994). Credit: Polsterer, Gieseke, and Igel (2015).

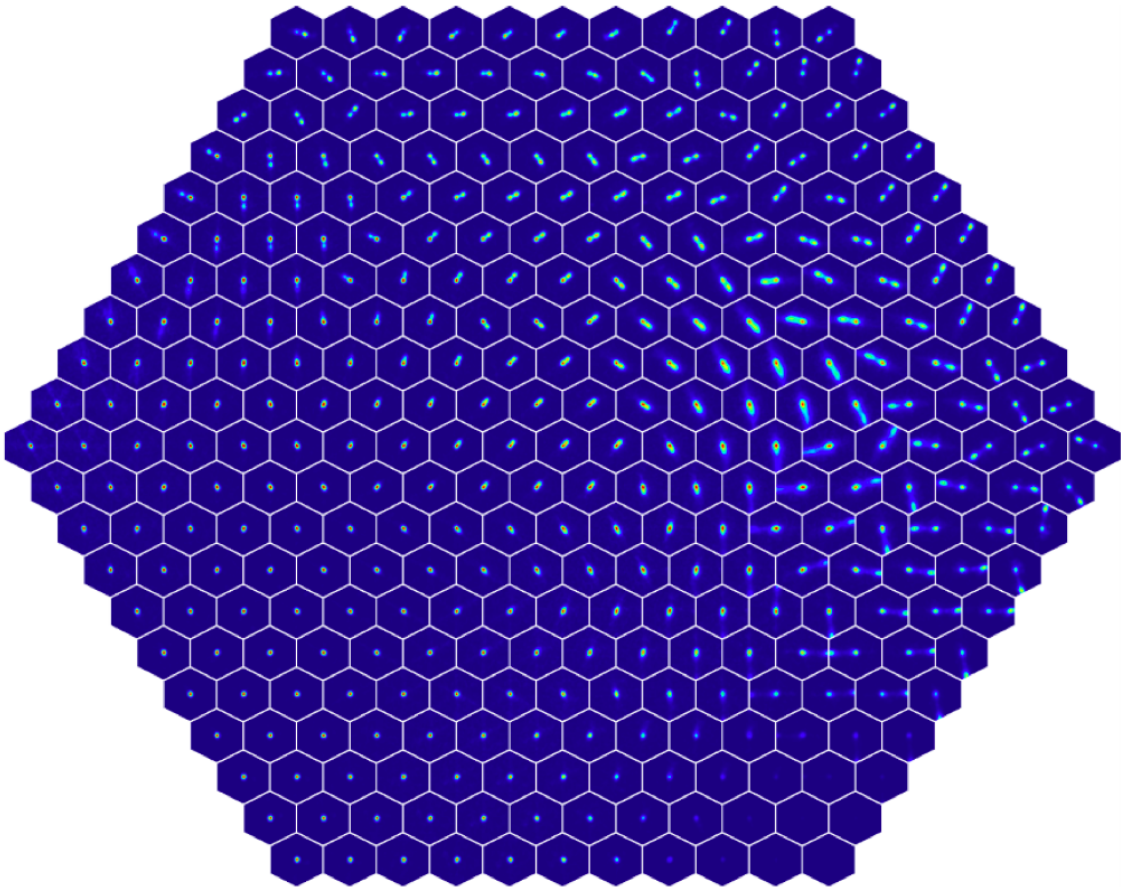

FIGURE 1.14: Hexagonal SOM of 331 nodes produced by PINK trained on 2×10^5 Radio Galaxy Zoo images. The image on each node represents a prototype of a class from the data. There are obvious groups of nodes showing similar characteristics across the map. Credit: Polsterer, Gieseke, and Igel (2015)

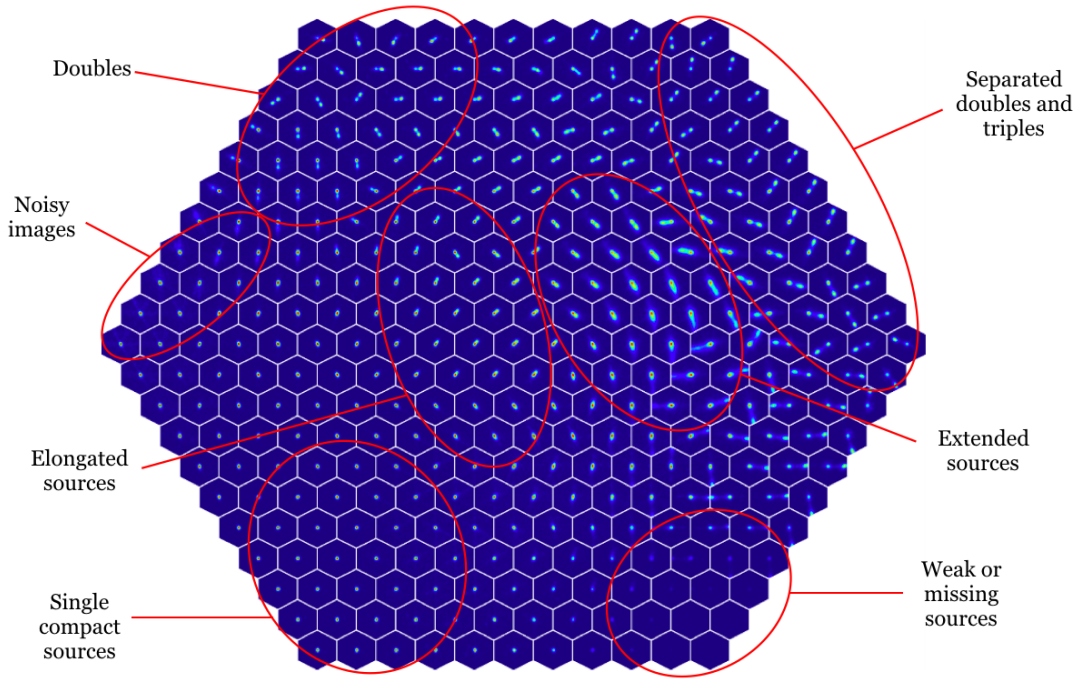

FIGURE 1.15: Hexagonal SOM produced by PINK trained on 2×10^5 Radio Galaxy Zoo images, as in 1.14. Groupings of nodes that display prototypes with similar characteristics are indicated. By classifying the 331 prototypes, the objects in each class are also classified as they are similar to the prototype. Credit: Polsterer, Gieseke, and Igel (2015)

Even though this appears to be a simple classification problem, it is made challenging by ML algorithms' poor spatial comparison capabilities. Humans are adept at this sort of comparison, performing various transformations on images to compare the key visual attributes without needing specific instruction. However, for ML algorithms, these transformations have to be applied to the input images, either before or during processing, to increase the accuracy of the similarity calculations and subsequently built model.

As a comparative example, consider the two images of a star in Figure 1.12. The stars are the same except they are at different locations in the image plane. A human will easily recognise that these are two images of the same object, however, a traditional ML algorithm will see these as completely different objects because they are not at the same location in both images. To reduce the effect image-plane differences may have on training and mapping a PINK SOM, the data is pre-processed ensuring the input images are of the same dimension, the target of interest is centred and the background noise removed.

During training and mapping, PINK calculates the similarity measure for each input for multiple rotations and flips of the image to reduce the impact image geometry may have on this metric. In Figure 1.13, the image flip and rotation transformations are visualised (Polsterer et al., 2016). The lowest similarity measure from the set for the different image realisations determines the winning node for the given input.

As an example to assess the usability and capability of PINK, Polsterer et al. (2016) tested the framework with data consisting of 2×10^5 images from Radio Galaxy Zoo (RGZ).[3] RGZ is a crowd-sourced citizen science project utilising thousands of volunteers to classify radio images of galaxies in the search for BHs. The project started in December 2013 and is ongoing.

The PINK SOM was set to have a hexagonal shape with 21 nodes across each diagonal resulting in a total of 331 nodes. The trained SOM is shown in Figure 1.14. Each node depicts a prototype of a class of galaxies from the data. Groups of prototypes with similar characteristics are found in the same locale; some of these have been identified in Figure 1.15.

In classifying the prototypes, all other images in the node's class are classified. This drastically reduces the number of objects to be inspected for classification from the size of the entire dataset (2×10^5 in the RGZ test case) to a few representative images (331 in Figure 1.14). Images with a large similarity measure to their class prototype are easily identified for further study. These are possibly interesting objects that are not well represented by the prototypes in the map.

The RGZ example exhibits the power of SOMs as a dimensionality reduction tool. It also demonstrates how using ML in certain science applications can save time and focus attention on objects of potentially greater interest.

1.3.3 Possible Applications of Machine Learning for the Event Horizon Telescope

In presenting the EHT results, Paper I reported the BH parameter constraints made from the observations. The emission region diameter was measured as 42 ± 3 μas, which is consistent with previous non-imaging constraints at 1.3 mm (Doeleman et al., 2012; Akiyama et al., 2015). For the adopted distance of $16.8 \pm$

[3] https://radio.galaxyzoo.org/

0.8 Mpc, the BH mass was found to be $M = 6.5 \pm 0.7 \times 10^9\, M_\odot$ which is in agreement with the stellar dynamics measurement in Gebhardt et al. (2011). Details of the image-plane and visibility-plane model fitting and parameter estimation techniques used can be found in Papers IV and VI.

The BH mass result was determined using both the visibility data (Paper VI) and the ring diameter calculated from the image domain (Paper IV). Because of the image quality, no other BH parameters could be reliably derived from the observations. Determining additional characteristics, such as the ring width and asymmetric deficit contrast, are motivating factors for upcoming work.

The potential GRMHD BH simulation parameter space is highly dimensional. ML techniques may provide useful to explore this problem. Below we present the work from two papers, namely van der Gucht et al. (2020) and Yao-Yu Lin et al. (2020), that considered using ML methods for GRMHD BH data analysis.

DEEP HORIZON; a Machine Learning Network that Recovers Accreting Black Hole Parameters

van der Gucht et al. (2020) built an ANN, DEEP HORIZON, with the purpose of recovering parameter values from BH GRMHD shadow images. The algorithm consisted of two networks: (1) a regression algorithm to recover the values of $M_{\rm BH}$, $\dot{M}$, $R_{\rm high}$, PA and i; and (2) a classification algorithm to recover the value of a_* which was a discrete variable in their data. They generated an image library of 10^5 GRMHD BH shadows in a SANE magnetism environment using the ray-tracing code RAPTOR (Bronzwaer et al., 2018).

DEEP HORIZON was able to recover the model parameters with a 95.9 % accuracy from the simulations. To mimic the EHT image reconstruction algorithm, the images were blurred with 2D Gaussian beams of 5 μas, 10 μas and 20 μas to ascertain the reliability of parameter recovery for images of a similar fidelity to the EHT images. Up to the 10 μas beam, DEEP HORIZON could reliably predict all parameters up to $3\,\sigma$.

With a 20 μas beam, however, the network could only report M and $\dot{M}$ with satisfactory accuracy. This is likely because these parameters are connected to the larger-scale features of the BH, specifically the size of the emission ring, which is consistently present. The PA and $R_{\rm high}$ parameter constraints were most affected by the convolution with the blurring kernels.

van der Gucht et al. (2020)'s analysis was also extended to BH simulations for a space-VLBI network at 690 GHz. Observations from such a telescope would not be impeded by atmospheric, weather and observation time restrictions that reduce the quality of the recoverable image. All parameters were recovered with high accuracy for these mock observations.

Although not directly comparable, the mean uncertainties on the mass measurement recovered by DEEP HORIZON are of the same order of magnitude as those of the EHT results (Paper VI). van der Gucht et al. (2020) concluded that such a method can be used to confirm EHT and future space-VLBI observation results given sufficient visibility coverage.

Feature Extraction on Synthetic Black Hole Images

Yao-Yu Lin et al. (2020) explored a similar ML technique as van der Gucht et al. (2020), using an Convolutional Neural Network (CNN) trained with 2×10^5 simulated BH GRMHD images generated with the IPOLE code (Mościbrodzka and Gammie, 2018). CNNs are particularly adept at processing image data which lends them to this problem. A full description of CNNs is beyond the scope of this thesis and is left up to the reader.

Yao-Yu Lin et al. (2020) only considered unblurred images, however, images from both MAD and SANE accretion flow models were included and the trained network was tested by providing images of BHs that had spin values that were not included in the training data. Two networks were trained; one classification algorithm to determine whether an image was MAD or SANE and one regression network to predict the spin. Feature maps were also analysed to determine what image features the network identified to predict the model spin value.

Once trained, Yao-Yu Lin et al. (2020) tested the network on three datasets. They found the network recovered the spin and magnetic flux classification from the image libraries with high accuracy. Notably, for one of the datasets, the MAD/SANE classifications were recovered with 99.89 % accuracy. On average, the values were recovered with 97.34 % accuracy for all three datasets.

For the spin values, the regression network achieves standard deviations of 0.013, 0.209 and 0.164 versus the truth values for the three datasets respectively. Although the network predicted the spin values with high accuracy, it favoured spin values that were included in the training data. This illustrates the impact training data can have on the output model.

Feature maps were used to determine which image characteristics the model used to predict spin values. Yao-Yu Lin et al. (2020) found the model focused on low surface brightness features of the images, specifically, the spiral structures in the accretion flow. These form from processes related to BH accretion and jet production which is believed to be caused by BH spin.

These two papers illustrate possible applications of ML to EHT science. However, the sizes of the datasets used were large even though the model parameter ranges were somewhat restricted. Empirical EHT data has an even larger number of impacting factors which include telescope, location, weather and other performance specific effects; in addition, the u, v space is sparsely sampled, and, because there are few antennas in the EHT array, each with a corruption of the incoming signals, calibration is not as robust as it can typically be for lower frequency radio interferometers.

In addition to the above, the emission around a BH itself evolves over time, on timescales of minutes for 10^6 $M_\odot$ BHs, and may change between observations on consecutive days. This, compounded with the large library of simulation data, is the motivation for exploring an ML technique that can reduce the scope of the parameter space sufficiently.

1.4 Project Objective and Thesis Layout

This work aims to conduct an exploratory study into the application of a particular class of ML algorithms to classify simulated GRMHD BH shadow images. We will test the ability of an unsupervised learning algorithm, a SOM, to distinguish differences between BH shadow images in an automated manner. We wish to make both qualitative and quantitative assessments on the ability of the resultant model to classify shadow images of a wide range of intrinsic geometric and thermodynamic parameters.

To achieve this, the algorithm PINK, a SOM that uses spatial characteristics to classify input, will be trained using a GRMHD BH simulation image library generated by the ray-tracing programme RAPTOR. The images display morphological differences based on simulation input parameters. Further, the images will be mapped to the SOM. An analysis of the distribution of parameters based on the models' mappings will conclude this work.

The thesis is organised as follows; in Chapter 2 we describe the GRMHD simulation dataset to be used and develop the method of analysis with example use of PINK on a crescent model dataset. Subsequently, we present the results of PINK for the BH GRMHD simulations and discuss the findings of these in Chapter 3. We conclude in Chapter 4 with a summary of the project.

Chapter 2

Method

In this chapter, we describe the two major components required in this thesis:

- a large library of GRMHD black hole images;
- the self-organising map algorithm PINK.

We then present the application of PINK to a parametric crescent model as a descriptive example ahead of GRMHD applications in Chapter 3.

2.1 GRMHD Image Library

A library of 10^5 BH model images was generated from GRMHD simulations, created with the Black Hole Accretion Code (BHAC, Porth et al., 2017) and processed using the ray-tracing code RAPTOR (Bronzwaer et al., 2018). This produced a library of images of the photon ring and accretion flow around a BH for different environment initialisations. A sample of the images contained in the final library can be seen in Figures 2.1 and 2.2. These show how different simulation parameters result in different BH morphology, discussed in more detail in the following section.

2.1.1 Model Parameters

The simulation input parameters are the mass of the BH, $M_{\rm BH}$, the accretion rate onto the BH, $\dot{M}$, the electron temperature prescription, $R_{\rm high}$ (e.g. Mościbrodzka, Falcke, and Shiokawa, 2016), the inclination, i, the viewing angle with respect to the BH spin axis, PA, and the dimensionless BH spin parameter, a_*. The images

TABLE 2.1: The GRMHD simulation parameters sampled ranges. The first five parameters were continuously sampled during the simulations, while a_* is set to one of the five values per simulation. $\dot{M}$ is sampled logarithmically and all other parameters were sampled linearly. Reproduced from van der Gucht et al. (2020).

Parameter	**Description**	**Range**
$M_{\rm BH}$	Mass of the central BH	$[2 \times 10^9 M_\odot, 8 \times 10^9 M_\odot]$
$\log_{10} \dot{M}$	Mass accretion rate	$[2 \times 10^{-6} M_\odot/\mathrm{yr}, 1 \times 10^{-2} M_\odot/\mathrm{yr}]$
$R_{\rm high}$	ion-electron temperature ratio	[1, 100]
i	inclination	[15°, 25°]
PA	position angle	[0°, 360°]
a_*	spin parameter	$0, \pm 0.5, \pm 0.9375$

were generated at an observing frequency of 230 GHz, the EHT observing frequency (Paper II), captured with a 0.1×0.1 milli-arcsecond^2 camera. The images have dimensions of 128×128 pixels (van der Gucht et al., 2020).

Within the simulation, all parameters are transformed to be dimensionless. The simulations are also scale free. To convert the parameters from code units to cgs units, the following are defined: the simulation length unit $\mathcal{L} = r_{\rm g}$ (see Equation 1.3); the simulation time unit $\mathcal{T} = r_{\rm g}/c$; and the simulation mass unit $\mathcal{M}$, which sets the density in the accretion flow. If $\dot{M}_{\rm sim}$ is the dimensionless accretion rate, then $\dot{M} = \dot{M}_{\rm sim}\, \mathcal{M}/\mathcal{T}$ (van der Gucht et al., 2020).

The simulation parameters were sampled from ranges as described in Table 2.1. $\dot{M}$ is sampled logarithmically due to the large parameter range it encompasses, while all other parameters are sampled linearly. a_* is the only discrete parameter in the simulations.

Each parameter affects the BH emission morphology in different ways and, here, we give a broad overview of these. $M_{\rm BH}$ alters the size of the photon ring; a higher mass BH will have a larger diameter photon ring than a lower mass BH. Illustrated in Figure 2.1, this follows from Equation 1.3 as the photon ring size is directly proportional to the BH mass. The flux observed is dependent on $\dot{M}$.

$R_{\rm high}$ characterises the region where the majority of the observed emission originates. A higher value for $R_{\rm high}$ indicates emission primarily originating from the jet, and a lower value indicates the dominance of disk emission. Figure 2.2

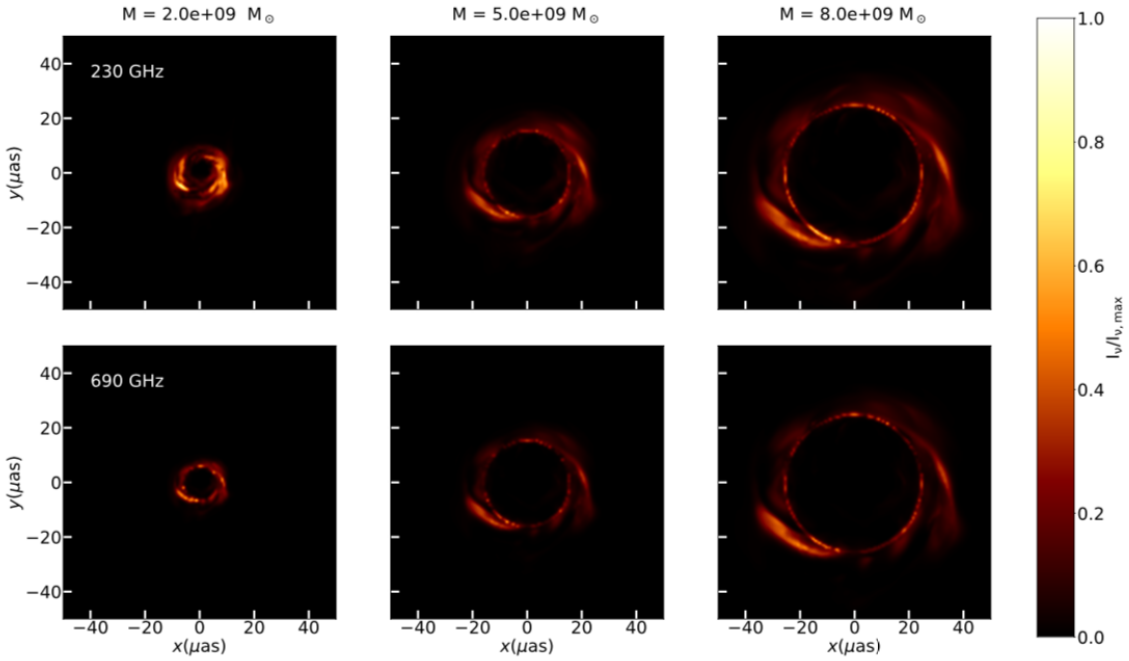

FIGURE 2.1: Snapshots from the GRMHD BH simulations. From left to right, $M_{\rm BH}$ increases as indicated. The top row shows images viewed at 230 GHz, the current observation frequency of the EHT. The second row depicts the same images observed at 690 GHz. At 690 GHz, the observing frequency simulated for the space-VLBI observations in van der Gucht et al. (2020). Credit: van der Gucht et al. (2020)

depicts how this influences the ring morphology; for low $R_{\rm high}$ (top row), there is extended emission outside the photon ring while for high $R_{\rm high}$ (bottom row) there is extended emission within the photon ring. This additional ring structure differentiates the models.

The spin parameter, a_*, indicates the magnitude of the BH's spin angular momentum and the orientation of the accretion flow concerning the BH spin, as defined in Equation 1.4. For disk motion in the same direction as the BH, a_* is positive. For motion in the opposite direction, a_* is negative.

Other parameters included in the simulations were the viewing angle, i, the angle between the observer and the accretion momentum vector, and the position angle, PA, the rotation of the image on the image plane. These parameters affect the orientation at which we view the BH and not the physical state of the GRMHD environment; they change the jet position and apparent asymmetry in the image. i in the simulations was restricted to the range given in Table 2.1 as there are observational constraints on this value for M87* (Walker et al., 2018).

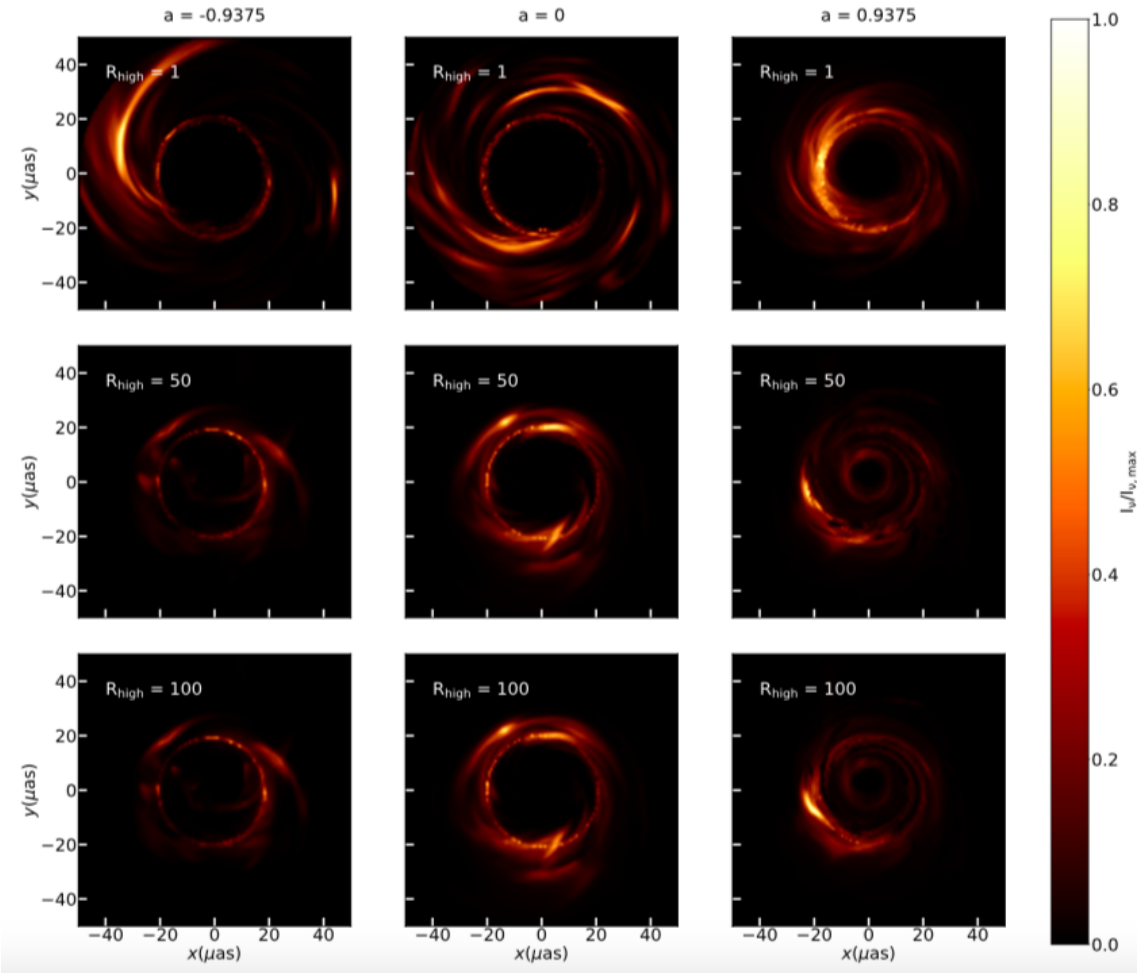

FIGURE 2.2: Snapshots from the GRMHD BH simulations generated at 230 GHz. From left to right, a_* increases from -0.9375 through zero to 0.9375. From top to bottom, R_{high} increases from 1, through 50 to 100. The shadows vary based on these parameters which are shown here. Credit: van der Gucht et al. (2020)

2.1.2 Simulations

Only SANE accretion flows were considered in the simulations. To capture the evolution of the accretion flow, five simulations were run, each initialised with one of the a_* values. The simulations were run until $t_{\text{final}} = 10^4\,\mathcal{T}$, each consisting of 10^3 snapshots captured at intervals of $10\,\mathcal{T}$.

For each spin value, the last 10^2 snapshots were used to compute the final images. These snapshots were selected as the accretion flow at this late stage in the simulations is well evolved and in a steady state. For each simulation's 10^2 snapshots, 2×10^2 images were computed by linearly sampling the parameters,

bar a_*, from the ranges in Table 2.1. These images made up the final dataset of 10^5 images.

As we have already discussed, different parameter values produce morphologically different BH shadows. With this image library, we want to explore to what extent certain BH shadow features are persistent for the parameter realisations. In turn, we want to identify any degenerate parameters based on the images' spatial characteristics and determine in what ways the parameter space can be reduced.

2.2 Parallelized rotation and flipping INvariant Kohonen maps

PINK is an unsupervised ML framework that employs a competitive training technique. It creates a compressed representation of data in the form of a 2D SOM based on the images' spatial characteristics. The algorithm employs the well-known similarity metric the Euclidean distance. The objective of the PINK framework is to provide a compact visual description of highly dimensional data. This allows for quick classification of multiple objects of the same class as the classification of a single representative prototype classifies all objects within that prototype class. It can also be used as a metric to determine outliers/anomalies by identifying objects that least fit their class prototype.

2.2.1 The PINK Self-Organising Map

A PINK SOM consists of a set of nodes, $\mathbf{p}_j \in P$, that exist in the latent space of the data. The nodes $P = \{\mathbf{p}_j = (\mathbf{w}_j, \mathbf{c}_j) \,|\, \mathbf{w}_j \in \mathbb{R}^d, \mathbf{c}_j \in \mathbb{N}, j = 1, ..., \mu_P\}$ map every weight vector, $\mathbf{w}_j$ to a map coordinate $\mathbf{c}_j$. Here, j is the size of the latent space.

The weight vector is the prototype that the node represents and the map coordinate gives an ordering of the nodes based on similarity to each other. The data analysed by pink is an image that is processed as a given set of n patterns.

To determine the weight-vectors, the map needs to be trained on a representative dataset. The weight vectors are initialised to some value and, iteratively t times the n patterns $\mathbf{y}_i \in \mathbb{R}^d$, with $i = 1, ..., n$, are applied to the map. Here, d and n are, respectively, the dimensions and size of the dataset.

By calculating the similarity measure $\Delta(\mathbf{y}, \mathbf{w}_j)$ between each pattern (image) $\mathbf{y}_i \in \mathbb{R}^d$ and weight $\mathbf{w}_j \in \mathbb{R}^d$ of every node $\mathbf{p}_j \in P$, the closest node or BMU $q(\mathbf{y} = \text{argmin}_{j=1,\ldots,\mu_P}\Delta(\mathbf{y}, \mathbf{w}_j))$ is determined. The nodes are then updated based on their distance to the BMU in the map $d(\mathbf{c}_{q(\mathbf{y})}, \mathbf{c}_j)$ and the number of training iterations t via a function $f(d(\mathbf{c}_{q(\mathbf{y})}, \mathbf{c}_j), t)$. The function f updates the weight vector $\mathbf{w}_j$ of a node $\mathbf{p}_j$ to a new value $\mathbf{w}'_j = \mathbf{w}_j + (\phi^{(j)}(\mathbf{y}) - \mathbf{w_j}).f(d(\mathbf{c}_{q(\mathbf{y})}, \mathbf{c}_j), t)$. $\phi^{(j)}$ is the standard Kohonen-map identity function and is used to align the coordinate systems of $\mathbf{y}$ and $\mathbf{w}_j$.

The function f consists of a distance-dependant element, which currently can be a Gaussian or Mexican Hat distance component in PINK, and simple linear damping based on the number of iterations t. Over the epochs, the weight-vectors converge to represent particular classes of objects from the data.

PINK assesses likeness based on spatial characteristics. Hence, to increase the fidelity of the training and mapping of the SOM, transformations of the input images are considered when determining a winning node. To achieve this, the algorithm calculates the Euclidean distances for all possible rotations and flips of objects to find the optimal value.

The difference between the image $\mathbf{y}_i$ and its best-fitting node $\mathbf{p}_{c_i}$ is a measure of how close the two values are. This is given by the quantization error (QE) defined as

$$QE = \frac{1}{n}\sum_{i=1}^{n} ||\mathbf{y}_i - \mathbf{w}_{c_i}|| \tag{2.1}$$

where n is, as previously stated, the number of patterns in the input image $\mathbf{y}_i$(Wandeto and Dresp, 2019). For an input image vector, the QE is the Euclidean distance between the image vector and the best matching node vector. The QE progression should converge to a minimum at the best matching node (Kohonen, Nieminen, and Honkela, 2009). The QE progression of the SOM converges over the training iterations.

Once training is complete, the weight vectors $\mathbf{w}_j$ are finalised. Images $\mathbf{y}$ can then be mapped to the prototypes P which give it a coordinate c in the map for the best matching node.

2.2.2 Map Parameters

A variety of parameters can be set when training and mapping with PINK. A full description of these, as well as descriptions of those not included here, can be found in the PINK documentation.[1] In Table 2.2, a selection of parameters relevant to this work are described.

Structurally, the SOM typically has either a hexagonal or quadratic shape. We chose to use a quadratic SOM for ease of reference in our outputs. The width and height of the SOM are set prior to any training which determines the total number of nodes in the SOM. The node dimensions are fixed; these cannot be bigger than the input images' dimensions.

In training, the image being considered can be rotated a number of times, from 4 (i.e. four 90° rotations) to 360 (i.e. 360 1° rotations) times. The number of rotations completed is traded off with processing time. It is recommended to keep the rotations at 360 for the best possible results, even if this increases computation time (Polsterer et al., 2016).

The number of training iterations completed is set by the user. To keep computation time to a minimum, it is best to set this low without having it negatively impact the result. Polsterer, Gieseke, and Igel (2015) found that increasing `-num-iter` > 4 did not have a significant impact on the SOM result. Given the spherical symmetric characteristics of our GRMHD BH simulation data, we expect this to hold true for our case.

The training and mapping of PINK are run separately. In training, the training dataset is the only input given and the fully trained SOM is the generated output. In mapping, the data to be mapped and the trained SOM are the inputs and the output is a similarity file that includes the similarity distances between the data and nodes. Convergence tests could be performed on subsets of the data, but that is not explored here.

Before using PINK on our GRMHD BH library, we tested it with a toy crescent model. Apart from the practical benefits, this exploration, presented in the following section, gave us a better understanding of the PINK framework. Here we present the results of that in a very similar way to what we will in the GRMHD case in Chapter 3.

[1] https://github.com/HITS-AIN/PINK

TABLE 2.2: Descriptions of a selection of PINK parameter inputs set in training and mapping the SOM.

Parameter	**Function from PINK Documentation**	**Input Values**
`-layout`	Layout of SOM	Quadratic, quadhex or hexagonal; default = quadratic
`-num-iter`	Number of training iterations	Integer value; default = 1
`-numrot`	Number of image rotations completed when calculating similarity measure	Integer value; 1 or a multiple of 4; default = 360
`-som-height`	Height dimension of SOM	Integer value; default = 10
`-som-width`	Width dimension of SOM	Integer value; default = 10
`-init`	Node weights initialisation	Zero, random, random with preferred direction, user set; default = zero
`-neuron-dimension`	Pixel dimension of nodes	Integer value; cannot be greater than the number of pixels in the input images; default = image-dimension$\times\sqrt{2}/2$
`-dist-func`	Distribution function for SOM update	Gaussian or Mexican Hat distance measure

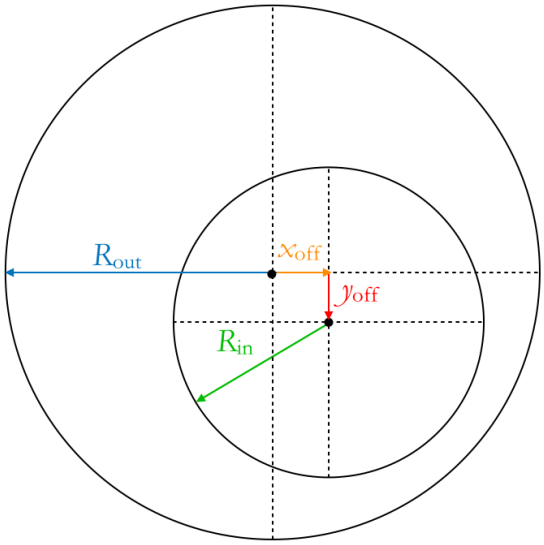

FIGURE 2.3: The geometry of the crescent model and corresponding parameters.

1	2	3	4	5	6	7	8
9	10	11	12	13	14	15	16
17	18	19	20	21	22	23	24
25	26	27	28	29	30	31	32
33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48
49	50	51	52	53	54	55	56
57	58	59	60	61	62	63	64

FIGURE 2.4: SOM node numbering reference. Nodes are numbered from 1 to 64 from left to right, top to bottom as if reading lines of text.

2.3 Crescent Toy Model

As discussed in Section 2.1, the emission around BHs varies observed as morphological differences between them. The differences are subtle and nonlinear, and so, to determine whether PINK is a suitable framework to explore GRMHD BH shadow images, we first investigated PINK's performance on a parameter crescent model.

The model used, a crescent, was chosen as it is a simple, well-understood model that loosely approximates the structure of the BH ring feature (Paper IV). We generated the dataset using the same model as described in Paper IV.

2.3.1 Crescent Model Parameters

A crescent is characterised by five parameters; R_{out}, the radius of the outer ring, R_{in}, the radius of the inner ring, x_{off}, the x-offset of the inner ring from centre, y_{off}, the y-offset of the inner ring from centre, and F, the integrated flux or brightness of the source. The inner ring has 0 flux. The geometry of this setup is shown in Figure 2.3.

We generated 10^5 crescent model images, uniformly sampling the parameters from the continuous ranges stated in Table 2.3. The images were set to have dimensions 100×100 pixels with the crescents centred on the outer ring centre.

TABLE 2.3: The crescent parameters and corresponding sample ranges. All parameters were uniformly sampled when generating the images.

Parameter	**Range**
F	[0, 40]
R_{out}	[25, 50]
R_{in}	[0, R_{out}-1]
x_{off}	[0, 10]
y_{off}	[0, 10]

2.3.2 PINK Crescent Self-Organising Map

We trained PINK on the 10^5 crescent images with parameters as set in Table 2.4. These settings reflect the setting used later with the GRMHD data.

The trained SOM is shown in Figure 2.5. For reference, nodes are numbered 1 to 64 from left to right, top to bottom as if reading lines of text. This identification scheme is depicted in Figure 2.4.

Following training, we mapped the 10^5 crescent images, the same set as was used to train the SOM, to the trained map. Figure 2.6 shows the number of images mapped to each node. Node 64 has the largest class size. This suggests its prototype represents a crescent that is common in the data.

The SOM in Figure 2.5 has easily distinguishable groups of crescents displaying similar characteristics. In Figure 2.7, we have indicated a few of these rudimentary clusters. We note that crescents with large outer radii are positioned along the right-hand side of the lattice; crescents with the smallest inner radii are grouped on a diagonal across the bottom left-hand corner of the map, and the crescents with the highest integrated fluxes are located along the top edge of the map.

To further our analysis, we used the image mappings and associated model parameters to determine to what degree the image clustering is present within the mapped parameter distributions. We present the results excluding those for

TABLE 2.4: The PINK parameters set when training and mapping the crescent data. These are the same as those set for the GRMHD data in Chapter 3

Parameter	Value
`-layout`	`quadratic`
`-num-iter`	`5`
`-numrot`	`360`
`-som-height`	`8`
`-som-width`	`8`
`-neuron-dimension`	`100`
`-init`	`random`
`-dist-func`	`Gaussian`

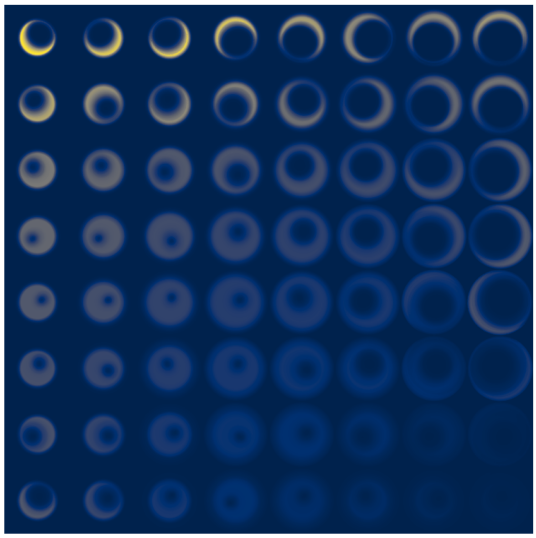

FIGURE 2.5: Trained quadratic PINK SOM for the 10^5 crescent models. Consisting of 64 nodes, each prototype represents a class of similar objects from the data.

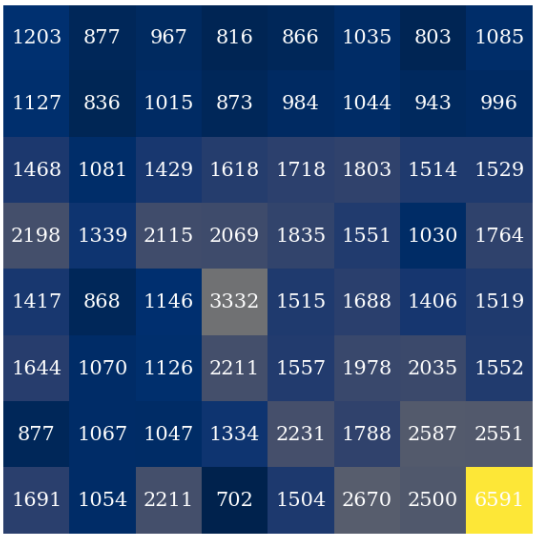

FIGURE 2.6: The number of crescent models mapped to each node of the SOM. Models are assigned to a node based on their similarity to the node's prototype.

y_{off} as they follow those of x_{off} given they are both geometric translations in the image plane.

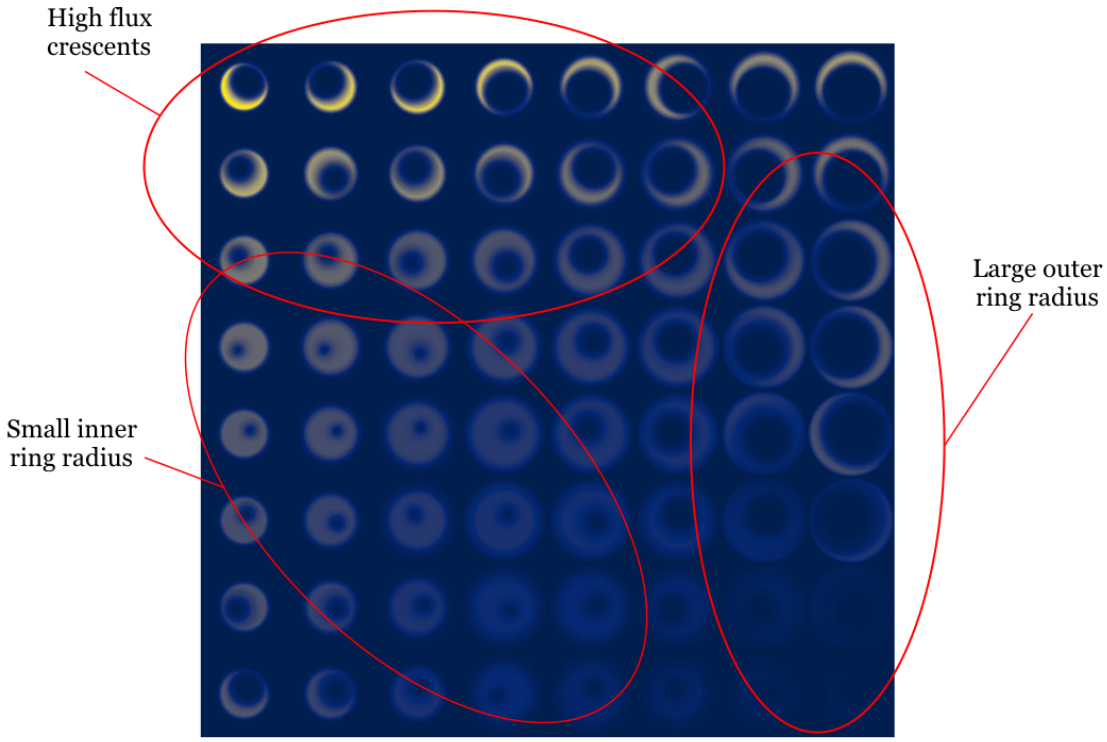

FIGURE 2.7: The SOM, as in Figure 2.5, with clusters of prototypes displaying similar characteristics identified.

2.3.3 Crescent Parameter Heatmaps

We used heatmaps to describe the prototypes' parameter distributions across the lattice. A heatmap is a colour-gradient representation of a variable in a 2D plane. The colours are representative of the variable values. The heatmaps of F, R_{out}, R_{in} and x_{off} are displayed in Figure 2.8. The value indicated on each node is the mode of the nodes' class distributions for each parameter.

The flux heatmap (2.8 (a)) shows the node values' increase in diagonal bands from the bottom right-hand corner to the top left-hand corner of the lattice. This structure follows what we identified in Figure 2.7.

For R_{out}, the heatmap (2.8 (b)) has a similar structure to that of F except it is mirrored across the counter diagonal. Here, the maximum valued nodes are clustered in the bottom right-hand corner of the lattice and the minimums are located in the first column.

The heatmap for R_{in} (2.8 (c)) has the most defined structure; similarly valued nodes are grouped in slanted bands orientated along the leading diagonal. The

(a) Flux, F

40.0	35.8	31.6	37.9	40.0	37.9	40.0	37.9
33.7	33.7	29.5	31.6	33.7	37.9	40.0	37.9
40.0	40.0	29.5	37.9	40.0	37.9	31.6	33.7
40.0	35.8	37.9	37.9	35.8	27.4	21.1	21.1
25.3	25.3	25.3	31.6	27.4	14.7	12.6	14.7
18.9	18.9	18.9	21.1	21.1	10.5	6.3	6.3
14.7	12.6	14.7	14.7	12.6	6.3	6.3	2.1
12.6	8.4	8.4	10.5	8.4	4.2	2.1	0.0

(b) Outer ring radius, R_{out}

26.3	27.6	30.3	31.6	34.2	35.5	36.8	39.5
26.3	28.9	31.6	32.9	35.5	39.5	42.1	42.1
25.0	30.3	32.9	35.5	39.5	43.4	46.1	46.1
25.0	31.6	35.5	39.5	43.4	48.7	50.0	47.4
25.0	30.3	36.8	44.7	47.4	48.7	50.0	50.0
25.0	30.3	34.2	50.0	48.7	34.2	50.0	50.0
26.3	28.9	32.9	36.8	50.0	43.4	50.0	48.7
27.6	30.3	30.3	34.2	38.2	28.9	27.6	50.0

(c) Inner ring radius, R_{in}

23.2	23.2	25.8	28.4	28.4	30.9	33.5	36.1
18.1	18.1	20.6	23.2	25.8	28.4	30.9	36.1
12.9	12.9	15.5	18.1	23.2	25.8	30.9	38.7
5.2	5.2	7.7	12.9	18.1	23.2	33.5	41.3
0.0	0.0	2.6	5.2	15.5	23.2	30.9	43.8
10.3	7.7	2.6	0.0	15.5	20.6	28.4	36.1
15.5	12.9	2.6	0.0	10.3	15.5	20.6	33.5
23.2	18.1	2.6	2.6	0.0	0.0	5.2	2.6

(d) x coordinate offset, x_{off}

5.3	8.4	4.7	2.6	6.8	6.3	3.7	5.3
7.4	8.9	9.5	6.8	7.9	8.9	10.0	9.5
4.2	1.1	4.7	3.2	6.8	9.5	7.9	7.9
8.4	6.3	3.7	6.3	1.1	2.1	3.7	1.1
0.0	2.1	7.9	6.8	5.3	2.6	3.7	0.0
10.0	5.3	6.3	0.5	5.8	8.4	9.5	2.6
10.0	9.5	4.7	5.3	8.9	6.8	8.4	8.4
6.8	2.1	4.2	2.1	3.2	6.8	9.5	2.6

FIGURE 2.8: The heatmaps for four of the crescent model parameters. Here, the heatmaps of the (a) flux, F, (b) outer ring radius, R_{out}, (c) inner ring radius, R_{in}, and (d) x coordinate offset, x_{off} are depicted. The value on each node is the parameter mode from the node's class of images.

band of the lowest valued nodes is located just above the bottom left-hand corner of the lattice. Again, this structure follows from what we previously identified in Figure 2.7 but here we see that this pattern repeats across the lattice.

(a) Flux, F (b) Inner ring radius, R_{in}

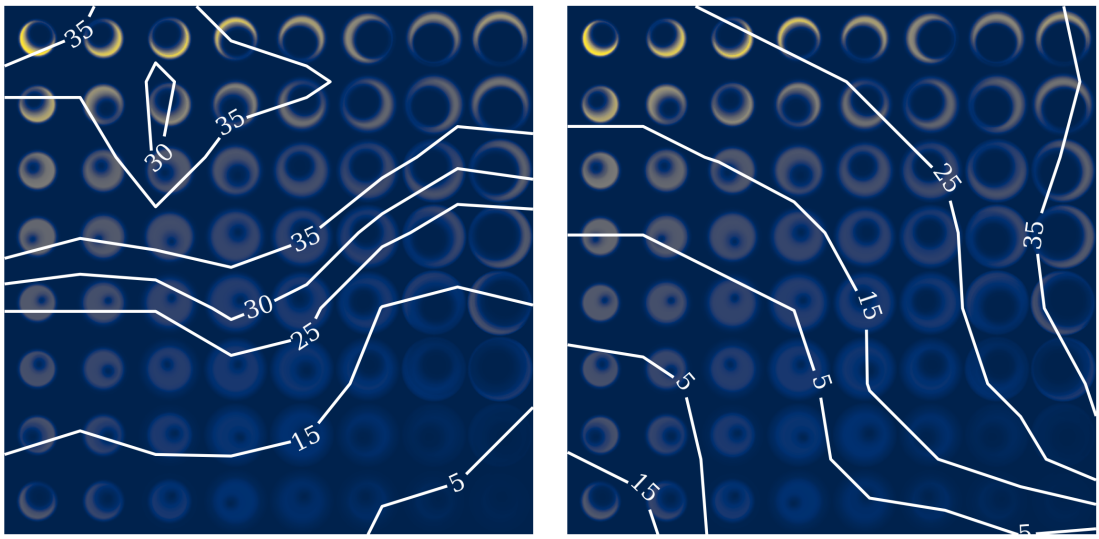

FIGURE 2.9: Contour maps of (a) flux and (b) R_{in} superimposed on the crescent SOM. This composite view helps to identify visual changes in prototypes and how they correspond to parameter changes of the nodes.

Finally, the x_{off} heatmap is shown in Figure 2.8 (d). There is no apparent structure within the distribution of values. This is expected given PINK's rotation and flip invariant similarity measure.

Interestingly, the heatmaps of both F and R_{out} have a small group of lower-valued nodes within their maximum node clusters. It is possible to discern this detail in the heatmaps as we concentrate on differences based on a singular parameter.

In Figure 2.9, we have plotted the contour maps of (a) flux and (b) R_{in} over the PINK crescent SOM. The contours follow from the heatmap values in Figure 2.8. These figures reinforce the observations made from the heatmaps. With the contours superimposed on the SOM, we have a composite view of prototypes differences and the corresponding parameter distribution across the map.

2.3.4 Parameter Distributions

The parameter modes of each node give a good intuition as to the distribution of parameters across the lattice. However, the distributions of the model parameters within each node class will vary.

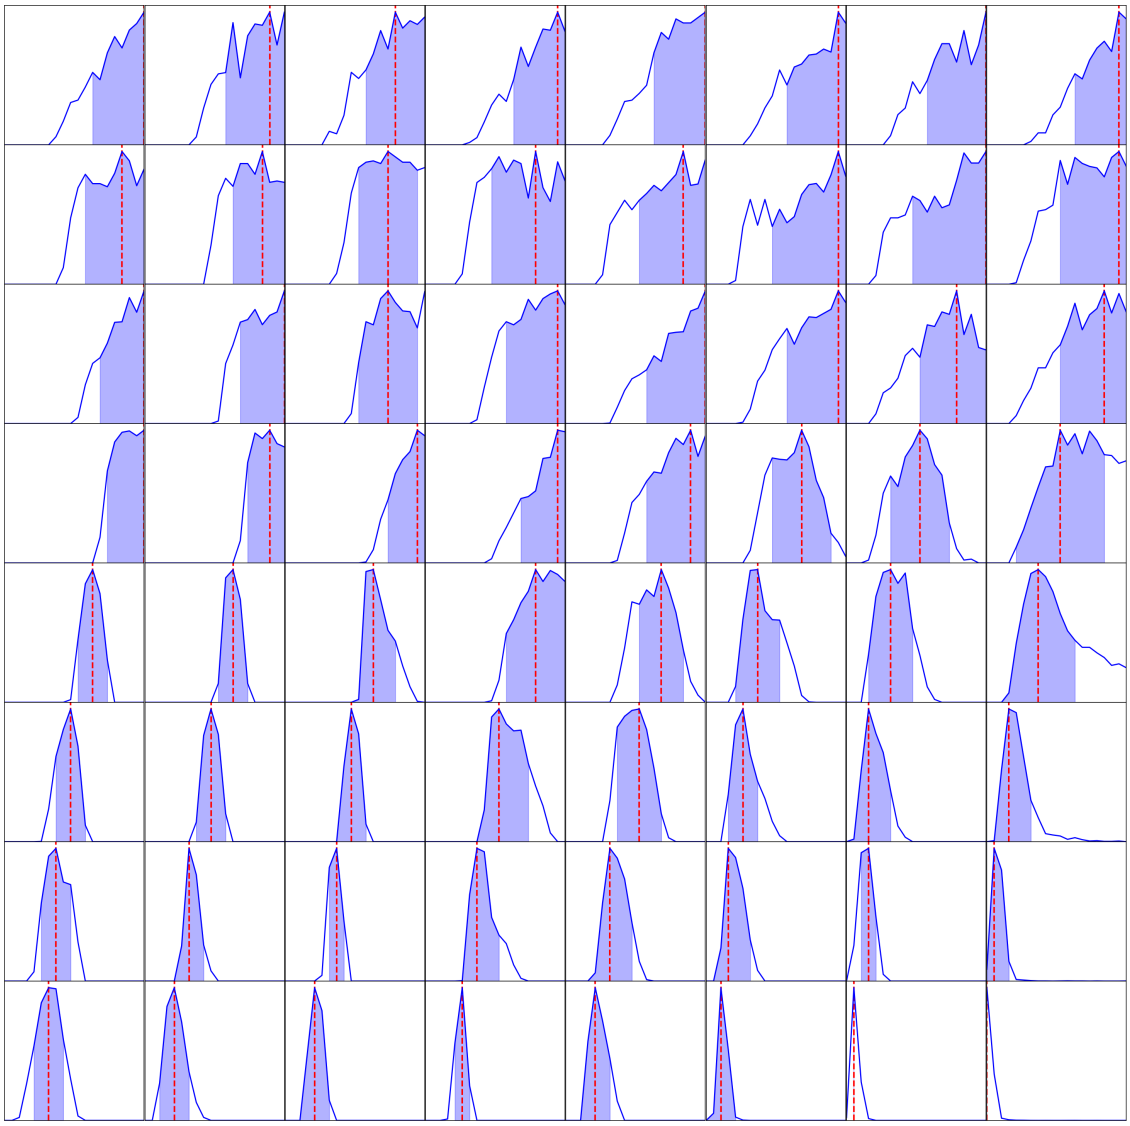

FIGURE 2.10: Distribution of mapped models' F across the SOM lattice. These distributions are created from the model parameters in each node's class. In each figure, the x-axis represents the entire parameter range; the value density is on the y-axis. As every distribution is normalised, the area under each curve equals 1. Since each node has roughly the same number of images, the relative densities are also captured by these distributions to a good approximation. The mode of the distribution is indicated with the red line. The blue shaded region represents the 68% confidence interval of the mode.

To dissect this, we plot the distributions of the model parameters within each node. We normalised each distribution to have an integral of 1 so comparisons between the distributions can be made despite the nodes having different class sizes.

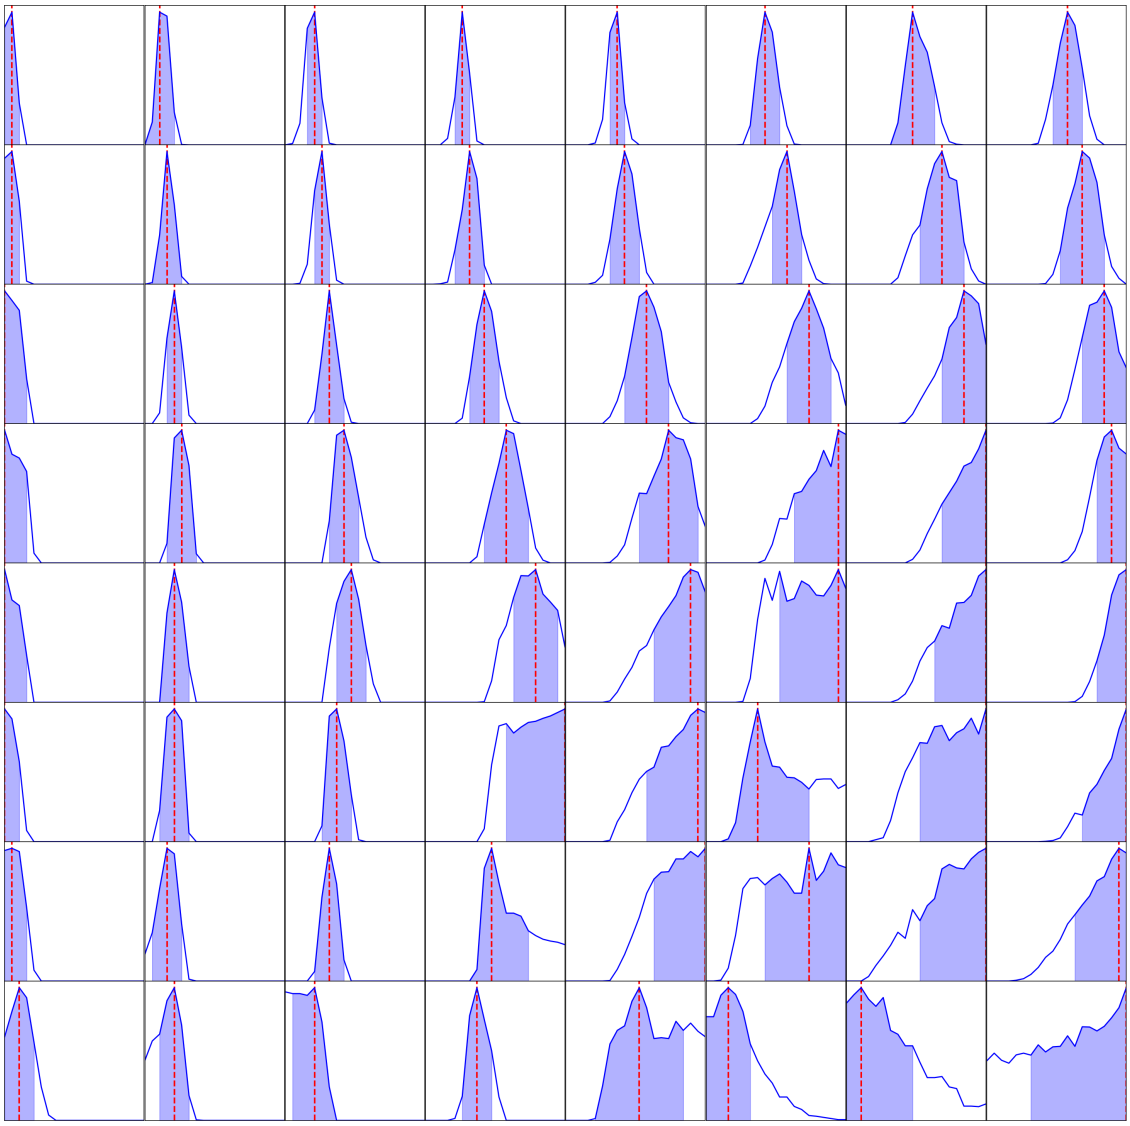

FIGURE 2.11: The same as Figure 2.10 but for R_{out}.

In the coming Figures 2.10 to 2.13, the x-axis of each distribution represents the entire range for that parameter. The blue line represents the full parameter distribution. The vertical red line indicates the mode for the distribution, or the most common value, and the blue shaded region represents the 68% confidence interval of the mode.

Flux, F

The flux distributions are shown in Figure 2.10. Distributions that peak in the lower region of the range, mainly those in the bottom three rows of the lattice,

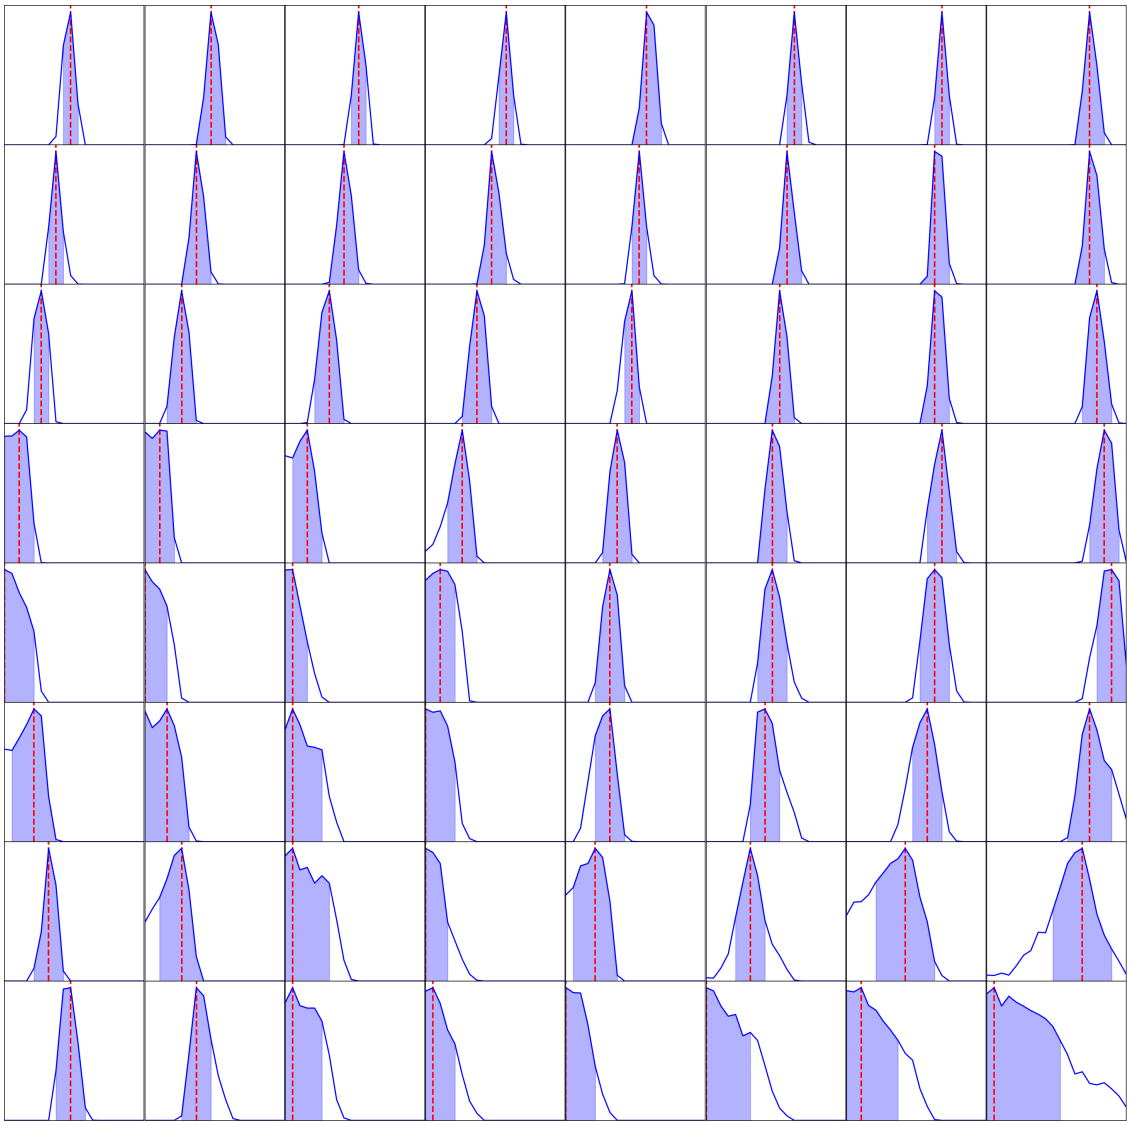

FIGURE 2.12: The same as Figure 2.10 but for R_{in}.

are narrowly spread in near-Gaussian distributions. Notably, nodes 63 and 64 are so narrowly distributed that their confidence intervals do not appear to extend beyond the mode of the distribution. It is important to note that node 64 contains all the zero flux (i.e. blank) images, explaining its large class size.

The flux distributions that have a peak near the maximum of the range are wider than those peaking in the lower part of the range. These distributions are notably broader and do not have a Gaussian distribution. The 68% confidence intervals are wider too.

Outer ring radius, R_{out}

In Figure 2.11, the distributions of R_{out} per node are displayed. Those in the first three columns and the top two rows have narrow distributions, indicated by their sharp peaks and narrow confidence intervals. These nodes all have peaks in the lower to mid part of the parameter range.

The distributions in the bottom right-hand corner of the lattice are wide and noisy in comparison. Rather than a sharp peak, these distributions have wider confidence intervals for the distribution mode; this is clearly shown in nodes 55 and 56. Node 64 includes models of R_{out} from the entire possible range, as expected as this class is made up of blank images.

Inner ring radius, R_{in}

The R_{in} distributions in Figure 2.12 are the narrowest in comparison to those of the two previous parameters. Bar node 64, the distributions have a clear peak with narrow confidence intervals.

The wide distribution seen in node 64 is a mapping effect; all zero-flux nodes were mapped to this node so models with varying R_{in} values are mapped to it.

Inner ring x-offset, x_{off}

Figure 2.13 shows the wide x_{off} distributions. All nodes have distributions that span the entire range and most are uniformly distributed. The confidence intervals reflect this with all occupying at least half of the entire parameter range. These distributions are as a result of PINK's rotation and flip invariant similarity measure. The crescent geometry, which in our model was affected by the inner ring offset, is not a dominant influence on the data mapping. The same distribution structure was seen for y_{off}.

Distribution maps overview

The F, R_{out} and R_{in} parameter distributions did show some degree of clustering within nodes. Considering the distributions of these parameters concurrently, there appears to be an inverse relationship between the width of the flux distributions and the R_{out} and R_{in} distributions.

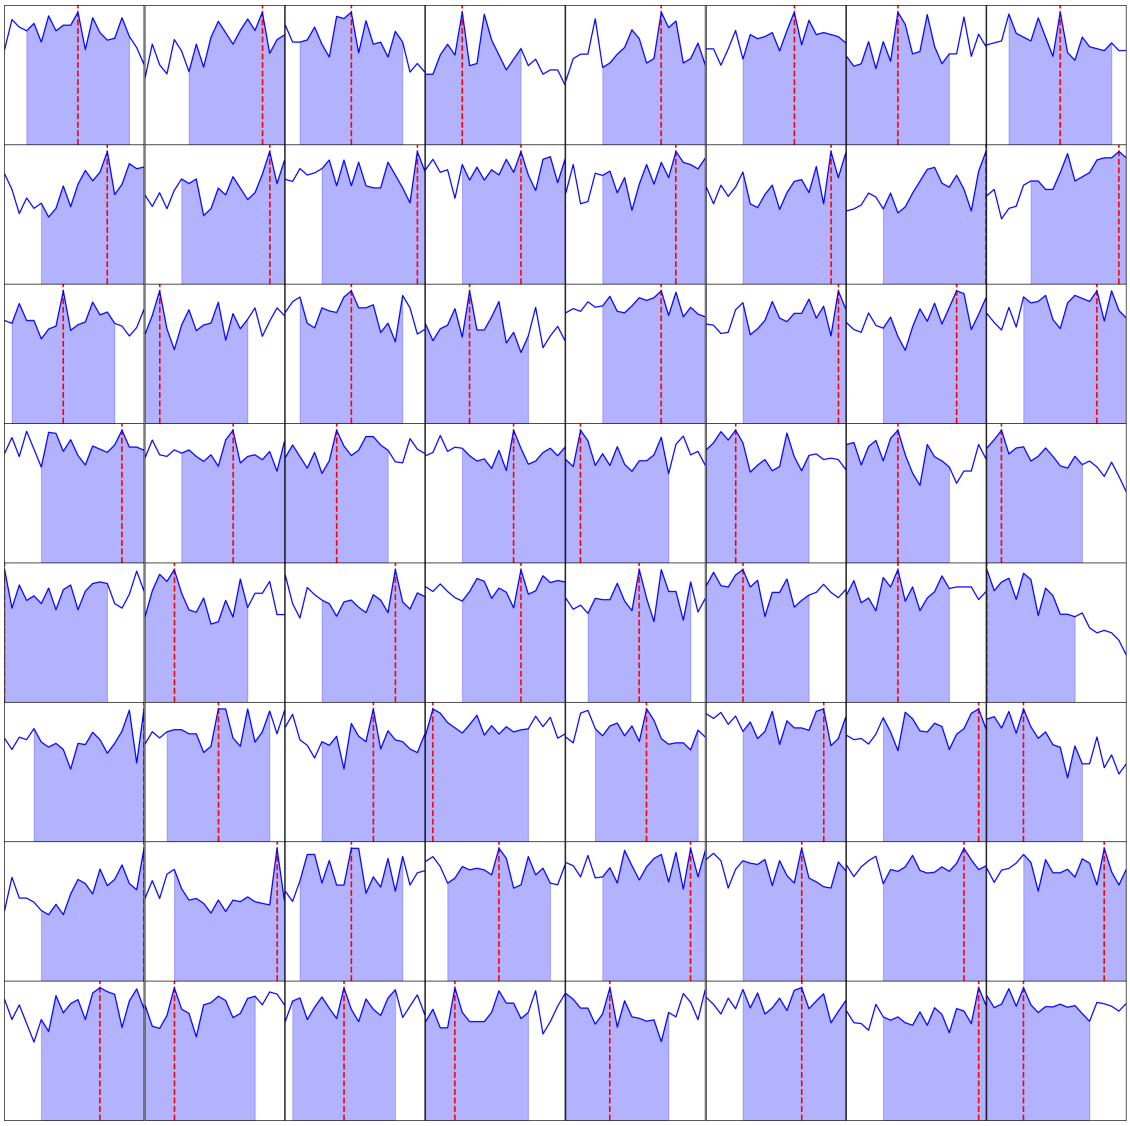

FIGURE 2.13: The same as Figure 2.10 but for x_{off}.

For node 64, the class is made up of almost exclusively zero-flux models. Given this, the crescents within the nodes class will have no observable image plane features, leading to all of these 'blank' images being mapped to node 64, irrespective of the other parameters. This is why for R_{out} and R_{in} the node 64 distributions are noticeably wider as compared to others for these parameters.

In nodes where both R_{out} and R_{in} are constrained, as in node 12, the corresponding flux distributions are wide and noisy. Given that, overall, there are a higher number of flux distributions with a widespread than for both R_{out} and R_{in}, we can infer that the network has identified the geometric features of a model as

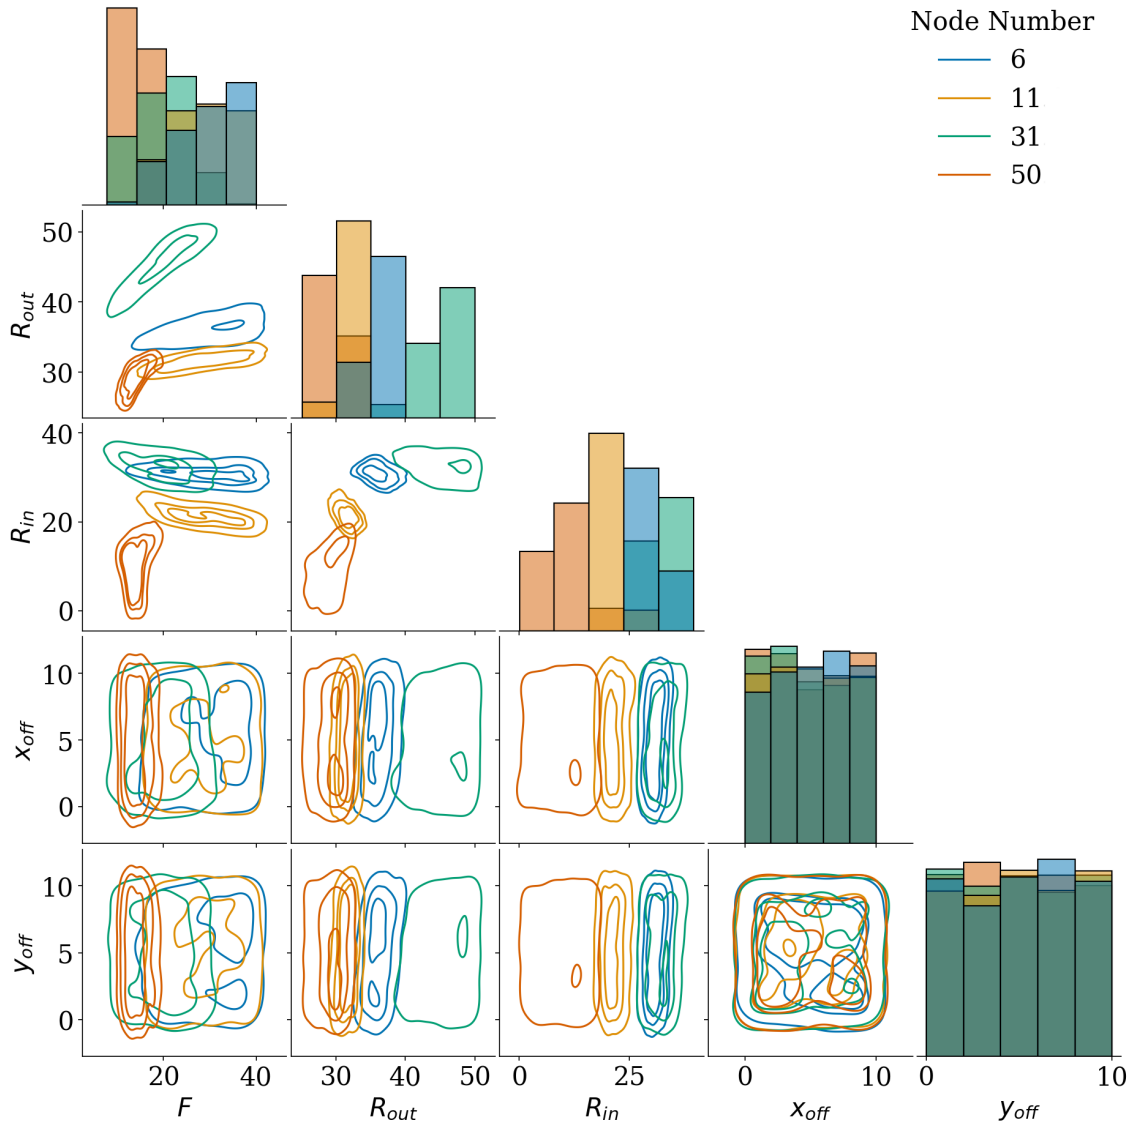

FIGURE 2.14: 2D-Joint parameter distributions of nodes 6, 11, 31 and 50 for the crescent model. Along the diagonal are the 1D parameter distributions. The 2D joint parameter distributions are in the lower triangle.

the dominant mapping determinant.

2.4 Further Analysis and Discussion

We plotted the 2D-joint parameter distributions for four quasi-randomly selected nodes, nodes 6, 11, 31 and 50, in Figure 2.14. These four nodes were chosen in part because they have similar class sizes (see Figure 2.6). Along the diagonal, the

1D parameter distributions are plotted, the same as in the previous distribution graphs. Below the diagonal, the joint parameter distributions are shown. By considering the joint distributions of four nodes concurrently, we can see if there are any trends or relationships between the parameter mapping across the lattice.

As expected, the joint distributions between all other parameters and both x_{off} and y_{off} show no correlation due to PINK's rotation/flip invariance. The x_{off} and y_{off} values of all the nodes span the entire possible range for the parameters. In particular, the joint distribution of the inner-offsets fill the entire plane.

The joint distributions of F, R_{out} and R_{in} show some structure. As previously noted, when the flux is well-constrained, as for node 50, the radii distributions extend across a larger part of the range than other nodes' distributions. In particular, node 50's R_{in} distribution spreads across half of the possible range.

However, the converse of the above is not necessarily true. Consider node 31; the R_{in} distribution spread is small but the flux and R_{out} distributions extend across a large part, over half for F and almost half for R_{out}, of their respective ranges.

For nodes 6 and 11, where the radii distributions are narrowly evidenced by the small area they occupy in the joint radii plot, the F distributions are wide. This suggests that, if both R_{out} and R_{in} have a small value range, the flux will have a large range. This suggests there may be a relationship between the flux and the radii of nodes, but this can only be inferred if the distributions of these parameters are known.

Finally, we consider the joint distribution of R_{out} and R_{in} for the four nodes. As previously noted, nodes 6 and 11 have narrow distributions for both parameters. Node 50 has a narrow R_{out} and a wide R_{in} distribution; node 31 has the opposite in terms of the radii distributions. For these two nodes, the narrow radius distribution exists at one end of the parameter ranges. It may be that, for the radii distributions, if one distribution is clustered at one extreme of a range the other will have a wide distribution. However, there is not enough information in Figure 2.14 to infer any general relationship between the joint radii distributions.

The above suggests that, for the crescent model used here, PINK's mapping is dominated by the crescent structure, i.e. the outer and inner radii, and then by the flux. This is only true if the flux is high enough to discern crescent features.

From this analysis of the toy model, we are satisfied with our understanding

of PINK and its performance on a model similar to a BH ring shadow. We apply the above methods to the GRMHD BH images, the results and discussion of which are found in Chapter 3. These will follow the same structure as presented here.

Chapter 3

Results and Discussion

In this chapter we present the primary results of this thesis separated into the three sections outlined below.

1. **Self-organising map lattice results** and their qualitative features and trends.
2. **Self-organising map correlations with intrinsic GRMHD parameters** which will connect image-domain features the algorithm has identified to any patterns, or lack thereof, on the SOM lattice structure.
3. **Statistical analysis**: the distribution of intrinsic parameter values within each node of the SOM lattice.

As in Section 2.3, the SOM nodes are numbered 1 to 64 from left to right, top to bottom as if reading lines of text (see Figure 2.4).

3.1 Self-Organising Map Lattice Results

Using the GRMHD image library, SOM algorithm and settings as described in Chapter 2, we generate the SOM presented in Figure 3.1. The following prototype features are immediately apparent:

1. the **size** of the emission ring changes across the lattice;
2. the **puffiness** or **thickness** of the emitting region varies between nodes;
3. some nodes display an **obvious two-sided jet** while, in others, only the **ring** structure is visible;
4. the shadow **central flux deficit** differ;
5. the **total flux** of the prototypes vary.

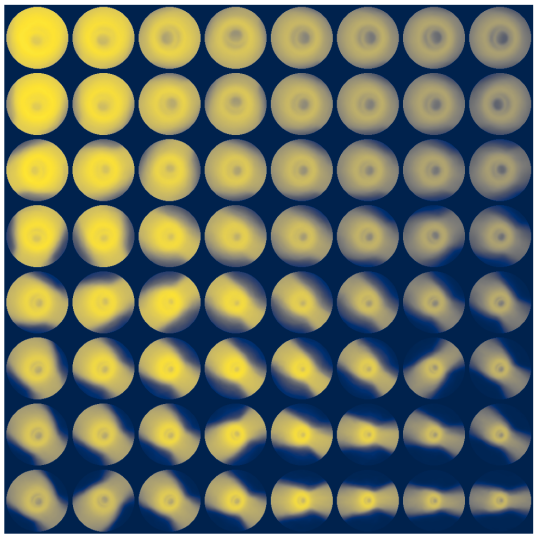

FIGURE 3.1: Trained PINK SOM of the 10^5 GRMHD BH shadow model images. Each node is a prototype of a class of models of the data. Morphological differences between nodes' prototypes are evident.

7600	3127	3054	2901	2562	2135	1861	1775
2668	2412	2192	1986	1725	1661	1343	1674
1882	1784	1697	1799	1824	1739	1230	1322
1677	1089	1192	790	1060	1662	1338	1138
1721	885	1113	993	643	444	763	2048
1634	875	828	765	854	985	1626	2258
1484	821	873	672	372	294	350	2096
1367	1418	1577	1796	438	828	973	2307

FIGURE 3.2: Number of GRMHD models mapped to each node of the SOM. This is the number of models in each node class and is the number of models in the dataset best represented by the corresponding prototype.

The number of models mapped to each node, shown in Figure 3.2, exhibits a similar structure to that for the crescent SOM in Figure 2.6. A corner node, node 1 in this case, has a significantly larger class size as compared to the others, indicating that the node's prototype may reflect a common morphology in the data.

We focus on the shadow which requires high contrast therefore we apply a pixel clip and pixel stretches to the SOM prototype images to better identify this feature. In Figure 3.3 we depict how the transformations change the prototype pixel distributions using node 32 as an illustrative example. The pixel transformations were not applied to the input data but the trained node prototypes. As these are the output of the SOM, the transformations did not affect the training or mapping.

The clip region was chosen based on the distribution maximum. We selected the pixels from 0.7 to 1.0 of the peak flux as this region contains the high-flux information from the BH which is associated with the ring feature.

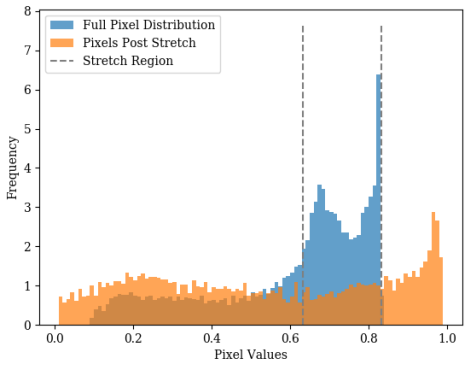

(a) Node 32 pixel distribution before (blue) and after (orange) the linear stretch was applied to the indicated stretch region.

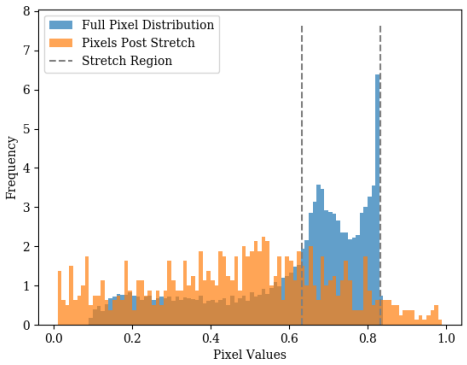

(b) Node 32 pixel distribution before (blue) and after (orange) the exponential transformation and linear stretch was applied to the indicated stretch region.

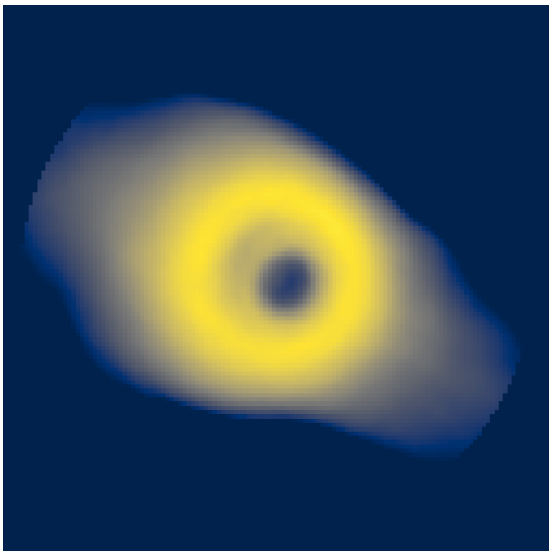

(c) Prototype for node 32 post the pixel clip and linear stretch transformation. The disk and jets of the image are visible.

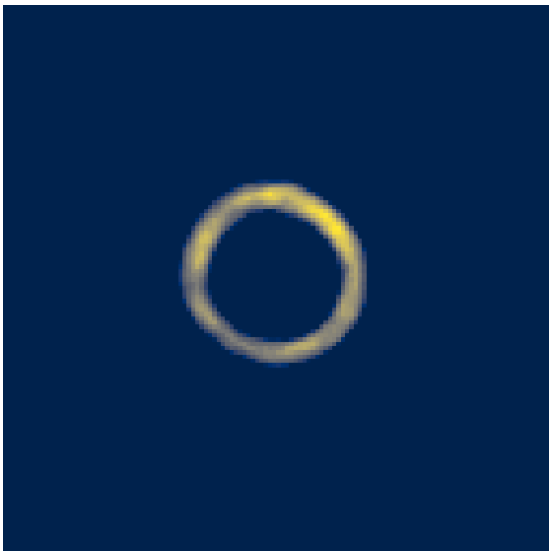

(d) Prototype for node 32 post the pixel clip, exponential transformation and linear stretch. Only the ring structure is visible.

FIGURE 3.3: The pixel transformations applied to node 32. The pixel distributions are in the top row and the resultant images in the bottom row.

We applied two different transformations to the clipped pixel region in isolation. In the first case, we did a simple linear stretch of the clipped region, extending the distribution across the entire original pixel flux range (Figure 3.3 (a)). In the second, we transformed the stretch region pixels as $f(p) = 10^p$, where p is the stretch region of the pixel distribution. We then, as before, stretched the

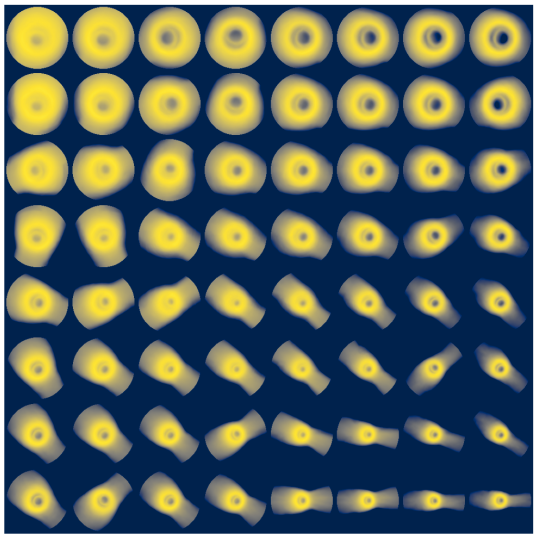

FIGURE 3.4: GRMHD PINK SOM with a linear pixel stretch applied. The pixel transformation makes the ring and jet structures more apparent.

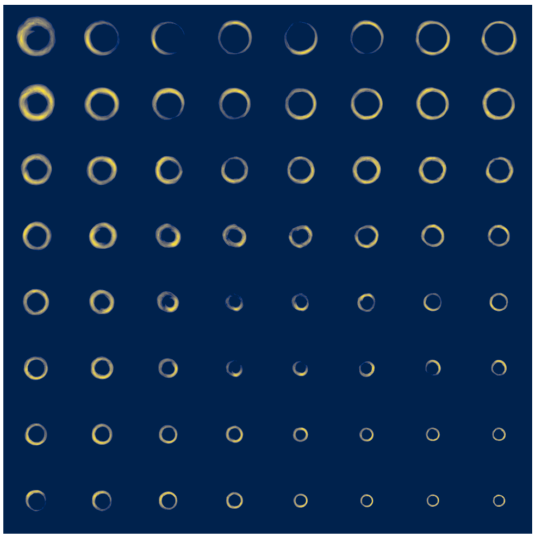

FIGURE 3.5: GRMHD PINK SOM with an exponential transformation and linear pixel stretch applied. The pixel transformation makes the ring structure more apparent.

transformed pixel distribution to extend across the original pixel flux range (Figure 3.3 (b)).

The linear stretch resulted in the image in Figure 3.3 (c). The exponential transformation (and linear stretch) resulted in the image in Figure 3.3 (d). The disk and jet structure is preserved in the linear case, while the exponentially-transformed image emphasises the shadow ring structure. The transformed SOMs are displayed in Figures 3.4 and 3.5 for the linear and exponential cases, respectively. Throughout the rest of this thesis, we will refer to these as the linear and exponential cases.

Following the same process as in Section 2.3, we begin our SOM analysis by visually identifying clusters of morphologically similar nodes. We completed this exercise for both the linearly and exponentially transformed SOMs.

In Figure 3.6 the linearly stretched map is shown. As indicated, the large, bright prototypes are located at the top of the lattice. Shadows displaying an obvious two-sided jet are grouped in the bottom half of the map; this is also where the smaller prototypes are, with the smallest being in the bottom right-hand corner of the map.

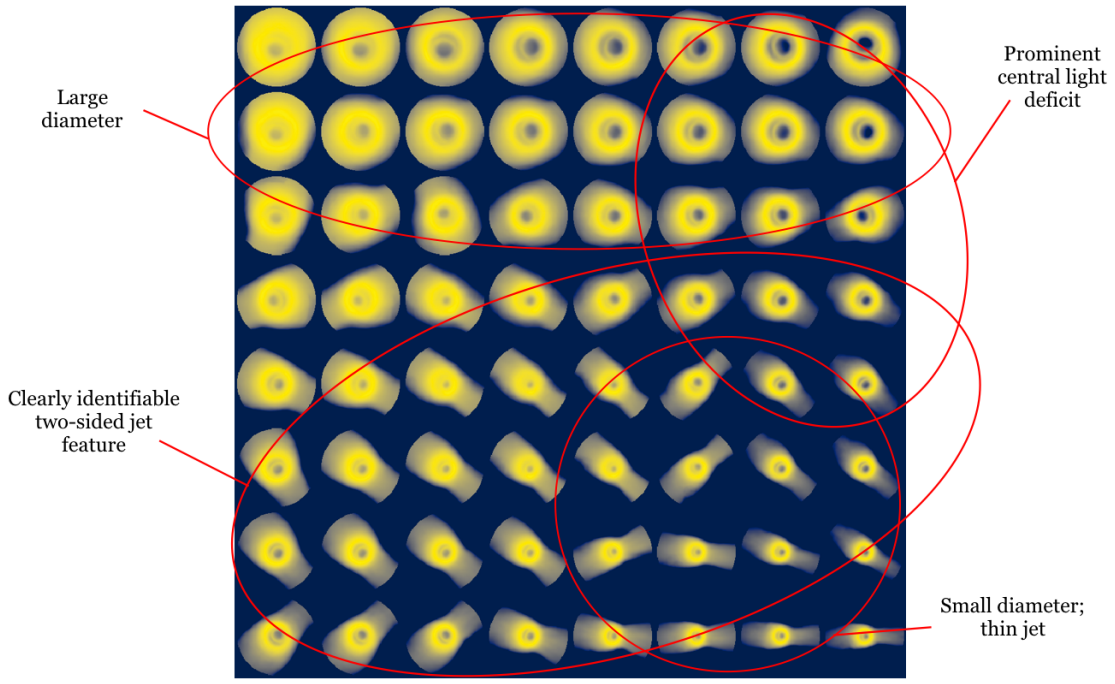

FIGURE 3.6: The linearly-stretched GRMHD SOM with clusters of prototypes displaying visibly similar characteristics identified.

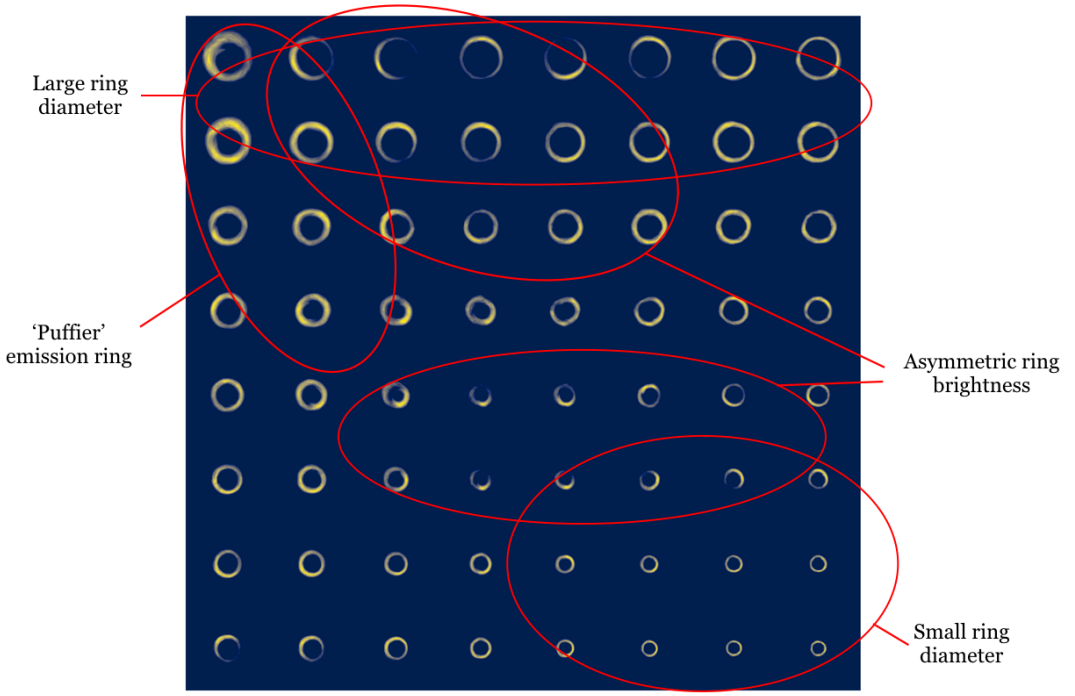

FIGURE 3.7: The exponentially-transformed GRMHD SOM with clusters of prototypes displaying visibly similar characteristics identified.

Considering the SOM as in Figure 3.7 of the exponentially transformed prototypes, the ring characteristics immediately discernible between the prototypes are

- ring **size**
- ring **'puffiness'**
- ring **asymmetry** or **azimutual structure**

with no jet features visible. In this figure, rings with a larger diameter are found in the first three rows of the lattice; those with a smaller diameter are in the bottom half of the map. Some rings, as in the top left-hand corner, are thicker or 'puffier'. All rings show some asymmetry, with this being more prominent in some.

The clusters we have identified here are rudimentary. They illustrate that PINK was able to cluster the models and generate representative prototypes, even if the prototypes do not display the detailed structure seen in the individual model images.

3.2 SOM Correlations with GRMHD parameters

In this section, we assess if the image-domain features map to intrinsic BH parameters used in the GRMHD simulation. This is a purely qualitative assessment, with a more detailed statistical analysis in the next section of this chapter.

In Figure 3.8, we have included the heatmaps of the four continuous GRMHD parameters: $M_{\rm BH}$ (3.8 (a)); $R_{\rm high}$ (3.8 (b)); i (3.8 (c)) and PA (3.8 (d)). In these heatmaps, the value displayed on each node is the parameter mode from the models' parameter distributions within the node class.

In Figure 3.8 (a) of the BH mass, the high valued nodes are located in the top row of the lattice; $M_{\rm BH}$ values decrease from the top-left node towards the bottom right-hand corner of the lattice. The minimum $M_{\rm BH}$ nodes are located in this bottom right-hand corner. This structure correlates with the prototype ring sizes we identified in Figure 3.6 and Figure 3.7. This indicates that the nodes' prototypes represent the $M_{\rm BH}$ of their class of models well.

The $R_{\rm high}$ heatmap (Figure 3.8 (b)) has high-valued nodes in a 'V' shape anchored in the bottom right-hand corner of the map. The largest cluster of maximum-valued nodes is in this corner. The low-valued nodes are grouped in the top left area of the lattice.

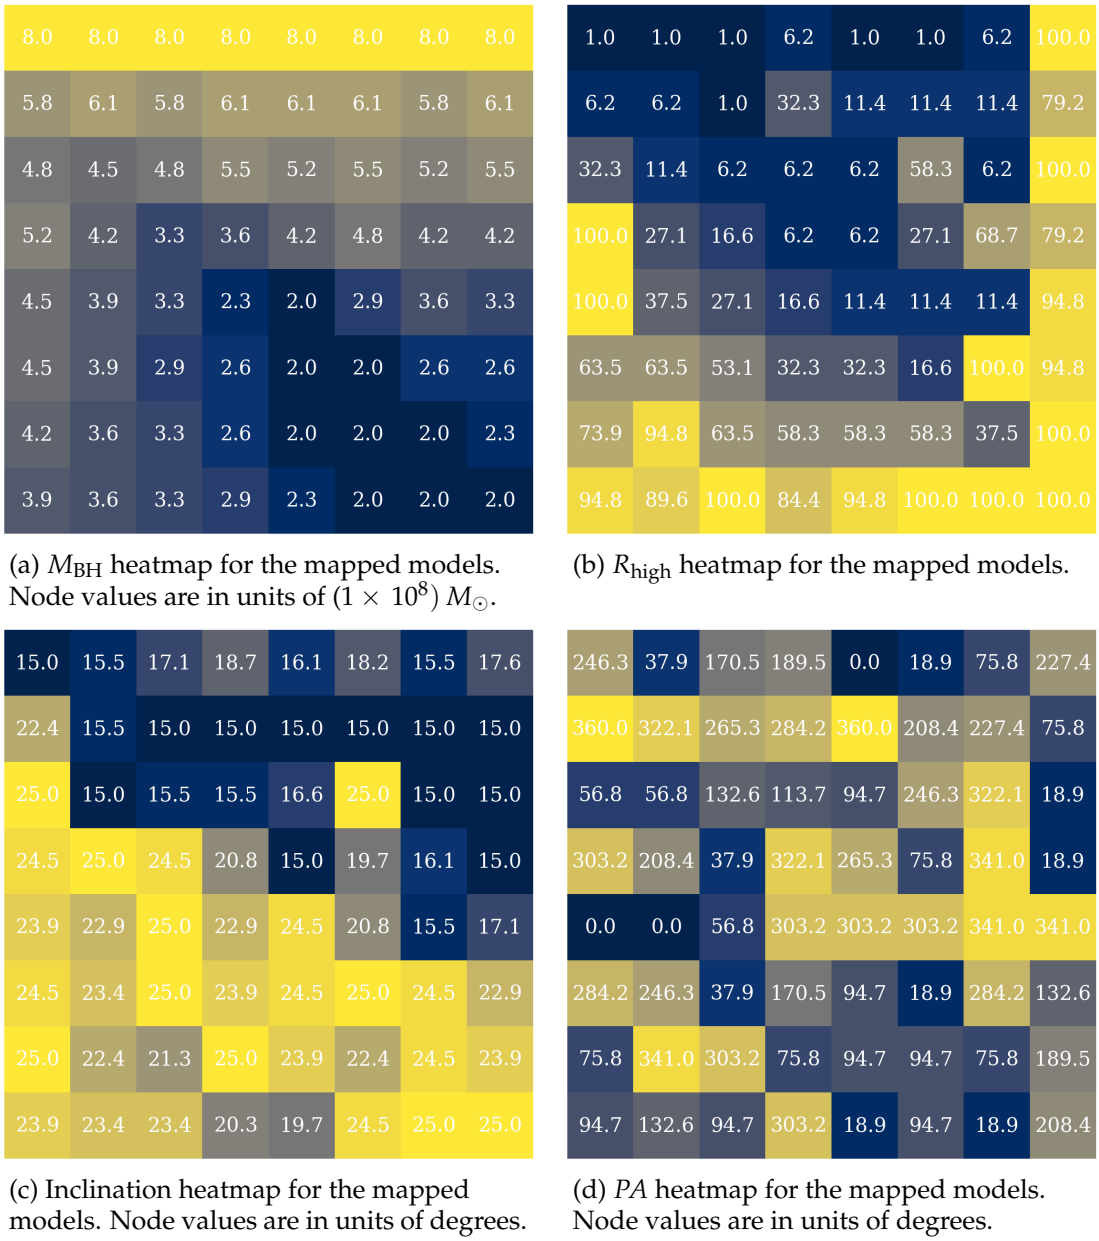

(a) M_{BH} heatmap for the mapped models. Node values are in units of $(1 \times 10^8)\, M_\odot$.

(b) R_{high} heatmap for the mapped models.

(c) Inclination heatmap for the mapped models. Node values are in units of degrees.

(d) *PA* heatmap for the mapped models. Node values are in units of degrees.

FIGURE 3.8: Heatmaps of the continuous GRMHD model parameters. The value displayed on each node is the peak of the parameter distribution, or mode of the distribution, of the models, mapped to the node.

There is a bifurcation of values across the leading diagonal in the inclination heatmap, Figure 3.8 (c). The maximum and minimum i values of the entire range are 15° and 25° respectively; both are present in the heatmap and, notably, node 18 and node 26, one of each end of the range, are adjacent. The nodes with the highest inclinations are found in the bottom triangle, while low inclination nodes

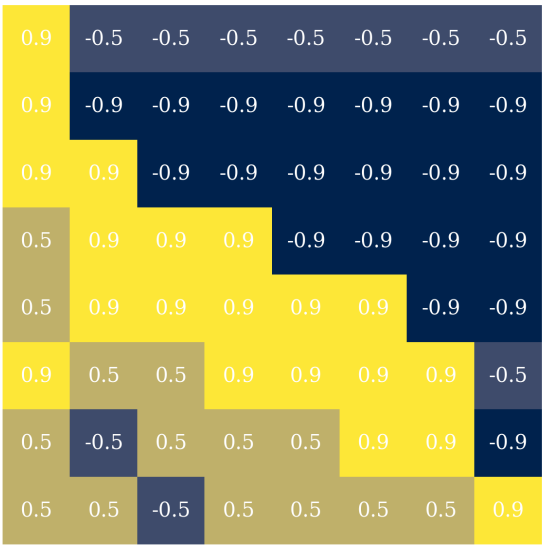

FIGURE 3.9: Heatmap for a_*. The value on each node is the mode spin value. There is a similar bifurcation structure to the one seen in the inclination heatmap. The minimum values are all in the top triangle of the lattice with the maximums distributed along the main diagonal. Nodes with R_{high} values of 0.5 are all in the lower triangle. Nodes of -0.5 appear mostly in the upper triangle and two are found in the lower triangle.

are in the upper triangle.

Figure 3.8 (d) for *PA* has no clusters of similar-valued nodes present. This behaviour is expected as PINK is rotation/flip invariant and, as *PA* affects the rotation of the BH shadow on the viewing plane, the modes' being randomly distributed indicates that the network was successful in this aspect.

Figure 3.9 shows the heatmap for a_*, the discrete, dimensionless spin parameter. This heatmap has a split of high and low values across the leading diagonal similar to that of the inclination. Interestingly, no nodes have a mode of 0. Those with a negative spin are mostly found in the upper triangle, with all minimum spins (-0.9375) in this region; the nodes with the largest (positive) spin values (+0.9375) are clustered along the main diagonal and all other positive valued nodes are located below the diagonal.

The M_{BH} heatmap indicates that the parameter values correlated well with the SOM. The a_* and inclination heatmaps have a similar structure with a split of values across the leading diagonal. This is a manifestation of the degeneracy seen in Figure 1.8. There is also no clear correlation between R_{high} and the prototypes.

Since the inclination in the simulations is constrained from observational data, its consequence on the prototype morphology is negligible. This follows from

Paper V; models did not morphologically differ significantly with small changes in i. PA does not influence PINK as the network is designed to ignore image-plane rotations.

(a) $M_{\rm BH}$ in units of $10^9\ M_\odot$

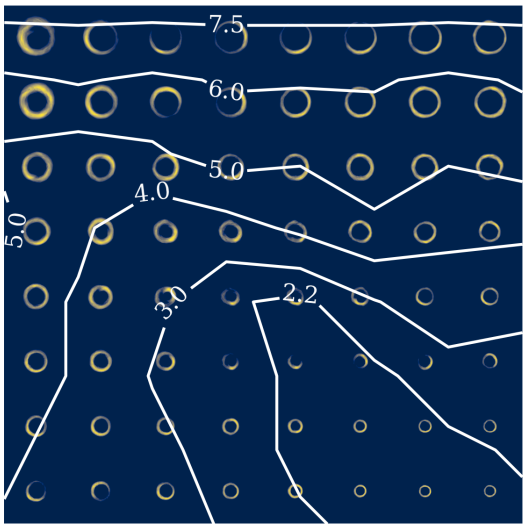

(b) $R_{\rm high}$

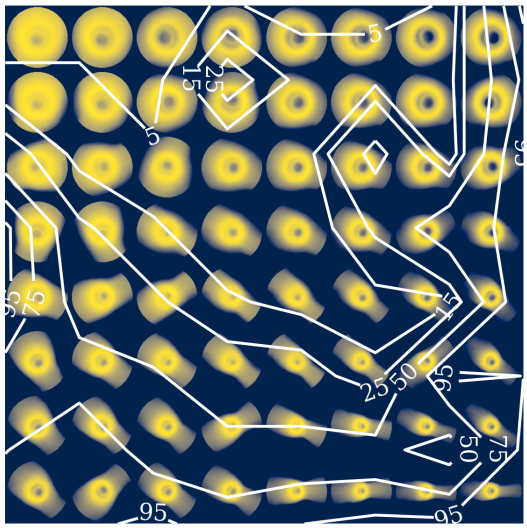

(c) Inclination

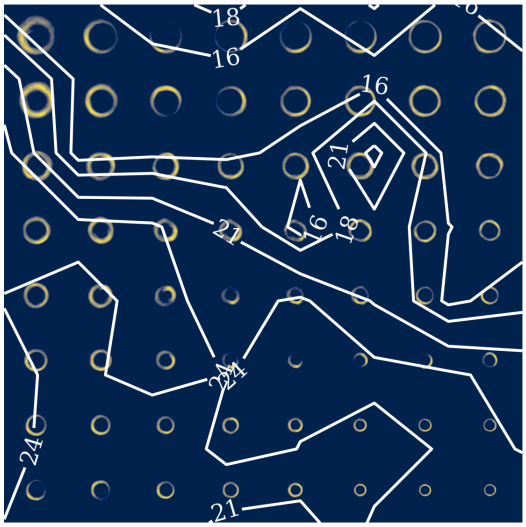

(d) a_*

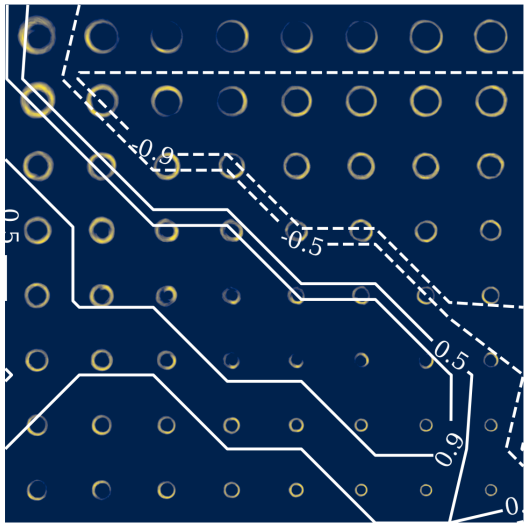

FIGURE 3.10: Contour maps of $M_{\rm BH}$, $R_{\rm high}$, i and a_* visualised over the GRMHD SOM. Those of $M_{\rm BH}$, i and a_* are plotted over the exponentially-transformed SOM, while the $R_{\rm high}$ contour map is shown over the linearly-stretched SOM.

To visualise the parameter distributions in conjunction with the SOM prototypes, we have plotted the contour maps of $M_{\rm BH}$, $R_{\rm high}$, i and a_* superimposed

onto the PINK SOM. These can be seen in Figure 3.10 (a-d). For all but $R_{\rm high}$, we have displayed the contours over the exponentially-transformed SOM as these parameters affect ring morphology the most. For $R_{\rm high}$, we have shown the contours over the linearly-stretched SOM as it is associated with the primary emission region between the disk and jet. We stress that this is simply a visual change, with no difference to the nodes and prototypes themselves.

As noted previously, $M_{\rm BH}$ correlates well with the ring size, which follows from the photon ring equation (Equation 1.3). For the other parameters, it is not immediately obvious how the parameter changes correspond to a visual difference in the prototypes changes when projected onto a 2D lattice.

3.3 Statistical Analysis

In this section, we perform a quantitative analysis of the distribution of intrinsic GRMHD BH parameters within each SOM node. In the following figures, we are assessing to what degree the parameters are clustered. This follows from the aims of this work; we seek to evaluate which GRMHD BH parameters are degenerate in terms of image morphology, i.e. the possible degeneracy between the individual parameters and node position. This can, in turn, reduce the parameter space which needs exploring when studying BHs.

We plot the parameter distributions for each node's class. The distributions have been normalised to allow for comparison between nodes despite their varying class sizes. The x-axes in the distribution graphs represent the full range for the given parameter. The distribution modes are indicated by a vertical red line; the shaded region under the distribution curves represent the 68% confidence interval of the mode. The confidence intervals give a sense of how degenerate the parameters are as a function of lattice position.

Black hole mass, $M_{\rm BH}$

In Figure 3.11 we have plotted the distributions of models' $M_{\rm BH}$ per node. Nodes 54 to 60 have narrow distributions that decrease sharply from a peak at the range minimum of $2 \times 10^9\, M_\odot$. Nodes that tend to have the larger $M_{\rm BH}$ models don't have as narrow distribution as the low mass nodes. However, they are still localised within the upper part of the range.

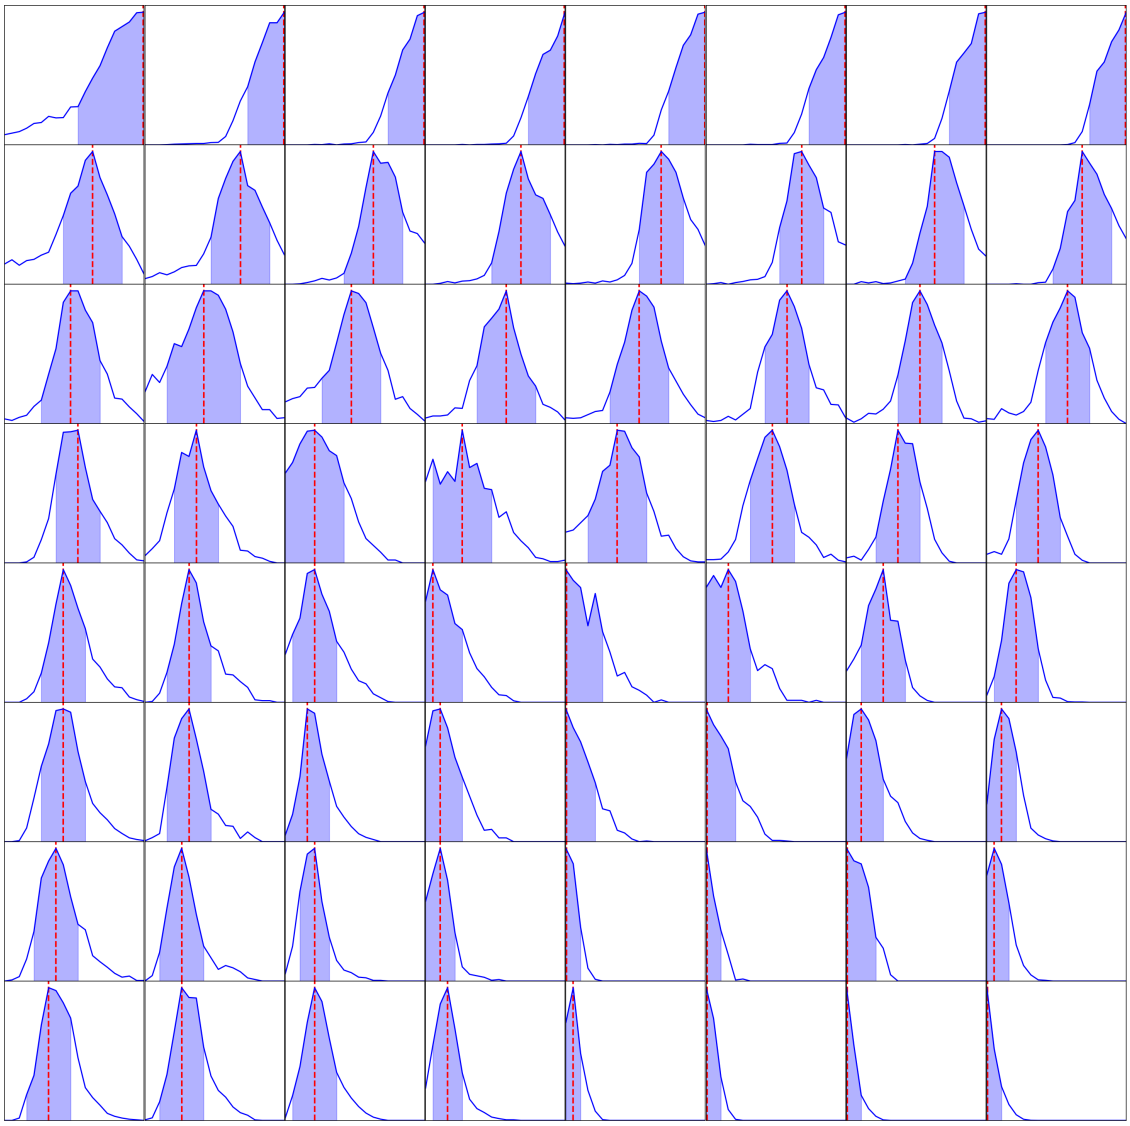

FIGURE 3.11: Distribution of mapped models' M_{BH} across the SOM lattice. These distributions are created from the model parameters in each node's class. In each figure, the x-axis represents the entire parameter range; the value density is on the y-axis. As every distribution is normalised, the area under each curve equals 1. The mode of the distribution is indicated with the red line. The blue shaded region represents the 68% confidence interval of the mode.

Nodes that peak in the middle of the range, for example, node 22, are symmetric and normally distributed. Distributions with a peak in the lower (upper) half of the range are right (left) skewed; node 43 (node 10) illustrates this. These trends are not due to the parameter limits but rather due to PINK's mapping.

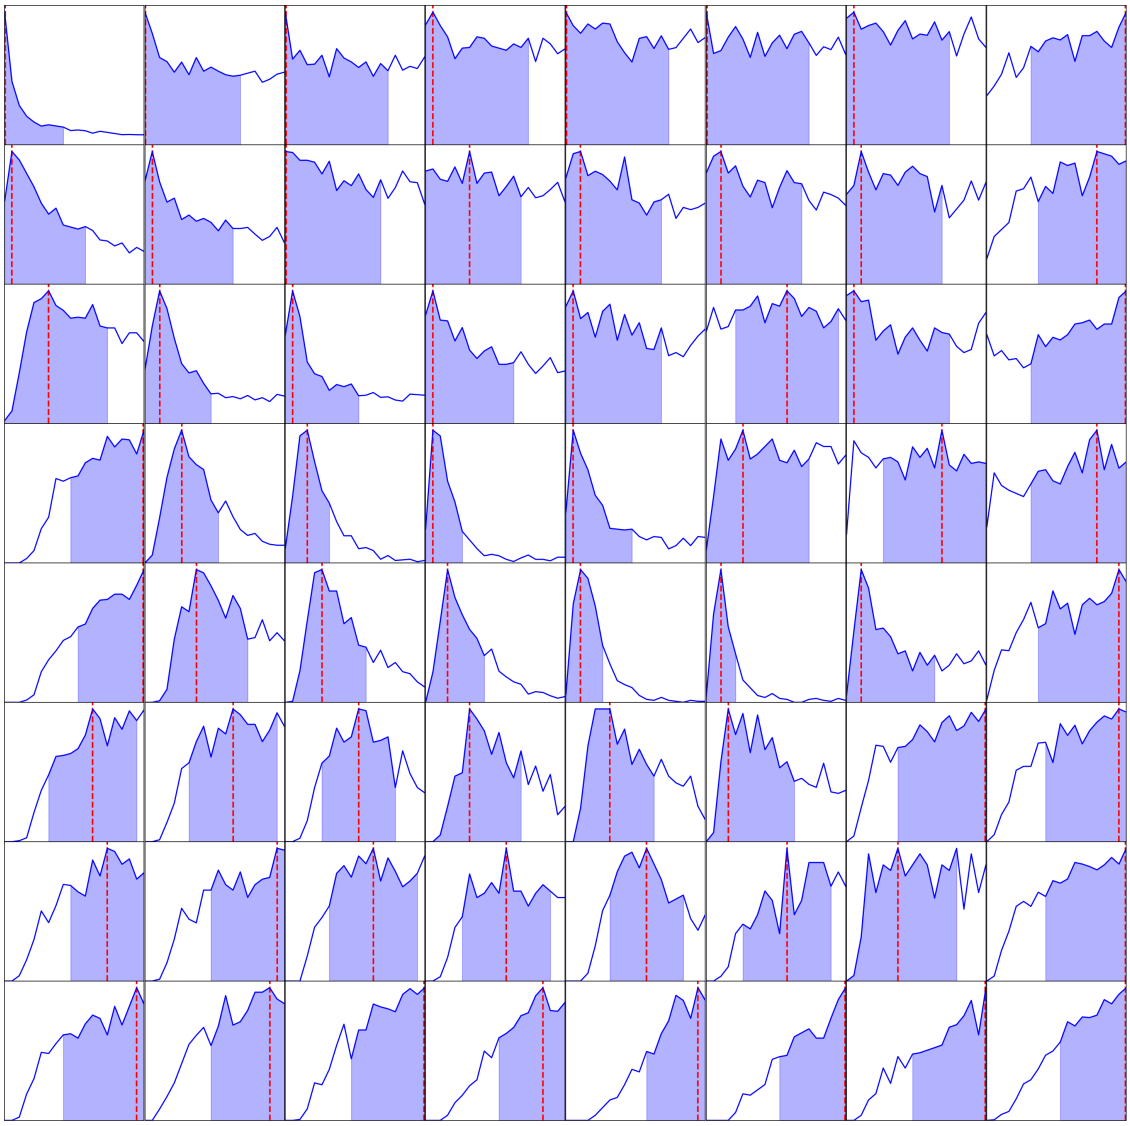

FIGURE 3.12: The same as Figure 3.11 but for R_{high}.

There is one node that behaves differently from the others: node 28. Its distribution has multiple local maxima and minima around the distribution peak. Despite this, the confidence interval is still similar in width to the other nodes.

In general, the distributions have clear peaks and narrow 68% confidence intervals. These tightly clustered distributions agree with the large correlation we noted previously between M_{BH} and ring diameter in the SOM and corresponding heatmap.

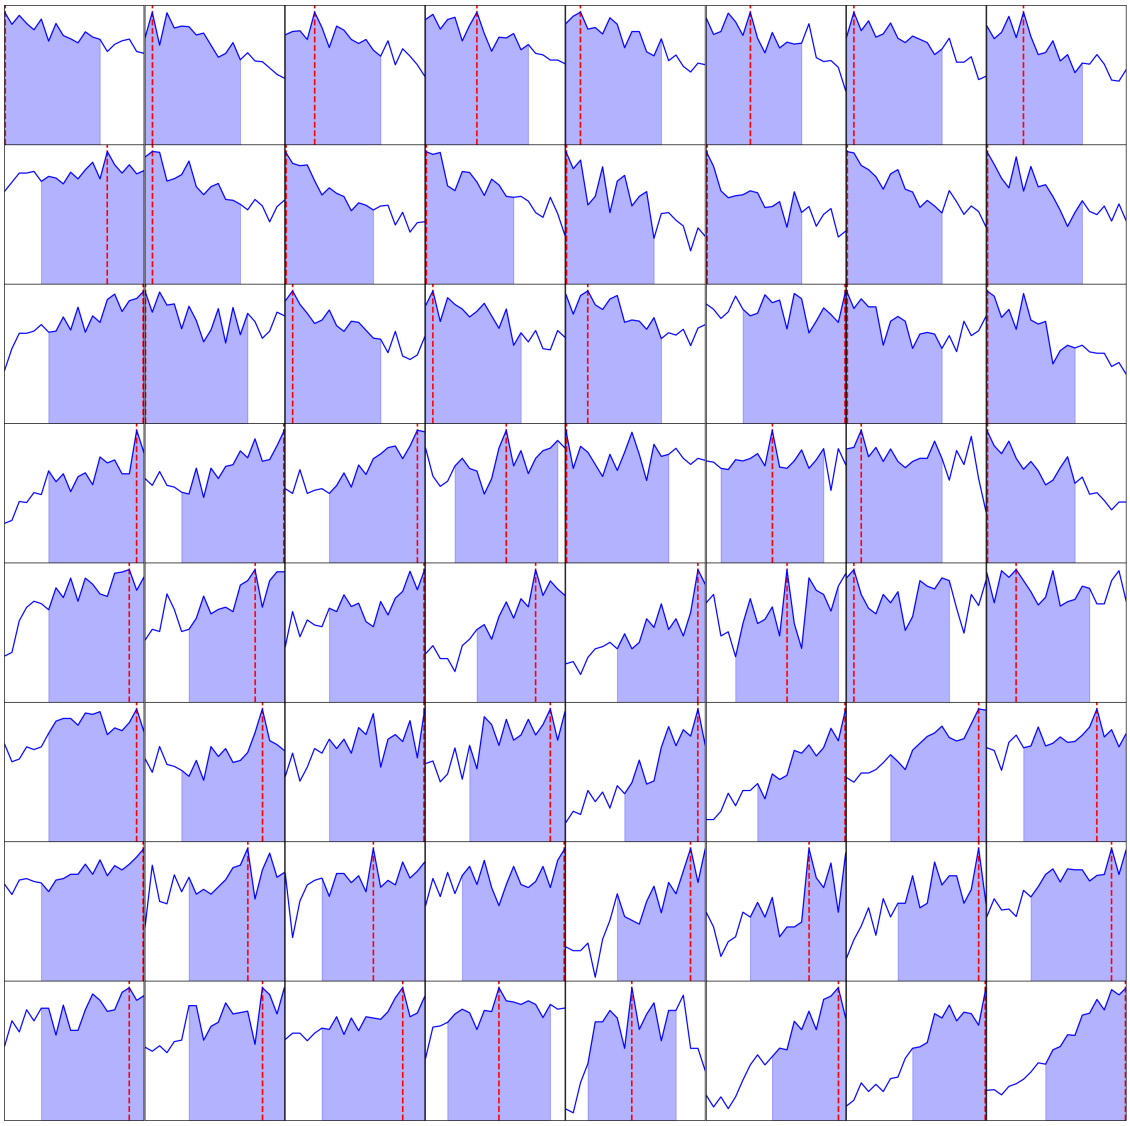

FIGURE 3.13: The same as Figure 3.11 but for inclination.

Electron temperature prescription, $R_{\mathbf{high}}$

The distributions of R_{high}, shown in Figure 3.12, show no symmetry. Nodes with peaks near or at the minimum of the range, e.g. nodes 1 and 38, have narrow distributions. In comparison, distributions that are concentrated in the upper half of the range, such as in the bottom row, have distributions that have a gentle incline towards the distribution maximum.

Distributions that peak near the middle of the range are widely dispersed.

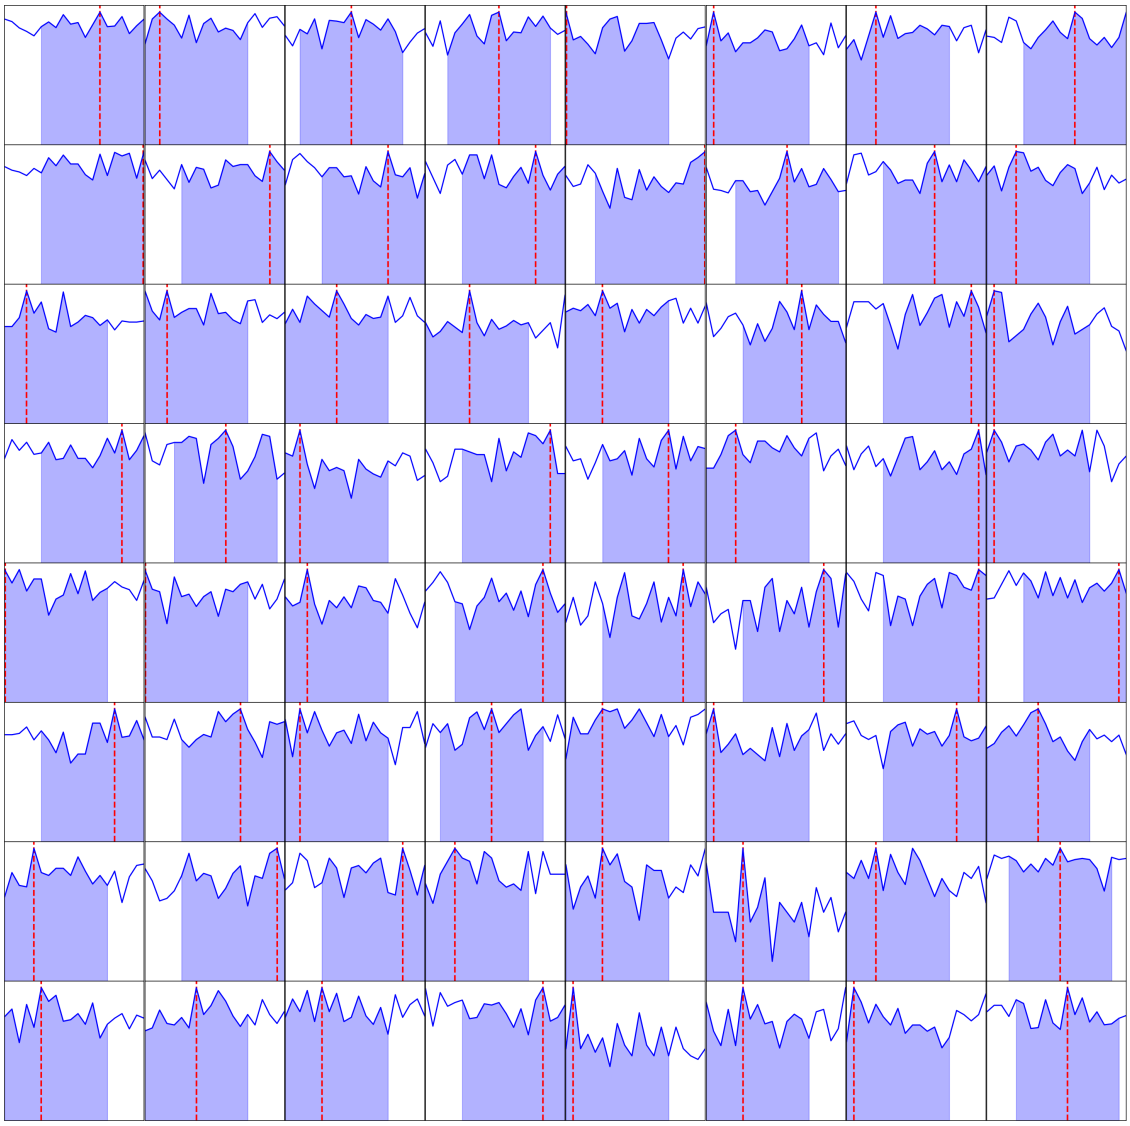

FIGURE 3.14: The same as Figure 3.11 but for *PA*

Some could even be considered as uniform such as in nodes 5 and 6. The confidence intervals for these nodes are wide, and the distribution modes do not necessarily describe anything meaningful. Consider node 7; it has a flat distribution with multiple maxima across the range that are of a similar height to the global peak. When considering the R_{high} heatmap (Figure 3.8 (b)) we understood node 7 to be a low-valued R_{high} node, but this view shows us that this is not the case and the parameter distribution must be taken into account when interpreting the SOM result and corresponding heatmaps.

Overall, R_{high} is not distributed across the SOM in small clusters as in the M_{BH}

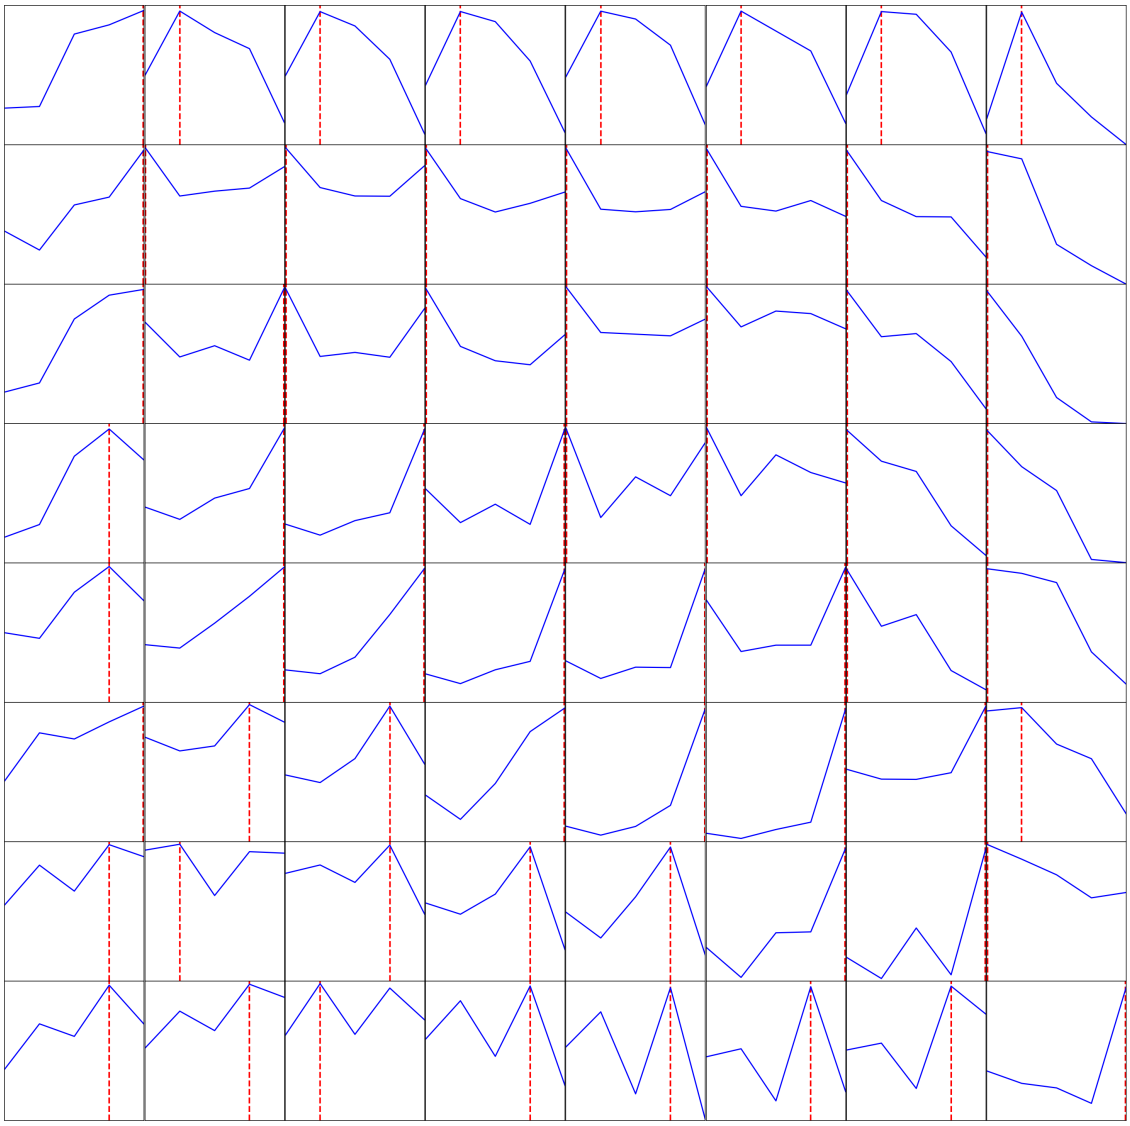

FIGURE 3.15: The same as Figure 3.11 but for a_*. As a_* is highly discretized with only five possible realisations it does not have a smooth structure like the other parameters.

distributions. Some of the distributions that peak at the minimum or maximum of the range are narrow, however, there are many nodes, like those in the top right of the map, that is widely dispersed. In these cases, it has proved critical to look at the node parameter distributions. For example, comparing nodes 7 and 8 based solely on their distribution modal values is not sufficient. This may make it difficult to see the differences between the prototypes.

Inclination, i

The inclination distributions, depicted in Figure 3.13, are widely dispersed. All contain models with i values from the entire range. The 68% confidence intervals are wide, with the narrowest of these occupying 53% of the range (node 62) and, the widest, 74% (node 48).

Some distributions have a slight concentration of models from one end of the range but are still not well constrained. As examples, nodes 8 and 64 have distributions more concentrated in the lower and upper range limits of the range, respectively, but both have i values well above zero across the entire range.

Many nodes display uniform i distributions. These distributions indicate that i has a small effect on ring morphology and, hence, is not a dominant PINK mapping factor.

Position angle, PA

We present the PA distributions in Figure 3.14. The distributions have a flat appearance with none containing values from a particular region of the range. The confidence intervals occupy a large portion of the range, between 73% and 79%, and have multiple local maxima and minima similar in height to the global values. This follows from PINK's flip/rotation insurance and the random distribution of values in PA's corresponding heatmap.

Dimensionless spin, a_*

Finally, Figure 3.15 shows the distributions of a_*; these look different to the other parameter distributions because a_* is a discrete and highly quantised variable. We have not indicated the 68% confidence interval because the small sample size made this an unreliable indication of degeneracy and/or difficult to interpret.

Distributions with a singular peak are characterised well by their mode, however, multiple nodes have two peaks of similar value (e.g. node 59). As with R_{high}, this may be why we can not see morphological differences between the nodes with different a_* modes since the values have a wide-spread distribution.

TABLE 3.1: Peaks and 68% confidence interval (CI) of M_{BH} in units of $(1 \times 10^9)\, M_\odot$, R_{high} and a_* for nodes 29, 37, 56 and 64.

Node	$M_{\mathbf{BH}}$	**68% CI**	$R_{\mathbf{high}}$	**68% CI**	a_*
Node 29	4.2	[2.95; 5.47]	6.2	[1; 48]	-0.9375
Node 37	2.0	[2.0; 3.58]	6.2	[1, 27]	0.9375
Node 56	2.3	[2.0; 2.95]	100.0	[43; 100]	-0.9375
Node 64	2.0	[2.0; 2.63]	100.0	[53; 100]	0.9375

Distribution maps overview

The parameter distributions have confirmed that a prototype's ring size and corresponding M_{BH} represents that of the nodes' classes well. This, however, is not necessarily true for the other parameters.

While some nodes have comparatively narrow R_{high}, i and a_* distributions, on the whole, the general lack of clustering means node classes aren't necessarily well represented by their prototype for these parameters.

3.4 Further Analysis and Discussion

In this section, we aim to investigate in more detail if there are in fact prototype differences between nodes with varying a_* and R_{high} distributions. We will look at a higher resolution subset of the parameter space, four node prototypes, to assess if morphological changes are corresponding to parameter differences that can be identified. We chose to consider four nodes to assess if the possible changes seen are consistent across the lattice. We use this as an illustrative example of more in-depth analyses that could be performed using our SOM framework and results, however, which are beyond the scope of this thesis.

To consider visual differences arising from changes in a_* and R_{high}, we zoom in on the four prototypes of nodes 29, 37, 56 and 64. In Table 3.1, we show the parameter values for these four specific nodes, excluding PA and i. We have chosen nodes that are all within a similar M_{BH} range, and have comparatively narrow a_* and R_{high} distributions.

We present the four prototypes, exponentially transformed and stretched, in Figure 3.16. The nodes with the linear stretch are shown in Figure 3.17. As indicated in the figures, the first column of nodes has low $R_{\rm high}$, the second column, high $R_{\rm high}$. The top row has the peak of a_* = -0.9375, the bottom row's peak is at a_* = +0.9375.

In these zoomed-in views, it is easy to identify morphological differences between prototypes of different $R_{\rm high}$ extremes. Those with small $R_{\rm high}$ have ring structures that have more detail, while the large $R_{\rm high}$ rings are thin with no prominent finer detail (Figure 3.16). In Figure 3.17, we see the images on the left are brighter as compared to those on the right. These morphological differences correlate with the GRMHD $R_{\rm high}$ parameter description, indicating that classification based on these features may be possible, given the other variables, in particular $M_{\rm BH}$, are constant and $R_{\rm high}$ is narrowly distributed.

A distinction between the models based on a_* is, however, not obvious. This is not surprising given that a_* affects the location of the bright region of the ring, i.e. the asymmetry of the ring, and ring diameter is not very sensitive to a_* (Johannsen and Psaltis, 2010; Paper V). PINK is also designed to train and map models irrespective of their geometry adding to the above.

We conclude this chapter by considering the joint parameter distributions for four nodes, nodes 8, 20, 33 and 60, in Figure 3.18. These nodes were chosen as they have similar class sizes (Figure 3.2) which, given the structure of the joint distributions plots, allows for direct comparison between the nodes' distributions. Along the leading diagonal, the 1D parameter distributions for each parameter are plotted (these are the same as the distributions presented in the previous section). For each parameter pair, the continuous joint parameter distributions are plotted in the lower triangle. We have excluded a_* from this plot as it is highly quantised. In Figure 3.18, we consider whether there are any general trends across the joint parameter distributions.

In Figure 3.18, each node is represented by a different colour. Hence, we can compare the joint parameter distributions between nodes. Besides $M_{\rm BH}$ having a large influence on the distributions, there is no overlap between the parameters 2D-distributions.

$M_{\rm BH}$ is distributed in narrow clusters, clearly seen in the joint distribution plots. However, this does not constrain the other parameter distributions. The i distributions extend across the entire possible range. The $R_{\rm high}$ distributions of

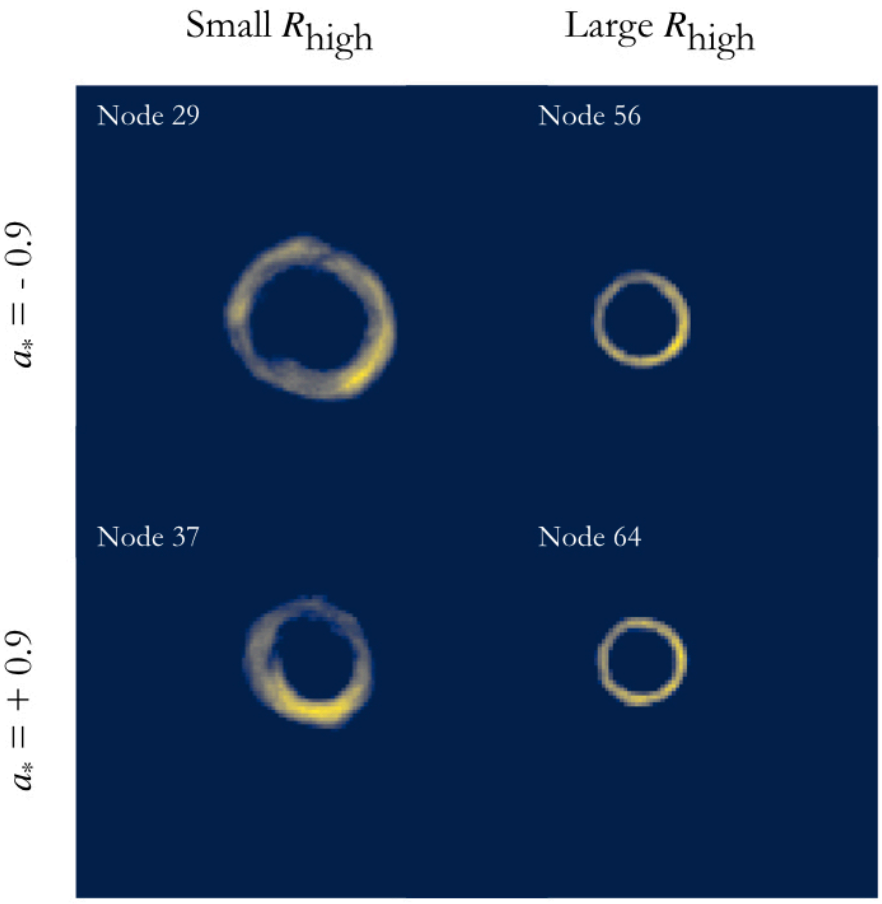

FIGURE 3.16: Four node prototypes with different extremes of a_* and $R_{\rm high}$. They have similar $M_{\rm BH}$ distributions. Shown here, nodes 29, 37, 56 and 64, with the exponential transformation and linear stretch applied.

nodes 8, 33 and 60 all peak near the maximum of the range but we cannot infer the node's mass distributions from this as the joint parameter distributions with $M_{\rm BH}$ shows clustering at the lower, mid and upper parts of the range.

The two parameters that are the most uniformly distributed, i and PA, have joint distributions that encompass the entire 2D plane. The same follows for the joint distributions of $R_{\rm high}$ with i and PA; knowing the value of one of these parameters does not indicate the value of the other, even for a narrow $R_{\rm high}$ distribution.

The lack of correlation between parameters indicates that, for the BH GRMHD models, using PINK as a dimensionality reduction technique, which clusters models based on spatial similarity, does not result in classes that can be reliably categorised based on the prototypes alone. The exception is $M_{\rm BH}$ which has a

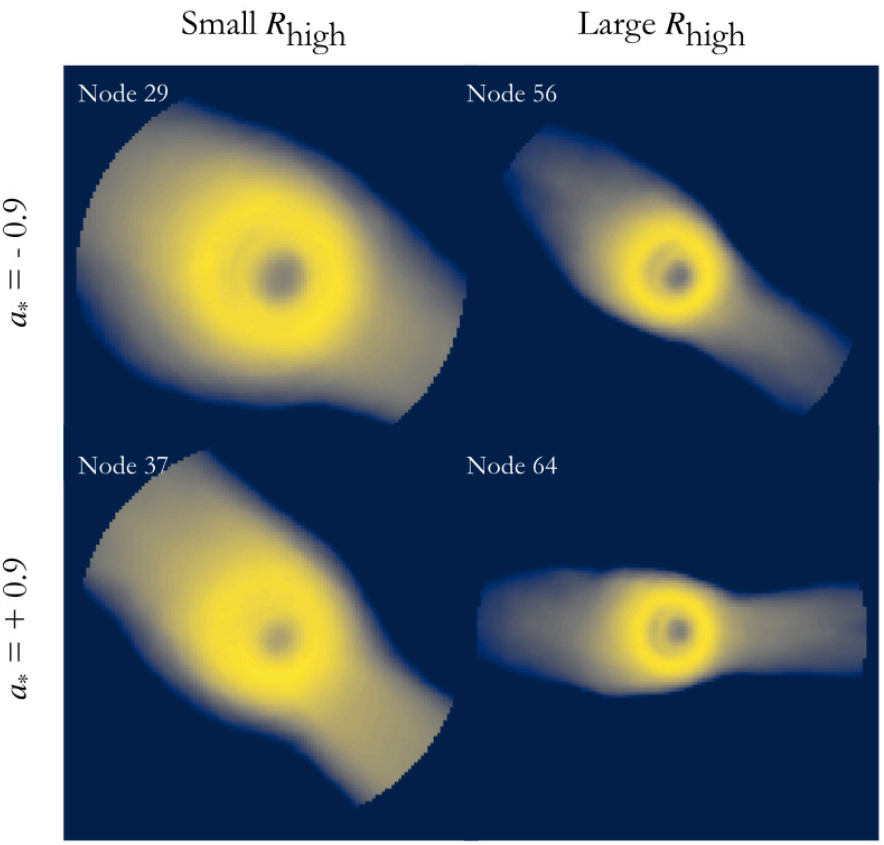

FIGURE 3.17: The same as Figure 3.16 but for the linear stretch.

dominant impact on object ring morphology.

The only other parameter that displayed some level clustering besides $M_{\rm BH}$ was $R_{\rm high}$, but the morphological effect this parameter has can only be identified in the prototypes if the $M_{\rm BH}$ distributions of nodes are in a similar range bracket and $R_{\rm high}$ has a narrow distribution.

PA impacts the geometry of the image, and hence, due to PINK's similarity measure, the characteristic it affects becomes redundant in the SOM training and mapping. a_* and i also affect the asymmetry in the ring and are sub-dominant in training and mapping.

This is highly relevant in the context of GRMHD simulations and emphasis the need to carefully consider the data for the SOM. The data structure itself has a large impact on the outcomes.

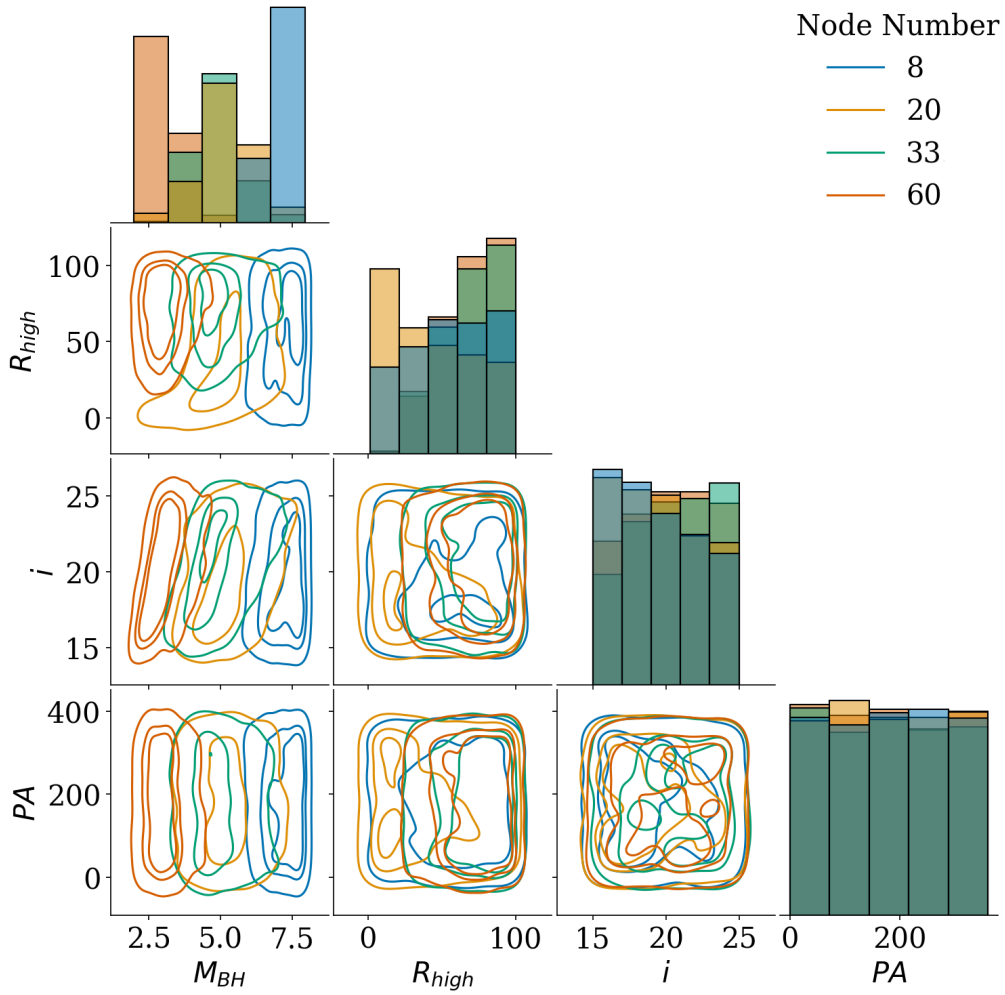

FIGURE 3.18: Joint parameter distributions of nodes 8, 20, 33 and 60 continuous GRMHD parameters. Along the diagonal are the 1D parameter distributions. The 2D joint parameter distributions are in the lower triangle.

3.5 Potential Future Applications

As a possible extension to this work, one may be able to select a family of GRMHD models that are most consistent with intrinsic BH data. A quantitative comparison between the data, including the constraints from other, independent pieces of observational information from, say, X-ray observations of M87*, and searching for parameter degeneracies that exist across the space may reduce the need to explore such large datasets by curating a set of models that are meaningfully different.

Chapter 4

Summary

4.1 Summary

The Event Horizon Telescope Collaboration released the first image of a BH shadow in 2019 (Paper I). The source, M87*, is an AGN that displays two highly Doppler-boosted jets of relativistic plasma emanating from its core. This has opened up a new era in studying SMBHs on event-horizon scales, revealing new insights and opportunities to explain the BH (a) jet launch, (b) shadow feature and (c) accretion flow characteristics.

There is a large suite of inputs required to study and model BH shadows. This space includes observations from multiple instruments at different wavelengths and scales; theoretic models and simulation output thereof; instrument-specific performance, which, for an interferometer such as the EHT, includes individual telescope performance, site-specific weather factors and other interferometric effects. Tools that effectively reduce this highly-dimensional parameter space can lead to more efficient data review, potentially making some approaches that require the full visibility dataset analysis more feasible.

Comparing theoretical model images with observational data is needed to extract meaning from the data. While approximations have been made, a rotating BH accreting from relativistic plasma has no simple analytical model solution. Instead, attention has shifted to numerical GRMHD simulations, which describe curved space-time, and are used to capture the accretion flow and BH shadow's evolution over time for a given set of initial conditions. However, these face the challenge of large data volume output, which are non-parametric and therefore expensive to confront with real data.

SOMs are a type of unsupervised learning algorithm that cluster data based on

similar characteristics. SOMs produce representations of (often) highly-dimensional parameter spaces in a 2D manner. This reduces the number of sources, for example, to examine from the size of the original dataset to the number of nodes in the map, with each node representing a prototype of a class from the data. Objects that are unlike their node class prototype are easily identifiable, so resources can be focused on anomalies of possibly interesting behaviour or phenomena.

In the astronomy context, SOMs have been used to, for example, classify radio maps from wide-field cm-wave radio telescopes with MeerKAT and LOFAR (Mostert et al., 2021). However, they have never been applied in the context of BH shadow images, which was the primary objective of this project.

The SOM algorithm utilised, PINK, classifies images based on image spatial-plane characteristics. The similarity measure it employs is both flip and rotation invariant to reduce the impact object geometry on the image plane has on training and mapping.

Before exploring the BH data, we validated PINK with a geometric model. We chose a crescent model, the same as used in Paper V, for its similarity to the BH shadow morphology observed (Paper I). We concluded that PINK created a SOM that satisfactorily represented the 10^5 dataset as 64 class prototypes, demonstrating its parameter-reduction ability and potential suitability to GRMHD BH image data exploration.

GRMHD BH simulations consider multiple parameters. Each alters the object morphology differently. Our data consisted of 10^5 GRMHD BH shadow images simulated in a SANE magnetism environment. The parameter subspace included $M_{\rm BH}$, $R_{\rm high}$, i, PA and a_* with ranges selected as relevant to M87*. Respectively, these influence: the size of the photon ring; ring flux based on the dominant emission region; asymmetry of the ring, rotation of the ring on the viewing plane and the location of the peak flux.

We proceeded by training an 8×8 PINK quadratic SOM with the image library. Subsequently, we obtained a mapping of the data to the SOM.

From the SOM, and corresponding model parameter distributions, we found:

- $M_{\rm BH}$ dominates the prototypes' location on the lattice and models' $M_{\rm BH}$ distributions are narrow;
- $R_{\rm high}$ influences the SOM mapping to a lesser degree than $M_{\rm BH}$, with some

R_{high} distributions displaying a broad clustering of model values while others have more uniform distributions;

- i distributions only display a slight concentration of certain model values for a few nodes;
- PA was uniformly distributed with no dominant peaks present;
- a_*, similar to i, has some nodes with a slight tendency towards a peak but most distributions are more chaotic.

The nodes' parameter distributions revealed that the SOM algorithm prioritised ring size as a class determining factor. The distributions of R_{high}, i, PA and a_* were widely dispersed except for some R_{high} instances. This is due to two main reasons: (1) the dominant impact M_{BH} has on ring morphology and (2) the parameters i, PA and a_* all affect the geometry and asymmetry of the ring and, due to PINK's flip and rotation invariant nature, were not clustered within image classes as these types of ring features become redundant within PINK.

Several visualisation and quantitative approaches were used to introduce the interdependence between all parameters. This implementation technique produced image classes that were predominantly determined by the BH ring size, and hence, M_{BH}. The many degeneracies between the image plane appearance and associated parameters imply that models with different parameters (other than M_{BH}) may appear morphologically similar.

This tool may be useful in analysing the morphological differences between models based on GRMHD parameters (other than M_{BH}) to determine which image-plane features are persistent. These features can be compared and evaluated for consistency for different M_{BH} values, reducing the parameter space needing exploration, especially in the case where visibility domain analyses follow, given their complex response.

The limitation of PINK as a tool in this context is the dominance of ring size as a class determining factor. The algorithm was unable to consistently differentiate images based on finer ring structure for models with M_{BH} in the range $2 \times 10^9 - 8 \times 10^9\, M_\odot$. The prototypes reflected this, having blurred rings without many substructures. This type of detail is needed to determine other physical quantities of interest for BHs.

4.1.1 Future Recommendations

The Event Horizon Telescope Collaboration recently presented work on the measured polarization of M87* (The Event Horizon Telescope Collaboration (2021a); The Event Horizon Telescope Collaboration (2021b)). They found that magnetic fields are dynamically important, favouring MAD models of the accretion flow for M87*. Exploring a MAD image library, while constraining the $M_{\rm BH}$ range based on previous results, may provide valuable insights into possible dependencies between the observed ring structure and polarimetry of M87*. The polarised intensity images, and specifically the configuration across the ring, may be a more discriminatory set of signatures to apply this SOM algorithm to.

Methods of the same type used in this work can be employed to explore the time-variability of Sgr A*. The changes in object morphology observed and corresponding parameter combinations that result in similar features can assist in determining the cause of the variability, especially given the much more tightly constrained mass of Sgr A* when compared with M87* (Gravity Collaboration et al., 2019).

This thesis only considered non-blurred BH images. However, intrinsic EHT observations are blurred due to a variety of factors. Future work should assess the effect blurring has on image-plane characteristics and parameter space reduction thereof.

Complex instrumental simulation pipelines, such as SYMBA for the EHT (Blecher et al., 2016; Janssen et al., 2019; Roelofs et al., 2020), can simulate the effects of actual observation conditions on data and calibrate the image in the same way (Janssen et al., 2019). These pipelines are often computationally expensive and, hence, processing large datasets through them is not feasible. If a set of representative prototypes can be found for a given dataset, processing these through the pipeline instead of the entire dataset will drastically reduce working time without affecting deductions that can be made from the pipeline output.

Bibliography

The Event Horizon Telescope Collaboration (2019a). "First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole". In: *ApJL* 875, p. 1. URL: https://iopscience.iop.org/article/10.3847/2041-8213/ab0ec7.

The Event Horizon Telescope Collaboration (2019b). "First M87 Event Horizon Telescope Results. II. Array and Instrumentation". In: *ApJL* 875, p. 2. URL: https://iopscience.iop.org/article/10.3847/2041-8213/ab0c96.

The Event Horizon Telescope Collaboration (2019c). "First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole". In: *ApJL* 875, p. 4. URL: https://iopscience.iop.org/article/10.3847/2041-8213/ab0e85.

The Event Horizon Telescope Collaboration (2019d). "First M87 Event Horizon Telescope Results. V. Physical Origin of the Asymmetric Ring". In: *ApJL* 875, p. 5. URL: https://iopscience.iop.org/article/10.3847/2041-8213/ab0f43.

The Event Horizon Telescope Collaboration (2019e). "First M87 Event Horizon Telescope Results. VI. The Shadow and Mass of the Central Black Hole". In: *ApJL* 875, p. 6. URL: https://iopscience.iop.org/article/10.3847/2041-8213/ab1141.

The Event Horizon Telescope Collaboration (2021a). "First M87 Event Horizon Telescope Results. VII. Polarization of the Ring". In: *ApJL* 910.L12, p. 48. URL: https://iopscience.iop.org/article/10.3847/2041-8213/abe71d.

The Event Horizon Telescope Collaboration (2021b). "First M87 Event Horizon Telescope Results. VIII. Magnetic Field Structure near The Event Horizon". In: *ApJL* 910.L13, p. 43. URL: https://iopscience.iop.org/article/10.3847/2041-8213/abe4de.

Abbott, B. P. et al. (Feb. 2016). "Observation of Gravitational Waves from a Binary Black Hole Merger". In: 116.6, 061102, p. 061102. DOI: 10.1103/PhysRevLett.116.061102. arXiv: 1602.03837 [gr-qc].

Akiyama, Kazunori et al. (July 2015). "230 GHz VLBI Observations of M87: Event-horizon-scale Structure during an Enhanced Very-high-energy γ-Ray State in 2012". In: 807.2, 150, p. 150. DOI: 10.1088/0004-637X/807/2/150. arXiv: 1505.03545 [astro-ph.HE].

Akiyama, Kazunori et al. (2015). "230 GHz VLBI OBSERVATIONS OF M87: EVENT-HORIZON-SCALE STRUCTURE DURING AN ENHANCED VERY-HIGH-ENERGY γ-RAY STATE IN 2012". In: *The Astrophysical Journal* 807.2, p. 150. DOI: 10.1088/0004-637x/807/2/150. URL: https://doi.org/10.1088/0004-637x/807/2/150.

Becker, Robert H., Richard L. White, and David J. Helfand (Jan. 1994). "The VLA's FIRST Survey". In: *Astronomical Data Analysis Software and Systems III*. Ed. by D. R. Crabtree, R. J. Hanisch, and J. Barnes. Vol. 61. Astronomical Society of the Pacific Conference Series, p. 165.

Bird, S. et al. (Dec. 2010). "The inner halo of M 87: a first direct view of the red-giant population". In: 524, A71, A71. DOI: 10.1051/0004-6361/201014876. arXiv: 1009.3202 [astro-ph.GA].

Bisnovatyi-Kogan, Gennady and Oleg Tsupko (July 2017). "Gravitational Lensing in Presence of Plasma: Strong Lens Systems, Black Hole Lensing and Shadow". In: *Universe* 3, p. 57. DOI: 10.3390/universe3030057.

Blakeslee, John P. et al. (Mar. 2009). "The ACS Fornax Cluster Survey. V. Measurement and Recalibration of Surface Brightness Fluctuations and a Precise Value of the Fornax-Virgo Relative Distance". In: 694.1, pp. 556–572. DOI: 10.1088/0004-637X/694/1/556. arXiv: 0901.1138 [astro-ph.CO].

Blandford, R. D. and D. G. Payne (June 1982). "Hydromagnetic flows from accretion disks and the production of radio jets." In: 199, pp. 883–903. DOI: 10.1093/mnras/199.4.883.

Blandford, R. D. and R. L. Znajek (May 1977). "Electromagnetic extraction of energy from Kerr black holes." In: 179, pp. 433–456. DOI: 10.1093/mnras/179.3.433.

Blandford, Roger, David Meier, and Anthony Readhead (2019). "Relativistic Jets from Active Galactic Nuclei". In: *Annual Review of Astronomy and Astrophysics* 57.1, 467–509. ISSN: 1545-4282. DOI: 10.1146/annurev-astro-081817-

051948. URL: http://dx.doi.org/10.1146/annurev-astro-081817-051948.

Blandford, Roger D. and Mitchell C. Begelman (Feb. 1999). "On the fate of gas accreting at a low rate on to a black hole". In: 303.1, pp. L1–L5. DOI: 10.1046/j.1365-8711.1999.02358.x. arXiv: astro-ph/9809083 [astro-ph].

Blecher, Tariq et al. (2016). "meqsilhouette: a mm-VLBI observation and signal corruption simulator". In: *Monthly Notices of the Royal Astronomical Society* 464.1, 143–151. ISSN: 1365-2966. DOI: 10.1093/mnras/stw2311. URL: http://dx.doi.org/10.1093/mnras/stw2311.

Bracewell, Ron (1965). *The Fourier Transform and its applications*.

Bronzwaer, T. et al. (May 2018). "RAPTOR. I. Time-dependent radiative transfer in arbitrary spacetimes". In: 613, A2, A2. DOI: 10.1051/0004-6361/201732149. arXiv: 1801.10452 [astro-ph.HE].

Bronzwaer, Thomas and Heino Falcke (2021). *The Nature of Black Hole Shadows*. arXiv: 2108.03966 [astro-ph.HE].

Cantiello, Michele et al. (Apr. 2018). "The Next Generation Virgo Cluster Survey (NGVS). XVIII. Measurement and Calibration of Surface Brightness Fluctuation Distances for Bright Galaxies in Virgo (and Beyond)". In: 856.2, 126, p. 126. DOI: 10.3847/1538-4357/aab043. arXiv: 1802.05526 [astro-ph.GA].

Cheung, C. C., D. E. Harris, and Ł. Stawarz (July 2007). "Superluminal Radio Features in the M87 Jet and the Site of Flaring TeV Gamma-Ray Emission". In: 663.2, pp. L65–L68. DOI: 10.1086/520510. arXiv: 0705.2448 [astro-ph].

Davelaar, J. et al. (Apr. 2018). "General relativistic magnetohydrodynamical κ-jet models for Sagittarius A*". In: 612, A34, A34. DOI: 10.1051/0004-6361/201732025. arXiv: 1712.02266 [astro-ph.HE].

Dexter, Jason, Jonathan C. McKinney, and Eric Agol (Apr. 2012). "The size of the jet launching region in M87". In: 421.2, pp. 1517–1528. DOI: 10.1111/j.1365-2966.2012.20409.x. arXiv: 1109.6011 [astro-ph.HE].

Doeleman, Sheperd S. et al. (Oct. 2012). "Jet-Launching Structure Resolved Near the Supermassive Black Hole in M87". In: *Science* 338.6105, p. 355. DOI: 10.1126/science.1224768. arXiv: 1210.6132 [astro-ph.HE].

Einstein, Albert (1915). "Die Feldgleichungen der Gravitation". In: *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin*, pp. 844–847.

Falcke, Heino, Fulvio Melia, and Eric Agol (Jan. 2000). "Viewing the Shadow of the Black Hole at the Galactic Center". In: 528.1, pp. L13–L16. DOI: 10.1086/312423. arXiv: astro-ph/9912263 [astro-ph].

Gebhardt, Karl and Jens Thomas (2009). "The black hole mass, stellar mass-to-light ratio, and dark halo in M87". In: *The Astrophysical Journal* 700.2, p. 1690.

Gebhardt, Karl et al. (2011). "The black hole mass in M87 from Gemini/NIFS adaptive optics observations". In: *The Astrophysical Journal* 729.2, p. 119.

Gebhardt, Karl et al. (Mar. 2011). "The Black Hole Mass in M87 from Gemini/NIFS Adaptive Optics Observations". In: 729.2, 119, p. 119. DOI: 10.1088/0004-637X/729/2/119. arXiv: 1101.1954 [astro-ph.CO].

Gillessen, S. et al. (Feb. 2009). "Monitoring Stellar Orbits Around the Massive Black Hole in the Galactic Center". In: 692.2, pp. 1075–1109. DOI: 10.1088/0004-637X/692/2/1075. arXiv: 0810.4674 [astro-ph].

Gravity Collaboration et al. (May 2019). "A geometric distance measurement to the Galactic center black hole with 0.3% uncertainty". In: 625, L10, p. L10. DOI: 10.1051/0004-6361/201935656. arXiv: 1904.05721 [astro-ph.GA].

Hilbert, David (1917). "Nachrichten von der Königlichen Gesellschaft der Wissenschaften zu Göttingen—Mathematisch-physikalische Klasse". In: *Weidmannsche Buchhandlung, Berlin*, pp. 53–76.

Ichimaru, S. (June 1977). "Bimodal behavior of accretion disks: theory and application to Cygnus X-1 transitions." In: 214, pp. 840–855. DOI: 10.1086/155314.

Janssen, M. et al. (2019). "rPICARD: A CASA-based calibration pipeline for VLBI data". In: *Astronomy Astrophysics* 626, A75. ISSN: 1432-0746. DOI: 10.1051/0004-6361/201935181. URL: http://dx.doi.org/10.1051/0004-6361/201935181.

Johannsen, Tim and Dimitrios Psaltis (July 2010). "Testing the No-hair Theorem with Observations in the Electromagnetic Spectrum. II. Black Hole Images". In: 718.1, pp. 446–454. DOI: 10.1088/0004-637X/718/1/446. arXiv: 1005.1931 [astro-ph.HE].

Johnson, Michael D. et al. (2020). "Universal interferometric signatures of a black hole’s photon ring". In: *Science Advances* 6.12, eaaz1310. DOI: 10.1126/sciadv.aaz1310.

Kerr, Roy P. (Sept. 1963). "Gravitational Field of a Spinning Mass as an Example of Algebraically Special Metrics". In: 11.5, pp. 237–238. DOI: 10.1103/PhysRevLett.11.237.

Kim, J. Y. et al. (Sept. 2018). "The limb-brightened jet of M87 down to the 7 Schwarzschild radii scale". In: 616, A188, A188. DOI: 10.1051/0004-6361/201832921. arXiv: 1805.02478 [astro-ph.GA].

Kohonen, Teuvo (1998). "The self-organizing map". In: *Neurocomputing* 21.1-3, pp. 1–6.

Kohonen, Teuvo, Ilari T. Nieminen, and Timo Honkela (2009). "On the Quantization Error in SOM vs. VQ: A Critical and Systematic Study". In: *Advances in Self-Organizing Maps*. Ed. by José C. Príncipe and Risto Miikkulainen. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 133–144. ISBN: 978-3-642-02397-2.

Kormendy, John and Douglas Richstone (Jan. 1995). "Inward Bound—The Search For Supermassive Black Holes In Galactic Nuclei". In: 33, p. 581. DOI: 10.1146/annurev.aa.33.090195.003053.

Livio, M. (Jan. 1997). "The Formation Of Astrophysical Jets". In: *IAU Colloq. 163: Accretion Phenomena and Related Outflows*. Ed. by D. T. Wickramasinghe, G. V. Bicknell, and L. Ferrario. Vol. 121. Astronomical Society of the Pacific Conference Series, p. 845.

Macchetto, Fo et al. (1997). "The supermassive black hole of M87 and the kinematics of its associated gaseous disk". In: *The Astrophysical Journal* 489.2, p. 579.

Miyoshi, Makoto et al. (Jan. 1995). "Evidence for a black hole from high rotation velocities in a sub-parsec region of NGC4258". In: 373.6510, pp. 127–129. DOI: 10.1038/373127a0.

Mościbrodzka, M. and C. F. Gammie (Mar. 2018). "IPOLE - semi-analytic scheme for relativistic polarized radiative transport". In: 475.1, pp. 43–54. DOI: 10.1093/mnras/stx3162. arXiv: 1712.03057 [astro-ph.HE].

Mościbrodzka, M. et al. (June 2017). "Faraday rotation in GRMHD simulations of the jet launching zone of M87". In: 468.2, pp. 2214–2221. DOI: 10.1093/mnras/stx587. arXiv: 1703.02390 [astro-ph.HE].

Mościbrodzka, Monika, Heino Falcke, and Hotaka Shiokawa (Feb. 2016). "General relativistic magnetohydrodynamical simulations of the jet in M 87". In: 586, A38, A38. DOI: 10.1051/0004-6361/201526630. arXiv: 1510.07243 [astro-ph.HE].

Mostert, Rafaël I. J. et al. (Jan. 2021). "Unveiling the rarest morphologies of the LOFAR Two-metre Sky Survey radio source population with self-organised maps". In: 645, A89, A89. DOI: 10.1051/0004-6361/202038500. arXiv: 2011.06001 [astro-ph.IM].

Narayan, Ramesh, Igor V. Igumenshchev, and Marek A. Abramowicz (Dec. 2003). "Magnetically Arrested Disk: an Energetically Efficient Accretion Flow". In: 55, pp. L69–L72. DOI: 10.1093/pasj/55.6.L69. arXiv: astro-ph/0305029 [astro-ph].

Narayan, Ramesh, Michael D. Johnson, and Charles F. Gammie (Nov. 2019). "The Shadow of a Spherically Accreting Black Hole". In: 885.2, L33, p. L33. DOI: 10.3847/2041-8213/ab518c. arXiv: 1910.02957 [astro-ph.HE].

Narayan, Ramesh et al. (Nov. 2012). "GRMHD simulations of magnetized advection-dominated accretion on a non-spinning black hole: role of outflows". In: 426.4, pp. 3241–3259. DOI: 10.1111/j.1365-2966.2012.22002.x. arXiv: 1206.1213 [astro-ph.HE].

Neugebauer, Gernot (2003). "The collapse to a black hole". In: *The Galactic Black Hole*. Ed. by Heino Falcke and Friedrich W. Hehl, pp. 72–94.

Owen, Frazer N., Jean A. Eilek, and Namir E. Kassim (Nov. 2000). "M87 at 90 Centimeters: A Different Picture". In: 543.2, pp. 611–619. DOI: 10.1086/317151. arXiv: astro-ph/0006150 [astro-ph].

Polsterer, K. L., F. Gieseke, and C. Igel (Sept. 2015). "Automatic Galaxy Classification via Machine Learning Techniques: Parallelized Rotation/Flipping INvariant Kohonen Maps (PINK)". In: *Astronomical Data Analysis Software an Systems XXIV (ADASS XXIV)*. Ed. by A. R. Taylor and E. Rosolowsky. Vol. 495. Astronomical Society of the Pacific Conference Series, p. 81.

Polsterer, K. L. et al. (2016). "Parallelized rotation and flipping INvariant Kohonen maps (PINK) on GPUs". English. In: *Proceedings of the 24th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2016)*. 24th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2016 ; Conference date: 27-04-2016 Through 29-04-2016. i6doc.com, pp. 405–410. ISBN: 978-2-87587-026-1.

Porth, Oliver et al. (2017). "The black hole accretion code". In: *Computational Astrophysics and Cosmology* 4.1. ISSN: 2197-7909. DOI: 10.1186/s40668-017-0020-2. URL: http://dx.doi.org/10.1186/s40668-017-0020-2.

Porth, Oliver et al. (Aug. 2019). "The Event Horizon General Relativistic Magnetohydrodynamic Code Comparison Project". In: 243.2, 26, p. 26. DOI: 10.3847/1538-4365/ab29fd. arXiv: 1904.04923 [astro-ph.HE].

Roelofs, F. et al. (2020). "SYMBA: An end-to-end VLBI synthetic data generation pipeline". In: *Astronomy Astrophysics* 636, A5. ISSN: 1432-0746. DOI: 10.1051/0004-6361/201936622. URL: http://dx.doi.org/10.1051/0004-6361/201936622.

Russell, Stuart J. and Peter Norvig (2021). *Artificial Intelligence: a modern approach*. 4th ed. Pearson.

Rybicki, George B. and Alan P. Lightman (1979). *Radiative processes in astrophysics*.

Sanders, D. B. et al. (Dec. 1989). "Continuum Energy Distributions of Quasars: Shapes and Origins". In: 347, p. 29. DOI: 10.1086/168094.

Schmidt, M. (Mar. 1963). "3C 273 : A Star-Like Object with Large Red-Shift". In: 197.4872, p. 1040. DOI: 10.1038/1971040a0.

Schwarzschild, Karl (Jan. 1916). "Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie". In: *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin*, pp. 189–196.

Shakura, N. I. and R. A. Sunyaev (June 1973). "Reprint of 1973A&A....24..337S. Black holes in binary systems. Observational appearance." In: 500, pp. 33–51.

Sun, Wei-Hsin and Matthew A. Malkan (Nov. 1989). "Fitting Improved Accretion Disk Models to the Multiwavelength Continua of Quasars and Active Galactic Nuclei". In: 346, p. 68. DOI: 10.1086/167986.

Synthesis Imaging in Radio Astronomy II (Jan. 1999). Vol. 180. Astronomical Society of the Pacific Conference Series.

Tchekhovskoy, Alexander, Ramesh Narayan, and Jonathan C. McKinney (Nov. 2011). "Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole". In: 418.1, pp. L79–L83. DOI: 10.1111/j.1745-3933.2011.01147.x. arXiv: 1108.0412 [astro-ph.HE].

Teo, Edward (2003). "Spherical Photon Orbits Around a Kerr Black Hole". In: *General Relativity and Gravitation* 35.11, pp. 1909–1926. DOI: 10.1023/A:1026286607562. URL: https://doi.org/10.1023/A:1026286607562.

Teuvo, Kohonen (2001). *Self-Organizing Maps*. 3rd ed. Springer-Verlag Berlin Heidelberg. DOI: 10.1007/978-3-642-56927-2.

van Cittert, P. H. (Jan. 1934). "Die Wahrscheinliche Schwingungsverteilung in Einer von Einer Lichtquelle Direkt Oder Mittels Einer Linse Beleuchteten Ebene". In: *Physica* 1.1, pp. 201–210. DOI: 10.1016/S0031-8914(34)90026-4.

van der Gucht, Jeffrey et al. (Apr. 2020). "Deep Horizon: A machine learning network that recovers accreting black hole parameters". In: 636, A94, A94. DOI: 10.1051/0004-6361/201937014. arXiv: 1910.13236 [astro-ph.HE].

Walker, R. Craig et al. (Mar. 2018). "The Structure and Dynamics of the Subparsec Jet in M87 Based on 50 VLBA Observations over 17 Years at 43 GHz". In: 855.2, 128, p. 128. DOI: 10.3847/1538-4357/aaafcc. arXiv: 1802.06166 [astro-ph.HE].

Walsh, Jonelle L et al. (2013). "The M87 black hole mass from gas-dynamical models of space telescope imaging spectrograph observations". In: *The Astrophysical Journal* 770.2, p. 86.

Wandeto, John and Birgitta Dresp (Aug. 2019). "The quantization error in a Self-Organizing Map as a contrast and color specific indicator of single-pixel change in large random patterns". In: *Neural Networks*. DOI: 10.1016/j.neunet.2019.08.014.

Yao-Yu Lin, Joshua et al. (July 2020). "Feature Extraction on Synthetic Black Hole Images". In: *arXiv e-prints*, arXiv:2007.00794, arXiv:2007.00794. arXiv: 2007.00794 [astro-ph.HE].

Zernike, F. (Aug. 1938). "The concept of degree of coherence and its application to optical problems". In: *Physica* 5.8, pp. 785–795. DOI: 10.1016/S0031-8914(38)80203-2.